

UC Santa Cruz

UC Santa Cruz Electronic Theses and Dissertations

Title

Essays On The Persistence Of The Forecast Bias Of Option Implied Volatility

Permalink

<https://escholarship.org/uc/item/5j25b8m5>

Author

Asensio, Ivan Oscar

Publication Date

2013

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
SANTA CRUZ

**ESSAYS ON THE PERSISTENCE OF THE FORECAST BIAS OF OPTION
IMPLIED VOLATILITY**

A dissertation submitted in partial satisfaction
of the requirements for the degree of

DOCTOR OF PHILOSOPHY

in

ECONOMICS

by

Ivan Oscar Asensio

September 2013

The Dissertation of Ivan Oscar Asensio is approved:

Professor Michael Hutchison, chair

Professor Joshua Aizenman

Professor Thomas Wu

Tyrus Miller
Vice Provost and Dean of Graduate Studies

Copyright © by

Ivan Oscar Asensio

2013

Table of Contents

Table of Contents	iii
Abstract	vi
Acknowledgements	viii
I Literature review	1
Forecast bias of the term structure of implied volatility	1
Forecast bias of implied volatility on ex post realized volatility	2
Forecast bias of model-free implied volatility	4
II The VIX-VIX futures puzzle	5
About the VIX	7
Pricing	7
Risk management	8
Index construction	8
About VIX futures	10
Pricing	11
The framework for evaluating the forecast bias of VIX futures	13
Theory	13
Constructing the VIX term structure	15
The testable equation for expectations hypothesis	16
Data	18
VIX, VIX futures, and the VIX term structure	18
About the curves associated with the VIX	19
Testing the forecast bias of the VIX term structure	20
Results	21
Subperiod analysis	22
What drives the forecast bias?	23
Open interest	24
Slow-moving capital	25
Performance of safe-haven assets	25
The cost of insuring against tail risks	26
Credit risk	27
Arbitrage existence and persistence	27
Replicating the VIX index	27
Arbitrage strategy, practical approach	29
Regression results	30

Closing comments	33
Appendix for Chapter II	35
VIX portfolios: Hedge or diversification?	35
Figures for Chapter II	40
Tables for Chapter II	52
III About the forecast bias of FX implied volatility	59
Foreign exchange markets	59
Liquidity	60
Data quality	60
Implied versus realized volatility	61
Measurement of realized volatility	62
Standard approach	62
Using high-frequency data	63
Data	64
Evaluating volatility forecast bias	65
Regression results	67
Analysis of errors	68
Measurement of financial openness	70
<i>De jure, de facto</i> , and hybrid measures of openness	70
Openness versus integration	71
The model	72
Background	72
Intuition about Black-Scholes	73
Production	75
Executing hedge trades across country borders	77
Equilibrium	78
Implications for prices	79
Analysis under no cross-border trading frictions	79
Analysis under symmetric cross-border trading frictions	80
Analysis under asymmetric cross-border trading frictions	80
The linkage between implied volatility forecast bias and integration	81
Financial integration index	83
About financial integration	83
Index construction	84
Validation of results	87
Robustness tests	88
Conclusion	89
Figures for Chapter III	92

Tables for Chapter III	94
References	101

Abstract

Essays on the persistence of the forecast bias of option implied volatility

by Ivan Oscar Asensio

Chapter I contains a literature review on the forecast bias of implied volatility based on the two fundamental questions addressed in the literature. Does implied volatility evolve as suggested by its term structure? And, is ex post realized volatility consistent with ex ante implied volatility forecasts? Chapters II and III contribute to this literature by focusing on explanations for and interpretations of forecast bias persistence. Some literature review sections which deal with specific applications of bias persistence were embedded in Chapters II and III, instead of the literature review in Chapter I, so as not to disrupt the flow of this dissertation.

The first strand investigates the evolution of implied volatility itself across time. Application of the expectations hypothesis suggests that the shape and slope of today's implied volatility term structure should be consistent with expected future changes of short-dated implied volatility. Analysis of the information content in VIX (CBOE Volatility Index) futures in Chapter II reveals a persistent forecast bias. I find that while there is evidence to uphold expectations hypothesis during the 2008-9 credit crisis period and before, in general ex ante forecasts of the future level of the VIX, implied by the VIX term structure, overshoot actual ex post changes, especially over shorter tenors. Strategies designed to profit from the bias reported have not been successful in eliminating the arbitrage opportunity. The forecast bias has increased as additional capital has flowed in, a result which does not support the slow-moving capital explanation of arbitrage persistence. Instead, I present evidence to suggest that the VIX-VIX Futures Puzzle is propagated by significant inflows of capital from non-professional investors via the proliferation of ETF (exchange-traded fund) offerings.

The second fundamental question addressed in this literature investigates the extent to which implied volatilities accurately forecast the ex-post realized volatility of underlying asset prices. In other words, are option premiums justified by the subsequent payouts? Chapter III is devoted to this question, based on an expansive dataset of implied volatility for over thirty currency pairs across developed and emerging economies. I report the forecast biasness and propose that

the magnitude of such may be used as a proxy for the degree of financial integration achieved for a particular country. I motivate this concept by furnishing a simple theoretical framework based on Dornbusch, Fischer, Samuelson (1977), based on the foundation that currencies with lower (greater) cross-border trading frictions would be expected to have options markets that exhibit lower (higher) levels of forecast bias. The transmission mechanism from cross-border currency trading frictions to FX implied volatility forecast bias is the hedging activity of currency option market-makers, based on the Black-Scholes approach to option-pricing. I conclude by introducing a new financial integration index which offers a number of benefits to existing approaches.

Acknowledgements

I express my deepest gratitude to my oral and dissertation committees, faculty advisors, and fellow students in my cohort. Your wisdom, direction, patience, and enthusiasm were instrumental in completing this work, and in my quest to transform ideas into scholarly research. Also, I appreciate the research support from the University of California Santa Cruz, the Cota-Robles Foundation, and the Sury Initiative for Global Finance and International Risk Management.

On a personal note, I dedicate this dissertation and the accomplishment it represents to Alicia Asensio, my wife. I share this with you, Julian Ivan, Sebastian Augusto, and Tammy Faye. You are the reason I get up every day.

Chapter I

Literature review

Implied volatility is the most important driver in the pricing of financial options. It represents the market's best guess at time t about future realized volatility over the period the option is active, $t + 1$. Implied volatility is determined not by a single formula or methodology, but by supply and demand dynamics. This is analogous to the manner in which, for instance, flood insurance rates are set. Factors such as the location, design and age of the structure, the recorded history of rains, earthquakes, and tsunamis for a particular region (with greater importance given to recent experience), and the cost of credit in the financial system, are all important in determining premiums. Ultimately, however, prices paid are determined by how much the insured are willing to pay.

The relevance and application of implied volatility measures to the valuation of assets, risk management, investment portfolio construction, accounting disclosure and monetary policy development highlight the importance of evaluating its forecast accuracy. Two fundamental questions have been addressed in the academic literature. The first investigates the evolution of implied volatility itself across time. The existence of a term structure, as it relates to option implied volatilities, implies that the market assigns different prices for different time horizons. Does implied volatility evolve according to term structure forecasts? Specifically, when the term structure is positively (negatively) sloped, does implied volatility at the short-end of the curve rise (fall) as much as predicted?

The second fundamental question addressed in the literature evaluates to what extent implied volatilities accurately forecast the ex post realized volatility of underlying asset prices. In other words, are option premiums incurred justified by the subsequent payouts on those options?

Forecast bias of the term structure of implied volatility

The existence of a term structure implies that the market places a different level of uncertainty about asset prices for different time horizons. In practice, this results in upward or downward

sloping term structures, seldom flat. Under the benchmark Black Scholes (BS) model, the concept of a term structure should not exist. The assumption is that implied volatility is fixed across option tenors. The divergence between theory and real-world application has been the motivation for vast research efforts involving finance theorists, behavioral economists, and non-academic researchers seeking trading profit opportunities.

The forecast accuracy of implied volatility term structure has been tested in the literature by application of the expectations hypothesis (EH). The intuition and tests are the same as have been applied originally to the term structure of interest rates¹. The theory suggests that the shape and slope of today's implied volatility term structure should link the long-end today with expected future changes of the short-end of the curve, or equivalently that the short-end today should be consistent with expected future changes of the long-end of the curve. Results when applied to financial options are mixed. Stein (1989) documents overreactions of the longer-dated option prices on the SPX index to changes in short-dated options. Campa and Chang (1995) develop a well accepted framework for testing EH, and their results uphold EH for a narrow set of currency option implied volatilities. Poteshman (2001) and Byoun et al. (2003) find the slope of the volatility term structure to have significant predictive ability for future implied volatility of the SPX. Mixon (2007) finds that by adjusting the implied volatility forecast by a volatility risk premium, the predictability along the term structure increases, however, not to the extent predicted by expectations hypothesis.

Forecast bias of implied volatility on ex post realized volatility

Poon and Granger (2003) offer an thorough survey of volatility forecasting. Based on 93 peer-reviewed research articles, the authors establish that option implied volatility contains a significant amount of information about future realized volatility, often providing more accurate forecasts than model-based efforts derived from time series models². Evidence is furnished for options on equities, broad-based price indices, interest rates, currencies, and commodities. Comparing across asset classes, it is reasonable to expect different levels of forecast accuracy

¹See Shiller (1979), Froot (1989) for an evaluation of EH for the term structure of interest rates.

²For surveys of such time series modeling techniques that use empirical data, see Campbell, Lo and MacKinlay (1997), Gouriéroux and Jasiak (2001), and Taylor (2005).

for options written on different assets. This is primarily due to differences in trading frictions. Because market-makers rely on executing transactions in spot, forward, and futures markets as a way to manage the risk in their options inventory portfolios, assets which trade in the presence of lower frictions (and thus greater liquidity, narrower bid-ask spreads, etc.) would be expected to have options markets which exhibit lower levels of forecast bias. Along these lines, the literature has focused on testing implied volatility for at-the-money (ATM) strikes or utilizing weighting schemes across different strikes which overweight the ATM strikes. Liquidity is superior.

Buraschi and Jackwerth (2001), Coval and Shumway (2001), Bakshi and Kapadia (2002), and Pan (2002) all report that ex ante estimates derived from implied volatility overshoot ex post realized volatility. The various explanations for the existence of the forecast bias fall under one of two categories: either the options market is inefficient for some reason³, or, the option pricing model is incorrect. With regard to the latter, while the empirical literature treats implied volatility as being exogenous, a few researchers have suggested that there is an element of endogeneity in prices related to model misspecification. Option pricing models used in practice, such as the Black-Scholes model, do not allow for a premium for bearing volatility risk. The crux of the argument is that volatility risk premium exists, and so prices derived by models that do not incorporate this important variable will automatically be biased. By applying the Heston (1993) model, researchers attempt to quantify the volatility risk premium and evaluate the extent to which the forecast bias is due to its omission. Benzoni (2001), Chernov (2001), and Mixon (2007) all find that forecast bias is reduced by incorporating this dimension.

Furthermore, the bias will persist only if it is difficult or impossible to construct and efficiently implement arbitrage strategies that will in time remove the market mispricing. Fleming (1998) finds material forecast bias in option prices on the S&P 100 index. Arbitrage strategies designed to profit from mispriced equity index options would involve actively managing the underlying basket of stocks in the index, an administratively demanding process. Alternatively, the arbitrage may be done indirectly via index futures, which trade with great liquidity. In both cases, the arbitrage requires replication of the options, which involves frequent trading and active management of the position which may discourage many potential arbitrageurs, and thus

³Some common explanations include the presence of trading frictions in hedging markets, liquidity frictions, and 'peso problems'.

allowing the bias to persist. Measurement of realized volatility is also a potential source of bias persistence. Poteshman (2000) finds that a more efficient volatility estimator based intra-day five minute returns removed over half of the bias present using daily data. Blair, Poon, and Taylor (2001) report up to a four-fold increase in r-squared coefficients when going from daily to high-frequency intra-day data.

Forecast bias of model-free implied volatility

The VIX is the most widely followed and well-known volatility product, representing the market's estimate of future volatility in the S&P 500 index (SPX) over a one-month period. The construction of the VIX and other volatility products is based on the development of model-free implied volatility measures. At a fundamental level, model-free implied volatility can be derived from a portfolio of plain vanilla options with strikes spanning the full range of possible outcomes for the underlying asset at expiry. The VIX, due to its construction, should possess better forecasting accuracy as compared to implied volatility estimates extracted from a single option⁴. The literature is mixed on this front. Becker, Clements and White (2006) examine whether the VIX contains any information relevant to future realized volatility beyond that reflected in model-based estimates based on empirical data. The authors conclude that the VIX does not add to the forecasting power of alternative approaches. Subsequently, Becker and Clements (2008) show that a combination of both the VIX and a model-based estimate is found to be superior to either method by itself. Konstantinidi et al. (2008) and Konstantinidi and Skiadopoulus (2011) demonstrate that VIX futures are predictable by their historical patterns, however the coefficients are too small to generate actual trading profits. Nossman and Wilhelmsson (2008) operate under a stochastic volatility framework in order to adjust actual VIX futures prices by a risk premium and find that EH cannot be rejected. Also they report that risk premium adjusted futures prices provide good forecasting ability of the VIX index with a 73% hit ratio. A potential limitation in their work however is that the researchers only use the near-term VIX futures contract with a maximum tenor of thirty days. In addition, their calibration of risk premium grows proportionally with time to maturity, with mean and upper bounds

⁴The literature on the forecast bias of options has primarily focused on implied volatility estimates based on individual options, at-the-money strikes are most common.

at two and four percent respectively for a thirty-day period. Extrapolating this to a six-month contract, which trades with ample volumes, would deem this an unrealistic metric. Bollerslev, Tauchen, and Zhou (2009) and Drechsler and Yaron (2011) study the relation between variance risk premia and equity risk premia and demonstrate that the VIX has significant forecasting ability for future SPX returns. Their models are based on the underlying VIX, determining that no additional information is being furnished by the term structure. Lu and Zhu (2010) study term structure dynamics and identify that a third factor, capturing curvature movements, is statistically significant. Shu and Zhang (2012) apply traditional linear Granger tests and find that VIX futures prices have predictive ability on the underlying VIX, however after exploring for non-linear relationships, the predictive ability goes away. Luo and Zhang (2012) propose a simple yet powerful two-factor stochastic volatility framework for the VIX, corresponding to the level and slope of the VIX term structure, which offers rich information content relative to historical volatility in forecasting future realized volatility. Ait-Sahalia, Karaman, and Mancini (2012) find that a significant jump risk component embedded in the VIX term structure contributes to a downward sloping term structure in quiet times, but an upward slope in turbulent times. In addition, the authors find that the expectations hypothesis, defined as the difference between forecasted and actual variance, does not hold, but biases and inefficiencies are modest for short horizons.

Chapter II

The VIX-VIX futures puzzle

Volatility products, which comprise a subset of derivative securities, have emerged over the last few decades due to the market's acceptance of volatility as being its own asset class. The traditional role of volatility as a singular input to standard option products is upgraded for this class of securities. For a vanilla option, volatility plays a supporting role in the valuation of contingent non-linear payouts. Its contribution to option value peaks for at-the-money (ATM) strikes but falls very quickly as the moneyness of the option deviates higher or lower⁵. Expressing a view on future volatility via standard options is thus an inefficient and ineffective proposition,

⁵The vega exposure, or price sensitivity to changes in volatility, of a standard vanilla option decreases as the strike of the option goes further in-the-money (ITM) or out-of-the-money (OTM).

as the contribution of volatility to the value of the derivative is heavily contingent on the level or return of the underlying asset. Volatility products, in contrast, grant the holder a well-defined exposure to volatility itself. There are two basic types, written primarily on major equity indices or currency prices⁶. On one hand, there are variance (volatility) swaps whose payoff is linked to the observed realized variance (volatility) over the contract period. Such instruments may be thought of as forward contracts on future realized variance (volatility). On the other, there are forward volatility agreements, which pay off according to changes in the implied volatility itself⁷. These are essentially forward contracts on future implied volatility. Carr and Lee (2009) provide an excellent and thorough review of volatility derivatives.

The VIX is the most widely followed and well-known volatility product, representing the market's estimate of future volatility in the S&P 500 index (SPX) over a one-month period. The VIX has been embraced as a risk management vehicle by investors, a barometer for risk aversion by financial markets participants, and an input to econometric specifications and robustness tests by academic researchers. A higher reading on the VIX corresponds to greater aversion to market risk, while lower readings are associated with rising risk appetite. Over the last few years the index has acquired global acceptance as the ultimate barometer for investor sentiment⁸.

The VIX is an index whose calculation is based on a set of option prices on the SPX across a wide variety of strikes⁹, thus a negative co-movement between the VIX and the SPX exists by construction. This is a stylistic fact. Figure II.1 tracks historical price data on the VIX and the SPX, from January 2006 to October 2012. The strong association between both price series is supported quantitatively by a correlation coefficient, based on daily changes, of -76% for the entire period. In addition, the average three-month rolling correlation is -83%. The negative number implies co-movements occur in the opposite direction.

- Insert Figure II.1 here -

Motivated by the fact that the VIX is not an investable asset, as exposure to the index is attained principally by trading VIX futures, the focus of Chapter II is to evaluate the information content of the VIX term structure. Tests of the expectations hypothesis on the VIX term

⁶The underlying instrument must trade with ample liquidity in order for construction of volatility derivatives to be possible.

⁷Implied volatility represents the market's estimate today about future realized volatility.

⁸The VIX index is commonly referred to as the market's 'fear gauge'.

⁹VIX calculation fully outlined in the next section.

structure are based on the risk-neutral formulation proposed by Campa and Chang (1995). An innovation proposed in this work involves construction of the VIX term structure using actual VIX futures market prices. The three key objectives are to document the forecast biasness and in the cases or time horizons over which forecast bias is present, address why it exists, and why it persists.

About the VIX

Pricing

The construction of the VIX and other volatility products is based on the development of model-free implied volatility measures. At the fundamental level, model-free implied volatility can be derived from a portfolio of plain vanilla options with strikes spanning the full range of possible outcomes for the underlying asset at expiry. The option prices in this approach are thus taken as given. Breeden and Lizenberg (1978), Demeterfi et al. (1999), Britten-Jones and Neuberger (2000), Carr and Wu (2006), Dupire (2006), Jiang and Tian (2007), Carr and Lee (2007) have described theoretical relationships between underlying option prices and model-free implied volatility, under the implicit assumption that the options on the underlying assets are priced properly.

A separate strand of literature running in parallel involves the use of alternative processes for asset prices, volatility dynamics and/or different option pricing models to price volatility products. This approach does not accept traded option prices on the SPX as given. Zhang and Zhu (2006), for instance, use the Heston stochastic volatility model to price VIX futures. Zhu and Zhang (2007) value VIX futures by applying a stochastic variance model to the evolution of the VIX itself and to deriving the term structure of forward variance. Lin (2007) uses an affine jump-diffusion model with jumps in both index and volatility processes to arrive at VIX futures prices. Sepp (2008b) applies a similar framework for calibration of both VIX futures and options on the VIX. Zhang and Huang (2010) highlight the importance of dynamics assumptions and parameter estimation by contrasting the results of a number of different approaches.

Risk management

Although not yet considered a stylistic fact, VIX-products are best utilized for purposes of diversification, and not necessarily to serve as true hedging vehicles. The difference is subtle but central to portfolio construction. Diversification involves decreasing portfolio variance, and thus portfolio risk, by adding elements that exhibit a low correlation to existing holdings. A hedge, by design, is intended to offset losses that may be incurred by exposure to a core asset position. This is can be said of futures and options on the SPX, for instance, which derive their value directly from SPX price changes.

This key distinction was made by Daigler and Rossi (2006) who report a significant diversification benefit comes from adding a long VIX position to an SPX portfolio. Szado (2009) evaluates the performance of overlaying a VIX portfolio on a base holding of stocks and bonds and finds a net reduction in aggregate risk of roughly one third. Delisle et al. (2010) go a step further and address the hedge inefficiency of VIX products by offering a VIX-replicating portfolio which outperforms a static buy and hold position in a portfolio specifically tied to the VIX.

In the Appendix section, I sketch out the basic theoretical argument for explaining the ineffectiveness of the VIX as a true hedge for long SPX holdings, highlighting instead the diversification benefits of the strategy.

Index construction

The VIX represents the conditional risk-neutral expectation of the volatility for the SPX index over the next calendar month

$$\sigma_t^{VIX} \cong E_t^{\mathbb{Q}}[\sigma_{SPX_{t \rightarrow t+1}}], \quad (1)$$

where σ^{VIX} is the estimate of volatility as quoted in annualized percentage terms according to standard market convention, and σ_{SPX} represents the realized volatility of the SPX, also expressed in annualized percentage terms, from time t to $t + 1$ months later. The methodology for arriving at VIX prices involves computing a weighted average of out-of-the-money (OTM)

option prices on the SPX across all strikes for the two nearby maturities¹⁰. The general formula used in the VIX calculation¹¹ is given by (2) and (3) below

$$\sigma_{p_j}^2 = \frac{2}{T_j} \sum_i \frac{\Delta K_i}{K_i^2} e^{rT} Q(K_i) - \frac{1}{T_j} \left[\frac{F_j}{K_0} - 1 \right]^2, \quad (2)$$

where $\sigma_{p_j}^2$ is the constructed variance of a portfolio of SPX options expiring at the two nearby maturities j and $j + 1$, T , expressed in years, is the common period over which all options in the calculation are active, F is the forward index level derived from coterminous index option prices¹², K_0 is the first strike below the forward, K_i is the strike price of the i^{th} out of the money option¹³, ΔK_i denotes the interval between strike prices, defined as $\Delta K_i = (K_{i+1} - K_i)/2$, r is the risk-free rate to expiration, $Q(K_i)$ is the midpoint between the bid-ask spread for each option with strike K_i , and r is the risk-free rate.

Table II.1 contains a collection of options on SPX that would be involved in a hypothetical calculation of the VIX as of 2-Nov-2012. In practice, option strikes are available for every five points on the SPX index, however for expositional purposes, the intervals used are fifteen and twenty-five SPX points apart for near and next-term maturities respectively.

- Insert Table II.1 here -

The salient information to be extracted from the table is as follows: 1) The strike of 1415 is the strike where the price difference between calls and puts for both maturities is smallest, 2) the strike of 1415 will be used to determine F and K_0 , which in turn determines the set of K_i , 3) the range of option strikes used for each maturity will vary as the calculation leaves out options for which the bid price is zero, and 4) the exact collection of options used will change in tandem with changes in the underlying price of the SPX as in-the-money (ITM) options are left out of the calculation.

The VIX, as quoted, is computed by deriving (2) for the near-term and the next-term maturities on SPX futures, indexed by j and $j + 1$ respectively¹⁴, and plugging below

¹⁰Maturities occur monthly. In eight out of twelve months in the year, VIX futures settle on the third Wednesday of each month, in the other four months, the futures expire on the fourth Wednesday of the month.

¹¹According to white paper on the VIX accessible on the CBOE website (www.cboe.com/micro/vix/vixwhite.pdf).

¹²This is equivalent to the strike price at which the price difference between the SPX call and put is smallest adjusted by e^{rT} (*call price - put price*).

¹³Note if a call option, then $K_i > K_0$ and if a put option, $K_i < K_0$.

¹⁴The near-term future is the 1st future, and the next-term future is the 2nd future. These terms will be used

$$\sigma_t^{VIX} = \sqrt{\left\{ T_j \sigma_{p_j}^2 \left[\frac{N_{T_{j+1}} - N_{30}}{N_{T_{j+1}} - N_{T_j}} \right] + T_{j+1} \sigma_{p_{j+1}}^2 \left[\frac{N_{30} - N_{T_j}}{N_{T_{j+1}} - N_{T_j}} \right] \right\} \frac{N_{365}}{N_{30}}}, \quad (3)$$

where N is the number of minutes for each referenced period. As suggested by the weighted average calculation in (3), VIX prices are updated continuously throughout the day to the very minute. Figure II.2 shows the output of this calculation on an intra-day basis.

