

Using Equity, Index and Commodity Options to Obtain Forward-Looking Measures of Equity and Commodity Betas, and Idiosyncratic Variance

Ehud I. Ronn

Department of Finance
University of Texas at Austin

July 2012

Revised: August 5, 2020

Acknowledgments: The author acknowledges with thanks the research assistance provided by Bomi Lee, Zack Liu, Xiang Kang, Natesh Arunachalam and Raghav Sarathy, while remaining solely responsible for any errors therein. The author acknowledges the helpful comments and suggestions of seminar participants at the University of Texas at Austin, Nanyang Business School, the University of York, University of Melbourne, the University of Auckland, the University of Canterbury, Otago University, the University of Wellington, Macquarie University, Oklahoma State University, Tsinghua University, Nankai University, Beihang University, Xiamen University and the University of Copenhagen,. Previous drafts of this paper were presented at the 2012 Energy Risk USA Conference, 2012 Global Derivatives USA conference, the 2012 Financial Research Association's "Early Ideas" Session, the 2013 ESCP Europe Business School Conference on Commodities and Financial Modeling, the Norwegian University of Science and Technology's 2013 Energy Economics and Finance Seminar, the 2014 EWGCFM Meeting & RCEM 2nd International Conference, the 2014 53rd EWGCFM Meeting & 2nd International Conference of the RCEM, the 2014 Energy Finance 2014 Conference, the 2015 32nd International Conference of the French Finance Association, the 2015 25th Pacific Basin Finance, Economics, Accounting and Management Conference, the 2015 10th Energy and Finance Conference, the 2015 26th Conference on Financial Economics and Accounting, the 2016 Center for Commodity Market Research Commodity Markets Conference, the 2016 Global Derivatives Trading & Risk Management Conference, the 2016 Wolfgang Pauli Institute (WPI) Conference on the Mathematics of Energy Markets, the 2016 SIAM Conference on Financial Mathematics and Engineering, the 2017 OptionMetrics Research Conference, the 2018 Auckland Centre for Financial Research Derivative Markets Conference and the 2019 Chicago Quantitative Alliance.

Author: Ehud I. Ronn

Address: Department of Finance
McCombs School of Business
University of Texas at Austin
2100 Speedway B6600
Austin, TX. 78712-1276

Tel.: (512) 471-5853
FAX: (512) 471-5073

E-mail: eronn@mail.utexas.edu

Abstract

This paper presents a parsimonious and theoretically-sound basis for extracting forward-looking measures of equity and commodity betas, and idiosyncratic variance.

Defining forward-looking betas and idiosyncratic variance as perturbations of historical estimates, we use the market prices of equity and index options under a single-factor market model to compute forward-looking term structures of equity betas and idiosyncratic variance. When applying the model to the market prices of options on oil company stocks and a market index, we are able to discern the market's perceptions regarding these oil companies' prospective beta, and hence signaling their future sensitivity to market changes. In turn, the prospective fraction of idiosyncratic variance relative to total variance provides a forward-looking market measure for onset of crises, when idiosyncratic risk fades relative to systematic, and complementing the information conveyed by VIX and the CBOE's equity implied correlation.

The model bears a natural extension to the joint use of options on equities and the futures price of oil. In so doing, we are able to discern a forward-looking oil beta. Such a beta, in conjunction with risk-neutral futures prices, gives rise to a CAPM-based forecast of oil prices.

Key Words: Implied volatilities, implied correlations and implied market betas

JEL Classification: G12 – Asset Pricing; G13 – Contingent Pricing

1 Introduction

The search for forward-looking indicators is a natural one in finance, as one of the primary roles of the discipline is to utilize market information to discern participants' views and expectations. This paper attempts to apply that principle to the analysis of implied volatility, implied correlation and implied betas, and their impact on investment analysis and practice.

Among the earlier to recognize the key role of markets as conveyors of information this was Nobel Laureate Friedrich A. von Hayek, who wrote (1945):

“Fundamentally, in a system in which the knowledge of the relevant facts is dispersed among many people, . . . [w]e must look at the price system as a mechanism for communicating information if we want to understand its real function.”

In a similar vein, Brealey, Myers and Allen (2006) write in their famous textbook:

“If [financial markets are] efficient, prices impound all available information. Therefore, if we can only learn to read the entrails,¹ security prices can tell us a lot about the future.”

That said, it should be duly noted, the notion of informationally-efficient markets does not imply omniscience: Thus, Paul A. Samuelson's (1966) famous quip:

“The stock market has forecast nine of the last five recessions.”

The key thrust of this paper is to document the “Message from Markets” as evidenced in the equity, index and commodity markets: Specifically, we seek to identify the “Message” that can be elicited from the concurrent observation of these markets and the derivatives written on these primary markets.

Equity implied volatility dates back to Latané and Rendleman (1976).² Implied correlations have been more challenging, originally requiring the simultaneous pric-

¹By way of explanation, Wikipedia provides an interpretation for their obscure reference: “In the religion of Ancient Rome, the practice of prophecy called for the inspection of the entrails [intestines] of sacrificed animals” to forecast the future.

²Some of the earlier literature referred to this volatility with the adjective “implicit.”

ing of individual *and* (relatively illiquid) spread options so as to permit isolation of the two vols and the implied correlation using the Margrabe (1978) formula. More recently, CBOE (2009) derived, and reports daily, an implied correlation index ρ_{Average} requiring only the implied vols of individual stocks and the market index

$$\rho_{\text{Average}} = \frac{\sigma_{\text{Index}}^2 - \sum_i w_i^2 \sigma_i^2}{2 \sum_i \sum_{j>i} w_i w_j \sigma_i \sigma_j}.$$

Implied equity betas are even more recent, and with several requiring alternative modifications on the “hybrid”-model use of option and historical data. Thus, French, Groth and Kolari (1983) use option-implied volatility with historical correlations. Siegel (1995) uses a hypothetical Margrabe-style exchange option to price “implicit betas.” Assuming the skewness of the idiosyncratic shock is zero, Chang, Christoffersen, Jacobs and Vainberg (2012) use option-implied volatility and skewness measures from out-of-the-money equity and index options to derive forward-looking betas. In turn, Buss and Vilkov (2012) use forward-looking information from option prices to estimate option-implied correlations, construct option-implied predictors of factor betas and find a monotonically increasing risk-return relation. Fouque and Kolman (2011) used a continuous-time CAPM with stochastic volatilities and forward-looking betas based on second and third risk neutral moments obtained from call option prices. D’Ecclesia et al. (2014) considered time-varying correlation between crude oil prices and stock markets (with structural breaks), with reported correlations being higher after 2004. Broadie, Chernov and Johannes (2007) use an affine jump-diffusion model to estimate risk premia using S&P futures options.