- Insert Figure II.2 here -

Although not perfect, the strength of the co-movement is evident.

About VIX futures

The VIX is an index, not an investable asset. The core method for attaining exposure to the VIX is via VIX futures, which began trading on the CBOE on 26-March-2004. VIX futures are essentially forward contracts on future implied volatility. There is however no cost of carry relationship between the VIX and VIX futures, as is standard between spot and futures prices of other exchange-traded assets¹⁵. Instead, by way of its construction, a position in VIX futures is an expression which links today's expected volatility to tomorrow's expected volatility,

$$\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = E_t^{\mathbb{Q}}[\sigma_{j,j \rightarrow j+1}^{VIX}], \quad (4)$$

where \mathcal{F} represents the conditional risk neutral expectation at time t of the VIX at the first future date of j . Note the VIX is always a one-month forecast of realized volatility, thus \mathcal{F} is a forecast of this one-month forecast. The tenor of \mathcal{F} , or the length of the forecasting period, will vary however. Combining equations (1) and (4) allows us to express futures prices as

$$\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = E_t^{\mathbb{Q}}[E_j^{\mathbb{Q}}[\sigma_{SPX_{j \rightarrow j+1}}]], \quad (5)$$

the expectation at time t , of the expectation at the time of the first future date j , of the realized volatility of the SPX index over the period j to $j + 1$. Figure II.3 illustrates the interaction interchangeably.

¹⁵Similarly, VIX futures prices do not contain elements related to insurance, storage, and transportation costs. This represents a departure from the extensive literature that deals with the forecast accuracy of futures prices for other assets such a commodities, currencies, interest rates.

between these concepts.

- Insert Figure II.3 here -

On 1-February-2012, time t , the closing price quotes, in annualized standard deviation terms, for the VIX and the 1st and second futures were $\sigma_t^{VIX} = 18.55$, $\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = 19.85$, and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX} = 22.05$ respectively. Note that σ_t^{VIX} is a forecast of realized volatility of the SPX over the immediate future period, while $\mathcal{F}_{t,j \rightarrow j+1}^{VIX}$ and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX}$ are forecasts of the future forecasts of realized volatility of the SPX. The horizontal dashed arrows in figure 3 represent the period over which such expectations apply.

The focus of this paper is not to what extent the VIX accurately forecasts the future realized volatility of the SPX, but rather the objective is to understand and characterize the evolution of the expectations. For illustrative purposes, a similar distinction can be made within the insurance industry. Expectations of floods in New York City is best captured by evaluating the price level and changes of flood insurance premiums. An analysis of the historical revenues earned from these premiums measured against subsequent payouts is also important, and indirectly contributes to the actuarial fair value of premiums, however, it does not capture the dynamic nature of the expectations of a flood on the part of end-users such as businesses and households.

Pricing

While VIX prices are derived from an exact calculation, VIX futures prices are ultimately determined by supply and demand dynamics in the market. The actuarial fair value a VIX futures contract can be determined from a synthetic calendar spread¹⁶ of SPX options bracketing the one-month period which starts at the futures expiration date, minus a term which estimates the risk-neutral variance estimate of VIX future. The derivation for this can be found in Carr and Wu (2006). Using the example in the previous section, the fair value of the next-term future is thus given by

$$\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX} = \sqrt{[\mathbb{P}_t - \hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}^2]}, \quad (6)$$

¹⁶A calendar spread is a trade involving the simultaneous sale and purchase of a pair of futures or vanilla options expiring on different dates. The legs of the spread vary only by maturity, as they are based on the same asset and notional amounts. The rationale for entering this trade is to take advantage of perceived value along the term structure, to partially finance long positions, or to cap the maximum loss of short positions.

where \mathbb{P} represents a portfolio of SPX options with long positions in out-of-the-money options expiring in $j + 2$ and short positions in out of-the-money options expiring in $j + 1$, and $\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}$ denotes an estimate of the cumulative variance of \mathcal{F} between t and $j + 2$. Using equation (2) allows expressing the first term on the right-hand-side as

$$\mathbb{P}_t = \sigma_{p_{j+2}}^2 - \sigma_{p_{j+1}}^2. \quad (7)$$

The second term, which is not known at time t , represents the concavity adjustment required since the static portfolio construction per equation (2) characterizes the variance profile, and does not sufficiently capture volatility, the square root of variance. In simple terms, if the VIX and VIX futures were quoted as variances and not annualized volatilities, this adjustment would not be required. Of course, expressing the VIX in volatility terms naturally increases the marketing appeal of the index, as volatility is easier to understand and interpret by market professionals, retail investors, and researchers¹⁷.

As mentioned, the term $\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}$, by definition an estimate of the variance of the underlying futures price, can never be known exactly at time t as it covers the period from t to expiration of the futures contract. Researchers have aimed to address this issue by applying different assumptions about asset and volatility dynamics in their construction of synthetic VIX futures curves using equations (2) and (3)¹⁸. Indirectly such calibrations, in one way or another, aim to estimate this unknown and dynamic parameter.

In this paper, I will work with actual VIX futures prices, thus, $\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+2}}$ is embedded in the price quote. Futures prices offer economically relevant information about the market's expectation of the variance profile of the futures contract throughout its life. In fact, the shorter the tenor, the more its pricing is impacted by this estimate of unrealized future variance, in relative terms. Another stylistic fact about volatility as reported in Poon and Granger (2003) is that forecasts of cumulative volatility become more accurate as the period of time over which volatility is generated grows due to cancellation of error and mean reversion dynamics.

Application to VIX futures fair value pricing per equation (6) would suggest

¹⁷There is such a product that is written on variances, but it is not very widely traded for reasons discussed. The S&P 500 3-month variance futures, which began trading on the CBOE on 18-May-2004, pay off according to the calculation of realized volatility over the period. The 12-month variance future, also launched in 2004, was delisted in 2011. Plans for a re-launch date of 4-October-2012 were postponed.

¹⁸As cited in section 2, see Zhang and Zhu (2006), Zhu and Zhang (2007), Lin (2007), Sepp (2008b), etc. for examples.

$$\text{Var}(\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+6}}) < \text{Var}(\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+3}}) < \text{Var}(\hat{\sigma}_{\mathcal{F}_{t \rightarrow j+1}}), \quad (8)$$

as the tenor of the futures contract decreases (increases), the uncertainty around the concavity adjustment rises (falls). This has two key implications. The first is that all else equal, shorter-dated futures should exhibit greater actual variance than longer-dated futures. This is a seemingly circular argument, but the intuition is simple. If the fair value pricing today is a function of projections about tomorrow's variance, and there is greater potential variability about this variance, then the fair value itself should exhibit greater variance. This is confirmed by the descriptive statistics in Table II.2 in a later section. The second implication is that an options market that develops on the actual VIX futures themselves should be characterized by downward sloping implied volatility curves. This is also the case, although options on VIX futures are not covered here.

The framework for evaluating the forecast bias of VIX futures

The expectations hypothesis will be applied to evaluate the forecast accuracy of VIX futures by constructing a VIX term structure according to Campa and Chang (1995). The focus is changes in expectations of future volatility, not expectations versus actual realized volatility which is a separate question. The distinction between both is addressed in the literature review in Chapter I.

Theory

Campa and Chang (1995) derive the relation for testing EH under three key assumptions: 1) volatility is stochastic, 2) the underlying asset and its volatility are uncorrelated, and 3) there is no volatility risk premium. The first assumption is a departure from the Black Scholes framework, and a well accepted concept within option pricing theory. As Hull and White (1987) demonstrate, the discrepancy between option prices using stochastic versus static volatility is independent of the level of volatility but increases with the time to maturity. The second assumption is a strong one, and generally not true in practice. A fall in asset prices is generally

associated with a rise in volatilities, as the demand for protection rises. On the other hand, rising asset prices results in less demand for protection and thus a decrease in implied volatilities. This is evident across equities, currencies, and other asset classes. Mixon (2007) tests EH using a similar formulation and points out that nonzero correlation would not present a problem provided there is no material wedge between the average expected volatility and the implied volatility for ATM options. Finally the last assumption is that volatility risk premium is zero. Campa and Chang do not test this assumption directly but instead present empirical evidence which suggests the assumption is plausible. They determine that even if a risk premium exists but not considered, the bias from omitted variables in their specification would contribute to further support of EH. Mixon (2007) expands on this work by defining volatility dynamics per Heston (1993) and proposes a formulation that is inclusive of a risk premium, thereby departing from Campa and Chang's risk neutral framework to an objective measure as suggested by the data. Although theoretically sound, his approach relies on estimates of future realized volatility which he furnishes using a GARCH (1,1) specification¹⁹.

In this paper, I will utilize Campa and Chang's risk neutral formulation to test EH. The rationale is as follows. For one, the VIX by construction does not constrain volatility to be constant. The index pools option prices for all the out-of-the-money strikes, thereby extracting all information conveyed by the short-end of the term structure and the implied volatility skew²⁰ per equations (2) and (3). With regard to the second assumption, I evaluate the implied volatility skew of options on the VIX index itself and based on Mixon (2007) confirm that there is no material wedge between the ATM strike, which is essentially the spot VIX, versus the average implied volatility of all ITM and OTM strikes for both calls and puts. Regarding the third assumption, as in Campa and Chang (1995), the omitted variable bias from using a specification that leaves out risk premiums would provide further support for EH. In other words, the regression coefficients are closer to unity, which contributes to forecast accuracy.

¹⁹Generalized Autoregressive Conditional Heteroskedasticity

²⁰Term structure refers to implied volatility differences for different maturities, while the skew refers to implied volatility differences for different strikes for a given tenor.

Constructing the VIX term structure

The innovation proposed in this paper involves testing EH based on the construction of the VIX term structure, which are essentially hypothetical variance swaps²¹, using actual VIX futures market prices. Nossman and Wilhelmsson (2008) test EH on actual VIX futures, but use only the near-term contract. As outlined previously, researchers have estimated the VIX term structure according to a model-free approach by replicating the CBOE formulas (2) and (3), many times under a wide variety of assumptions about asset and volatility dynamics, as well as alternative models for option pricing. Ait-Sahalia, Karaman, and Mancini (2012) generate the VIX term structure in the presence of jumps for the underlying asset for instance.

Whereas the standard VIX represents the risk neutral expectation of the realized volatility of the SPX over a one-month period per equation (1), the VIX term structure represents long-dated VIX contracts which are estimates of the realized volatility over a period of more than one month. For example, the long-dated VIX contract which expires at the next-term or second VIX future date, $j + 1$, is defined as

$$\sigma_{t,t \rightarrow j+1}^{VIX} \cong E_t^{\mathbb{Q}}[\sigma_{SPX_{t \rightarrow j+1}}]. \quad (9)$$

Then applying the definition of expectations hypothesis allows constructing the variance swap which expires at the next term date as the sum of the VIX plus the near-term VIX future

$$E_t \left[T_{j+1} \left(\sigma_{t,t \rightarrow j+1}^{VIX} \right)^2 \right] = E_t \left[T_j \left(\sigma_{t,t \rightarrow j}^{VIX} \right)^2 \right] + E_t \left[(T_{j+1} - T_j) \left(\mathcal{F}_{t,j \rightarrow j+1}^{VIX} \right)^2 \right], \quad (10)$$

where T is the time period expressed in years ending at the time specified by the subscript j , j represents the near-term or 1st future expiry date, $j + 1$ represents the next-term future expiry date, $j \rightarrow j + 1$ implies that the future fixes at j but yields a forecast of realized volatility on the SPX to $j + 1$, and t denotes today. The left hand side of (10) constitutes the long-dated VIX contract or variance swap, while the right hand side contains elements that are available per market pricing of the VIX and VIX futures. It is important to note that the calculation for the long-dated VIX contract which expires at the next future date does not contain the VIX future that expires at the next future date. It contains only the VIX and the VIX future that expires

²¹The terms variance swap and long-dated VIX contract will be used interchangeably, both represent hypothetical instruments that extend VIX beyond a one month tenor. The VIX at the various future tenors constitutes its term structure. This is not the same concept as the VIX futures curve.

at the near-future date. Figure 3 illustrates the intuition behind this one period lag. I proceed this way and construct VIX term structure out seven months. The VIX as quoted represents the first point on the VIX term structure.

Figure II.4 shows the resulting VIX term structure based on the sample quotes from an earlier section.

- Insert Figure II.4 here -

The constructed VIX term structure will be lower than the VIX futures curve when the slope of the VIX futures curve is positive. Forecasts of cumulative volatility become more accurate as the period of time over which volatility is generated grows due to cancellation of error and mean reversion dynamics. Thus a six-month long-dated VIX contract or variance swap represents the expected variance in the SPX over a six-month period, while the six-month VIX future represents the estimate today for the VIX index, or equivalently the thirty-day estimate of the variance in the SPX, six months from now. The latter naturally carries greater uncertainty, and thus a higher volatility price.

The testable equation for expectations hypothesis

Using the results of Hull and White (1987), the derivation by Campa and Chang of the testable equation for EH starts with establishing that the appropriate price for an option quoted at time $t = 0$ equals the Black-Scholes price, evaluated at the average variance over the life of the option

$$C^{HW} = E[C^{BS}(\bar{\sigma}_{0,t}^2)], \quad (11)$$

where C^{HW} and C^{BS} denote the prices under Hull and White and Black-Scholes respectively. It is well known, however, that the Black-Scholes model evaluated at the average variance overprices ATM options, and this overpricing increases as the time to maturity grows. And thus, for an ATM option, the concavity in σ implies

$$C^{HW} = \theta_t C^{BS} E[(\bar{\sigma}_{0,t}^2)], \quad (12)$$

where θ_t corrects for the mispricing, is less than one, decreases as time to maturity grows. The concavity is small for ATM options, and thus (12) can be expressed as

$$\theta_t C^{BS} E[(\bar{\sigma}_{0,t}^2)] \cong C^{BS} \left[\theta_t^2 E(\bar{\sigma}_{0,t}^2) \right], \quad (13)$$

relating the average variance over the life of the option to the implied Black-Scholes variance. Applying the law of iterated expectations on the definition of EH given in (13), Campa and Chang show the relation for current and future expected ATM implied volatilities, expressed in variances or volatility squared

$$\sigma_{0,km}^2 = \left(\frac{1}{k} \right) E_0 \left[\sum_{i=0}^{k-1} \sigma_{i,m,(i+1)m}^2 \right] \left(\frac{\theta_{km}}{\theta_m} \right)^2, \quad (14)$$

where m is the number of months until expiration for the short-dated option, k as the number of periods of length m , and θ is the concavity adjustment for a given tenor. Equation 14 says that the current volatility quote equals the average of the current and expected future short-dated volatility quotes. In other words, the slope of the term structure is informative about where implied volatility will be in the future. Next, I apply the same simplifying assumption, $\theta_{km}/\theta_m = 1$, as in Campa and Chang (1995). There are pros and cons to this assumption. The benefit is that the testable equation becomes more streamlined and intuitive. The cons are that the omission of terms results in the introduction of omitted variable bias. If the covariances between the term structure level and the term structure slope are negative (positive), then the omitted variable bias results in regression coefficients (for tests of expectations hypothesis) that are biased downward (upward). A downward bias for coefficients is acceptable, while an upward bias is not. Expectations hypothesis is upheld in this context if coefficients are larger (and closer to 1.0). Thus failure to reject the hypothesis test using a specification that is biased downward is conservative, but if the specification is upwardly biased, then this may be a false positive. I proceed with this specification, as the covariances aforementioned are negative²².

Since the level of volatility follows a near unit-root process, equation (14) is tested in terms of long-short spread rather than using variance levels directly. Then, subtracting the current short-dated option variance, $(\sigma^{VIX})^2$, to both sides yields the testable equation for the expectations

²²The time series of the seven-month point on the term structure is negatively correlated to the term structure slope between one and seven months, and so on for the other points on the term structure.

hypothesis on the VIX term structure, returning to the notation used in this paper, yields

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} [(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2] = \alpha_0 + \beta_0 [(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2] + \sum_{i=1}^{k-1} u_i, \quad (15)$$

where u_i represent the expectational errors. It should be noted from equation (15) that the VIX futures as quoted, \mathcal{F} , enter the regression as the constructed long-dated VIX contracts, or variance swaps, per the calculation in equation (10).

Data

The dataset for this paper contains daily closing prices for the SPX, the VIX, and VIX futures from January 2006 to November 2012. The original source of this data is the CBOE, however, they were gathered through the Bloomberg System. In addition, daily data for spot prices and volumes for ETF offerings linked to the VIX will also be used.

VIX, VIX futures, and the VIX term structure

A single reading per day is available for the VIX index. The frequency for the analysis will be weekly, however. Weekly data is selected as it helps reduce the forecast bias introduced by including overlapping forecasts in the analysis, which produces serially correlated errors. In addition, each day there is a strip of VIX futures contracts that expires at set dates in the future. The near or 1st future expires in the same month, the next or 2nd future expires in the following calendar month and so on. When VIX futures were launched on 26-March-2004, there were a total of four futures contracts trading on any given day. Today there are total of nine futures contracts, with the longest maturity being more than one year out. In order to utilize as much history as possible in this analysis, I include the first six months of maturities only, and focus on tenors two, three, and six. This allows merging the older data with the newer data, without having to rely too much on interpolation and extrapolation to fill missing maturities. Table II.2 contains descriptive statistics on the VIX, VIX futures, and the resulting VIX term structure spanning the period 2006-2012.

- **Insert Table II.2 here** -

The average figures suggest there is a tendency for the VIX futures curve to be upward sloping. In addition, the ranges of prices for the VIX as well as the various points along the VIX futures curve and VIX term structures convey that the volatility of the prices themselves decreases as the time to maturity increases, and this is confirmed by the standard deviation readings. The intuition for this was discussed in the VIX futures pricing section. Essentially as the tenor of the futures contract rises, the uncertainty around the concavity adjustment should fall. Skewness readings are in line with implied volatilities in other asset classes. Kurtosis figures indicate fat tails for the underlying VIX, but not for the set of futures. The short-end of the resulting VIX term structure displays fat tails.

Figure II.5 tracks the spread between six-month VIX futures (left axis) and the underlying VIX (right axis) for the entire sample period.

- **Insert Figure II.5 here** -

A few observations worth noting. For one, roughly 75% of the time, the spread between the long-end future and the VIX is positive. Also, the spread went steeply negative during the height of the 2008 credit crisis and then materially negative in the latter half of 2011 when Euro zone crisis fears peaked. And three, the average spread when positive was substantially greater post-crisis, as opposed to pre-crisis.

About the curves associated with the VIX

Figure II.6 is a snapshot of the various curves associated with the VIX as of the close of business on 3-November-2010, the day the second round of quantitative easing (QE2) was announced in the United States.

- **Insert Figure II.6 here** -

The VIX index closed at 19.56, while the 6-month VIX future closed above 25, displaying a significant premium. As discussed, when the VIX futures curve is upward sloping, the constructed VIX term structure or variance swap curve, will generally be lower. Also depicted in the

figure are the term structures for ATM and 25-delta OTM implied volatilities for SPX options. SPX options are the building blocks of the VIX index. It is common for the term structure of ATM volatilities on underlying SPX options to be lower than the VIX term structure across tenors, as the latter prices in the skew that is prevalent for OTM options on all financial assets. By construction, the VIX gives the holder exposure to the full set of OTM options on SPX at any one particular time, and thus the skew is reflected in the higher term structure.

Regarding the relationship between the VIX term structure and the term structure of 25-delta OTM SPX options, we would expect the latter to be higher on average. The payoff of a variance swap such as the VIX is convex in volatility. This means that an investor who is long a variance swap will benefit from boosted gains and discounted losses, a phenomenon which is amplified when volatility skew is steep. Thus, the fair strike of a variance swap is often in line with the implied volatility of 40-delta SPX puts, which is lower than that of 25-delta puts in the presence of skew.

Testing the forecast bias of the VIX term structure

Expectations hypothesis will be tested to address the following research question. When the VIX term structure is positively (negatively) sloped, does the VIX subsequently rise (fall) as much as predicted? Specifically, equation (15) evaluates the long-short spread of the constructed VIX term structure at time $t = 0$ against subsequent changes in the VIX over two, three, four, five, six and seven-month periods. The long-short spread is an unbiased predictor of the future movements in the VIX if $\alpha = 0$ and $\beta = 1$. This would imply that for every unit of variance in the long-short spread of today, we would expect one unit of variance in subsequent movements in the VIX index. Rejections of this joint hypothesis test will be reported at the 5% level.

In addition, while it is useful to have unbiased forecasts, biasness and predictive power are two separate concepts. As Poon and Granger (2003) point out, a biased forecast can have predictive power, but an unbiased forecast is useless if forecast errors are always big. For this reason, I also run a hypothesis test to evaluate the extent to which $\beta = 1$. Rejection of a hypothesis test based on this single coefficient would not resolve the issue of biasness, as a joint hypothesis test would be required, however the test is important from the standpoint of highlighting potential arbitrage opportunities. I also run a third test for investigating whether

$\beta = 0$. Rejecting this hypothesis would confirm that the VIX term structure contains valuable information content.

Results

Table II.3 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately. The regressions are based on weekly readings from January 2006 to February 2012. Weekly data is selected as it helps reduce the forecast bias introduced by including overlapping forecasts in the analysis, which produces serially correlated errors. The serial correlation is also addressed by applying error correction techniques outlined in Newey and West (1987)²³.

- Insert Table II.3 here -

The joint hypothesis test results dictate that the long-short spread does predict the direction of the subsequent move in the VIX correctly, since the beta coefficients are positive, but for the most part, not to the extent implied by the expectation hypothesis, since they are all less than one. The coefficient reading of 0.802 for the seven-month point on the VIX term structure implies that for every 10 units in today's long-short spread, a rise in the VIX index of 8.02 units is expected over the next seven months. The beta coefficients for all points on the curve, with the exception of the 7-month, are all significantly less than one at the 5% level, implying that VIX futures are consistently overpriced. The forecast bias increases substantially with shorter tenors. This is a feature also prevalent in the results for Campa and Chang (1995) and Mixon (2007), confirming a well-known stylistic fact about financial options. Long-dated options offer better value to hedgers and speculators as opposed to short-dated options. The results of the second hypothesis test, which evaluates whether $\beta = 1$ individually, does not reveal any new information over the joint test. The results of the third hypothesis test, which evaluates whether $\beta = 0$, establishes that there is no information content in the VIX term structure over the short-part of the curve only (at two and three-month tenors).

²³The lags for this error correction technique will be selected according to the point on the term structure being evaluated. The longer the tenor, the greater the lag.

Subperiod analysis

To control for the time period effect, I run the hypothesis tests for three disjoint periods surrounding the 2008-9 credit crisis period: pre, during, and post. Table II.4 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately for the pre-crisis subperiod.

- Insert Table II.4 here -

With exception of the seven-month point on the curve, the beta coefficients are closer to one in comparison to those of the entire sample, and the improvement is greatest at the short-end. In addition, the standard errors are larger in this subperiod, contributing to joint hypothesis test results which suggest no forecast bias for the four through seven month points on the curve. During this subperiod, the VIX averaged 12.8%, the six-month long-short spread of the VIX term structure averaged 1.6%, and the six-month VIX future averaged 2.4% above the underlying VIX index. The term structure was consistently forecasting an appreciation in the VIX, and the forecast was upheld for the term structure at months four, five, six and seven.

Table II.5 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately for the subperiod that covers the credit crisis of 2008-9.

- Insert Table II.5 here -

This was a highly volatile period. Financial option prices rose across asset classes, commensurate with increased market uncertainty. The average level of the VIX index was 31%, more than double the pre-crisis period, and peaked at over 70% following the Lehman bankruptcy in the fourth quarter of 2008. Also the term structure was inverted for a majority of this period, a common feature that arises during times of financial distress. The average spread between the 6-month VIX future and the VIX was -3.9%. Results show beta coefficients are higher across the board, in fact, for the seven-month point on the term structure, it is above one, which suggests subsequent move in the VIX was greater than forecast by the term structure. The hypothesis test results demonstrate no forecast bias for the five, six, and seven points on the curve. The intuition from these results suggests that options are well-priced during periods of extreme volatility, despite the higher premiums paid.

Table II.6 contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately for the subperiod following the credit crisis of 2008-9.

- Insert Table II.6 here -

Post-crisis, there are rejections of the expectations hypothesis across the board. The beta coefficients are further away from one as compared to the other subperiods, in fact, they are negative for the two and three-month points on the curve. A negative coefficient implies that a positive slope on the term structure was associated with subsequent falls in short-dated volatility forecasts. The average reading on the VIX was 24.3% for this period, roughly a quarter below the crisis period average. However, the slope of the term structure was consistently positive and steep. The average long-short spread of the VIX term structure at the six month point was 2.9% and the average spread between the 6-month VIX future and the VIX was 3.8%. The divergence between current and future expectations is an important phenomenon. The overall level of risk as determined from the VIX was lower during the post-crisis period, but the expectations for the future level of the VIX to rise as extracted from the term structure slope were among the highest levels registered. The hypothesis tests determined that the expectations of subsequent rise in the VIX based on the positively sloped term structures did not manifest.

To summarize the results of the tests, I will highlight the some important observations. First, VIX futures are consistently overpriced relative to the subsequent moves in the underlying VIX index. Second, the forecast bias increases substantially with shorter tenors. Third, the forecast bias is smallest during periods of extreme volatility. Four, deviations from expectations hypothesis are greatest during the post-crisis period.

What drives the forecast bias?

Now that the existence of the forecast bias of the VIX term structure has been documented, characterized and quantified, the focus shifts to addressing why it exists and why it persists. A number of possible factors might influence the size of the forecast bias over time. This section includes a discussion of each, and a description of the specific variables that will be used to establish and evaluate the linkage quantitatively in the next section.

Open interest

The VIX itself is not a tradeable asset. It is not possible to speculate or invest in the index directly. Exposure to the VIX is attained by synthetically replicating the index using portfolios of individual options on the S&P 500 according to the methodology outlined by the CBOE, trading futures and options whose payoff is determined by future levels of the VIX, or via ETF's or managed portfolios offered by the investment management industry. These methods are listed in order of decreasing difficulty of implementation.

Over the years following the credit crisis of 2008-9, the proliferation of ETF offerings involving the VIX has opened the flood gates and capital has poured into the strategy. Today an investor does not have to replicate the index or trade VIX futures or options on the VIX to get exposure to the index, both without question more difficult avenues. Instead, an investor may buy or sell individual ETF shares, via a number of online portals, at very low transaction costs. VIX-related ETF's, along with managed funds offerings, have quickly grown to a market capitalization estimated at \$3-5 billion between 2009 and 2012. The market capitalizations of three of the larger funds, according to Bloomberg data, are as follows: 1) the iPath S&P 500 VIX short-term futures fund with \$1.56 billion in assets (ticker VXX), 2) the iPath S&P 500 VIX medium term fund with \$116m in assets (ticker VXZ), and the iPath S&P 500 VIX dynamic fund with \$309m in assets (ticker XVZ), as of November 2012.

The transmission mechanism from ETF's to VIX futures markets is simple. In order to provide individual investors with a risk profile that tracks the underlying VIX index, professional money managers must trade futures and, to a lesser degree, options on the VIX. Mass buying of VIX ETF's then translates into large amounts of capital flowing into futures markets. A number of different portfolio offerings are available, based on varying strategies. The VXX ETF, the largest by market capitalization, offers exposure to the short-end of the term structure and thus a risk profile that most closely mimics that of the underlying VIX index²⁴. The XVZ ETF, in contrast, offers exposure to the entire term structure depending on perceived value at the discretion of the money manager. Figure II.7 tracks the open interest in VIX futures along with the market capitalization of the VXX ETF.

²⁴The descriptive statistics in table II.2 demonstrate that the distribution of the short-end of the term structure more closely matches the distribution of the underlying VIX index, as compared to the longer-end of the curve.

- Insert Figure II.7 here -

Open interest for VIX futures markets averaged approximately 50,000 contracts prior to 2009. After 2009, there was a sharp rise in open interest, no doubt related to the proliferation of VIX ETF offerings, as seen in figure II.7. Today, ETF's represent the main vehicle for accessing the VIX, and inflows into ETF's implies inflows into the market for VIX futures. Open interest may play a key role in explaining the magnitude of the forecast bias in the VIX term structure, thus, I include these data in the regression analysis. The source of open interest data is the CBOE, however, the data was gathered through Bloomberg.

Slow-moving capital

Application of the theory of slow-moving capital, as suggested by Gromb and Vayanos (2002), Brunnermeier and Pedersen (2009), Duffie (2010), Ashcraft et. al (2010) and others, stresses that slow-moving capital may play a key role in propagating mispricing in financial markets. Is it plausible that changes in the availability of capital that could potentially be directed towards arbitrage strategies contributes to changes in the forecast bias in the VIX term structure? I follow Matthias, Longstaff, and Lustig (2013) in testing for this by including changes in total global hedge fund net asset values as one of the variables in the regression. The ticker of the specific time series from Bloomberg is HFRXGL Index²⁵, however, the original source of this data is Hedge Fund Research Inc..

Performance of safe-haven assets

A number of investments have emerged as potential safe havens following the 2008-9 credit crisis: US treasuries, gold and other precious metals, agricultural commodities, the Japanese yen, the Swiss franc, the Chinese renminbi, land, real estate, etc.. Some have fared better than others²⁶,

²⁵The HFRX Global Hedge Fund Index is designed to be representative of the overall composition of the hedge fund universe. It is comprised of all eligible hedge fund strategies; including but not limited to convertible arbitrage, distressed securities, equity hedge, equity market neutral, event driven, macro, merger arbitrage, and relative value arbitrage. The strategies are asset weighted based on the distribution of assets in the hedge fund industry.

²⁶Currency liberalization has helped the Chinese renminbi hold its value despite expectations of a soft/hard landing.

some are more accessible than others²⁷, while some have lost their safe haven status altogether²⁸. For the regression analysis, I will use the price of gold as representative of this asset type. The outperformance of safe haven assets, itself a manifestation of fear-driven buying, may have an impact on the mispricing of volatility products since it would imply investor interest in risk management. The Bloomberg index for these data is XAU Curncy.

The cost of insuring against tail risks

Interest in protecting financial interests against tail risks or 'black swan' events is another byproduct of the 2008-9 credit crisis. Note, this is not the same as the previous factor. Investing in safe haven assets is motivated by the return of principal, not necessarily the return on principal, and thus this typically involves a reallocation of capital from one source (higher yielding riskier asset) to another (lower-yielding safer asset). Insuring against tail risks, however, implies continuing to hold the higher-yielding riskier asset, but removing the impact of disaster outcomes via hedging vehicles. The relative cost of insuring against tail risks is a variable that captures this important element of investor attitudes and behavior. My initial conjecture is that rising interest in hedging tail risks exacerbates the forecast bias in the VIX term structure.

As a suitable proxy, I will include the 25-delta 6-month option implied volatility risk reversals for the USDBRL exchange rate. Expressed in units of annualized volatility per annum, this time series tracks the differential between out-of-the-money puts and out-of-the-money calls on the Brazilian real (a high-yielding emerging market currency widely considered a risky asset). The greater the differential, the greater the interest in buying protection against a fall in the real. This is associated with rising aversion to risk. The Bloomberg index for these data is USDBRL25R6M Curncy.

²⁷Hard assets such as forestry, land, and real estate are in general more difficult to procure as investments, manage, and value.

²⁸The Swiss National Bank pegged the franc to the euro on 6-September-2011 in response to excessive inflows and currency strength which it declared to be "a threat to the economy". The level of the peg represented over a 20% depreciation of the franc from recent highs.

Credit risk

Seeking arbitrage profits generally involves capital intensive strategies²⁹. If the cost of credit rises either too quickly or too far above cost of funds, the speculative capital available for such pursuits will decrease. Swap spreads represent the yield differential, in basis points, for swap contracts versus treasuries. Swap contracts are the benchmark instrument for pricing loans in the private sector, while treasuries represent the US government's cost of borrowing. Divergences between the two are associated with rising levels of credit risk in the financial system. Because the tenor of the trades involving VIX futures are under 1-year, I will include the on-the-run 12-month swap spreads in the regression, available through Bloomberg ticker USSP2 Index.

Arbitrage existence and persistence

The premise behind financial markets arbitrage is that asset mispricing or systematic forecast biases may be exploited by simultaneously buying and selling two different versions of the same risk profile, at two different prices. The underpriced (overpriced) asset is bought (sold), and the risk is offset by selling (buying) the other. Barring any unforeseen circumstances³⁰, the price differential is the profit. In other words, it is a hedged bet offering lower risk and lower return. I will construct an arbitrage strategy based on the forecast bias documented that is implementable in practice, and evaluate the extent to which changes in the magnitude of the arbitrage profits are impacted by each of the factors discussed. Note the arbitrage profits are a proxy for the forecast bias³¹. The proxy is necessary for evaluating its existence and persistence in a practical setting.

Replicating the VIX index

The textbook approach for constructing an arbitrage strategy which aims to profit from the forecast bias discussed in this paper would involve systematic selling of VIX futures against a long position in a VIX-replicating portfolio constructed from underlying options on the S&P

²⁹This will be illustrated in the next section which outline the construction of the arbitrage strategy involving the VIX term structure.

³⁰Such as financial fraud, insolvency, illiquidity.

³¹The terms may be used interchangeably when discussing the results and intuition.

500 equity index. As characterized previously, ex ante forecasts of the future level of the VIX extracted from VIX futures consistently overshoot ex post realizations. VIX futures would be sold, and the risk would be hedged by buying the synthetically constructed VIX. The premium earned from the former would be greater than the cost of replicating the index, and the difference would be profit. Although mathematically precise, this avenue is not easily implemented in practice. Even if implementation is attempted, the performance slippage between the VIX index and the VIX-replicating portfolio may become large, thereby eating into the expected return from the arbitrage.

There are 3 elements that challenge the feasibility of constructing a VIX-replicating portfolio: the volume of trades, the frequency required, and the high transaction costs for low-delta options. Strict construction of the definition of the VIX would involve holding portfolios of hundreds, even thousands of options, at any one particular time. Table II.1 contains a collection of options on SPX that would be involved in a single hypothetical calculation of the VIX as of 2-Nov-2012. In practice, option strikes are available for every five points on the SPX index, however for expositional purposes, the intervals used are fifteen and twenty-five SPX points apart for near and next-term maturities respectively.

Once the portfolio is constructed, price fluctuations of the SPX index would require rebalancing the VIX-replicating portfolio multiple times a day. The greater the volatility of the SPX index, the more rebalancing required, resulting in thousands of individual transactions a day, often times in just an hour. Developing a computer algorithm for executing the trades is necessary.

Finally, the theoretical calculation of the VIX uses mid prices. In financial markets, the mid price is the price between the best price of the selling dealer's offer or ask price and the best price of the buying dealer's bid price. Many times it is simply the average or midpoint of the current bid and ask prices being quoted by the dealer. Constructing and managing a VIX-replicating portfolio represents an important deviation from theory, as it is not possible to transact at mid prices. Options would be bought at the dealers offer price (higher than the mid), and options would be sold at the seller's bid price (lower than the mid price). Each trade executed conceivably presents a departure away from the benchmark VIX index. In fact, such transaction costs increase disproportionately for deep out-of-the-money options.

It is possible to deviate from strict construction of the VIX definition outlined by the CBOE.

Delisle et al. (2010) offer an alternative VIX-replicating portfolio which would significantly reduce the administration requirements. However, the tradeoff of such an approach is that it would undoubtedly introduce an additional layer of performance slippage between theory, the VIX index, and practice, the VIX-replicating portfolio.

Arbitrage strategy, practical approach

In contrast, the results of the hypothesis tests can be used to construct an arbitrage strategy that may be feasibly implemented in practice. Table II.3 establishes that the forecast bias is greatest at the short-end of the VIX term structure, and narrowest at the long-end. Tables II.4-6 confirm this across subperiods. The foundation for the arbitrage based on these results involves systematic selling the VIX term structure at the short-end where futures are generally overpriced (leg 1), and hedging this position by establishing a long exposure, equal in notional, to the long-end of the VIX term structure, where the expectations hypothesis is upheld, implying no forecast bias (leg 2). The log return of leg 1 is given by the following

$$r_{leg1} = \sum_{j=1}^2 \ln \left(\frac{\mathcal{F}_{0,j \rightarrow j+1}^{VIX}}{\sigma_{j \rightarrow j+1}^{VIX}} \right), \quad (16)$$

where $\mathcal{F}_{0,1 \rightarrow 2}^{VIX}$ and $\mathcal{F}_{0,2 \rightarrow 3}^{VIX}$ represent the near and next futures contracts that will be shorted at time $t = 0$, the payouts of which will be determined by the prevailing level of the VIX one and two months later, denoted by $\sigma_{1 \rightarrow 2}^{VIX}$ and $\sigma_{2 \rightarrow 3}^{VIX}$ respectively. Leg 1 is expected to generate profits on average.

Similarly, the log return of leg 2 is given by the following

$$r_{leg2} = \sum_{j=1}^7 \ln \left(\frac{\mathcal{F}_{0,j \rightarrow j+1}^{VIX}}{\sigma_{j \rightarrow j+1}^{VIX}} \right), \quad (17)$$

where $\mathcal{F}_{0,j \rightarrow j+1}^{VIX}$ for $j = 1 \dots 7$ represent the long positions in VIX futures established at time $t = 0$, whose payouts will be determined according to the prevailing level of the monthly fixes in the VIX index starting with the near futures expiry and ending at the 7-month expiry date. Leg 2 is also expected to generate profits on average, albeit smaller than leg 1, according to the lower beta coefficients from Table 3.

The profit to the arbitrageur is the net payout from simultaneously buying leg 1 and selling leg 2 at the respective ratios given by

$$r_{arb} = \frac{r_{leg1}}{2} - \frac{r_{leg2}}{7}, \quad (18)$$

where the gains from being long leg 1 are expected to exceed losses from shorting leg 2. The trade is replicated weekly. Changes in the size of the arbitrage profit, which is an expression of the forecast bias identified through the various tests of the expectations hypothesis, will be the dependent variable in the regression.

Regression results

I explore the contribution of the factors that might influence the size of the forecast bias over time by regressing weekly changes in the realized profit from the arbitrage on weekly changes in the explanatory variables. The regression is carried out twice. The first uses the entire sample of data. The second evaluates only the post-crisis period. Table II.7 reports the results.

- Insert Table II.7 here -

Starting with the full period analysis, the results indicate that the forecast bias is affected only by the capital flow variable. The sign of the coefficient is particularly illuminating. It suggests that rises in global hedge fund net asset values widen the arbitrage available, itself a representation of the forecast bias. The coefficient estimate indicates that a 10% increase in hedge fund capital in the system would be associated with a 16.9% rise in arbitrage profits. On the surface, it would seem this result does not provide support for the slow-moving capital hypothesis. However, it is important to note that this regressor includes both capital inflows and capital appreciation. This is the definition of net asset value. In other words, it cannot be unequivocally determined that a rise in net asset value directly translates into flows into a particular arbitrage strategy. Rises in net asset value imply both inflows and capital gains.

The subperiod regression indicates that the forecast bias is affected by both the open interest and the capital flow variables, significant at the 5% level. The fit as compared to the full period analysis is superior, based on the improved R-squared reading of 12.1%. The coefficients are

both positive, suggesting that capital inflows into the greater market in general via hedge funds and specifically into VIX futures markets, in concert, exacerbate the forecast bias in the VIX term structure. Once again, the opposite situation would be expected according to the slow-moving capital hypothesis. Valuable insight is gained from differentiating between the parties involved in these two types of capital allocation. In addition, the interpretation of the capital flow variable must be expanded.

Open interest factor implies inflows from non-professional investors

The transmission mechanism from ETF's to VIX futures markets has been established. In order to provide individual investors with a risk profile that tracks the underlying VIX index, professional money managers offering such ETF's must trade futures and, to a lesser degree, options on the VIX. Mass buying of VIX ETF's then translates into large amounts of capital flowing into futures markets. This is confirmed by the open interest variable. Open interest for VIX futures markets averaged approximately 50,000 contracts, prior to 2009 and prior to the launch of VIX ETF funds. After 2009, there was a sharp rise in open interest, multiples above previous levels, no doubt related to the proliferation of VIX ETF offerings following the credit crisis of 2008-9.

With that as background, it can be said that the open interest variable used in the regression is largely a reflection of increases or decreases in VIX ETF volumes. Historically, the bulk of ETF usage has come from non-professional investors, as opposed to hedge funds, institutional and professional investors³². Thus, the open interest factor is a reflection of inflows of non-professional investors, which are in general, less sophisticated than professional investors in a relative sense. This is passive capital in search of buy and hold strategies, as opposed to arbitrage opportunities.

³²According to the annual study of the United States investment management market conducted by Greenwich Associates during 2010, ETF usage amongst United States pension funds, endowments, and foundations grew to approximately 14%.

Capital flow factor implies inflows of more sophisticated capital, as well as capital appreciation

The capital flow variable used in the regression is proxied by the HFRX Global Hedge Fund Index, a strategy-weighted representation of the total net asset value of the hedge fund industry comprised of all eligible hedge fund strategies including but not limited to convertible arbitrage, distressed securities, equity hedge, equity market neutral, event driven, macro, merger arbitrage, and relative value arbitrage. This undoubtedly reflects more sophisticated capital which could in theory be readily deployed to exploiting arbitrage opportunities as they arise. The hedge fund investor may not be sophisticated, but full discretion for asset allocation is given to the hedge fund manager. As previously stated, net asset value reflects increases in inflows as well as capital appreciation. Due to this confounder, it can be established that the open interest factor is more direct measure of capital flows into VIX-related strategies.

Intuition

The positive coefficients for the open interest and the capital flow variables for the regression following the credit crisis of 2008-9 imply that increases in the availability of capital contribute to the biasness of the VIX term structure. The open interest variable reflects inflows of capital from non-professional investors, which are generally less sophisticated, while the capital flow variable reflects inflows of more sophisticated capital, in addition to capital appreciation. Both are happening in tandem. It is the former, the open interest variable, that can be interpreted to be evidence against the slow-moving capital hypothesis, as it more clearly reflects inflows into VIX-related strategies. This phenomenon is illustrated in greater detail in Figure II.8.

- Insert Figure II.8 here -

I construct a series based on the 2-month VIX term structure tests of the expectations hypothesis. The points on the time series represent the beta coefficients for rolling 1-year regressions, per equation 15. A decrease in the coefficient away from one implies a larger forecast bias. The series is then plotted against the market capitalization for the VXX ETF, the largest offering from this family of funds. The link between the size of the forecast bias and the size of the VIX fund is evident, as previously established from tests of the expectations hypothesis.

Lastly, the positive coefficient for the capital flow variable may also imply that as hedge funds perform well, so does the market in general. The combination of a bullish cycle for risk assets and more inflows into VIX futures (as suggested by the open interest variable) presents the perfect storm for a widening of the forecast bias in the VIX term structure. As the S&P 500 rises, the VIX index falls. If this is happening at the same time as strong inflows are heading into the market for VIX futures, this will result in a steepening of the VIX term structure and thus higher forecasts for the future level of the VIX. These forecasts will not manifest as long as the bull cycle persists, however.

Closing comments

The focus of this paper is first to evaluate the information content of VIX futures prices, the core method of attaining exposure to the VIX index. There are two main objectives. The first is to identify, characterize, and quantify the forecast bias of the VIX term structure, which is constructed from VIX futures prices. There are four key findings along these lines.

First, VIX futures are consistently overpriced relative to the subsequent moves in the underlying VIX index. Second, the forecast bias increases substantially with shorter tenors. Third, the forecast bias is smallest during periods of extreme volatility. Fourth, deviations from expectations hypothesis are greatest during the post-crisis period. These findings describe the VIX-VIX Futures Puzzle.

The second objective, once the forecast bias has been documented, is to shed light on why the forecast bias exists and why it persists. I first identify a number of factors that might influence the size of the forecast bias over time: futures open interest, hedge fund capital flows, performance of safe haven assets, the costs of insuring against tail risks, and the amount of credit risk in the financial system. I then construct an arbitrage strategy which aims to profit from the forecast bias I identify, and regress weekly realized changes in the arbitrage profits versus changes in the factors. Note the size of the arbitrage profits are a proxy for the size of the forecast bias. The results of this regression for the subperiod following the 2008-9 credit crisis suggest that capital inflows into VIX futures (the open interest factor) and into the greater market in general via hedge funds (the capital flows factor), in concert, exacerbate the forecast bias in the VIX term structure. This is a particularly surprising result, as the opposite situation

would be suggested by application of the slow-moving capital hypothesis. The persistence of the arbitrage is related to key differences between the parties involved in the capital allocations. The open interest variable reflects inflows from non-professional investors, which are generally less sophisticated than professional investors, while the capital flow variable reflects inflows of more sophisticated capital, as well as capital appreciation. As the VIX has become more accessible to the average investor, this has created distortions to VIX futures markets, especially for short-dated tenors which are most actively used for management of ETF funds. Lastly, the positive coefficient for the capital flow variable may also imply that as hedge funds perform well, so does the market in general. The combination of a bullish cycle for risk assets and more inflows into VIX futures (as suggested by the open interest variable) presents the perfect storm for a widening of the forecast bias in the VIX term structure.

Future research will be dedicated to evaluating alternative explanations of the VIX-VIX Futures puzzle, including the development of the idea that the forecast bias in the VIX term structure may be attributed to investor appetite for paying a certain premium to VIX products because of the diversification benefits offered. Another interesting avenue for research involves the development of a new formulation for testing EH for the VIX and other volatility products, deviating from a risk-neutral framework. In addition, while I focused on tests of the expectations hypothesis applicable to the term structure of implied volatility, it should be noted that VIX futures are more of a hybrid product. By construction, the VIX is determined from underlying option prices, however, VIX futures themselves offer linear returns, a feature of futures and forward contracts. Methods for testing the expectations hypothesis in futures markets were not directly applied here since the VIX is not an investable asset such as gold, oil, and other commodities where cost of carry relationships, storage costs, transportation costs, and insurance premiums impact futures pricing. Nonetheless a new formulation combining the elements of both futures and options is worth exploring.

Appendix for Chapter II

VIX portfolios: Hedge or diversification?

The objective is to establish that capital allocations to VIX portfolios provide diversification, not direct offsets aimed at covering losses, a feature typically associated with hedging activity. This class of investor desires a risk profile that is directly tied to market sentiment and not necessarily to underlying equity prices. There is a common misconception that a capital allocation to a portfolio constructed from VIX products would provide a suitable hedge to a long equity position given the strength of the co-movement between the VIX and the SPX on a mark-to-market basis. Figure II.9, however, shows a strong counterargument.

- Insert Figure II.9 here -

A dollar invested in the S&P 500 index in February 2011 would have been worth about the same in February 2012, with minimal divergence from the benchmark starting point. A dollar invested in VXX, an exchange traded fund (ETF) which takes long positions in VIX futures, would also have been worth a dollar at the end of the period but its value would have, in contrast, oscillated significantly over this time period, falling over thirty percent in the summer of 2011 and then nearly doubling in the fourth quarter of 2011 when the Euro zone sovereign debt crisis dominated headlines. The combined risk profile of these two assets would not constitute prudent hedging activity in a traditional sense. A hedge by definition should offset losses, not enhance returns or exacerbate volatility as is evident in this case.

Figure II.10 shows the rolling three-month correlation of changes between the SPX and the VIX, and the realized standard deviation of each.