In contrast, our approach is a distinctly different hybrid model: We use historical betas as well as historical idiosyncratic variances as inputs in order to obtain the *perturbations* or *adjustment factors* (to these historical measures) implied by observed equity and index option prices. It is important to note that, in so doing, the information forwarded by option prices is the information *incremental* to what is observed in the historical estimates. By using at-the-money (ATM) options, our approach obviates the need to compute skewness, an issue that can be especially acute in shorter-dated options where the vol skew is known to be more pronounced. Moreover, by utilizing the entire term structure of volatilities (for each stock and

the S&P), we are able to compute a term structure of betas and idiosyncratic variances out to the most-distant option expiration date. The inclusion of time-to-maturity terms is designed to lessen any bias that would otherwise permeate the intercept terms.

Finally, it is instructive to compare and contrast the oil-company equity betas reported here with those obtained by Chang, Christoffersen, Jacobs and Vainberg (CCJV) (2012) and Buss and Vilkov (2012). CCJV generally report far-lower betas (0.33 to 0.80), but that may well be due to the earlier 1996 – 2004 time period covered in their sample: As is well-known, crude-oil prices exhibited a far-lower correlation with the S&P 500 prior to the Great Recession. Whether the heightened correlation exhibited more recently is due to the disparate perspectives of “financialization of the energy industry” or to “integrated capital markets” remains an issue to be resolved. Buss and Vilkov (2012), in turn, report a marginally higher implied beta relative to the historical beta. One of the findings of the current paper is that the relationship between historical and implied is quite sensitive to the time period being analyzed — specifically, whether we are in “crisis” or calm mode.

There are three sets of questions we seek to address in this paper:

1. Using information in the Prices of Equity, Index and Commodity Options, we seek to extract forward-Looking equity betas. This permits us to address two questions. The first is, What is the prospective equity beta? The interest in such a beta is long-standing, having in part to do with the ability of the CAPM to constitute an empirically-consistent asset pricing model.

A second question of some interest pertains to this: Is there a forward-looking *term structure* of equity betas? The answer to this question could have important implications for capital budgeting as well as investments.

2. Financial economists have long had an interest in the dichotomy between systematic and idiosyncratic variances. Our interest in this paper focuses on the change in emphasis in the advent of a financial crisis: Thus, since systematic risk dominates during crises, can we discern when prospective idiosyncratic variance is decreasing?
3. The third set of questions pertains to the matter of forward-looking commodity

(crude-oil) betas. Here, the first issue is whether we can discern, in the sign of forward-looking betas, indicators of supply- or demand-side crises?

And the second question here pertains to forecasts of crude oil prices. While there is an abundance of prognosticators, we consider a *financial-economics* approach to forecasting spot prices simultaneously using several forward-looking measures, including implied beta, Sharpe Ratio and implied volatility, to infer a CAPM-based forecast of oil prices.

Inter alia, this paper also addresses the phenomenon oft-cited as the “financialization of commodity [or oil] markets” — e.g., Hamilton and Wu (2012), Henderson, Pearson and Wang (2012), Singleton (2012), and Basak and Pavlova (2013). Whereas those and other papers take a pejorative view, in part reflective of the term “financialization,” that same phenomenon is addressed in this paper as “integrated financial markets,” an event both natural and unexceptional. It is noteworthy the terms authors use frequently reflect the writers’ malignant or benign perspective of the event.

The paper now proceeds as follows. Considering the format of standard equity options, Section 2 demonstrates the empirical methodology by presenting a one- and two-factor model for equity options and its corresponding econometric specification. Section 3 details the data requisite for the testing of these two models, whereas Section 4 proceeds to report the corresponding empirical results. Section 5 uses an analogous one-factor model to relate oil futures contracts to the S&P 500 Index equity market, producing forward-looking measures that permit explicit testing of one-year-ahead forecasts of the spot price of oil. Section 6 concludes.

2 Using Equity Options to Obtain Forward-Looking Equity Betas and Forward-Looking Idiosyncratic Variance

2.1 The One-Factor Model

Consider the one-factor market-model equation,

$$R_i = a_i + \beta_{i, \text{SPX}} R_{\text{SPX}} + e_i \tag{1}$$

where we assume $\text{Corr}(R_{\text{SPX}}, e_i) = 0$. Applying the variance operator to (1) yields

$$\Sigma_i^2 = \beta_{i,\text{SPX}}^2 \sigma_m^2 + \sigma_i^2, \quad (2)$$

where

$\Sigma_i^2 \equiv \text{Var}(R_i)$, the variance of the return on stock i

$\sigma_m^2 \equiv \text{Var}(R_{\text{SPX}})$, the variance of the return on the S&P 500 market index

$\sigma_i^2 \equiv \text{Var}(e_i)$, the idiosyncratic variance

Eq. (2) applies to *historical* $\{\widehat{\Sigma}_i, \widehat{\beta}_{i,\text{SPX}}, \widehat{\sigma}_i\}$ data with $\widehat{\beta}_{i,\text{SPX}} \equiv \text{Cov}(R_i, R_{\text{SPX}}) / \text{Var}(R_{\text{SPX}})$ over some specified time interval (such as a 60-day moving window).³