- Insert Figure II.10 here -

As established, the two indices are well correlated. Correlation however does not tell the whole story, and a high correlation, even if persistent, does not ensure that gains and losses on the core investment will be sufficiently offset by gains and losses on the hedge. The optimal hedge ratio may be changing despite the stability of the correlation.

The time-varying nature of optimal hedge ratios has been addressed extensively in the literature. There is no debate about the fact that optimal hedge ratios change across time, the controversy centers around whether or not there is merit in implementing dynamic versus static approaches in the context of portfolio construction. Figlewski (1984), Lypny (1988), and Baillie and Myers (1989) find that a dynamic hedge strategy outperforms the time invariant hedge, while Smirlock (1985) and Cecchetti et al. (1988) find that time invariant hedge ratios perform better than the dynamic approach.

I will sketch out the basic theoretical argument for explaining the ineffectiveness of the VIX as a true hedge for a long position in the SPX. Despite the strong correlation, the ratio of the volatilities is the source of the ineffectiveness of the hedge. Consider a two-asset portfolio comprised of long positions in the SPX and VIX indices³³ whose combined variance is given by

$$\sigma_{SPX,VIX}^2 = w_{SPX}^2 \sigma_{SPX}^2 + w_{VIX}^2 \sigma_{VIX}^2 + 2\rho w_{SPX} w_{VIX} \sigma_{SPX} \sigma_{VIX}, \quad (19)$$

where w represents the weights of each asset, σ^2 are the respective variances, and ρ is the correlation between the two assets. The two assets by definition offer linear returns. Suppose w_{SPX} is normalized at 1.0. The exercise involves selecting w_{VIX} such that the combined portfolio variance reduction is maximized, and thus we compute the partial derivative of equation (19)

$$\frac{\partial \sigma_{SPX,VIX}^2}{\partial w_{VIX}} = 2w_{VIX} \sigma_{VIX}^2 + 2\rho \sigma_{SPX} \sigma_{VIX}. \quad (20)$$

The optimal allocation to the VIX is then found by setting (20) equal to zero and solving for w_{VIX}

$$w_{VIX} = -\rho \frac{\sigma_{SPX}}{\sigma_{VIX}}. \quad (21)$$

Equation (21) says that the optimal allocation to the VIX should be proportional to the ratio of the volatilities of both elements in the portfolio. Figure II.11 shows this ratio historically, based on the rolling 3-month calculations displayed in figure II.10.

- Insert Figure II.11 here -

³³For expositional purposes, I will assume the VIX is directly investable. The results should hold if actual market instruments are used.

Working with the assumption that $\rho = 1$, a reading of 0.10 as in 2006 would suggest the optimal allocation to the VIX should be 10 cents for every dollar invested in the SPX, in other words, equal to the ratio depicted in figure 6 since w_{SPX} is normalized at 1.0. During the credit crisis of 2008, this optimal hedge ratio would have risen close to five times 2006 levels. The intuition for this is as follows. By construction, the VIX is designed to maintain a constant price sensitivity to the implied volatility of the SPX over a one-month period. This is achieved by constructing the index as being a weighted average of all OTM options for the two nearby expiries, centered around the prevailing SPX futures curve, at any point in time³⁴. Thus, changes in the value of the underlying SPX index will change the hypothetical portfolio of options that determine the value of the VIX. As the SPX drops, the hedger would ideally want to continue to hold options struck at the higher strikes. The VIX, however, does not offer this risk profile, as it is only defined by the full set of out-of-the-money (OTM) options. Once an option goes in-the-money (ITM), it will no longer be reflected in the price of the VIX. Instead, the VIX holder would synthetically own a greater concentration of the lower strikes in tandem with the fall in the SPX index, and none of the higher ITM strikes that would insulate the fall in equity prices. Maintaining an optimal hedge ratio would then require upsizing w_{VIX} as spot SPX is dropping. Along similar lines, as SPX rises, maintaining an optimal hedge ratio would require reducing w_{VIX} . This is seen clearly confirmed by figure 11, the optimal volatility-reducing holding of the VIX, w_{VIX} , would have fallen dramatically following the post 2008-9 credit-crisis bottom of the SPX in March 2009.

Furthermore we can establish that a dramatic shift in the ratio of volatilities has the same impact as that of a significant drop in correlation from the standpoint of hedge effectiveness, defined as

$$\mathcal{H}_{VIX} = -\frac{\sigma_{SPX,VIX}^2 - \sigma_{SPX}^2}{\sigma_{SPX}^2}, \quad (22)$$

where \mathcal{H} takes on a maximum value of 1.0 denoting perfect effectiveness. Combining (19) and (21) arrives at the expression for hedge effectiveness

$$\mathcal{H}_{VIX} = \rho. \quad (23)$$

³⁴Refer to table 1 for illustration.

In other words, if the hedge ratio is constructed according to the optimal ratios outlined in (21), then \mathcal{H} is purely a function of the correlation. Suppose however that due to a change in the ratio of the volatilities as depicted in figure 11, the portfolio is hedged according to the less optimal hedge ratio given by

$$w_{VIX} = -\frac{1}{3}\rho\frac{\sigma_{SPX}}{\sigma_{VIX}}, \quad (24)$$

which says that despite the rise in ratio of volatilities, the allocation to the VIX, w_{VIX} , was not changed. The hedge effectiveness for this less optimal portfolio would be given by

$$\mathcal{H} = \frac{1}{18}\rho^2. \quad (25)$$

Figure II.12 contains a plot of the efficiency under the optimal and sub-optimal hedge ratios against all possible values of ρ .

- Insert Figure II.12 here -

It is clear that even at perfect correlation levels, the hedge efficiency is low under a sub-optimal hedge ratio. Otherwise stated, the impact on hedge efficiency of a portfolio of two assets that exhibit a high correlation to one another but constructed at a suboptimal hedge ratio is similar to that of a portfolio of assets that exhibit a low correlation, even if this portfolio is constructed according to the optimal hedge ratio. The rapidly changing ratio of the volatilities is the root of hedge inefficiency. Low hedge efficiency is associated with diversification, not hedging.

A true hedge from the standpoint of offsetting gains and losses on an underlying core investment in S&P 500 stocks would involve SPX futures or options. Futures and options span the full spectrum of payoff outcomes. The former offers linear returns, full downside protection, does not require upfront premium, but offers no upside potential. The resulting risk profile for hedging via SPX futures is equivalent to exiting the investment in the SPX altogether. On the other hand, options on the SPX offer asymmetric returns, full protection and full upside participation, but require initial premium which averages five to ten percent per annum. VIX portfolios are generally marketed as combining the attractive elements of both SPX futures and options: linear returns, no upfront premium and upside potential. There is, however, no free lunch in

financial markets and thus the holder must give up having full downside protection. In addition, the value of the VIX portfolio itself may fall in tandem with rises in equity prices. Investor incentives with regard to allocations to VIX portfolios are thus associated with diversification.

Figures for Chapter II

Figure II.1. Historical prices for VIX and SPX. This figure tracks historical price data on the VIX and the SPX, from January 2006 to October 2012. The strong association between price series is supported quantitatively by a correlation coefficient, based on daily changes, of -76% for the entire period. In addition, the average three-month rolling correlation is -83%. The negative number implies co-movements occur in the opposite direction.




Figure II.2. Intra-day prices on VIX and SPX on 2-Nov-2012. This figure tracks historical price data on the VIX and the SPX on an intra-day basis. Although not perfect, the strength of the co-movement is evident.




Figure II.3. VIX, VIX futures, and expectations of SPX realized volatility illustrated. This figure illustrates the interaction between the VIX, VIX futures, and expectations of SPX realized volatility. On 1-February-2012, time t , the closing price quotes, in annualized standard deviation terms, for the VIX and the 1st and second futures were $\sigma_t^{VIX} = 18.55$, $\mathcal{F}_{t,j \rightarrow j+1}^{VIX} = 19.85$, and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX} = 22.05$ respectively. Note that σ_t^{VIX} is a forecast of realized volatility of the SPX over the immediate future period, while $\mathcal{F}_{t,j \rightarrow j+1}^{VIX}$ and $\mathcal{F}_{t,j+1 \rightarrow j+2}^{VIX}$ are forecasts of the future forecasts of realized volatility of the SPX. The horizontal dashed arrows represent the period over which such expectations apply.




Figure II.4. VIX, VIX futures, and resulting VIX term structure. Figure shows the resulting VIX term structure based on the sample quotes from an earlier section. The constructed VIX term structure will be lower than the VIX futures curve when the slope of the VIX futures curve is positive. Forecasts of cumulative volatility become more accurate as the period of time over which volatility is generated grows due to cancellation of error and mean reversion dynamics. Thus a six-month long-dated VIX contract or variance swap represents the expected variance in the SPX over a six-month period, while the six-month VIX future represents the estimate today for the VIX index, or equivalently the thirty-day estimate of the variance in the SPX, six months from now. The latter naturally carries greater uncertainty, and thus a higher volatility price.




Figure II.5. Spread between six month VIX futures versus the VIX. Figure tracks the spread between six-month VIX futures (left axis) and the underlying VIX (right axis) for the entire sample period. A few observations worth noting. For one, roughly 75% of the time, the spread between the long-end future and the VIX is positive. Also, the spread went steeply negative during the height of the 2008 credit crisis and then materially negative in the latter half of 2011 when Euro zone crisis fears peaked. And three, the average spread when positive was substantially greater post-crisis, as opposed to pre-crisis.




Figure II.6. VIX futures curve, VIX term structure, and SPX implied volatility term structure. Figure is a snapshot of the various curves associated with the VIX as of the close of business on 3-November-2010, the day the second round of quantitative easing (QE2) was announced in the United States. The VIX index closed at 19.56, while the 6-month VIX future closed above 25, displaying a significant premium. As discussed, when the VIX futures curve is upward sloping, the constructed VIX term structure or variance swap curve, will generally be lower. Also depicted in the figure are the term structures for ATM and 25-delta OTM implied volatilities for SPX options. SPX options are the building blocks of the VIX index. It is common for the term structure of ATM volatilities on underlying SPX options to be lower than the VIX term structure across tenors, as the latter prices in the skew that is prevalent for OTM options on all financial assets. By construction, the VIX gives the holder exposure to the full set of OTM options on SPX at any one particular time, and thus the skew is reflected in the higher term structure. Regarding the relationship between the VIX term structure and the term structure of 25-delta OTM SPX options, we would expect the latter to be higher on average. The payoff of a variance swap such as the VIX is convex in volatility. This means that an investor who is long a variance swap will benefit from boosted gains and discounted losses, a phenomenon which is amplified when volatility skew is steep. Thus, the fair strike of a variance swap is often in line with the implied volatility of 40-delta SPX puts, which is lower than that of 25-delta puts in the presence of skew.




Figure II.7. Open interest in VIX futures versus market capitalization on VXX ETF. Figure tracks the open interest in VIX futures along with the market capitalization of the VXX ETF. Open interest averaged approximately 50,000 contracts prior to 2009. After 2009, there was a sharp rise in open interest, no doubt related to the proliferation of VIX ETF offerings. Today, ETF's represent the main vehicle for accessing the VIX, and inflows into ETF's implies inflows into the market for VIX futures.



Figure II.8. Rolling beta for 2-1 regression versus VXX ETF fund market capitalization. Figure shows that the forecast bias in the VIX term structure has increased following the 2008-9 credit crisis. The points on the time series represent the beta coefficients for rolling 1-year regressions, per equation 15. A decrease in the coefficient away from one implies a larger forecast bias. The series is then plotted against the market capitalization for the VXX ETF, the largest offering. The link between the size of the forecast bias and the size of the VIX fund is evident.




Figure II.9. Performance of a dollar invested in the SPX index versus a dollar invested in a VIX ETF. Figure shows that a dollar invested in the S&P 500 index in February 2011 would have been worth about the same in February 2012, with minimal divergence from the benchmark starting point. A dollar invested in VXX, an exchange traded fund (ETF) which takes long positions in VIX futures, would also have been worth a dollar at the end of the period but its value would have, in contrast, oscillated significantly over this time period, falling over thirty percent in the summer of 2011 and then nearly doubling in the fourth quarter of 2011 when the Euro zone sovereign debt crisis dominated headlines. The combined risk profile of these two assets would not constitute prudent hedging activity in a traditional sense. A hedge by definition should offset losses, not enhance returns or exacerbate volatility as is evident in this case.




Figure II.10. Rolling 3-month correlation and realized standard deviation for SPX and VIX. Figure shows the rolling three-month correlation of changes between the SPX and the VIX, and the realized standard deviation of each. As established, the two indices are well correlated.




Figure II.11. Ratio of 3-month realized volatility of SPX over the realized volatility of the VIX. In a two-asset portfolio made of up an underlying investment in the SPX and a hedging vehicle such as the VIX, the optimal allocation to the VIX should be proportional to the ratio of the volatilities of both elements in the portfolio. Figure shows this ratio historically, based on the rolling 3-month calculations of the volatilities.





Figure II.12. Hedge efficiency at different levels of correlation. Figure contains a plot of the efficiency under the optimal and sub-optimal hedge ratios against all possible values of the correlation coefficient. It is clear that even at perfect correlation levels, the hedge efficiency is low under a sub-optimal hedge ratio. Otherwise stated, the impact on hedge efficiency of a portfolio of two assets that exhibit a high correlation to one another but constructed at a suboptimal hedge ratio is similar to that of a portfolio of assets that exhibit a low correlation, even if this portfolio is constructed according to the optimal hedge ratio.



Tables for Chapter II

Table II.1

Collection of SPX options involved in σ_{VIX} calculation*

Table contains a collection of options on SPX that would be involved in a hypothetical calculation of the VIX as of 2-Nov-2012. In practice, option strikes are available for every five points on the SPX index, however for expositional purposes, the intervals used are fifteen and twenty-five SPX points apart for near and next-term maturities respectively. The salient information to be extracted from the table is as follows: 1) The strike of 1415 is the strike where the price difference between calls and puts for both maturities is smallest, 2) the strike of 1415 will be used to determine F and K_0 , which in turn determines the set of K_i , 3) the range of option strikes used for each maturity will vary as the calculation leaves out options for which the bid price is zero, and 4) the exact collection of options used will change in tandem with changes in the underlying price of the SPX as in-the-money (ITM) options are left out of the calculation.

Date: 2-Nov-2012			SPX range: 1414-1433			σ^{VIX} range: 16.10-17.59		
Near term expiry 17-Nov-2012			Near term expiry 19-Dec-2012					
$F_1 = 1414.5$			$F_2 = 1412.25$					
Strike	Call / Put	Mid price	Strike	Call / Put	Mid price			
1235	P	0.20	1115	P	0.50			
1250	P	0.40	1140	P	0.85			
1265	P	0.60	1165	P	1.40			
1280	P	0.80	1190	P	1.85			
1295	P	1.00	1215	P	2.25			
1325	P	1.25	1240	P	3.30			
1310	P	1.50	1265	P	4.50			
1340	P	2.25	1290	P	6.00			
1355	P	3.45	1315	P	8.50			
1370	P	5.00	1340	P	12.0			
1385	P	8.00	1365	P	16.5			
1400	P	12.0	1390	P	23.3			
1415	P	17.5	1415	P	32.8			
1415	C	17.0	1415	C	31.0			
1430	C	11.0	1440	C	17.0			
1445	C	6.00	1465	C	8.50			
1460	C	3.00	1490	C	4.00			
1475	C	1.50	1515	C	1.80			
1490	C	0.80	1540	C	0.80			
1505	C	0.50	1565	C	0.60			
1520	C	0.40	1590	C	0.40			
1535	C	0.30	1615	C	0.30			
1550	C	0.20	1640	C	0.10			

* Calculation leaves out options for which the bid price is zero

Table II.2

Descriptive statistics for VIX, VIX futures, and the resulting term structure (2006-2012)

Table contains descriptive statistics on the VIX, VIX futures, and the resulting VIX term structure spanning the period 2006-2012. The average figures suggest there is a tendency for the VIX futures curve to be upward sloping. In addition, the ranges of prices for the VIX as well as the various points along the VIX futures curve and VIX term structures convey that the volatility of the prices themselves decreases as the time to maturity increases, and this is confirmed by the standard deviation readings. The intuition for this was discussed in the VIX futures pricing section. Essentially as the tenor of the futures contract rises, the uncertainty around the convexity adjustment should fall. Skewness readings are in line with implied volatilities in other asset classes. Kurtosis figures indicate fat tails for the underlying VIX, but not for the set of futures. The short-end of the resulting VIX term structure displays fat tails.

Futures						
	Near future	Next future	3rd month	4th month	5th month	6th month
	$\mathcal{F}_{0,1 \rightarrow 2}^{VIX}$	$\mathcal{F}_{0,2 \rightarrow 3}^{VIX}$	$\mathcal{F}_{0,3 \rightarrow 4}^{VIX}$	$\mathcal{F}_{0,4 \rightarrow 5}^{VIX}$	$\mathcal{F}_{0,5 \rightarrow 6}^{VIX}$	$\mathcal{F}_{0,6 \rightarrow 7}^{VIX}$
Ave	23.7	24.2	24.4	24.5	24.6	24.6
Max	63.7	56.9	52.8	48.5	46.4	44.4
Min	10.5	11.9	12.8	13.3	13.7	14.1
SD	10.1	9.01	8.27	7.76	7.42	7.13
Skew	1.37	0.99	0.74	0.53	0.39	0.29
Kurt	2.25	1.19	0.59	0.0	-0.32	-0.49

Term structure (Constructed from futures prices)							
	VIX Index	2nd month	3rd month	4th month	5th month	6th month	7th month
	$\sigma_{0 \rightarrow 1}^{VIX}$	$\sigma_{0 \rightarrow 2}^{VIX}$	$\sigma_{0 \rightarrow 3}^{VIX}$	$\sigma_{0 \rightarrow 4}^{VIX}$	$\sigma_{0 \rightarrow 5}^{VIX}$	$\sigma_{0 \rightarrow 6}^{VIX}$	$\sigma_{0 \rightarrow 7}^{VIX}$
Ave	23.4	24.8	24.7	24.6	24.6	24.6	24.6
Max	74.3	78.7	70.6	65.7	62.1	59.3	57.1
Min	9.9	10.7	11.6	12.1	12.5	12.7	12.9
SD	11.3	11.0	10.0	9.42	9.00	8.67	8.41
Skew	1.76	1.51	1.28	1.12	0.99	0.89	0.80
Kurt	4.08	3.21	2.23	1.70	1.28	0.96	0.71

Table II.3

Evaluating the forecast biasness of the VIX term structure

Full sample: 4-January-2006 to 1-February-2012 (318 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	-0.002	-0.002*	-0.002*	-0.002*	-0.001*	0.000*
(s.e.)	(0.010)	(0.009)	(0.008)	(0.006)	(0.004)	(0.002)
β_0	0.802	0.713*	0.593*	0.446*	0.236*	0.011*
(s.e.)	(0.117)	(0.133)	(0.147)	(0.165)	(0.187)	(0.081)
R-squared	0.301	0.242	0.174	0.107	0.037	0.0002
N	318					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.802	0.713*	0.593*	0.446*	0.236*	0.011*
(s.e.)	(0.117)	(0.133)	(0.147)	(0.165)	(0.187)	(0.081)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.802*	0.713*	0.593*	0.446*	0.236	0.011
(s.e.)	(0.117)	(0.133)	(0.147)	(0.165)	(0.187)	(0.081)

[†] Tests the 7-month point on the VIX term structure versus changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table II.4

Evaluating the forecast biasness of the VIX term structure

Pre-crisis period: January-2006 to December-2007 (104 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately for the pre-crisis subperiod.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	0.002	0.002	0.001	0.003	-0.002*	0.000*
(s.e.)	(0.003)	(0.003)	(0.002)	(0.002)	(0.001)	(0.001)
β_0	0.585	0.639	0.701	0.677	0.527*	0.186*
(s.e.)	(0.226)	(0.213)	(0.179)	(0.134)	(0.110)	(0.107)
R-squared	0.089	0.116	0.152	0.171	0.139	0.052
N	104					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.585	0.639	0.701	0.677*	0.527*	0.186*
(s.e.)	(0.226)	(0.213)	(0.179)	(0.134)	(0.110)	(0.107)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.585*	0.639*	0.701*	0.677*	0.527*	0.186
(s.e.)	(0.226)	(0.213)	(0.179)	(0.134)	(0.110)	(0.107)

[†] Tests the 7-month point on the VIX term structure versus changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table II.5

Evaluating the forecast biasness of the VIX term structure

Crisis period: January-2008 to April-2009 (66 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately for the subperiod that covers the credit crisis of 2008-9.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	0.040	0.032	0.024	0.016*	0.008*	0.003*
(s.e.)	(0.036)	(0.035)	(0.031)	(0.024)	(0.015)	(0.007)
β_0	1.12	0.990	0.828	0.628*	0.353*	0.075*
(s.e.)	(0.124)	(0.127)	(0.127)	(0.124)	(0.156)	(0.087)
R-squared	0.485	0.393	0.293	0.189	0.076	0.009
N	66					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	1.12	0.990	0.828	0.628*	0.353*	0.075*
(s.e.)	(0.124)	(0.127)	(0.127)	(0.124)	(0.156)	(0.087)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	1.12*	0.990*	0.828*	0.628*	0.353*	0.075
(s.e.)	(0.124)	(0.127)	(0.127)	(0.124)	(0.156)	(0.087)

† Tests the 7-month point on the VIX term structure versus changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table II.6

Evaluating the forecast biasness of the VIX term structure

Post-crisis period: April-2009 to February-2012 (149 weekly observations)

Table contains the results for the joint hypothesis test $\alpha = 0$ and $\beta = 1$, and the individual tests for each coefficient separately for the subperiod following the credit crisis of 2008-9.

$$\left(\frac{1}{k}\right) \sum_{i=1}^{k-1} \left[(\sigma_{i,i \rightarrow i+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] = \alpha_0 + \beta_0 \left[(\sigma_{0,0 \rightarrow k+j}^{VIX})^2 - (\sigma_{0,0 \rightarrow j}^{VIX})^2 \right] + \sum_{i=1}^{k-1} u_i$$

Joint null hypothesis $\alpha = 0$ and $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
α_0	-0.018*	-0.015*	-0.015*	-0.008*	-0.004*	0.000*
(s.e.)	(0.008)	(0.008)	(0.007)	(0.006)	(0.004)	(0.002)
β_0	0.764*	0.620*	0.434*	0.221*	-0.003*	-0.160*
(s.e.)	(0.223)	(0.218)	(0.217)	(0.216)	(0.162)	(0.071)
R-squared	0.151	0.101	0.053	0.016	0.000	0.000
N	148					

Single null hypothesis $\beta = 1$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.764	0.620	0.434*	0.221*	-0.003*	-0.160*
(s.e.)	(0.223)	(0.218)	(0.217)	(0.216)	(0.162)	(0.071)

Single null hypothesis $\beta = 0$

	7-1 [†]	6-1	5-1	4-1	3-1	2-1
β_0	0.764*	0.620*	0.434*	0.221	-0.003	-0.160*
(s.e.)	(0.223)	(0.218)	(0.217)	(0.216)	(0.162)	(0.071)

† Tests the 7-month point on the VIX term structure versus changes in the VIX index (1-month)

* Indicates rejection at the 5% level

Table II.7

Regression results of weekly changes in realized arbitrage profits on changes in open interest, capital flow, safe haven performance, tail risk hedge costs, and credit risk factors

Table contains the regression results of weekly changes in realized arbitrage profits on changes in open interest, capital flow, safe haven performance, tail risk hedge costs, and credit risk factors. Arbitrage profits are estimated by selling the 2-month term structure and simultaneously buying the 7-month term structure in equal lots. The open interest variable describes the total open interest for VIX futures on the CBOE. Capital flows are estimated from total global hedge fund net asset values, as tracked by Hedge Fund Research Inc.. Safe have performance is proxied by the price of gold spot. Tail risk hedge costs are estimated by the relative cost of out-of-the-money puts on the Brazilian real versus the US dollar. Changes in swap spreads for 1-year tenor are used to represent changes in credit risk in the financial system.