Relationship (2) also holds *prospectively*, that is, to implied vols $\{\Sigma_i, \sigma_i\}$ extracted from option prices on the individual equities i and the market portfolio $m \equiv \text{SPX}$. Now consider a specification that explicitly models the relationship between historical estimates $\{\widehat{\beta}_{i,\text{SPX}}, \widehat{\sigma}_i\}$ and forward-looking ones $\{\beta_{i,\text{SPX}}, \sigma_i\}$. In theory, the difference between historical and *ex-ante* statistics arises from two sources:

1. The information set. The historical returns, variances and covariances are due to a specific realization of uncertainty. That information set is (part of) investors' perceptions of the future, but investors may and presumably do consider other sources of information in forming expectations of the future.
2. A risk premium, aka the "market price of volatility risk." Technically, implied vols are risk-neutral expectations of future realized vols, but as is well-known, there can be a non-zero market price of volatility risk that separates the statistical from the risk-neutral expectations. While the literature is not unanimous, most researchers have found risk-neutral implied vols exceed their statistical-expectations counterparts: See Jackwerth and Rubinstein (1996), Pan (2002), Bakshi and Kapadia (2003a, 2003b), Low and Zhang (2005), Doran and Ronn (2008) and Bollerslev, Gibson and Zhou (2011). Since both the LHS and RHS of eq. (2) pertain to risk-neutral expectations, the forward-looking betas we obtain here are risk-neutral.

³Throughout this paper, variables with a "carat" $\widehat{}$ denote historical estimates.

The econometric model we posit is one which utilizes *cross-sectional daily tests* while allowing for a term structure of betas: In this test we employ a cross-sectional analysis at a given point in time, taking into explicit consideration the entire term structure of betas. Since equity implied vols are provided on a daily basis out to 24 mos. maturities, we fill in the missing observations by assuming forward vols are constant between observable expiration dates.⁴

The hybrid portion of the model links the historical estimates $\{\hat{\beta}_i, \hat{\sigma}_i\}$ to their forward-looking analogues $\{\beta_{iT}, \sigma_{iT}\}$ via a linear additive (3) correction: At any date t ,

$$\begin{cases} \beta_{itT} &= \hat{\beta}_{it} + \alpha_{1t} + \alpha_{2t}T \\ \sigma_{itT} &= \hat{\sigma}_{it} + \alpha_{3t} + \alpha_{4t}T \end{cases} \quad (3)$$

where the coefficients $\{\alpha_{1t}, \alpha_{2t}, \alpha_{3t}, \alpha_{4t}\}$ are stock- i independent. Accordingly, the non-linear optimization performed to estimate the four coefficients is:

$$\min_{\{\alpha_{1t}, \alpha_{2t}, \alpha_{3t}, \alpha_{4t}\}} \sum_i \sum_T \left[\Sigma_{iT}^2 - (\hat{\beta}_{it} + \alpha_{1t} + \alpha_{2t}T)^2 \sigma_{mT}^2 - (\hat{\sigma}_{it} + \alpha_{3t} + \alpha_{4t}T)^2 \right]^2 \quad (4)$$

For each stock i on date t we have (at most) $T - 4 = 20$ d.f. Although we naturally retain the i -dependence of the historical estimates $\{\hat{\beta}_{i,\text{SPX}}, \hat{\sigma}_i\}$, to increase the number of degrees of freedom, we make the simplifying assumptions the estimated coefficients $\{\alpha_{1t}, \alpha_{2t}, \alpha_{3t}, \alpha_{4t}\}$ are all i -independent: This increases the number of degrees of freedom by the number of stocks $N = 8$ to $24N - 4 = 188$.

The coefficients $\{\alpha_{1t}, \alpha_{2t}, \alpha_{3t}, \alpha_{4t}\}$ should be interpreted as the information option prices provide *above and beyond* the historical information incorporated into

⁴Using the principle variance is additive across maturities — since S&P returns are uncorrelated across time — for both individual stocks and market index, the algorithm is the following. Assume we observe implied vols for monthly maturities 1, 2, 3 and 6 mos.

For the maturity $m = 4$, we solve for σ_4 using the observable σ_3 and σ_6 :

$$(4/12)\sigma_4^2 = (3/12)\sigma_3^2 + [\sigma_6^2(6/12) - \sigma_3^2(3/12)]/3.$$

We repeat this for successive maturities out to 24 mos. for the S&P, and 18 mos. for those stocks whose actively-traded options do not extend to 24 mos.

the historical estimates of beta $\widehat{\beta}_{it}$ and idiosyncratic risk $\widehat{\sigma}_{it}$. There are at least two reasons we might be interested whether an analysis of the type eq. (9) hybrid model we are considering here gives rise to meaning/interesting results:

1. Obviously, across the entire market, the weighted average beta for any maturity T is by definition 1.0. However, that does not imply a unit beta for any specific sub-section of stocks. Thus, the question is whether the beta time-to-maturity coefficient α_{4t} is non-zero for the oil equities under consideration here.
2. By virtue of the forward-looking nature of implied vols, the prices of options provide a contemporaneous “Message from Markets.” In this case, we can examine a measure of forward-looking *idiosyncratic* risk, with respect to whatever information these contain *above and beyond* the comparable historical estimate.

2.2 The Two-Factor Model: Joint Equity and Crude-Oil Betas

As a link to the latter part of this paper, which focuses on oil futures contracts, this two-factor model enriches the previous by explicitly incorporating a dependence on the price of oil, where the latter is represented by the array of crude-oil futures prices out to a two-year maturity. Again suppressing the time subscript t , the two-factor model is:

$$R_i = a_{iT} + b_{imT}R_m + b_{i,CLT}R_{CLT} + e_T, \quad (5)$$

where

b_{imT} = Partial regression coefficient for the stock i on m , where m is the S&P 500 Index

$b_{i,CLT}$ = Partial regression coefficient for stock i on CL_T

We now clarify the dependence of these parameters on the time-to-maturity dimension T , we note that for each and every stock i , we can apply the variance

operator for each maturity T to obtain:⁵

$$\Sigma_{iT}^2 = b_{imT}^2 \sigma_{mT}^2 + 2 b_{imT} b_{iT, CL} \rho_{m, CL, T} \sigma_{mT} \sigma_{CL, T} + b_{i, CL, T}^2 \sigma_{CL, T}^2 + \sigma_{iT}^2 \quad (6)$$

where

b_{imT} = First partial correlation coefficient for stock i for maturity T

$b_{i, CL, T}$ = Second partial correlation coefficient for stock i for maturity T

Σ_{iT} = Implied vol on stock i for maturity T

σ_{mT} = Implied vol on stock market for maturity T

$\rho_{m, CL, T}$ = Corr between market and CL for maturity T

σ_{iT} = Idiosyncratic vol for stock i maturity T

In order to be able empirically to address the issue of forward-looking crude-oil betas, we multiply and divide the second expression on the RHS by σ_{mT}^2 :

$$\begin{aligned} \Sigma_{iT}^2 &= b_{imT}^2 \sigma_{mT}^2 + 2 b_{imT} b_{iT, CL} \frac{\rho_{m, CL, T} \sigma_{CL} \sigma_{mT}}{\sigma_{mT}^2} \sigma_{mT}^2 + b_{i, CL, T}^2 \sigma_{CL, T}^2 + \sigma_{iT}^2 \\ &\equiv b_{imT}^2 \sigma_{mT}^2 + 2 b_{imT} b_{iT, CL} \beta_{CL, T} \sigma_{mT}^2 + b_{i, CL, T}^2 \sigma_{CL, T}^2 + \sigma_{iT}^2 \\ &= (b_{imT}^2 + 2 b_{imT} b_{iT, CL} \beta_{CL, T}) \sigma_{mT}^2 + b_{i, CL, T}^2 \sigma_{CL, T}^2 + \sigma_{iT}^2 \end{aligned} \quad (7)$$

where $\beta_{CL, T}$ is the standard beta of the T -maturity CL contract on the SPX, $\beta_{CL, T} \equiv \text{Cov}(\text{CL}_T, \text{SPX}) / \sigma_{mT}$.

While there is much we can analyze here, our primary focus of analysis in the two-factor model will be the forward-looking values of $\beta_{CL, T}$ and σ_{iT} . In so doing, we will seek to determine the degree to which forward-looking oil betas are different from their historical counterparts, and the degree to which forward-looking idiosyncratic risk differs from its historical. As before, the hybrid portion of the model links the historical estimates $\{\hat{b}_{imT}, \hat{b}_{iT, CL}, \hat{\beta}_{CL, T}, \hat{\sigma}_{iT}\}$ to their forward-looking analogues $\{b_{imT}, b_{iT, CL}, \beta_{CL, T}, \sigma_{mT}\}$ via a linear additive (3) correction: Thus, at any

⁵In the empirical testing of this model, we will utilize values of T from 1/12 to 2 yrs.

date t ,

$$\begin{cases} b_{t,imT} &= \hat{b}_{t,im,T} + \alpha_{5t} \\ b_{t,iT,CL} &= \hat{b}_{i,T,CL} + \alpha_{6t} \\ \beta_{t,CL,T} &= \hat{\beta}_{t,CL,T} + \alpha_{7t} \\ \sigma_{itT} &= \hat{\sigma}_{t,iT} + \alpha_{8t} \end{cases} \quad (8)$$

where the coefficients $\{\alpha_{5t}, \alpha_{6t}, \alpha_{7t}, \alpha_{8t}\}$ are stock- i independent. In sharp contrast to (3), it is important to note (8) no longer uses an affine form for the maturity date T . The rationale is parsimony in the numbers of parameters to be estimated: Since all the parameters estimated in (8) include dependency on T , it is no longer “necessary” to include an explicit affine term for T .

Accordingly, the non-linear optimization performed to estimate the seven coefficients $\boldsymbol{\alpha} \equiv \{\alpha_{5t}, \alpha_{6t}, \alpha_{7t}, \alpha_{8t}\}$ is:

$$\min_{\{\boldsymbol{\alpha}\}} \sum_i \sum_T \left[\Sigma_{iT}^2 - b_{imT} (b_{imT} + 2 b_{iT,CL} \beta_{CL,T}) \sigma_{mT}^2 - b_{i,CL,T}^2 \sigma_{CL,T}^2 - \sigma_{iT}^2 \right]^2. \quad (9)$$

3 Data for the One- and Two-Factor Models

The empirical results reported here cover the period Oct. 1, 2007 – June 30, 2017. Covering the period from the pre-Great Recession S&P stock-market high (Oct. 9, 2007), the data include the Great Recession and its recovery, the period of price increase associated with the “Arab Spring” 12/1/10 – 4/1/11,⁶ and the precipitous decline in oil prices in 2014. $N = 8$ stocks are utilized. They are the stocks included in Bloomberg’s BUSOILP Index: Anadarko Petroleum Corp., Apache Corp., ConocoPhillips Co., Chevron Corp., Hess Corp., Occidental Petroleum Corp., Marathon

⁶The choice of the “terminal date” April 1, 2011 was driven by a strong financial-market indicator that these political events were no longer of first-order importance to the oil markets. This will be more fully described below.

Oil Corp. and Exxon Mobil Corp.⁷

We need specify the empirical proxies for the three sets of historical variables $\{\widehat{\Sigma}_{it}, \widehat{\beta}_{i, \text{SPX}, t}, \widehat{\sigma}_{it}\}$ as well as the nine ATM implied vols on the S&P 500 and the individual stocks i , $\{\sigma_{mt}, \Sigma_{it}\}$:

1. The historical estimates $\{\widehat{\Sigma}_{it}, \widehat{\beta}_{i, \text{SPX}, t}, \widehat{\sigma}_{it}\}$ on each stock i are obtained from 60-day rolling regressions of eq. (1) culminating on date t .
2. The nine ATM implied vols for maturities $T \leq 2$ $\{\sigma_{mtT}, \Sigma_{itT}\}$ are obtained from ATM options on the S&P 500 Index and the eight individual BUSOILP stocks as observed on date t .