Full sample: January-2006 to February-2012		
Explanatory variable	Regression coefficient	T-statistic
Open interest	0.039	1.16
Capital flow	1.693	2.22*
Safe haven performance	0.144	1.09
Tail risk hedge costs	0.016	0.38
Credit risk	-0.013	-0.55
F	5.94**	
R-squared	0.053	
N	317	
Post-crisis sample: April-2009 to 1-February-2012		
Explanatory variable	Regression coefficient	T-statistic
Open interest	0.112	1.95*
Capital flow	3.02	2.40*
Safe haven performance	-0.181	-1.61
Tail risk hedge costs	-0.050	-0.46
Credit risk	-0.0005	-0.02
F	18.1**	
R-squared	0.121	
N	147	

* Indicates significance at the 5% level based on Newey-West standard errors

** Indicates significance at the 1% level based on Newey-West standard errors

Chapter III

About the forecast bias of FX implied volatility

Implied volatility is a projection or forecast of the realized volatility that will be exhibited by asset prices in the future. Chapter II was devoted to evaluating the evolution of implied volatility itself across time, based on its term structure. This Chapter will be devoted to the second fundamental question addressed in the volatility forecasting literature. I investigate the extent to which implied volatility accurately forecasts ex-post realized volatility. In other words, are option premiums justified by the subsequent payouts? I will report the forecast biasness of foreign exchange (FX) implied volatility using an expansive dataset of over thirty currency pairs across developed and emerging economies, and evaluate the information content of the persistence of the biases present. In particular, I propose that the magnitude of forecast bias may be used as a proxy for the degree of financial integration achieved for a particular country. I motivate this concept by furnishing a simple theoretical framework based on Dornbusch, Fischer, Samuelson (1977), which is based on the premise that currencies with lower (greater) cross-border trading frictions, would be expected to have options markets that exhibit lower (higher) levels of forecast bias. The transmission mechanism from cross-border currency trading frictions to FX implied volatility forecast bias is the hedging activity of currency option market-makers, based on the Black-Scholes approach to option-pricing. I conclude by introducing a new financial integration index which offers a number of benefits to existing approaches, primarily as it relates to incorporating the most recent information into current readings³⁵.

Foreign exchange markets

Foreign exchange is a suitable asset type for a study that aims to analyze differences in the forecast accuracy of implied volatility across individual assets, in this case currencies, for two reasons: liquidity and data availability and quality.

³⁵Measures of openness or integration that are based on the annual IMF AREAER report for instance will reflect information from the previous year, as opposed to the current year.

Liquidity

Foreign exchange (FX) represents the largest and most liquid market in the world. According to the latest Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity released by the Bank of International Settlements (BIS), the daily turnover in FX markets is approximately four trillion US dollars (USD). This includes spot trades and all derivative transactions executed over-the-counter (OTC), that is, through banks and broker-dealers and not through a central exchange. As Poon and Granger (2003) point out in their review of volatility forecasting in financial markets, it is reasonable to expect different levels of forecast accuracy for options written on different assets. Option market-makers must hedge the risk in their inventory portfolios by executing transactions in spot, forward, and/or futures markets. Thus, assets with lower (greater) trading frictions, would be expected to have options markets that exhibit lower (higher) levels of forecast bias. Applying this concept across the FX arena suggests that the implied volatility of fully convertible currencies that trade free of restrictions should be less biased with regard to forecast accuracy than implied volatility for currencies that are not fully convertible or trade according to restrictions. The transmission vehicle from currency trading frictions to implied volatility forecast bias is naturally the hedging requirements of market-makers. The model section provides a theoretical framework that explains this linkage.

Data quality

Greater liquidity also implies availability of higher quality data. Testing of the forecast bias of implied volatility, which is exogenous in this literature, relies on the measurement of realized volatility using underlying asset prices. This is generally the dependent variable in the formulation, calibrated over the period the option is active. This is also a potential source of bias persistence. It is widely accepted that measurement according to the textbook approach of using a single observation per day contributes to forecast bias. Poteshman (2000) finds that a more efficient volatility estimator based intra-day five minute returns removed over half of the bias present using daily data. Blair, Poon, and Taylor (2001) report up to a four-fold increase in R-squared coefficients when going from daily to high-frequency intra-day data. I follow this guidance and use intra-day data in this analysis, which is available on a tick by tick basis through Bloomberg. Price quotes for exchange rates are captured continuously from the major broker-

dealers 24 hours a day, over weekends and all global holidays. The aggregation methodology for price quotes ensures the integrity of the data³⁶.

Implied versus realized volatility

Implied volatility is the most important driver in the pricing of financial options. It represents the market's best guess at time t about future realized volatility over the period the option is active, $t + 1$. Implied volatility is determined not by a single formula or methodology, but by supply and demand dynamics. This is analogous to the manner in which, for instance, flood insurance rates are set. Factors such as the location, design and age of the structure, the recorded history of rains, earthquakes, and tsunamis for a particular region (with greater importance given to recent experience), and the cost of credit in the financial system, are all important in determining premiums. Ultimately, however, prices paid are determined by how much the insured are willing to pay.

Figure III.1 tracks the implied volatility (at time t) versus realized volatility (at time $t + 1$) for options on the EURUSD (number of US dollars per 1 euro) exchange rate.

- Insert Figure III.1 here -

Figure III.2 tracks both volatilities for the USDCNY (number of Chinese renminbi per 1 US dollar) exchange rate.

- Insert Figure III.2 here -

For any given date in which both lines have the same value, implied volatility equals ex post realized volatility. If implied volatility is higher (lower) than future realized volatility, the forecast is biased upward (downward). It is visually evident that over the period depicted, implied volatility tracked ex post realized volatility quite closely for the EURUSD exchange rate, while consistently overshooting in the case of USDCNY.

³⁶Prices for major broker-dealers are captured simultaneously, any outliers are thrown out, and the computed average represents the single global market price. These quotes are available to all market participants. The process ensures price transparency and consistency are maintained.

Measurement of realized volatility

Standard approach

As established, implied volatility is an exogenous variable, while realized volatility is a latent variable. The standard textbook formulation for measurement of realized volatility uses one observation per day in

$$X = \sqrt{\frac{260}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2}, \quad (26)$$

where X is the annualized realized volatility, x are daily price quotes (generally taken at the close of business), and 260 is the annualization factor representing the number of business days in one year³⁷. This approach is problematic for a number of reasons. For one, currency options are quoted on a 365-day basis, thus representing a mismatch in day count. For instance, an option quoted and dealt on Friday for expiry the following Monday has three days of time value, comprised of one business day and two-weekend days. The calculation of realized volatility according to equation (26) over this time period would pick up a single observation, the change in price from close of business Friday to close of business Monday. Three days of time value versus one day of calibration represents a source of potential systematic bias. Another issue that arises due to the low granularity of daily spot data is that while developed market currencies like the euro, the British pound, and the Swiss franc trade 24 hours a day, many Emerging Market currencies trade only during local business hours. Thus, using only a single observation per day vastly under-represents actual price action. Further, once-per-day readings are typically gathered at 5pm Eastern Standard Time (EST), the close of business in New York. Many Asian currency markets for instance close much earlier in the day according to EST (at close of business in Asian time zones). The price reflected on the screen following the close of business for Asian and even European time zones, however, may have gravitated away from the actual closing price due to positioning, rebalancing of overnight order books, and supply and demand conditions in markets for non-deliverable contracts which trade globally without restriction. Using a single data point per day would not incorporate this valuable information.

³⁷See Chapter 11 in textbook “Options, Futures & Other Derivatives” by John C. Hull.

Using high-frequency data

In order to correct for deficiencies in the standard calibration method, I use intra-day exchange rate data as opposed to daily. This follows the work of Blair, Poon, and Taylor (2001), Poteshman (2000), and others who have found that increasing the frequency of data is effective in reducing forecast bias. There is a tradeoff to consider, naturally. Using minute-by-minute data for instance would capture price changes that are small in magnitude, which may not constitute actual market transactions but instead may simply reflect changes in the bid or ask prices posted by dealers. Such price changes, although recorded by the Bloomberg System, would not constitute volatility in the exchange rate. Thus, price moves that occur within the bid-ask spread are not included in the calculation as they would potentially introduce systematic bias. Additional filtering is done by using market price data captured every 30 minutes. A 30-minute time interval strikes an appropriate balance between added granularity and reduction of measurement bias. The revised standard deviation formula used in this analysis is thus given by

$$X = \sqrt{\frac{365}{T} \sum_{i=1}^n \left[\ln \left(\frac{S_i}{S_{i-1}} \right) \right]^2} \quad (27)$$

where T is the number of calendar days in observation period, S is the spot rate, and n is the number of 30-minute intervals in T .

In addition, this calibration approach has the added benefit that it more closely mimics the way market-makers manage the gamma risk in currency option portfolios³⁸. As spot exchange rates move, the writer of the option must buy and sell the underlying currency, or derivatives whose payoff is determined by the currency³⁹, according to changes in the expected payout of the option. Such rebalancing, known as delta-hedging, generally occurs multiple times a day in actual practice. This is the basic idea behind the Black-Scholes pricing option pricing model, and thus, a volatility measurement methodology that uses multiple observations per day is more consistent with this process. See Bakshi and Nikunj (2002) for an empirical review of delta-hedge practices by market-makers.

³⁸Gamma is the rate of change of delta, the approximate probability that the option will be in-the-money at expiry. The greater the gamma, the more trades the market-maker will have to execute in order to manage the risk of loss from the option.

³⁹Options on currencies that are not freely convertible are generally hedged via non-deliverable contracts, as spot is restricted.

Data

Currency spot transactions, forwards and options trade either directly between counterparties, a medium referred to as over-the-counter (OTC), or in centralized exchanges. The former comprises over 90% of the total daily turnover as reported by the latest BIS Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity. As a result, I use OTC data for both currency prices and implied volatilities in this study. The complete dataset is available from December 2007 to April 2013⁴⁰. All spot prices are quoted versus the USD, the base currency. As discussed, for realized volatility calculations, I will use spot prices captured every 30-minutes (48 observations per day, 336 observations per week, etc.). Implied volatilities are quoted in annualized percentage terms across the curve. The benchmark ‘on the run’ tenors are 1-week, 2-week, 1-month, 2-month, 3-month, 6-month, 9-month, 1-year, and then annually to 5-years for major currencies. Longer tenors are available, but trade with limited liquidity. For this study, I will use 1-month options. This tenor is the most actively traded, has the greatest liquidity and the narrowest spreads. The source of this data is Bloomberg. Table III.1 describes the set of currencies in the study.

- Insert Table III.1 here -

Table III.1.A contains descriptive statistics for monthly snapshots of implied volatility and the computed realized volatility for each currency pair. I use monthly data so that there are no overlapping observations. This helps address the potential impact of the biases that typically arise in time series analysis from using overlapping observations including inefficient estimators, serially correlated errors, artificially small standard errors⁴¹. In addition, I carry out unit root tests to confirm the volatility series are stationary AR(1) with drift. Justification for describing the evolution of volatility as a random process with drift, which itself is determined by the rate at which volatility reverts to its long-run average, is supported by Heston (1993).

- Insert Table III.1A here -

⁴⁰Spot data and implied volatility on a daily basis is available for longer, however, the higher-frequency intra-day spot data that is needed for realized volatility calculations is only available starting December 2007.

⁴¹See Harri and Brorsen (2009).

A few observations are worth noting. The differential between the average implied and realized volatility figures is narrowest for the G10 set of currencies. The differentials are, in general, much wider for currencies across LATAM, ASIA, and EEMEA, suggesting that option prices may be higher than justified by ex post realizations⁴². One notable example of such is the Argentine peso. The average 1-month implied volatility registered at 12.5%, while the average 1-month realized volatility over the period the option is active was 2.2%. Furthermore, generally speaking, the standard deviation of realized volatility is lower or on par with the standard deviation of implied volatility. Exceptions include the currencies of Argentina, Chile, India, Korea, Philippines, and Taiwan. This is illuminating as these countries have all been associated with low levels of financial integration. The fact that option prices for this subset exhibit lower standard deviation suggests there is an additional element, potentially non-random, which is embedded in prices⁴³. Lastly, the Dickey-Fuller tests reject the null hypothesis of a unit root in favor of the stationary with drift alternative. The test statistic is more negative than the critical value, and most cases are rejected at the 5% level.

It is important to note that all options included in this study trade globally, despite convertibility restrictions for many of the currencies involved. For example, currencies such as ARS, PEN, KRW, and PHP cannot be traded on a spot basis unless parties are domiciled within the borders of the country. Generally central bank approval is required to do so, along with additional administrative requirements imposed by authorities. Despite this, however, there is an active market for options written on these currencies outside the borders. These options trade without restriction on a non-deliverable basis which means they are cash-settled in dollars at expiry. Global corporations and fund managers use them for hedging and risk-taking purposes. No party takes delivery of the actual currency due to restrictions however.

Evaluating volatility forecast bias

The regression-based method generally used in the volatility forecasting literature for examining the forecast accuracy of implied volatility today on ex post realized volatilities of financial market

⁴²This will be formalized by the regression results in the next section.

⁴³This will be confirmed in the model section.

asset prices, according to the thorough review by Poon and Granger (2003), is characterized by:

$$X_{t+1} = \alpha + \beta \hat{X}_t + u_{t+1} , \quad (28)$$

where \hat{X}_t is implied volatility (the key driver of the price of the option) at time t , X_{t+1} is realized volatility over $t + 1$, the period over which the option is active, α is the intercept, and u are the i.i.d. errors. One important note is that a variation of this equation has been widely used in the international finance literature for testing the forecasting power of the currency forward premium on future exchange rate returns as reported by Hodrick (1987). Fama (1984) used the following regression

$$\Delta s_{t+1} = b_0 + b_1(f_t - s_t) + e_{t+1} , \quad (29)$$

where Δ is the backward difference operator, f_t is the forward rate, s_t is the spot rate, and e are the i.i.d. errors. If predictable excess currency returns are zero then $E_t(s_{t+1}) = f_t$ and $b_1 = 1$. This simple test represents one of the unsolved puzzles in the field, as typically b_1 are significantly less than one and even negative⁴⁴. Fama brings structure to the problem and explains the forecast bias by determining that predictable foreign exchange returns are in essence 'foreign exchange risk premiums'. In other words, investors are compensated for seeking returns outside domestic markets. The second leading explanation involves the presence of large systematic forecast errors.

Returning to equation 28, Poon and Granger establish that the prediction is unbiased only if $\alpha = 0$ and $\beta = 1$. This describes the situation where one unit of implied volatility translates into one unit of realized volatility, or equivalently, one dollar spent on currency options translates into one dollar in at-expiry payouts. The ex ante prediction undershoots ex post actual if $\alpha > 0$ and $\beta = 1$ or $\alpha = 0$ and $\beta > 1$. A dollar allocated to options translates into more than a dollar of subsequent payouts. One would not expect this situation to hold over the long run, as it would imply that the broker-dealer is running an inventory portfolio of options at a loss. The ex ante prediction overshoots ex post actual if $\alpha < 0$ and $\beta = 1$ or $\alpha = 0$ and $\beta < 1$.

⁴⁴A negative number describes the situation where an investor would buy a foreign bond that yields more than the domestic bond at time t , and the investor would be rewarded with an appreciation of the foreign currency versus the domestic at $t + 1$, in addition to the enhanced yield. This is the gist of the profitable carry trade in global financial markets.

This is the case where options pay out less than premiums paid. In the case where $\alpha > 0$ and $\beta < 1$, the most common scenario in this literature, ex ante prediction in general undershoots in low volatility environments and overshoots in high volatility environments⁴⁵. The fourth case is where $\alpha < 0$ and $\beta > 1$, which implies that in typically ex ante prediction overshoots in low volatility environments and undershoots in high volatility environments⁴⁶. Buraschi and Jackwerth (2001), Coval and Shumway (2001), Bakshi and Kapadia (2002), and Pan (2002) all report that ex ante estimates derived from implied volatility overshoot ex post realized volatility. The various explanations for the existence of the forecast bias fall under one of two categories: either the options market is inefficient⁴⁷, or, the option pricing model is incorrect. In this paper, I argue that the presence of cross-border trading frictions, which impacts the ability of market-makers to efficiently hedge the risk in option inventory holdings, translates into biased currency option prices. This will be addressed in the model section.

Regression results

Consistent with the volatility forecasting literature, I will carry out regressions, per equation 28, based on disjoint monthly observations for reasons discussed in the data section. I evaluate the implied volatility, \hat{X}_t , for 1-month options on the set of currencies registered on the last business day of each month versus the realized volatility, X_{t+1} , over the month the option is active. Individual hypothesis tests for $\alpha = 0$ and $\beta = 1$ are carried out, as well as the joint test. Rejections reported at the 5% level. In addition, the Box-Pierce Q-test statistic evaluates the null hypothesis that the residuals from the ordinary least-squares (OLS) regression are not serially correlated. The first twelve lags are considered. All results are documented in Table III.2.

- Insert Table III.2 here -

⁴⁵For example, suppose $\beta = 1.20$ and $\alpha = -0.015$. In a low volatility environment such that the observed $\hat{X} = .05$, the predicted value for ex post realized volatility would be lower. However, in a high volatility environment such that the observed $\hat{X} = .10$, the predicted value for ex post realized volatility would instead be higher.

⁴⁶For example, suppose $\beta = 0.80$ and $\alpha = 0.015$. In a low volatility environment such that the observed $\hat{X} = .05$, the predicted value for ex post realized volatility would be higher. However, in a high volatility environment such that the observed $\hat{X} = .10$, the predicted value for ex post realized volatility would instead be lower.

⁴⁷Some common explanations include the presence of trading frictions in hedging markets, liquidity frictions, and the 'peso problem'.

The prediction is unbiased only if $\alpha = 0$ and $\beta = 1$. In general, this is more prevalent for the currencies of G10 and EEMEA countries, although the latter display comparatively low R-squared coefficients. There are more rejections of the hypothesis tests for Latin American and Asian currencies, the two exceptions being the Mexican peso and the Hong Kong dollar. In addition, the direction of the deviation, whether positive or negative, has important implications. Starting with the set of G10 currencies, with the exception of JPY, these test results can be characterized by the combination, $\alpha < 0$ and $\beta > 1$. Only for AUD, GBP, and NZD, however, are the deviations significant at the 5% level. For this set of currencies, ex ante prediction generally overshoots in low volatility environments and undershoots in high volatility environments. In other words, the projected ex post realized volatility from smaller values of \hat{X}_t is impacted more by the negative α coefficient than by the positive β coefficient, while the projections from larger values of \hat{X}_t are impacted more by the positive β coefficient than by the negative α coefficient. Moreover, rejections for the currencies of LATAM are characterized by the combination, $\alpha > 0$ and $\beta < 1$, implying that currency options are generally underpriced in low volatility environments and overpriced in high volatility environments. Similarly as before, this means that projected ex post realized volatility from smaller values of \hat{X}_t is impacted more by the negative β coefficient while the projections from larger values of \hat{X}_t are impacted more by the positive α coefficient. Lastly, rejections of the hypothesis tests across the Asian region are characterized by negative or depressed β coefficients, while at the same time for all but two currencies, the α coefficient is not significantly different than zero. This is the case where option prices are consistently biased upward, or equivalently, the premiums paid over the entire sample exceeded the subsequent option payouts.

Analysis of errors

There is some evidence of serial correlation amongst the errors, as noted from the Box-Pierce Q-test statistics. The first twelve lags are considered, representing a year's worth of observations. The length was selected to ensure robustness⁴⁸, but kept to a number that is not large relative to the total number of observations in the sample⁴⁹. Recall, I have ruled out one of the primary

⁴⁸Serial correlation tested beyond the first lag gives a more complete picture as opposed to Durbin-Watson tests which only consider the first lag.

⁴⁹If the lag is too large, this would render the results useless.

sources of systematic errors by working with non-overlapping observations. Serially correlated errors, which are present in roughly a third of the currency tests, may result in depressed standard errors. In addition, the residuals themselves may contain information that can be used to improve the linear fit. A third problem is the relative size of the errors themselves, as Fama (1984) finds is a contributor to the artificially small b_1 coefficients in equation 29. Large standard errors also have the potential to render even unbiased forecasts, as suggested by the alpha and beta coefficients in (28), to be of little practical use due to amount of variance explained.

With regard to the standard errors, I proceed along two fronts. In the next section, I explore the presence of omitted variable bias which may account for some of the non-randomness of the errors from using the specification in (28) to test the forecast bias of currency option implied volatility. In addition, I will evaluate the relative size of the residual errors according to methods outlined by Blair, Poon, and Taylor (2001). The authors propose the following formulation for measuring explanatory power

$$\mathbb{P} = 1 - \frac{\sum (X_t - \tilde{X}_t)^2}{\sum (X_t - \mu_X)^2}, \quad (30)$$

where X is the dependent variable, implied volatility, and \tilde{X} is the predicted value per (28). The ratio \mathbb{P} compares the sum of the squared prediction errors with the sum of the squared deviations of X . Values closer to one imply that prediction errors are small in relation to the variability of the dependent variable. This is a desirable characteristic of a well-behaved forecasting model. In contrast, a negative value for \mathbb{P} implies forecast errors have a greater amount of variation than the actual variation in X , an undesirable condition. Table III.3 contains the results of the test that evaluates the relative size of the forecast errors.

- Insert Table III.3 here -

About half of the readings fall in the range $\mathbb{P} = [0.40, 0.60]$. This implies that the variation in the residual errors in such cases is roughly half the size of the variation in the variable being forecasted. This result does not suggest that the size of the forecast errors impacts the explanatory power of the model. There is a large subset of currencies that registered readings of less than 0.30. In such cases, the size of the forecast errors may be cause for concern, as the magnitude of the errors is approaching, although not exceeding, the magnitude of the variance in

the independent variable. There were no cases where \mathbb{P} was negative, which would unequivocally mean that the forecast errors too large. In the next section, I propose that the presence of cross-border trading frictions may be contributing to the systematic nature of errors.

Measurement of financial openness

Financial openness refers to the degree to which capital is allowed to flow freely in and out of a country's borders. The post 2008-9 crisis period has seen great proliferation in the methods used by central banks to restrict openness as authorities struggle with multiple mandates: reduce leverage in the system, promote growth, and manage the exchange rate. A decade ago this was a problem reserved for emerging economies, today restricting openness is a high priority agenda item of both emerging and developed economies. In addition, the frequency of the adoption of new measures or the removal of existing measures depending on market conditions has risen materially. The benefits of adopting anti-openness measures, however, remains difficult to ascertain as reported by Dooley (1997), Glick, Guo, and Hutchison (2006), Magud and Reinhart (2006), and Habermeier, Kokenyne, and Baba (2011). Measurement of openness itself is a challenge.

De jure, de facto, and hybrid measures of openness

De jure and *de facto* measures have been developed to characterize, quantify, compare and contrast the degree of financial openness for a particular country. *De jure* methods, proposed by Grilli and Milesi-Ferretti (1995), Quinn (1997, 2003), Mody and Murshid (2005), Glick and Hutchison (2005), Chinn and Ito (2006, 2008), and others, are based on known laws and regulatory filings. *De facto* measures, in contrast, are derived retrospectively using empirical time series data. De Gregorio (1998), Cheung et al. (2003), Lane and Milesi-Ferreti (2007) have developed *de facto* measures.

In one way or another, *de jure* methods use IMF's Annual Report on Exchange Arrangements and Exchange Restrictions (AREAER). Although this introduces a degree of objectivity, there are limitations. Private sector transactions influence and may run counter to a country's regulation efforts. Also, a material time lags exist between the adoption of new regulation and

the IMF's annual report, which suggests that *de jure* openness measures available at any given time do not reflect the latest action by authorities. Finally, *de jure* openness readings, by design, change with the introduction of new capital controls, not necessarily with the adoption of more controls of the same type. Progress is made on this front by Rodriguez and Wu (2013) who introduce a set of capital control and prudential FX measure indices which aim to capture the intensity of such policies.