4 Empirical Results for the One- and Two-Factor Models

Deferring the results for oil futures section to Section 5, the initial empirical results pertain to the time-series of the estimated parameters for betas, α_{1t} and α_{2t} , and the corresponding parameters for idiosyncratic volatility α_{3t} and α_{4t} .

4.1 Results — α_{1t} and α_{2t}

The time-series of α_{1t} addresses the issue of whether the option-beta is materially different from the historical beta. In turn, the estimated coefficients of α_{2t} speaks to the issue of whether there is, at time t , a statistically-significant upward- or downward-sloping term-structure of Betas in oil stocks.⁸

⁷The Anadarko Petroleum Corp., a member of the BUSOILP eight, owned a 25% interest in Deepwater Horizon, the oil rig operated by BP PLC in the Gulf of Mexico. As a consequence of the catastrophic sinking of that oil rig on April 20, 2010 Anadarko stock underperformed the BUSOILP index by as much as 42% by June 30, 2010. To eliminate the dependence on that unique event, APC was removed from the analysis effective April 1, 2010.

⁸The statistical significance of the estimated weighted-average of α_{1t} and α_{2t} was tested using the variance of the inverse-variance weighted average of each coefficient given by their respective $1 / \sum_i 1 / \sigma_i$.

Table 1 — Time-Series Analysis of α_{1t} and α_{2t}
 Data Period: Oct. 1, 2007 – June 30, 2017
 Number of Observations: $N = 2414$

	α_{1t}	α_{2t}
<u>Oct. 1, 2007 – June 30, 2017; $N = 2414$</u>		
Percentage Statistically Non-Zero	0.753	0.541
No. of Positive Coefficients		
No. of Positive Coefficients + No. of Negative Coefficient	0.458	0.292
Inverse-Variance $\frac{1/\hat{\sigma}_{\alpha t}^2}{\sum_t (1/\hat{\sigma}_{\alpha t}^2)}$ Weighted Average	0.118**	-.0048**
<u>Oct. '07 – Dec. '10; $N = 779$</u>		
Percentage Statistically Non-Zero	0.810	0.633
No. of Positive Coefficients		
No. of Positive Coefficients + No. of Negative Coefficient	0.575	0.371
Inverse-Variance $\frac{1/\hat{\sigma}_{\alpha t}^2}{\sum_t (1/\hat{\sigma}_{\alpha t}^2)}$ Weighted Average	0.143**	-.0036**

** : Significant at 1%

In weighing the import of these results, it should be borne in mind there are two possible reasons for a non-zero value to the estimated parameter values α_{1t} and α_{2t} (and, indeed, α_{3t} and α_{4t}). The first is that option prices incorporate newly-arrived information that is not within the historical $\hat{\beta}_{i,SPX,t}$ and $\hat{\sigma}_{it}$ on each stock i obtained from 60-day rolling regressions. And the second explanation is that there are differences between the subjective P and risk-neutral Q probability measures. Finally, it is critical to note that, across the entire stock-market (encompassing, but obviously not limited to, the eight oil-sector stocks analyzed here), the weighted average value of historical and forward-looking betas should both be unity — implying that, across all stocks, the weighted average of α_{1t} and α_{2t} is zero.

It is of interest to pursue further the negativity of α_{2t} results reported in Table 1. Of some interest is the empirical examination of the “Term Structure of Equity Betas.” In the following, we examine the time-series of statistically-significant estimates of α_{2t} during two periods:

1. The period immediately at the onset of the financial crisis, and the subsequent recovery
2. The period surrounding the “Arab Spring” Dec. 2010 – April 2011

We consider the empirical relationship between the time-series of the estimated α_1 and α_2 in $\beta_{itT} = \hat{\beta}_{it} + \alpha_{1t} + \alpha_{2t}T$: A large magnitude of α_1 is indicative of a sharp change from the historical value. Accordingly, we examine for the period 10/1/2007 – 6/30/2017, the empirical time-series of the correlations $\text{Corr}(\alpha_1, \alpha_2)$:

Table 2 — Time-Series of the Empirical Correlations of α_{1t} and α_{2t}

Time Period	Highlights of Period	$\text{Corr}(\alpha_1, \alpha_2)$
	Entire Period	−0.177
2007-11-27 to 2008-07-31	Great Recession I	−0.511
2008-08-01 to 2009-09-30	Great Recession II	−0.162
2009-09-01 to 2010-12-31	Time Before Arab Spring	−0.207
2011-01-01 to 2011-04-30	Arab Spring	−0.029
2011-05-01 to 2014-06-19	Stable Time Before Oil Crash	0.072
2014-06-20 to 2016-02-29	Oil Crash	−0.514
2016-03-01 to 2017-06-30	Post-Oil Crash Period	−0.543

The negative empirical correlation over virtually the entire period is indicative of a mean-reversion of equity beta towards an historical value.

4.2 Results — α_{3t} and α_{8t}

In this section, we focus on the forward-looking information implicit in the intercept term for the measures of idiosyncratic risks α_{3t} presented in the one-factor model (3) and α_{8t} of the two-factor model (8).

By way of background, a timeline of the major events in Sep. 2008 includes:

1. Sep. 7: The Federal Government takes over Fannie Mae and Freddie Mac
2. Sep. 15: Lehman Brothers files for bankruptcy
3. Sep. 16: The Fed bails out insurance giant AIG

In addition to these news reports and the concurrent level/return of equity prices, there are two forward-looking market indicators for the advent of the financial crisis:

1. The level of VIX, the implied vol on the S&P 500 Index. Naturally, such an index rises during times of financial, economic or political crises. While the S&P index bottomed out during the financial crisis on March 3, 2009, VIX peaked on Nov. 20, 2008 at a level of 80.9%.⁹
2. The level of ICJ, the implied correlation in the S&P 500 Index. Assuming an identical correlation across all stocks, the implied correlation reported by the CBOE is:¹⁰

$$\rho_{\text{Average}} = \frac{\sigma_{\text{Index}}^2 - \sum_i w_i^2 \sigma_i^2}{2 \sum_i \sum_{j>i} w_i w_j \sigma_i \sigma_j}. \quad (11)$$

⁹The all-time high for equity implied vol occurred on Black Monday Oct. 19, 1987, when VIX's S&P 100 predecessor VXO spiked to 150.2%.