Priced-based *de facto* measures are based on the rationale that true integration of capital markets should be reflected in common prices of similar financial instruments across national borders. Isolating the linkage between measurable variables and openness is a challenge, however. In addition, deviations from Uncovered Interest Parity (UIP) or International Capital Asset Pricing Model (ICAPM) may be attributed to other factors such as the Peso-Problem Hypothesis, not necessarily restrictions to the capital account. Quantity-based *de facto* measures are based on actual capital flows data, arguably a more direct approach, although such data is highly volatile, subject to measurement errors, revisions, and time lags.

Hybrid measures combine elements of both *de jure* and *de facto* methods. The KOF Index of Globalization (Dreher 2006) offers a measure of openness which combines laws and regulatory actions with economic flows of goods, capital, and services. Gwartney, Lawson, and Hall (2011) publish the Economic Freedom of the World index which measures the degree to which policies and institutions of countries are supportive of economic freedom. One particular part, or subindex, specifically measures the extent to which a wide variety of restraints affect international exchange including tariffs, quotas, hidden administrative restraints, and exchange and capital controls.

There are pros and cons to each approach, and the degree of usefulness varies by application.

Openness versus integration

Often in the capital controls literature, the terms financial openness and financial integration are used interchangeably, or the clear distinction between them is not made explicit. Financial openness is a concept which tracks the degree of ease or difficulty for inflows or outflows to cross the borders of a particular country. Many of the openness measures discussed previously associate changes in capital control policy to changes in openness. However, a change in policy,

which may result in greater openness, does not ensure success in achieving greater financial integration, one of three key central bank policy objectives that encompass the macroeconomic policy Trilemma. Aizenman, Chinn and Ito (2010) discuss integration as being an end objective which, if targeted, creates a tradeoff in achieving monetary autonomy and exchange rate stability. Financial openness, then, can be thought of as the means to the end goal, financial integration. See Le (2000) for a theoretical treatment of the difference between and interaction of both concepts.

The model

As discussed, option market-makers must hedge the risk in their inventory portfolios by executing transactions in spot, forward, and/or futures markets. In this section, I provide a framework that demonstrates that currencies with lower (greater) trading frictions, would be expected to have options markets that exhibit lower (higher) levels of forecast bias. The transmission vehicle from currency trading frictions to FX implied volatility forecast bias is the hedging requirements of market-makers. I propose that the magnitude of forecast bias may be used as a proxy for the degree of financial integration achieved for a particular country. The premise is that hedging activity is more efficient in integrated markets.

Background

The risk of a financial option to its writer (market-maker hereafter) is managed by maintaining a dynamic position in the underlying asset throughout the life of the option. This concept was introduced by Black and Scholes (1977) who derived a stochastic partial differential equation, now called the Black–Scholes equation, which governs the price of the option over time. The key idea was that the option risk could be perfectly hedged, or replicated⁵⁰, by buying and selling the underlying asset based on the expected payout of the option, or in other words, the probability of exercise at expiry. For decades now, market-makers have followed variations of this approach to manage options portfolios across asset classes. The prices of financial options (or equivalently implied volatility hereafter) are thus intimately linked to the hedging activity

⁵⁰Hedging or replicating an option are mechanically the same operation, one is simply the reverse of the other.

of market-makers.

The framework presented which addresses country financial integration in the context of implied volatility biasness involves a Dornbusch, Fischer, and Samuelson (1977) two-country model. There is a currency option market-maker in each country who must execute individual currency trades (hedge trades hereafter), in either the home or foreign country, in order to manage the risk in having sold a currency option to a corporation or end-user. The option is written on the exchange rate between the currency of the home and the currency of the foreign country. Exogenous differences between countries are described by the productivity parameters, the endowed supply of foreign exchange liquidity, and the cost of executing a hedge trade which comes from two sources (the cost of rebalancing the hedge and the cross-border trading frictions). The individual hedge trades executed over the life of the option make up a composite portfolio by way of a CES aggregator. The price index for this portfolio represents the total cost of hedging the option risk (ex post), which should be equal the premium paid at inception by the corporation (ex ante according to Black-Scholes) assuming option prices are unbiased⁵¹.

Intuition about Black-Scholes

Suppose a US-based corporation plans to buy an asset from a European entity currently valued at ten million euros. The acquisition will occur in one year, thus a rise in the value of the euro versus the US dollar implies the US-based corporation will need more US dollars to render payment in euros to the overseas entity. This is considered undesirable risk. As a result, the corporation buys an option from a market-maker which gives the corporation the right to buy ten million euro in exchange for US dollars at 1.30 (number of dollars per 1 euro), equal to the the spot exchange rate at the time the option is purchased, for a premium of five hundred thousand euros. The option is active for one year. Note the option does not obligate the corporation to buy euro at 1.30. For whatever reason if the deal does not close, there is no residual obligation to the corporation.

The market-maker, on the other hand, is at risk to deliver euros at the price of 1.30 to the corporation. If the euro has strengthened versus the dollar on the date the option expires, the

⁵¹Poon and Granger (2003) offer an thorough survey of volatility forecasting. Based on 93 peer-reviewed research articles, the authors establish that option implied volatility contains a significant amount of information about future realized volatility, often providing more accurate forecasts than model-based efforts derived from time series models.

benefit enjoyed by the corporation from exercising the option presents a risk to the market-maker. In contrast if the euro has depreciated, the market-maker would not be at risk since the option would expire worthless and the corporation can simply buy ten million euros on the open market for less than originally budgeted. The premium paid for the option is a sunk cost in this case. The at-expiry payoffs of this zero-sum arrangement can be described by $\mathcal{P}_c = -\mathcal{P}_m$, or expanding each side further

$$\text{Max}\{N(S_T - K), 0\} - c = -[-[\text{Max}\{N(S_T - K), 0\} - c]], \quad (31)$$

where \mathcal{P}_c is the net payoff in US dollars to the corporation from the purchased option inclusive of premium, \mathcal{P}_m is the net payoff in US dollars to the market-maker for having sold the option inclusive of premium, N is fixed amount of euros that may be purchased under the terms of the option, S_T is the spot exchange rate at expiry, K is the strike rate of the option, and c is the initial premium paid.

Conceptually speaking, the market-maker may employ a strategy for hedging the aforementioned risk which involves buying ten million euros at the time the option is sold to the corporation, at the rate of 1.30, the same rate as the strike rate of the option. This holding of euro will be referred to as 'the Hedge'. The market-maker, although obligated to deliver euros to the corporation at the predetermined rate of 1.30, would have secured the ten million euros at the same rate of 1.30 and thus would be insulated from losses that would be incurred from a strengthening in the euro. In this case, the market-maker keeps the five hundred thousand premium as profit. However, under the same hedging strategy, if the euro were to fall in value over the year the option is active, the market-maker would not need to deliver euros to the corporation, as the option would expire worthless. There would be no risk to the market-maker from the option. Note the Hedge was not utilized. This ten million euro purchased at inception would have fallen in value in US dollar terms and thus must be unwound at a loss. If the fall in value exceeds five hundred thousand euros, the initial premium earned for selling the option, then the market-maker would experience a net loss on the combined position (initial premium minus the cost of unwinding the Hedge). Constructing the Hedge from a single purchase of euros is thus an ineffective strategy for protecting the market-maker from losses.

Along these lines, a more suitable strategy would involve adjusting the Hedge according to

the evolution of the exchange rate over the life of the option. This implies starting with a core holding of euros which at inception would be tied to the expected probability that the option will have intrinsic value at expiry. Using the above example, because the strike rate of the option is equal to the prevailing spot rate at the time the option is sold⁵², under the assumption that exchange rates are a random walk, the market-maker would estimate there is a 50% probability that the exchange rate will be above 1.30 at expiry, and 50% that it will be below 1.30. Thus, the Hedge at inception would involve holding five million euros, half the amount of euros that must be delivered according to the terms of the option. From this initial starting point, the market-maker will increase the size of the Hedge as the euro becomes stronger versus the US dollar, and similarly will decrease the Hedge as the euro becomes weaker, as euros will not be needed if the option expires worthless. The end objective for the market-maker is to hold ten million euros if the option has value at expiry (exchange rate will be above 1.30), or hold no euros if the option expires worthless (exchange rate will be below 1.30). The process of buying more euros when the euro rises in value, and selling euros when the euro falls in value represents the cost of hedging the option risk⁵³. In the presence of transaction costs for the individual hedge trades, the cost of maintain the Hedge rises.

Production

There are two countries, the home country, $i = 1$, and the foreign, $i = 2$. Each country is endowed with supply of foreign exchange liquidity, \mathcal{L}_i , which is not mobile across country borders. Only two currencies are involved, the home currency and the foreign currency, and thus the model is based on one single exchange rate. In each country there is a continuum of hedge trades possible belonging to the unit interval indexed by $x \in [0, 1]$. The technology available to each country for carrying out the exchange of one currency versus another is described by

$$y_i(x) = z_i(x)^{-\theta} \ell_i(x), \quad (32)$$

⁵²For expositional purposes I am assuming the forward curve for the exchange rate between euros and US dollars is flat or zero across all tenors.

⁵³Note, the strategy of buying an asset following a rise and then subsequently selling it following a fall (and vice versa) generates trading losses by construction.

where $z_i(x)^{-\theta}$ represents productivity and $\ell_i(x)$ is the amount of liquidity required to carry out an individual hedge trades. With regard to productivity, $z_i(x)$ is an independent draw from an exponential distribution with parameter λ_i . In particular, if $\lambda_i > \lambda_j$, then exchanging one currency for another is a more efficient process in country i . A larger θ implies more variation in productivity, which is the same for both countries by assumption.

Since a particular hedge trade can be executed in either country, the hedge trades can instead be identified by their vector of cost draws $z = (z_1, z_2)$, and thus equation 32 can be expressed as a function of z in

$$y_i(z) = z_i^{-\theta} \ell_i(z). \quad (33)$$

The currency option market-maker must execute hedge trades in order to manage the risk in a currency option that was sold to an end-user. The option is written on the exchange rate between the currency of the home country and the currency of the foreign country. The risk of an option to its writer is managed by maintaining a dynamic position in the underlying exchange rate based on the expected payout on the option, or in other words, maintenance of the Hedge. The option market-maker may access the liquidity of the foreign exchange market of the home or the foreign country. The collection of hedge trades necessary to maintain the Hedge from inception to maturity, is described by

$$Q_i = \left[\int q_i(z)^{\frac{\eta-1}{\eta}} \varphi(z) dz \right]^{\frac{\eta}{\eta-1}}, \quad (34)$$

where η is the elasticity of substitution between any two hedge trades, $q_i(z)$ is the quantity of hedge trades used by the option trader in country i , and $\varphi(z)$ is the joint density of the cost draws across countries. The price index in country i associated with this production is given by

$$P_i = \left[\int p_i(z)^{1-\eta} \varphi(z) dz \right]^{\frac{1}{1-\eta}}. \quad (35)$$

Equation 35 represents the premium, or equivalently, the implied volatility for the option. It follows from the option trader's cost minimization that the first order condition can be expressed as

$$q_i = \left(\frac{p_i}{P_i \varphi(z)} \right)^{-\eta} Q_i, \quad (36)$$

which describes the downward sloping demand curve for hedge trades.

Executing hedge trades across country borders

Liquidity markets are competitive. The liquidity required to execute individual hedge trades incurs a cost to the market-maker that is equal to the marginal product, s_i . This is known as the hedge rebalancing cost, which specifically arises from the process of buying currencies after a rise and selling after a fall⁵⁴. In addition, there may be an additional cost, denoted by τ_{ij} , which arises when a market-maker in country i accesses the foreign exchange market in country j in the presence of trading frictions. Along these lines, I assume unity implies there are no trading frictions for executing a trade within country borders, $\tau_{11} = \tau_{22} = 1$. Thus it follows that $\tau_{12} \geq 1$ and $\tau_{21} \geq 1$.

The market-maker will execute an individual hedge trade in the country that can deliver the lowest price. Hence, the price in country i to execute any hedge trade z is $p_i(z) = \min[p_{i1}(z), p_{i2}(z)]$. When a market-maker in the home country accesses the liquidity in the home country, the cost to execute a single hedge trade x is given by $p_{11}(x) = s_1/z_1(x)^{-\theta}$. When the market-maker in the home country accesses the liquidity in the foreign country, the cost to execute a single hedge trade x is $p_{12}(x) = s_2\tau_{12}/z_2(x)^{-\theta}$. Along these lines, hedge trades will be executed in country 1 if $p_{11}(x) \leq p_{12}(x) \Leftrightarrow s_1/z_1(x)^{-\theta} \leq s_2\tau_{12}/z_2(x)^{-\theta}$. Solving this equation yields a value \bar{x}_1 such that all hedge trades $x \in [0, \bar{x}_1]$ are executed in country 1. Similarly, it can be analogously derived that a value \bar{x}_2 exists such that all hedge trades in $x \in [\bar{x}_2, 1]$ are executed in country 2.

All hedge trades along the continuum can conceivably be executed in either country. If $\tau_{12} = \tau_{21} = 1$, then $\bar{x}_1 = \bar{x}_2$, half the trades are executed in country 1 and half in country 2, implying complete specialization and full financial integration⁵⁵. However, if cross-border trading frictions exist, $\bar{x}_1 \neq \bar{x}_2$. The market-maker in country 1 will execute all hedge trades

⁵⁴This is described in the previous intuition section.

⁵⁵Financial markets are integrated when the law of one price holds. In this model, when the prices of identical hedge trades are the same in both countries, the market-makers take advantage of the combined liquidity in both countries to fulfill hedge rebalancing needs.

in $x \in [\bar{x}_1, 1]$ in country 2 and the market-maker in country 2 will execute all hedge trades in $x \in [0, \bar{x}_2]$ in country 1. Thus, hedge trades in the range $[\bar{x}_2, \bar{x}_1]$ are only executed in the country of domicile, at a greater cost as compared to the case where there are no cross-border trading frictions. The larger the cross-border trading friction, the greater the range of hedge trades that must be executed domestically. This situation implies country 1 is less integrated from the global economy, or equivalently country 2, and vice versa.

In this model, price equalization for the collection of hedge trades necessary for maintenance of the Hedge per equation 35 is a working assumption. In other words, the price of an option on a given exchange rate is the same in every country⁵⁶. Integration, thus, is defined by the condition that cross-border trading frictions force a market-maker to execute a portion of necessary hedge trades domestically, at a higher price, in the presence of trading frictions. The greater the cross-border trading frictions, the more the market-makers are forced to deal domestically, and thus the less integrated countries 1 and 2 are.

Equilibrium

Equilibrium is characterized by the trade balance condition $s_1 \mathcal{L}_1 \pi_{12} = s_2 \mathcal{L}_2 \pi_{21}$, where π_{ij} is the proportion of hedge trades executed in country j by the market-maker in country i . Equilibrium also implies that the breakdown of hedge trades for the market-maker in country 1 can be described by $\pi_{11} + \pi_{12} = 1$, and similarly for the market-maker in country 2, $\pi_{22} + \pi_{21} = 1$. Incorporating the pricing rule that the price for executing a hedge trade in country i is $p_i(z) = \min[p_{i1}(z), p_{i2}(z)]$, yields the expression $\pi_{12} = Pr\{p_{12}(z) < p_{11}(z)\}$. In other words, by the law of large numbers π_{12} can be interpreted as describing the probability that the price in country 2 offers the market-maker in country 1 a lower price as opposed to prices domestically. After some rearranging and applying some basic properties of the exponential distribution⁵⁷, the proportions can be expressed as

$$\pi_{ij} = \frac{1}{1 + \left(\frac{s_i}{s_j}\right)^{-1/\theta} \tau_{ij}^{1/\theta} \left(\frac{\lambda_i}{\lambda_j}\right)}, \quad (37)$$

⁵⁶This is consistent with pricing in over-the-counter markets, where broker-dealers set the prices of currency options through a single global network.

⁵⁷If $u \sim \exp(\mu)$, then for any $k > 0$, $ku \sim \exp(\mu/k)$, and if $u_1 \sim \exp(\mu_1)$ and $u_2 \sim \exp(\mu_2)$, then $Pr\{u_1 \leq u_2\} = \frac{\mu_1}{\mu_1 + \mu_2}$.

recalling that z_i has an exponential distribution with parameter λ_i .

Implications for prices

The model proposed by Leland (1985) is an extension to Black-Scholes which incorporates transaction costs⁵⁸. It is Leland's framework which I will specifically follow. The option hedging strategy depends on the size of transaction costs (which are represented here by p_i and determined from two components, the hedge rebalancing cost, s_i , and the cross-border friction, τ_{ij}) and the time period between hedge portfolio revision (which determines the total number of hedge trades executed throughout the life of the option, represented by ℓ_i , the amount of liquidity required for maintenance of the Hedge). The price index for the collection of hedge trades expressed by equation 35 is essentially an average of three subintervals: trades that can be executed only in country 1, trades that can be executed only in country 2, and trades that can be executed in both countries. An analytic expression for the price index can be derived by exploiting the same basic properties of the exponential distribution. This expression for the relative price indices for countries 1 and 2 is

$$\frac{P_1}{P_2} = \left(\frac{1 + \left(\frac{s_2}{s_1}\right)^{-1/\theta} \tau_{12}^{-1/\theta} \left(\frac{\lambda_2}{\lambda_1}\right)}{\tau_{21}^{-1/\theta} + \left(\frac{s_2}{s_1}\right)^{-1/\theta} \left(\frac{\lambda_2}{\lambda_1}\right)} \right)^{-\theta}. \quad (38)$$

Analysis under no cross-border trading frictions

By design, this model assumes price equalization for the composite hedge trade portfolio, or equivalently option prices, such that $P_1 = P_2$. As discussed, option prices charged to end-users are the same in country 1, as in country 2. If there are no cross-border trading frictions, denoted by $\tau_{12} = \tau_{21} = 1$, then $\pi_{12} = \pi_{21} = 0.5$, the market-maker in country 1 executes half of the hedge trades at home and half in the foreign country, and vice versa. This implies full integration, regardless of the equilibrium marginal product ratio s_1/s_2 .

⁵⁸Transaction costs invalidate the Black-Scholes argument for option replication, since continuous rebalancing of the Hedge implies infinite trading, and thus no upper limit on the total cost of strategy. Otherwise stated, the 'pure' Black-Scholes strategy holds only in the limiting case of zero transaction costs.

Analysis under symmetric cross-border trading frictions

Suppose countries 1 and 2 are symmetric in the following exogenous variables such that $\mathcal{L}_1 = \mathcal{L}_2$, $\lambda_1 = \lambda_2$, and $\theta = 0.5$. The rationale for such an assumption is that the individual currency hedge trades that must be executed in order to hedge the risk of having written the option involve a single exchange rate only. Thus a market-maker in country 1 would face the same liquidity and productivity for executing hedge trades in country 1 as the market-maker in country 2 in his or her own country. The point of interest is naturally exploring the impact of cross-border trading frictions (the market-maker in country 1 accessing country 2 and vice versa). Suppose that $\tau_{12} = \tau_{21} = 1.5$. Recall that by assumption $P_1/P_2 = 1$. As a result, marginal costs are equalized, $s_1 = s_2$, but the level of integration for countries 1 and 2 as result of cross-border trading frictions is lower. From equation 37, $\pi_{12} = \pi_{21} = 0.31$. The market-maker in country 1 executes only 31 percent of hedge trades in country 2, and the market-maker in country 2 executes only 31 percent of hedge trades in country 1. This implies approximately 38 percent of all hedge trades are executed domestically due to cross-border frictions, a departure from integration.

In addition, a key point to note is that although prices are the same across countries, the presence of cross-border trading frictions increases the overall level of prices within the borders of each country.

Analysis under asymmetric cross-border trading frictions

Under the same assumptions about symmetry as before, $\mathcal{L}_1 = \mathcal{L}_2$, $\lambda_1 = \lambda_2$, and $\theta = 0.5$, suppose that $\tau_{12} > \tau_{21} = 1.0$. This situation describes the involvement, for instance, of a market-maker in the United States (Country 1) and one in Korea (Country 2). The former incurs cross-border trading frictions, $\tau_{12} > 1.0$, in order to execute hedge trades in Korea from the United States⁵⁹. Cross-border trading frictions may arise from the introduction of capital controls, convertibility restrictions, and other anti-openness policies. However the market-maker in Korea would not face the same frictions if accessing the liquidity in the United States. In this case, $s_2 > s_1$, and $\pi_{21} < \pi_{12} < 0.50$. Similarly, the overall level of prices in each country is higher in the presence

⁵⁹Although foreigners are prohibited from buying and holding the won currency, a market-maker in the United States is able to access currency hedging vehicles that settle on a non-deliverable basis, not requiring convertibility of the currency.

of cross-border frictions. It is illuminating to note that even though one country does not have cross-border frictions, the presence of trading frictions in the other country implies a deviation from full integration.

The linkage between implied volatility forecast bias and integration

Leland (1985) offers an intuitive explanation in his conclusion which I will paraphrase for application to this context:

Inclusive of added transaction costs from cross-border trading frictions, the net price for executing a hedge trade to buy a particular currency for the market-maker is slightly higher than the price without the cross-border trading frictions. Similarly the net price for executing a hedge trade to sell a particular currency is slightly lower. This accentuation of up and down movements of the currency prices can be modeled as if the volatility of the actual currency prices was higher. This implies upward biased implied volatility prices, X_t .

Starting with the case of no cross-border trading frictions, option prices, denoted by P_i , are determined from the hedge rebalancing cost, s_i , and the trade balance condition. As discussed previously, the forecast accuracy of option prices, according to the review of volatility forecasting by Poon and Granger (2003), can be tested empirically. Letting $P_i = \hat{X}_i$, recall the regression-based method generally used for examining the forecast accuracy of implied volatility (or equivalently, option prices) today on ex post realized volatilities (or equivalently, option payouts) of financial market asset prices is characterized by:

$$X_{t+1} = \alpha + \beta \hat{X}_t + u_{t+1}, \quad (39)$$

where \hat{X}_t is implied volatility at time t , X_{t+1} is realized volatility over $t + 1$, the period over which the option is active, α is the intercept, and u are the i.i.d. errors. The prediction is unbiased only if $\alpha = 0$ and $\beta = 1$.

In the presence of cross-border trading frictions, $\tau_{12} > 1$ and $\tau_{21} \geq 1$ or $\tau_{12} \geq 1$ and $\tau_{21} > 1$, implied volatility can be decomposed as follows $\hat{X} = \hat{X}' + \kappa$, where κ represents the differential between prices with and without cross-border frictions, all else equal. Note, by assumption,

$X_1/X_2 = X'_1/X'_2$. The true model, inclusive of frictions can be expressed as

$$X_{t+1} = \alpha' + \beta' \hat{X}_t' + \gamma \kappa_t + u'_{t+1} . \quad (40)$$

I am unable to test equation 40 directly since there is no empirical representation for κ . Thus, I will test the forecast bias of implied volatility using equation 39 and correct for omitted variable bias.

There are two necessary conditions, which in this case are met, which allows for proceeding with (39) as opposed to the true model in (40). First, the cross-border trading frictions κ are a determinant of realized volatility, X . Specifically, a larger κ would be expected to be associated with a smaller X . Cross-border frictions which arise as a result of capital controls, convertibility restrictions, and other anti-openness policies work to dampen cross-border flows, and this works to diminish realized volatility in the respective currency. This is confirmed by the descriptive statistics in Table III.I. For example, the mean values for realized volatility, X , across Asia (a region well known to implement capital control and anti-openness policies), are smaller generally speaking than the mean values for realized volatilities of developed economy currencies. The second condition requires that κ and \hat{X} be correlated. In this case, it is expected that cross-border frictions and implied volatility (or option prices) are positively related. The rationale is that it should be more expensive to ensure risks related to currencies for countries which impose anti-openness regulation. This is once again confirmed by Table III.I. Implied volatility for the currencies of LATAM, for instance, are higher in general once again compared to those of the currencies of developed economies.