¹⁰Under appropriate assumptions, the implied correlation can easily be derived. For any portfolio, the variance property holds by definition:

$$\sigma_{\text{Index}}^2 = \sum_i w_i^2 \sigma_i^2 + 2 \sum_{i=1}^N \sum_{j>i}^N w_i w_j \sigma_i \sigma_j \rho_{ij}, \quad (10)$$

where

σ_i = volatility of asset i

w_i = weight of asset i in the index

ρ_{ij} = correlation coefficient between assets i and j

If we now set ρ equal across all assets i and j , we can solve for the ρ_{Average} in eq. (11) in the text.

ρ_{Average} is computed using implied vols σ_i extracted from “SPX options prices, together with the prices of options on the 50 largest stocks in the S&P 500 Index” (<http://www.cboe.com/micro/impliedcorrelation/>). Whereas ICJ was as low as .428 on Oct. 10, 2008, it reached its financial-crisis high of 0.721 on Oct. 29, 2008.

The reason within-market correlations spike during a market crash is well-known: During a market decline, especially when it is sharp, equities’ systematic risk dominates their unsystematic portions.

3. Although the following is not an equity-based measure, it is nevertheless of interest to report what was happening concurrently in the *bond market*: The Merrill Lynch C0A0 Corp. Bond Index peaked on Oct. 30, 2008 at a yield of 9.00%.

The values of the two indicia, VIX and ICJ, are reported in the two figures, 1 and 2. These form the backdrop and benchmarks for the information set available in the time-series of the coefficient α_{3t} , the values of which are designed to answer the question: By signaling forward-looking lower idiosyncratic risk lower than its historical value — i.e., $\alpha_{3t} < 0$ — when do equity options signal the arrival of a financial crisis?

The data plotted in Fig. 3 displays α_{3t} over the period 11/29/07 – 7/31/09. In this time-series, we can observe several episodes of negative α_{3t} 's: Dec. 4, 2007, Feb. 8, 2008, Aug. 21, 2008, Sep. 19, 2008, Nov. 10, 2008, Dec. 19, 2008 and Jan. 8, 2009. Comparable results are depicted for α_{8t} in Fig. 4.

In comparison, during the period of the financial crisis, VIX exceeded 30% for the first time on Sep. 15, 2008, whereas the implied-correlation value exceeded 0.5 on Oct. 16, 2008.

Of substantial interest is the behavior of α_{3t} on the *recovery phase* of the financial recession. As is well-known, VIX peaked on 11/20/08, and the S&P bottomed out on 3/9/09. In terms of α_{3t} , the first large positive numbers are evidenced on 2/9/09, six weeks after the VIX peak but a month before the S&P trough.

Revisiting the above results from a slightly different perspective, as a final test we examine the relative informativeness of historical and forward-looking idiosyncratic variances. To do so, consider the two measures denoted Stat_1 and Stat_2 :

$$\text{Stat}_{1t} = \frac{1}{8} \sum_{i=1}^8 \left(\frac{\hat{\sigma}_{it}}{\widehat{\Sigma}_{it}} \right)^2 \quad (12)$$

$$\text{Stat}_{2t} = \frac{1}{8} \sum_{i=1}^8 \left(\frac{\hat{\sigma}_{it} + \alpha_{3t}}{\Sigma_{it, T=2/12}} \right)^2 \quad (13)$$

Stat_{1t} should be interpreted as the average fraction of historical idiosyncratic variance relative to historical total variance, whereas Stat_{2t} computes the numerator as a forward-looking measure relative to the two-mo. implied vol $\Sigma_{it, T=2/12}$.

In light of the similarity between the two Stat measures, we focus on the latter in Table 3 below. The table presents the timeline for the Great Recession:

Table 3 — Financial Timeline for Great Recession and Key Dates
Identified by Current Model: Forward-Looking Idiosyncratic Variance

Date	Event
Oct. 9, 2007	Stock market at pre-Recession peak
Sep. 7, 2008	Federal Government takes over Fannie Mae and Freddie Mac
Sep. 12, 2008	Forward-looking $\text{Stat}_{2t} \equiv \frac{1}{8} \sum_{i=1}^8 \left(\frac{\hat{\sigma}_{it} + \alpha_{3t}}{\sum_{it, T=2/12}} \right)^2$ declined sharply
Sep. 15, 2008	Lehman Brothers files for bankruptcy
Sep. 16, 2008	Fed bails out insurance giant AIG
Oct. 30, 2008	Merrill Lynch C0A0 Corp. Bond yield peaks at 9.00%
Nov. 20, 2008	VIX peaks at 80.9%
Feb. 9, 2009	First sign of recovery: <i>Positive</i> α_{3t}
March 9, 2009	S&P at Recession trough
March 20, 2009	S&P Implied Correl peaks at 0.748

The date of Sep. 12, 2008 was well into the onset of the Great Recession, as evidenced by the events which took place one week earlier. But of greater significance are the first buds of the recovery displayed on Feb. 9, 2009, when α_{3t} turned significantly positive.

We conclude this section by computing the correlations between α_{3t} , α_{8t} and VIX over different subperiods of interest.

Table 4 — Correlation Matrix: Shocks to One- α_{3t} and Two-Factor α_{8t} Idiosyncratic Risks with VIX

Time Period	Event	Corr: One-Factor α_{3t} with Two-Factor	Corr: One-Factor α_{3t} with VIX	Corr: Two-Factor α_{8t} with VIX
Entire Data Period: 10/1/07 – 6/30/17		0.5311	−0.1038	−0.1425
2007-10-01 to 2008-07-31	Great Recession I	0.6564	−0.1358	−0.1496
2008-08-01 to 2009-02-28	Great Recession II	0.4202	−0.0692	−0.2517
2009-03-01 to 2010-12-31	Time before Arab Spring	0.5834	0.1470	0.1430
2011-01-01 to 2011-04-30	Arab Spring	0.3572	−0.0152	−0.0198
2011-05-01 to 2014-06-19	Stable Time before Oil Crash	0.4977	−0.1011	−0.2089
2014-06-20 to 2016-02-29	Oil Crash	0.4460	−0.0031	−0.4196
2016-03-01 to 2017-06-30	After Oil Crash	0.6314	−0.1756	−0.3494

The three concluding columns of Table 4 convey the important information: For every sub-period, $\text{Corr}(\alpha_{3t}, \alpha_{8t})$ were always positive. Moreover, with the sole exception occurring during the relatively-stable period preceding the onset of the “Arab Spring,” $\text{Corr}(\alpha_{3t}, \text{VIX})$ and $\text{Corr}(\alpha_{8t}, \text{VIX})$ were negative, indicating a forward-looking increase in idiosyncratic variance was coincident with a decline in VIX, both of course indicating a more pacific time period.

5 The One-Factor Oil Model: Crude-Oil Betas

5.1 The One-Factor Oil Model

By analogy to the one-factor equity model, we now consider a one-factor model for the different maturities of the oil futures market. Accordingly, for varying maturities T , apply a market-model to returns on crude-oil futures contracts,

$$r_T = a_T + \beta_T R_{\text{SPX}} + e_T \quad (14)$$

$$\Sigma_T^2 = \beta_T^2 \sigma_{mT}^2 + \sigma_T^2 \quad (15)$$

where

$\Sigma_T^2 \equiv \text{Var}(r_T)$, the variance of the return on crude-oil futures contract of maturity T

$\beta_T \equiv \text{Cov}(R_T, R_{\text{SPX}}) / \text{Var}(R_{\text{SPX}})$, market beta of oil futures contract of maturity T

$\sigma_{mT}^2 \equiv \text{Var}(R_{\text{SPX}, T})$, the variance of the return on the S&P 500 market index to expiration date T

$\sigma_{iT}^2 \equiv \text{Var}(e_{iT})$, the idiosyncratic variance

As before, the hybrid portion of the model links the historical estimates $\{\hat{\beta}_i, \hat{\sigma}_i\}$ to their forward-looking analogues $\{\beta_i, \sigma_i\}$ via an additive (16) correction: Pooling all maturities $T \leq 2$ yrs., at any date $t \in [10/1/2007 - 6/30/2017]$,

$$\begin{cases} \beta_{Tt} = \hat{\beta}_{Tt} + \alpha_{9t} \\ \sigma_{Tt} = \hat{\sigma}_{Tt} + \alpha_{10,t} \end{cases} \quad (16)$$

It might be usefully noted the absence of a time-to-maturity regressor T is attributable to the import of parsimony: Because the variables $\hat{\beta}_{Tt}$ and $\hat{\sigma}_{Tt}$ already contain T -dependence, we assign only the intercept terms α_{9t} and $\alpha_{10,t}$.

For this case, on any date- t the minimization performed to estimate the two coefficients is:

$$\min_{\{\alpha_{9t}, \alpha_{10,t}\}} \sum_T \left[\Sigma_{tT}^2 - (\hat{\beta}_{tT} + \alpha_{9t})^2 \sigma_{mT}^2 - (\hat{\sigma}_{tT} + \alpha_{10,t})^2 \right]^2. \quad (17)$$

Whereas much of the focus in the one-factor equity model was on the informativeness of forward-looking idiosyncratic variances, in this one-factor model we are keenly interested in the intercept term α_{9t} , as it might signify substantial change in the prevailing oil regime.

5.2 The Stylized Facts Regarding the Correlation of the Spot Price of Oil with the S&P 500

By way of characterizing the correlation of spot oil prices with the S&P 500 from Oct. 1, 2007, Fig. 5 is instructive. The figure plots the 45-day moving historical correlation between the rates of return on the prompt-month oil futures contract and the S&P 500 stock price index.

With the important exception of early 2011, most of the time since the second half of 2008, oil contracts have exhibited positive comovement with equity markets. It is useful succinctly to summarize “demand-side” vs. “supply-side” effects in the following Table 5:

Table 5 —
Distinguishing between Demand- and Supply-Side Shocks on Oil Prices

	Demand-Side Shocks	Supply-Side Shocks
Definition	Arise from a change in the economy’s demand for oil products, usually associated with either growth or recessionary conditions	Are attributable to constraints on oil supplies reaching consumer markets, due to either hurricanes in the Gulf of Mexico, geopolitical crises in the Middle East or spiking oil prices outpacing economic growth
Correlation(Oil Prices, S&P 500)	Positive correlation	Negative correlation

5.3 Data for the One-Factor Oil Model

Analogous to the data for the one- and two-factor equity models, the data here covers the same date range of Oct. 1, 2007 – June 30, 2017. The prices we observe are those for WTI oil-futures contracts out to a 24-mos. maturity, as well as option prices (i.e., implied vols) for options as far out as two-yrs. expiration. As was the case for option prices in the equity model, we use the principle variance is additive across maturities and const. forward vol to interpolate for any missing expiration dates.

5.4 Empirical Results for the Forward-Looking Betas' α_{9t} and a Contrast with the Two-Factor Equity Model's α_{7t}

Figs. 6 and 7 graph the time-series of α_{9t} during the two critical periods of the Great Recession and the “Arab Spring.” With respect to the Great Recession, it is clear α_{9t} turned significantly negative in Sep. 2008: While oil prices were still rising until July 3, 2008, the correlation with equity markets was negative. Then, as the historical correlation began to change sign, the forward-looking beta was significantly more positive, indicating a descent into the demand-side shock of the Great Recession.

The “Arab Spring” began with the immolation in Tunisia on Dec. 17, 2010. However, it was not until the third week of Jan. 2011 that the forward-looking oil beta turned negative in harmony with a supply-side shock. From Figs. 5 and 7, it is of note that, in terms of the oil markets, the “Arab Spring” appeared to conclude April 5, 2011. While the political analysis of that is well beyond the scope of this paper, the onset of these political events was in North Africa, which does not produce substantial oil for export. Moreover, oil markets signaled they were significantly becalmed by the recognition there would not be contagion to the Persian Gulf, where substantial oil is actually produced.