The direction of the omitted variable bias can be determined from both the sign of $Corr(\kappa, X)$ and the sign of γ . The combination of $Corr(\kappa, X) < 0$ and $\gamma > 0$ implies κ is positively related to the error term u' and thus β , in the testable equation 39, is biased downward. Cross-border trading frictions, $\tau_{12} > 1$ and $\tau_{21} \geq 1$ or $\tau_{12} \geq 1$ and $\tau_{21} > 1$, lead to depressed values for β . Option prices are thus biased upward, and this has been demonstrated to imply a departure from full financial integration. With regard to the error term, if the omitted variable were uncorrelated with the explanatory variables, then the impact from its omission would be captured by the intercept, α , and the error terms. This is not the case, however, as the omitted variable I argue is correlated to \hat{X} . However, the correlation may not be perfect and thus it is likely that the

intercept is also affected, along with the β coefficient. Note this rationale is consistent with the regression test outlined by equation 39 which is commonly used in the volatility forecasting literature. The forecast bias is evaluated according to the joint test, $\alpha = 0$ and $\beta = 1$, thus the impact from the omission of the variable manifests in both the regressor coefficient and the intercept term.

Financial integration index

About financial integration

Financial markets are integrated when the law of one price holds. Assets generating identical cash flows should command the same return, regardless of the domicile of the seller and of the asset holder. Integration can then be measured by comparing the returns of identical assets between one country and another. The concept is simple, however, identifying the assets for comparison is widely considered to be a difficult task.

It is established in a report prepared by the EU Commission by Adams and others (2002) that existing indicators of financial integration fall into four broad categories: 1) indicators of credit and bond market integration⁶⁰, 2) indicators of stock market integration⁶¹, 3) indicators of integration based on economic decisions of households and firms⁶², and 4) indicators of institutional differences that may induce financial market segmentation⁶³. Each approach offers advantages and disadvantages related to the set of countries covered, the amount of history available, the frequency and lag of the updates, etc. Naturally, the degree of usefulness for the various measures of integration varies by application. In addition, changes to the financial and economic landscape may render even the most widely followed and accepted measures to be temporarily useless. For instance, during the credit crisis of 2008-9, the lack of liquidity of US dollars globally on safe-haven buying caused massive deviations from covered interest rate

⁶⁰Deviations from covered interest rate parity (CIP) are the most commonly used measures of integration from this type.

⁶¹The notion of home bias is related to the degree financial integration. Home bias should disappear when financial markets are perfectly integrated. By the same token, the presence of home bias reveals lack of financial integration.

⁶²A fall in the correlation between national saving and investment signals an increase in global financial integration for instance.

⁶³This type of financial integration measurement considers the effects of different legal and institutional frameworks on the degree of financial market integration. As an example, the index of regulation of dispute resolution, proposed by Djankov et al. (2001), measures the extent to which legal procedures differ from informal dispute resolution.

parity (CIP) involving countries that are considered fully integrated by a number of measures (such as Canada, Japan, and core Euro zone nations). In fact, deviations to CIP have persisted for years in some cases.

The framework presented in this paper falls into category 1 above. The basic thesis is that the degree and efficiency of cross-border banking activity is expected to increase under greater integration. In other words, the elimination of barriers to international capital flows, together with a general relaxation of barriers to entry, can induce increases in cross-border credit flows, the share of loans extended by foreign (or out-of-state) banks, or as suggested in this paper, the ability of market-makers to access the liquidity in foreign exchange markets across borders for the purposes of hedging the risk in options inventory holdings. Several studies have used such indicators to investigate the reaction of credit markets to the lifting of regulations on segmentation such as the state-level branching restriction in the United States⁶⁴. In another study, Berger, Kashyap and Scalise (1995) document that the share of financial assets controlled by out-of-state holding companies increased from 2.1 percent in 1979 to 27.9 percent in 1994 when the US branching restrictions were removed. Petersen and Rajan (2000) report similar results, showing that US firms are choosing to borrow from increasingly distant banks.

Index construction

The degree of implied volatility forecast bias for the set of currencies spanning G10, LATAM, ASIA, and EMEA countries was evaluated from the regression results using the formulation described by equation 39. Recall the null for the joint hypothesis test of the coefficients in (39) was

$$H_0 : \alpha = 0 \text{ and } \beta = 1 . \tag{41}$$

I use this information to construct a financial integration index, FI_j , which is composed of a simple aggregation of the joint significance of the coefficients per

$$FI_j = F(r, n - k - 1) , \tag{42}$$

⁶⁴See Boyd and Gertler (1993), and Jayaratne and Strahan (1996).

where j refers to each respective currency in the study and $F(r, n - k - 1)$ is the F-statistic with $r = 2$ restrictions, one for each coefficient being tested, and $n - k - 1 = 63$ degrees of freedom for n observations and k independent variables⁶⁵. The F-test is used to determine if a group of variables are jointly significant statistically. In this context, low and insignificant F-statistics imply unbiased volatility forecasts according to the joint hypothesis test. Specifically this means the residuals under the restricted model are not significantly greater than the residuals assuming the variables are removed. This situation is associated with greater financial integration as described in the model section. However, if the residuals are a significantly larger with the aforementioned restrictions as opposed to without, the F-statistic will be larger, thus implying the fit worsened by adding the restrictions and this represents greater evidence for rejecting the null hypothesis in (41). This describes cases where implied volatility forecast biasness exists, and this is associated with lower degrees of financial integration. Table III.4 contains a ranking of FI_j readings for 2013.

- Insert Table III.4 here -

Values below 3.15 for FI_j do not lead to rejections of the null hypothesis at the 5% confidence level. This includes currencies such as the Swiss franc, Canadian dollar, and the Japanese yen, currencies of countries that are generally considered to be well integrated to the world economy according to other leading measures of openness or integration⁶⁶. This was also the case for currencies not generally associated with high degrees of openness such as PLN and ZAR. There is evidence to suggest that Poland may be a lot more integrated than previously believed. All restrictions on currency transactions were removed in 2001. In addition, Polish government passed a law in 2012 that brought all monetary regulations into compliance with European Union (EU) standards. The same law also removed all restrictions on capital flows between Poland and EU member states⁶⁷. With regard to South Africa, the trend since 1994 has been for an easing of exchange controls. According to the same source, HSBC FX Strategy, in October 2009 officials released a Medium Term Budget Policy Statement which incorporated a number of measures for the relaxing of foreign exchange controls including: 1) increased the limit on the amount

⁶⁵This index generates a reading for each country, based on the implied volatility forecast biasness of its currency versus the US dollar. Recall that all currency prices are quoted versus the US dollar. The implicit assumption is thus that the US is fully integrated to the world economy.

⁶⁶See table III.5 in the 'Validation of results' section.

⁶⁷Source: HSBC Emerging Market Currency Guide (2013), written by HSBC FX Strategy.

of capital that domestic institutions can send overseas, 2) removed the mandatory conversion of foreign currency proceeds to rand received by export companies, 3) allowed South African companies to open foreign bank accounts, 4) and increased the limits on domestic individuals' ownership of foreign investments. These legal and regulatory changes have, according to FI_j readings, contributed to greater financial integration. The next section covers three examples of alternative measures that register much different readings for both Poland and South Africa. The underlying reasoning behind this is related to the considerable timing lag between the latest updates available for the alternative measures and the present. I will expand on this point in the next section.

In contrast, higher readings were registered for the Chinese renminbi, the Philippine peso, the Malaysian ringgit, the Indian rupee, and the Argentine peso. These are the currencies of countries which have been associated with a history of engagement in anti-openness policy and regulation. One surprise reading was that of Hong Kong. The F-statistic was significant at even the 1% level, reflecting that implied volatility forecast bias exists, suggesting a departure from full integration. There are a number of plausible explanations. For one, a closer inspection of the individual t-tests for coefficients α and β reveal that deviations from 0 and 1 are not statistically significant. In other words, the individual tests statistics suggest no forecast bias, while the joint test does. The disconnect may be attributable to the low, near-zero, levels for the dependent variable X in (39). From Table III.1.A, the mean and standard deviation for HKD is 0.6% and 0.4%, respectively. This implies that the unrestricted model in the calculation of the F-statistic would likely register very low residual errors, and thus adding any restrictions would cause the errors to grow substantially on a relative basis. This in turn leads to a high, significant reading for the F-test, even though the individual tests for each coefficient are not significant. An intuitive explanation of this phenomenon involves the proliferation of trading of the Chinese renminbi in the offshore market. Hong Kong was the first, and is today the largest, trading center for deliverable renminbi outside the Mainland. In essence, Hong Kong is the hub for both the Hong Kong dollar and the deliverable renminbi (trading acronym CNH), thus, cross-border frictions that still exist within the borders of China may be manifesting in the forecast bias of option prices for HKD.

Validation of results

I validate the FI_j readings by correlating the results in Table III.4 with the latest readings for three openness or integration measures cited previously⁶⁸: the Chinn-Ito index (KAOPEN) by Chinn and Ito (2006, 2008), the KOF Index of Globalization (KOF) by Dreher (2006), and the Economic Freedom of the World index (EFW) by Gwartney, Lawson, and Hall (2011). The results are compiled in Table III.5.

- Insert Table III.5 here -

The correlation coefficients of FI_j to KAOPEN, KOF, and EFW are 56%, 49%, and 40%⁶⁹, respectively. An important detail to consider is that the latest readings available as of Q2 2013 for KAOPEN⁷⁰, KOF⁷¹, and EFW⁷² reflect data from 2011, 2010, and 2010, respectively. The time lags are considerable, ranging between two and three years despite efforts made by the authors to furnish annual publications.

The correlation analysis serves as evidence that the FI_j index is capturing, via a *de facto* approach, some of the same information with regard to financial openness and integration being captured by the *de jure* and hybrid indices. Naturally the set of countries covered by FI_j is much narrower than the set of currencies covered under other indices. With 32 currencies pairs, the country coverage for the FI_j index approaches 50 if all the countries in the Euro zone are assigned the FI_j value for the EUR. In addition, the United States would also be assigned the maximum integration value by assumption. Note, the US dollar (USD) is the base currency for all option implied volatilities.

However, a key advantage of the FI_j index is that the data for generating FX implied volatility forecast bias measures is available without time lags. This is not the case for KAOPEN,

⁶⁸Recall the point was made earlier that the terms are commonly interchangeably in the literature. Although I am of the view that these are two separate concepts. Financial openness can be thought of as the means to the end goal, financial integration. See Le (2000) for a theoretical treatment of the difference between and interaction of both concepts.

⁶⁹Correlation calculation removes the ARS data point, as the outlier reading for the F-statistic distorts the calculation due to properties of leverage. Also the figures reported are the absolute values of the respective correlation coefficients, as the sign does not matter in this case.

⁷⁰From http://web.pdx.edu/~ito/Chinn-Ito_website.htm, the Chinn-Ito Index as of 4/24/2013 is updated through 2011.

⁷¹From <http://globalization.kof.ethz.ch>, 2013 KOF Index of Globalization rankings are based on raw data for the year 2010. The Economic Globalization subindex is used for comparison.

⁷²From <http://www.freetheworld.com/2012/EFW2012-complete.pdf>, the Economic Freedom of the World: 2012 Annual Report rankings are from 2010. The Freedom to Trade Internationally subindex is used for comparison.

KOF, EFW, and other measures. This feature is important when a central bank embarks on a new anti-openness agenda, after not being very active prior to this change. The previous section cited the cases of Poland and South Africa as examples of the potential pitfalls that arise due to the time lags. Another example is Peru. Anti-openness measures announced in 2010 included⁷³: 1) Increased capital gains taxes for non-residents investments, 2) New 30% income tax for settlement of derivative contract with offshore banks, 3) Increased reserve requirements for local banks for new credit lines and foreign currency transactions, 4) New bank limits on net FX positions, 5) Cap on private pension fund overseas investments.

From Table III.5, the openness readings for Peru for FI_j , KAOPEN, KOF, EFW are 34.1, 2.4390, 73.27, 8.60, respectively. For the country mix in this study, these readings correspond to percentiles 32%, 100%, 55%, 93%, respectively. Note 100% implies fully open borders. While the FI_j and KOF indices responded to the active central bank activity that took place in 2010 based on readings that reflect a material deviation from financial openness and integration, KAOPEN and EFW register readings that suggest otherwise. This again may be attributable to the sizable timing lags in acquiring the data necessary for proper updates of the respective measures.

Robustness tests

The robustness tests involve an assessment of the alternative explanations for the forecast bias of FX implied volatility. Three plausible drivers that can be tested are size of the market, liquidity, and the existence of dual-markets.

The size of the market will be approximated by two variables, the approximate daily turnover for the spot market for each currency and the daily turnover for forwards, or non-deliverable forwards (NDF's) for currencies that are not convertible, for each currency. A market-maker will use a combination of spot, forwards, and NDF's to hedge the risk in options portfolios. Liquidity will be represented by the bid-ask spread for a typical FX option transaction (institutional, not retail). The greater the liquidity, the smaller the bid-ask spread. Existence of dual-markets is a binary variable which takes on a value of 1 if forwards and options trade on a non-deliverable basis outside a respective country's borders due to convertibility restrictions of the currency

⁷³Not an exhaustive list. Source HSBC FX Strategy.

(thereby implying the existence of dual-markets), or a value of 0 if forwards and options trade on a deliverable basis anywhere in the world. There are three sources for this data: the latest BIS Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity, HSBC's Emerging Markets Currency Guide 2013, and the HSBC FX Trading desk in New York (the source for information on bid-ask spreads for options on developed economy currencies).

Because the robustness data is available only on an annual basis, it is not possible to incorporate it into the tests of implied volatility forecast bias, as the frequency of the data used for the regressions carried out is monthly⁷⁴. Nonetheless as a robustness check, I will regress the FI_j values for 2013 with each of variables described above. Table III.6 shows the data and results of this analysis.

- Insert Table III.6 here -

High R-squared coefficients would be cause for concern as this would imply that FI_j index values could be easily replicated by any of the robustness variables including size of the market, liquidity, and the existence of dual-markets. This is not the case. R-squared coefficients are 4% for the regression on both of the size variables and 3% for the regression on the liquidity variable. The R-squared for the regression on the existence of dual-markets variable is material at 33%, however, the lack of granularity makes this variable inadequate for a measurement of financial openness or integration. In other words, the degree of integration would only take on two values.

Conclusion

Chapter III makes two contributions literature of volatility forecasting.

As established, measurement of realized volatility has been documented as a potential source of implied volatility forecast bias persistence by Blair, Poon, and Taylor (2001), Poteshman (2000), and others. It is well documented that using daily data is problematic, on the other hand, using high-frequency intra-day data is more computationally expensive, and introduces other potential sources of bias such as those that arise from the aggregation of small returns

⁷⁴Residuals should be orthogonal to the robustness variables, however this cannot be tested using the formulation in equation 39.

that occur within bid-ask spread levels⁷⁵. I propose an improvement to the calibration approach for realized volatility for exchange rates that uses 30-minute returns, with additional filtering that ignores insignificant moves, and demonstrate that implied volatilities on a set of liquid, convertible, and freely traded currencies are free of forecast bias (as compared to ex post realized volatilities) over a multi-year period (2007-2013).

Second, I propose that the magnitude of forecast bias may be used as a proxy for the degree of financial integration achieved for a particular country. I motivate this concept by furnishing a simple theoretical framework based on Dornbusch, Fischer, Samuelson (1977), which is based on the premise that currencies with lower (greater) cross-border trading frictions, would be expected to have options markets that exhibit lower (higher) levels of forecast bias. The transmission mechanism from cross-border currency trading frictions to FX implied volatility forecast bias is the hedging activity of currency option market-makers, based on the Black-Scholes approach to option-pricing. Based on this, I construct a financial integration index, FI_j , from a simple aggregation of the joint significance of the α and β coefficients in the formulation widely used in the literature according to Poon and Granger (2003). Values below 3.15 for $FI_j = F(r, n - k - 1)$ do not lead to rejections of the null hypothesis at the 5% confidence level. This includes currencies such as the Swiss franc, Canadian dollar, and the Japanese yen, currencies of countries that are generally considered to be well integrated to the world economy according to other leading measures of openness or integration. In contrast, higher readings were registered for the Chinese renminbi, the Philippine peso, the Malaysian ringgit, the Indian rupee, and the Argentine peso. These are the currencies of countries which have been associated with a history of engagement in anti-openness policy and regulation, resulting in a departure from full integration. I cite a few instances, Poland, South Africa, and Peru, where FI_j readings disagree with previous work, and argue that significant time lags in a number of the alternative approaches, namely the Chinn-Ito index (KAOPEN) by Chinn and Ito (2006, 2008), the KOF Index of Globalization (KOF) by Dreher (2006), and the Economic Freedom of the World index (EFW) by Gwartney, Lawson, and Hall (2011), are the source of the disagreement, a phenomenon which favors the index I propose.

This framework has a number of benefits. For one, it does not restrict measurement of openness or integration to the adoption or removal of capital controls. This is important as central

⁷⁵This problem was addressed in this paper.

banks have more tools at their disposal to achieve objectives including prudential FX measures and monetary policy, today increasingly unconventional. In addition, the data for generating FX implied volatility forecast bias measures is available without time lags. In contrast, the latest IMF AREAER report reflects information from the previous year, as opposed to the current year. Thus, openness measures that rely on this data will be lagged by at least this period of time, likely longer. Finally, the framework involving FX implied volatility can be used to furnish estimates of integration today, as well as future changes to the degree of integration. This is an attractive dimension, as the impact of central bank policy changes, or announcement of policy changes, is generally not immediate or easily quantifiable. This framework aggregates all available information and generates one single reading. Finally, the availability of higher frequency readings is also possible.

Figures for Chapter III

Figure III.1. Implied versus realized volatility for EURUSD. This figure tracks the implied volatility (at time t) versus realized volatility (at time $t + 1$) for options on the the EURUSD (number of US dollars per 1 euro) exchange rate.





Figure III.2. Implied versus realized volatility for USDCNY. This figure tracks the implied volatility (at time t) versus realized volatility (at time $t + 1$) for options on the the USDCNY (number of renminbi per 1 US dollar) exchange rate.



Tables for Chapter III

Table III.1

Set of currencies

Table contains the set of currencies in the study for the period December 2007 to April 2013. All currency quotes are versus the US dollar (USD).

Subgroup	Currency	Trading Acronym
G10*	Australian dollar	AUD
	Canadian dollar	CAD
	Swiss franc	CHF
	Euro	EUR
	Great Britain pound	GBP
	Japanese yen	JPY
	Norwegian kroner	NOK
	New Zealand dollar	NZD
	Swedish kroner	SEK
LATAM**	Argentine peso	ARS
	Brazilian real	BRL
	Chilean peso	CLP
	Colombian peso	COP
	Mexican peso	MXN
	Peruvian sol	PEN
ASIA	Chinese renminbi	CNY
	Hong Kong dollar	HKD
	Indonesian rupiah	IDR
	Indian rupee	INR
	Korean won	KRW
	Malaysian ringgit	MYR
	Philippine peso	PHP
	Singapore dollar	SGD
	Thai baht	THB
Taiwan dollar	TWD	
EEMEA***	Czech koruna	CZK
	Hungarian forint	HUF
	Israeli shekel	ILS
	Polish zloty	PLN
	Russian ruble	RUB
	Turkish lira	TRY
	South African rand	ZAR

* G10 - Group of 10 industrialized nations

** LATAM- Latin America

*** EEMEA - Eastern Europe, Middle East, Africa

Table III.1.A

Descriptive statistics for implied and realized volatility

Table contains descriptive statistics for implied and realized volatility for each currency pair, as well as unit root tests to confirm the volatility series are stationary AR(1) with drift.

<i>N=64</i>	Implied Volatility			Realized Volatility		
	Mean (SD) %	Max - Min %	Dickey-Fuller	Mean (SD) %	Max - Min %	Dickey-Fuller
AUD	14.4 (6.0)	41.6 - 6.5	-2.46**	15.6 (11.3)	75.0 - 2.8	-2.87**
CAD	11.5 (4.6)	27.1 - 5.1	-1.95**	11.5 (6.3)	31.9 - 2.3	-2.05**
CHF	12.2 (3.3)	23.5 - 6.9	-2.74**	12.5 (5.5)	28.2 - 3.0	-2.68**
EUR	12.3 (3.9)	27.5 - 7.1	-2.56**	11.6 (5.3)	28.4 - 3.4	-2.24**
GBP	11.1 (4.4)	28.2 - 5.1	-2.00**	10.7 (6.1)	30.5 - 2.2	-1.83**
JPY	12.3 (4.1)	32.3 - 7.0	-2.98**	11.6 (6.1)	38.1 - 2.9	-3.63**
NOK	14.9 (4.6)	31.0 - 7.9	-2.22**	15.2 (7.2)	39.0 - 4.8	-2.17**
NZD	15.2 (5.4)	37.3 - 7.5	-2.55**	16.3 (9.4)	61.1 - 5.1	-2.75**
SEK	15.2 (4.7)	30.4 - 9.1	-2.22**	15.7 (7.4)	36.7 - 5.9	-2.12**
ARS	12.5 (12.4)	71.7 - 5.0	-3.11**	2.2 (1.2)	6.6 - 0.7	-4.78**
BRL	16.0 (7.2)	44.1 - 6.4	-2.33**	12.2 (9.3)	53.8 - 1.5	-3.33**
CLP	14.4 (5.7)	34.4 - 5.4	-2.61**	7.6 (4.4)	21.8 - 0.8	-3.33**
COP	16.0 (6.6)	35.2 - 6.6	-2.36**	8.3 (5.3)	24.4 - 1.4	-3.86**
MXN	14.0 (7.1)	39.8 - 5.7	-2.41**	12.4 (10.2)	71.9 - 2.4	-4.25**
PEN	8.7 (4.3)	21.2 - 4.0	-1.40*	4.8 (4.9)	25.6 - 0.4	-4.77**
CNY	2.6 (1.2)	6.0 - 0.8	-3.33**	1.5 (1.0)	6.1 - 0.2	-4.04**
HKD	0.8 (0.2)	1.5 - 0.4	-3.53**	0.6 (0.4)	1.8 - 0.1	-3.54**
IDR	12.8 (9.4)	50.0 - 5.2	-2.32**	9.0 (11.3)	66.4 - 0.7	-3.30**
INR	10.9 (4.3)	29.7 - 5.4	-3.02**	4.2 (2.0)	10.2 - 1.3	-3.97**
KRW	15.4 (11.4)	63.2 - 5.2	-2.25**	11.6 (5.3)	28.4 - 3.4	-3.03**
MYR	8.5 (2.3)	15.0 - 4.9	-3.58**	4.0 (1.9)	10.3 - 0.7	-4.40**
PHP	9.5 (4.3)	26.8 - 4.3	-2.36**	4.4 (2.5)	12.5 - 0.4	-5.23**
SGD	7.2 (2.4)	14.0 - 3.4	-2.86**	5.3 (3.0)	12.9 - 0.7	-2.63**
THB	8.1 (2.8)	18.0 - 4.6	-4.40**	4.6 (2.9)	16.0 - 0.0	-5.00**
TWD	6.5 (2.0)	12.3 - 3.5	-3.33**	2.6 (1.4)	5.7 - 0.0	-4.08**
CZK	16.5 (6.4)	37.9 - 9.9	-2.04**	16.2 (8.0)	50.1 - 5.5	-3.18**
HUF	19.7 (6.6)	41.4 - 12.3	-2.59**	21.0 (8.8)	53.9 - 8.6	-2.76**
ILS	10.3 (3.2)	19.5 - 6.6	-1.76**	8.4 (4.3)	20.0 - 2.6	-2.96**
PLN	19.3 (7.7)	46.3 - 11.0	-2.28**	19.0 (10.1)	47.5 - 6.7	-2.56**
RUB	12.1 (4.8)	27.1 - 5.8	-2.94**	8.2 (6.8)	55.4 - 2.4	-5.55**
TRY	13.6 (5.7)	38.7 - 5.2	-2.87**	9.8 (6.6)	41.5 - 1.3	-3.07**
ZAR	18.9 (6.7)	55.3 - 12.0	-3.65**	17.9 (10.2)	65.5 - 6.3	-3.88**

Table III.2

Tests of forecast bias of FX implied volatility

Implied volatility, \hat{X}_t , for 1-month currency options is evaluated versus the realized volatility, X_{t+1} , over the month the option is active. Individual hypothesis tests for $\alpha = 0$ and $\beta = 1$ are carried out, as well as the joint test. Serial correlation test results also reported.

$X_{t+1} = \alpha + \beta\hat{X}_t + u_{t+1}$						
$N=65$		α	β	F-statistic for joint test	R-squared	Box-Pierce Q-statistic (12 lags)
G10	AUD	-0.050*	1.422*	4.36 \circ	0.571	10.9
	CAD	-0.019	1.155	1.36	0.694	25.8†
	CHF	-0.015	1.135	0.40	0.445	28.3†
	EUR	-0.011	1.025	1.91	0.568	25.5†
	GBP	-0.030*	1.225*	4.28 \circ	0.769	10.0
	JPY	-0.007	0.996	0.94	0.450	9.0
	NOK	-0.030	1.219	1.66	0.609	18.2
	NZD	-0.043	1.348*	4.01 \circ	0.604	13.8
	SEK	-0.030	1.225	1.98	0.609	23.8†
LATAM	ARS	0.020*	0.018*	618 \circ	0.037	53.3†
	BRL	-0.036*	0.982	13.5 \circ	0.585	8.40
	CLP	-0.001	0.535*	169 \circ	0.478	28.6†
	COP	0.002	0.497*	131 \circ	0.380	15.2
	MXN	0.022	0.729	2.53	0.257	7.72
	PEN	0.004	0.495*	34.1 \circ	0.193	24.4†
ASIA	CNY	0.004	0.401*	79.4 \circ	0.234	37.0†
	HKD	-0.002	0.976	19.4 \circ	0.392	35.6†
	IDR	-0.004	0.725*	9.34 \circ	0.383	20.8†
	INR	0.011*	0.293*	772 \circ	0.419	8.60
	KRW	0.004	0.524*	63.6 \circ	0.482	13.8
	MYR	0.000	0.470*	318 \circ	0.337	17.7
	PHP	0.013*	0.318*	309 \circ	0.333	11.1
	SGD	-0.014	0.922	28.9 \circ	0.522	12.3
	THB	0.022*	0.299*	62.7 \circ	0.078	19.0
	TWD	0.003	0.367*	375 \circ	0.285	15.6
EEMEA	CZK	0.042	0.719*	2.56	0.329	28.8†
	HUF	0.045	0.824	1.63	0.381	27.7†
	ILS	-0.016	0.961	1.47	0.505	27.1†
	PLN	0.012	0.918	0.36	0.519	19.2
	RUB	-0.019	0.828	17.9 \circ	0.345	10.3
	TRY	-0.012	0.806	22.8 \circ	0.488	11.3
	ZAR	0.002	0.931	0.74	0.381	8.36

* denotes significance at 5% level for individual tests, \circ 5% level for joint test

† denotes evidence of serially correlated errors at 5% level

Table III.3**Tests of the size of the forecast errors per Blair, Poon, and Taylor (2001)**

Table III.3 tests for the relative size of the forecast errors. About half of the readings fall in the range $\mathbb{P} = [0.40, 0.60]$, a result which does not suggest that the size of the forecast errors impacts the explanatory power of the model. There is a large subset of currencies that registered readings of less than 0.30. Such forecast errors may be cause for concern, as the magnitude of the errors is approaching, although not exceeding, the magnitude of the variance in the independent variable. There were no cases where \mathbb{P} was negative, which would unequivocally mean that the forecast errors too large.

		$\mathbb{P} = 1 - \frac{\sum (X_t - \bar{X}_t)^2}{\sum (X_t - \mu_X)^2}$
<i>N=65</i>		\mathbb{P}
G10	AUD	0.571
	CAD	0.694
	CHF	0.445
	EUR	0.568
	GBP	0.769
	JPY	0.450
	NOK	0.608
	NZD	0.604
	SEK	0.609
LATAM	ARS	0.037
	BRL	0.585
	CLP	0.478
	COP	0.380
	MXN	0.257
	PEN	0.193
ASIA	CNY	0.234
	HKD	0.392
	IDR	0.383
	INR	0.419
	KRW	0.482
	MYR	0.337
	PHP	0.333
	SGD	0.522
	THB	0.078
	TWD	0.285
EEMEA	CZK	0.329
	HUF	0.381
	ILS	0.505
	PLN	0.519
	RUB	0.345
	TRY	0.488
	ZAR	0.381

Table III.4**Financial integration readings based on FX implied volatility forecast bias (2013)**

Table III.4 contains a ranking of FI_j readings for all currencies in the study. In this context, low and insignificant F-statistics imply unbiased implied volatility forecasts according to the joint hypothesis test. This is associated with greater financial integration. Along these lines, larger values for the F-statistics represent greater evidence in favor of a rejection of the null. This is the case where implied volatility forecast biasness exists, and this is associated with lower degrees of financial integration.

$$FI_j = F(r, n - k - 1)$$

Currency	FT_j
CHF	0.4
PLN	0.4
ZAR	0.7
JPY	0.9
CAD	1.4
ILS	1.5
HUF	1.6
NOK	1.7
EUR	1.9
SEK	2.0
MXN	2.5
CZK	2.6
NZD	4.0
GBP	4.3
AUD	4.4
IDR	9.3
BRL	13.5
RUB	17.9
HKD	19.4
TRY	22.8
SGD	28.9
PEN	34.1
THB	62.7
KRW	63.6
CNY	79.4
COP	131
CLP	169
PHP	309
MYR	318
TWD	375
INR	772
ARS	6180

$k = 2$
 $n - k - 1 = 63$

Table III.5

Side-by-side comparison of openness readings

Table III.5 contains a side-by-side comparison of the latest readings for the FI_j index for all currencies in the study versus the latest readings for the Chinn-Ito index (KAOPEN) by Chinn and Ito (2006, 2008), the KOF Index of Globalization (KOF) by Dreher (2006), and the Economic Freedom of the World index (EFW) by Gwartney, Lawson, and Hall (2011).

	FI_j^*	$KAOPEN^\dagger$	KOF^\ddagger	EFW^\ddagger
	Correlation ^o	56%	49%	40%
CHF	0.4	2.439	78.4	7.14
PLN	0.4	0.064	72.8	7.27
ZAR	0.7	-1.169	64.6	7.16
JPY	0.9	2.439	44.0	7.16
CAD	1.4	2.439	75.8	7.50
ILS	1.5	2.439	78.8	7.95
HUF	1.6	2.439	89.6	7.92
NOK	1.7	2.439	73.9	6.98
EUR	1.9	2.439	71.6	7.65
SEK	2.0	2.439	88.7	7.88
MXN	2.5	1.120	59.9	6.74
CZK	2.6	2.439	85.1	7.77
NZD	4.0	2.439	80.6	8.45
GBP	4.3	2.439	78.0	8.48
AUD	4.4	1.120	76.4	7.28
IDR	9.3	-0.113	58.5	6.74
BRL	13.5	-0.113	52.4	7.13
RUB	17.9	0.678	55.6	6.08
HKD	19.4	2.439	97.6	9.02
TRY	22.8	0.064	56.1	7.33
SGD	28.9	2.439	97.6	9.05
PEN	34.1	2.439	73.3	8.60
THB	62.7	-1.169	68.4	6.63
KRW	63.6	0.942	59.6	7.21
CNY	79.4	-1.169	51.1	6.57
COP	131	-0.113	52.8	6.82
CLP	169	1.384	82.5	8.20
PHP	309	-1.169	53.3	6.69
MYR	318	-1.169	78.8	7.25
TWD	375	-1.169	51.1	7.55
INR	772	-1.169	42.7	6.28
ARS	6180	-0.808	42.6	6.24

* *de facto index (latest 2013)*, † *de jure index (last updated 2011)*

‡ *hybrid indices (contain raw data from 2010)*

o *Correlation calculation ignores the ARS data point as a way to address leverage*

Table III.6

Robustness tests for FI_j index

Table III.6 contains FI_j for 2013 values along with the latest snapshot of the robustness variables: size of market, liquidity, existence of dual markets. Reported are the R-squared coefficients from regressions of the 2013 values of FI_j on each of the robustness variables.

	FI_j	Size of market I (spot in \$bill)	Size of market II (fwd / NDF \$ bill)	Liquidity (bid-ask in %)	Existence of dual markets (1=Yes, 0=No)
	R-squared	4%	4%	3%	33%
CHF	3.390	60.0	20.0	0.10	0
CAD	3.015	4.0	4.5	0.50	0
SEK	2.602	3.0	8.0	1.00	0
GBP	2.584	210.0	68.0	0.10	0
HKD	2.276	67.3	21.8	0.10	0
ILS	2.144	2.0	2.5	0.40	0
EUR	2.064	2.5	4.0	0.70	0
NOK	1.842	9.0	2.5	0.20	0
AUD	1.200	400.0	132.0	0.10	0
PLN	1.162	17.0	5.0	0.20	0
SGD	1.132	15.0	8.0	0.40	0
NZD	1.019	1.6	6.0	0.50	0
CZK	0.711	23.0	7.5	0.20	0
JPY	0.519	130.0	43.0	0.10	0
IDR	-0.113	92.0	30.0	0.20	0
ZAR	-0.632	0.5	0.7	1.00	1
CNY	-0.670	7.0	3.0	0.40	1
TRY	-0.727	40.0	12.0	1.00	0
HUF	-1.317	4.0	4.0	0.10	0
PEN	-1.484	8.5	1.3	0.50	0
RUB	-1.519	5.0	5.0	0.50	0
MXN	-1.722	1.6	0.2	2.50	1
BRL	-1.751	0.9	0.7	1.25	1
COP	-2.073	10.0	5.0	0.50	1
CLP	-2.216	30.0	3.0	0.20	1
TWD	-2.801	1.5	0.5	1.00	1
MYR	-3.025	2.3	1.0	1.00	1
KRW	-3.266	1.0	1.0	1.00	1
ARS	-4.858	1.7	1.8	1.50	1
PHP	-4.878	1.5	3.5	0.50	1
INR	-5.760	4.0	4.5	0.40	1
THB	-6.490	0.3	0.1	6.50	1

Source: BIS Triennial Central Bank Survey of Foreign Exchange and Derivatives Market Activity, HSBC's Emerging Markets Currency Guide 2013, HSBC FX Trading desk in New York

References

- [1] Adam, K., Jappelli, T., Menichini, A., Padula, M., Pagano, M. (2002). Analyse, compare, and apply alternative indicators and monitoring methodologies to measure the evolution of capital market integration in the European Union. Centre for Studies in Economics and Finance (CSEF) and the Department of Economics and Statistics University of Salerno.
- [2] Ait-Sahalia, Y., Karaman, M., Mancini, L. (2012). The term structure of variance swaps, risk premia and the expectations hypothesis. Working paper. Princeton University, NBER, Swiss Finance Institute, University of Zurich, and EPFL.
- [3] Aizenman, J., Chinn, M., Ito, H. (2010). Assessing the emerging global financial architecture: Measuring the Trilemma's configurations over time. NBER Working Paper No. 14533.
- [4] Bakshi, G., Kapadia, N. (2002). Delta hedged gains and the negative market volatility risk premium. *Review of Financial Studies* 16, 2, 527-566.
- [5] Baillie, R., Myers, R. (1989). Modeling commodity price distributions and estimating the optimal futures hedge. Unpublished manuscript. Department of Economics, Michigan State University.
- [6] Becker, R., & Clements, A. (2008). Are combination forecasts of S&P 500 volatility statistically superior? *International Journal of Forecasting*, 24, 122- 133.
- [7] Becker, R., Clements, A. & White, S. (2006). Does implied volatility provide any information beyond that captured in model-based volatility forecasts? *Journal of Banking and Finance*.
- [8] Benzoni, L. (2001). Pricing options under stochastic volatility: An empirical investigation. Working Paper. Carlson School of Management, Minneapolis.
- [9] Berger, Allen N., Anil K. Kashyap, Scalise, J. (1995). The transformation of the U.S. banking industry: What a long, strange trip it's been. *Brookings Papers on Economic Activity* 2, 55-218.
- [10] Black, F., Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* 81 (3): 637-654.
- [11] Blair, B., Poon, S., Taylor, S. (2001). Forecasting S&P 100 volatility. The incremental information content of implied volatilities and high frequency index returns. *Journal of Econometrics* 105, pp. 5-26.
- [12] Bollerslev, T., Tauchen, G., Zhou, H. (2009). Expected stock returns and variance risk premia. *Review of Financial Studies* 22, 4463.
- [13] Boyd, John H., Gertler, M. (1993). U.S. commercial banking: Trends, cycles, and policy. *NBER Macroeconomics Annual*, 319-368.
- [14] Breeden, D., Litzenberger, R. (1978). Prices of state-contingent claims implicit in options prices. *Journal of Business* 51, 621-651.
- [15] Britten-Jones, M., Neuberger, A. (2000). Option prices, implied price processes, and stochastic volatility. *Journal of Finance* 55, 839-866.
- [16] Brunnermeier, M., Pedersen, L. (2009). Market liquidity and funding liquidity. *Review of Financial Studies* 22, 2201-2238.

- [17] Buraschi, A., Jackwerth, J. (2001). The price of a smile: Hedging and spanning in option markets. *Review of Financial Studies* 14, 2, 495-527.
- [18] Campa, J., Chang, K. (1995). Testing the expectations hypothesis on the term structure of volatilities in foreign exchange options. *Journal of Finance* 50, No. 2.
- [19] Campbell, J.Y., Lo, A.W. & MacKinlay, A.G. (1997). *The Econometrics of Financial Markets*. Princeton University Press. Princeton NJ.
- [20] Carr, P., Lee, R. (2009). Volatility derivatives. *Annual Review of Financial Economics* 1, 319-339.
- [21] Carr, P., Wu, L. (2006). A tale of two indices. *Journal of Derivatives* 13, 13-29.
- [22] Chernov, M. (2000). Implied volatility as forecasts of future volatility, the market risk premia, and returns variability. Working Paper. Columbia Business School.
- [23] Cheung, Y. W., Chinn, M. D., & Fujii, E. (2003). Chinese economies in global context: the integration process and its determinants. *Journal of the Japanese and International Economies*.
- [24] Chinn, M., Ito, H. (2006). What matters for financial development? Capital controls, institutions, and interactions. *Journal of Development Economics* 81, 163-192.
- [25] Chinn, M., Ito, H. (2008). A new measure of financial openness. *Journal of Comparative Policy Analysis* 10, 309-322.
- [26] Coval, J., Shumway, T. (2001). Expected option returns. *Journal of Finance* 56, 3, 983-1009.
- [27] Daigler, R., Rossi, L. (2006). A portfolio of stocks and volatility. *Journal of Investing*, Summer, 99-106.
- [28] Delisle, J., Doran, J., Krieger, K. (2010). Volatility as an asset class: Holding VIX in a portfolio. Unpublished manuscript.
- [29] De Gregorio, J. (1998). Financial integration, financial development and economic growth. Chile: Department of Industrial Engineering, Universidad de Chile, unpublished manuscript.
- [30] Demeterfi, K., Derman, E., Kamal, M. (1999). A guide to volatility and variance swaps. *Journal of Derivatives* 6, 9-32.
- [31] Dooley, M. (1997). A survey of the literature on controls over international capital transactions. *IMF Staff Papers* 43, 639-687.
- [32] Dornbusch, R., Fischer, S. and Samuelson, P.A. (1977). Comparative advantage, trade and payments in a ricardian model with a continuum of goods. *American Economic Review*, 67, 823-839.
- [33] Dreher, A. (2006). Does globalization affect growth? Evidence from a new Index of Globalization. *Applied Economics* 38, 10: 1091-1110.
- [34] Dreschsler, I., Yaron, A. (2011). What's vol got to do with it? *Review of Financial Studies* 24, 1.
- [35] Duffie, D., Singleton, K. (1997). An econometric model of the term structure of interest rate swap yields. *Journal of Finance* 52, 1287-1321.

- [36] Dupire, B. (2006). Model free results on volatility derivatives. Working paper. Bloomberg, NY: SAMSI.
- [37] Fama, E. (1984). The information in the term structure. *Journal of Financial Economics* 13, 509-528.
- [38] Fama, E. (1984). Forward and spot exchange rates. *Journal of Monetary Economics* 14, 319-338.
- [39] Figlewski, S. (1986). Hedging with financial futures for institutional investors. Ballinger Publishing. Cambridge, MA.
- [40] Fleming, J. (1998). The quality of market volatility forecasts implied by S&P 100 index option prices. *Journal of Empirical Finance* 5, 317-345.
- [41] Froot, K. (1989). New hope for the Expectations Hypothesis of the term structure of interest rates. *Journal of Finance*. Vol. XLIV, No. 2.
- [42] Glick, R., Hutchison, M. (2005). Capital controls and exchange rate instability in developing economies. *Journal of International Money and Finance* 24, 387-412.
- [43] Glick, R., Guo, X., Hutchison, M. (2006). Currency crisis, capital-account liberalization, and selection bias. *Review of Economics and Statistics* 88, 698-714.
- [44] Gouriéroux C. & Jasiak J. (2001). *Financial Econometrics*. Princeton University Press. Princeton NJ.
- [45] Grilli, V., Milesi-Ferretti, G. (1995). Economic effects and structural determinants of capital controls. *Staff Papers - International Monetary Fund* Vol. 42, No. 3, pp. 517-551.
- [46] Gromb, D., Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics* 66, 361-407.
- [47] Gwartney, J., Lawson, R., Hall, J. (2011). *Economic freedom of the world: 2011 Annual report*. Fraser Institute.
- [48] Habermeier, K., Kokenyne, A., Baba, C. (2011). The effectiveness of capital controls and prudential policies in managing large inflows. *IMF Staff Discussion Note* 11/14.
- [49] Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, Volume 6, number 2, pp. 327-343.
- [50] Hodrick, R. J. (1987). *The empirical evidence on the efficiency of forward and futures foreign exchange markets*. Hardwood Academic Publishers. Chur, Switzerland.
- [51] Huang, Y., Zhang, J. (2010). The relation among SPX options, variance futures, and VIX futures.
- [52] Hull, J., White, A. (1987). The pricing of options on assets with stochastic volatility. *Journal of Finance* 42, 281-300.
- [53] Jayaratne, Jith, Strahan, P. (1996). The finance-growth nexus: Evidence from bank branch deregulation. *Quarterly Journal of Economics* 111 (4), August, 639-670.
- [54] Jiang, G., Tian, Y. (2007). Extracting model-free volatility from option prices: An examination of the VIX index. *Journal of Derivatives* 14(3), 1-26.

- [55] Konstantinidi, E., Skiadopoulos, G. (2011). Are VIX futures prices predicable? An empirical investigation. *International Journal of Forecasting*, forthcoming.
- [56] Konstantinidi, E., Skiadopoulos, G., Tzagkaraki, E. (2008). Can the evolution of implied volatility be forecasted? Evidence from European and US implied volatility indices. *Journal of Banking and Finance* 32, 2401-2411.
- [57] Lamoureux, C.G., Lastrapes, W. (1993). Forecasting stock return variance: Understanding stochastic implied volatility. *Review of Financial Studies*, 6, 293– 326.
- [58] Lane, P. R. and Milesi-Ferretti, G. (2006). The external wealth of nations mark II: Revised and extended estimates of foreign assets and liabilities, 1970-2004. IMF Working Paper 06/69.
- [59] Le, H. (2000). Financial openness and integration. Asia Pacific School of Economics and Management working papers. Development issues, 00/4.
- [60] Leland, H. (1985). Option pricing and replication with transactions costs. *The Journal of Finance*, Vol. 40, No. 5 (Dec., 1985), 1283-1301.
- [61] Lin, Y. (2007). Pricing VIX futures: Evidence from integrated physical and risk-neutral probability measures. *Journal of Futures Markets* 27, 1175-1217.
- [62] Lu, Z., Zhu, Y. (2010). Volatility components: The term structure dynamics of VIX futures. *Journal of Futures Markets* 30, No. 3, 230-256.
- [63] Lypny, G. (1988). Hedging foreign exchange risk with currency futures: Portfolio effects. *Journal of Futures Markets* 8, 703-715.
- [64] Magud, N., Reinhart, C. (2006). Capital controls: An evaluation. NBER Working Paper Series, no. 11973.
- [65] Matthias F., Longstaff, F., and Lustig H. (2013). The TIPS-Treasury Bond Puzzle. *Journal of Finance*, Forthcoming. DOI: 10.1111/jofi.12032.
- [66] Mixon, S. (2007). The implied volatility term structure of stock index options. *Journal of Empirical Finance* 14, No. 3. 333-354.
- [67] Mody, A., Antu, P. M., (2005). Growing up with capital flows. *Journal of International Economics*, Vol. 65, No. 1 (January), 249–66.
- [68] Newey, W., West, K. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review* 28, 777-787.
- [69] Nossman, M., Wilhelmsson, A. (2008). Is the VIX futures market able to predict the VIX Index? A test of the expectations hypothesis. *Journal of Alternative Investments*.
- [70] Pan, J. (2002). The jump risk premia implicit in options: Evidence from an integrated time-series study. *Journal of Financial Economics* 63, 1, 3-50.
- [71] Petersen, Mitch, Rajan, R. (2000). Does distance still matter? The information revolution in small business lending. NBER Working Paper no. 7685, May, Cambridge MA.
- [72] Poon, S., Granger, C. (2003). Forecasting volatility in financial markets: A review. *Journal of Economic Literature*. Vol XLI, 478-539.
- [73] Poteshman, A. (2000). Forecasting future volatility from option prices. University of Illinois Urbana-Champaign, Working Paper.

- [74] Poteshman, A. (2001). Underreaction, overreaction, and increasing misreaction to information in the options market. *Journal of Finance* 56, 851-876.
- [75] Rodriguez, M., Wu, T. (2013) The effect of capital controls and prudential FX measures on options-implied exchange rate stability. Working paper. Federal Reserve Bank of San Francisco, University of California Santa Cruz.
- [76] Quinn, D. (1997). The correlates of change in international financial regulation. *American Political Science Review* 91(3): 531-551.
- [77] Quinn, D. (2003). Capital account liberalization and financial globalization, 1890 – 1999: A synoptic view. *International Journal of Finance and Economics*, Vol. 8, No. 3, p. 189-204.
- [78] Sepp, A. (2008b). Pricing options on realized variance in Heston model with jumps in returns and volatility. *Journal of Computational Finance* 11(4), 33-70.
- [79] Shu, J., Zhang, J. (2012). Causality in the VIX futures markets. *Journal of Futures Markets* 32, No. 1, 24-46.
- [80] Stein, J. (1989). Overreactions in the options markets. *Journal of Finance* 44, 1011-1023.
- [81] Szado, E. (2009). VIX futures and options - A case study of portfolio diversification during the 2008 financial crisis. Working paper, University of Massachusetts, Amherst.
- [82] Taylor, S. J. (2005). *Asset Price Dynamics, Volatility, and Prediction*. Princeton University Press. Princeton NJ.
- [83] Vasilellis, G., Nigal M. (1996). Forecasting volatility for portfolio selection. *Journal of Business Finance and Accounting*, 23, 1, pp. 125–143.
- [84] Zhang, J., Huang, Y. (2010). The CBOE S&P 500 three-month variance futures. *Journal of Futures Markets* 30(1), 48-70.
- [85] Zhang, J., Zhu, Y. (2006). VIX futures. *Journal of Futures Markets* 26, 521-531.
- [86] Zhu, Y., Zhang, J. (2007). Variance term structure and VIX futures pricing. *International Journal of Theoretical and Applied Finance* 10(1), 111-127.