We conclude this section by tabulating in Table 6 below the correlation between the term α_{9t} the two-factor equity model's α_{7t} :

Table 6 —
Correlations of Perturbations $\{\alpha_7, \alpha_9\}$ in the One- and Two-Factor Models

Time Period	Event	Corr (α_7, α_9)
	Entire Period	0.071
2007-11-27 to 2008-07-31	Great Recession	0.277
2008-08-01 to 2009-09-30		0.127
2009-10-01 to 2010-12-31	Time Before Arab Spring	-0.016
2011-01-01 to 2011-04-30	Arab Spring	0.215
2011-05-01 to 2014-06-19	Stable Time Before Oil Crash	0.256
2014-06-20 to 2016-02-29	Oil Crash	-0.117
2016-03-01 to 2017-06-30	After Oil Crash	-0.100

The correlations between α_{7t} and α_{9t} are not strong nor are they uniformly positive. That may be a consequence of the highly non-linear aspect of the two-factor equity model.

5.5 Using Forward-Looking Betas in the CAPM to Forecast the Spot Price of Oil

As previously noted, there is a plethora of oil-price forecasters. However, just as CCJV (2012) demonstrated some success in testing the CAPM with forward-looking betas, our model here offers the challenge of using forward-looking betas in the CAPM to forecast the spot price of oil.

As is well-known, in an informationally-efficient market, futures prices represent the risk-neutral expectations of spot prices:

$$E^*(S_T | I_t) = F,$$

where

$E^*(\cdot | I_t)$ = Risk-neutral expectations at time t , based on the information set I_t

S_T = Spot price at some future date T

F = Futures price at date t

The point here is that futures prices are assumed to incorporate all information I_t traders believe is relevant in setting prices. What futures prices cannot contain is the risk premium:

$$E(S_T | I_t) = F + \text{Oil Risk Premium},$$

where $E(\cdot | I_t)$ = standard (physical measure) expectations at time t . Thus, the challenge we wish to address using the CAPM is to provide a *forward-looking* risk premium, meaning one based on prospective, rather than merely retrospective/historical, statistics.

The specification of a CAPM approach to the Commodity Market Price of Risk is straightforward. Let

μ_i = Expected return on maturity i

μ_M = Expected return on the market portfolio

r = Riskfree rate of interest

In that case,

$$\begin{aligned}
\mu_i &= \beta_i (\mu_M - r) \\
&= \frac{\text{Cov}(R_i, R_M)}{\text{Var}(R_M)} (\mu_M - r) \\
&= \frac{\rho_i \sigma_i \sigma_M}{\sigma_M^2} (\mu_M - r) \\
&= \frac{\rho_i \sigma_i}{\sigma_M} (\mu_M - r) \\
&= \rho_i \sigma_i \frac{\mu_M - r}{\sigma_M}
\end{aligned} \tag{18}$$

Combining the CAPM with the oil futures markets yields:

$$\mu_{it} = \rho_{it} \sigma_{it} \frac{\mu_M - r}{\sigma_M} \equiv \rho_{it} \sigma_{it} \lambda_t \tag{19}$$

With respect to futures contract of maturity i ,

$$\begin{aligned}
E(F_{iT}) &\equiv F_{i0} \exp\{\mu_{iT}T\} \\
&= F_{i0} \exp\{\rho_{it} \sigma_{it} \lambda_t T\} \\
\implies \frac{1}{T} \ln \left[\frac{E(F_{iT})}{F_{i0}} \right] &= \rho_{it} \sigma_{it} \lambda_t
\end{aligned}$$

$$\begin{array}{l}
\text{Annualized Expected} \\
\text{Futures Price Change}
\end{array}
\equiv \rho_{it} \left(\begin{array}{l} \text{Current CL}_i \\ \text{Implied Vol} \end{array} \right) \left(\begin{array}{l} \text{Current Stock Market} \\ \text{Sharpe Ratio} \end{array} \right)$$

The clear implication of eq. (19) is this: When $\rho_{it} < 0$ — say, because of a geopolitical crisis — the resulting $F_{i0} > E(F_{iT})$ reflects the intuitive notion of a risk premium attributable to concerns over oil supplies reaching consumer markets. In words, a negative correlation implies an upward-bias of the futures price relative to its expected-spot price (under the physical measure) counterpart.

The advent of a perceived supply-side crisis generally gives rise to numerous references in the lay press to an oil-price surcharge. For example, “The possibility

that there might be a disruption in oil supply at some time in 2012 as Iran retaliates has, I think, permanently embedded a \$10 to \$20 premium in the price of oil,” according to the chief global economist at the Economic Outlook Group cited in the 12/28/11 issue of the *New York Times*. The approach proffered here thus provides a quantitative measure for this risk premium.¹¹ Moreover, the procedure detailed here provided a forward-looking correlation, and as such may be superior to the historical correlation between crude-oil futures and the S&P.

Generally, when we test forecasts in finance, we make the plausible assumption expectations materialize on average (over long periods of time), standard tests call for confronting models with realized returns. However, when available — such as those provided by Bloomberg — *expectations* can provide an alternate test. Specifically, the Bloomberg system reports the oil-price forecasts reported by professionals in the oil sector. While we will subsequently test these in depth, it is useful here to demonstrate the impact of demand- and supply-side crises on the bias of these forecasts relative to the corresponding futures contracts. To see this, consider Figs. 8 and 9.

Fig. 8 presents these price forecasts on a “normal” date, Feb. 15, 2016, when the oil beta was positive. In contrast, Fig. 9 presents these on March 2, 2011, the heart of the financial “Arab Spring” in 2011. The key line to observe is “Diff (Median – Current),” which computes the difference between the median of the price forecasts and the corresponding current futures contract.

¹¹ To be fair, such references are typically made with respect to *spot prices*, but the theory suggests they be reflected in the prices of futures contracts.

While not pertaining directly to the model’s forward-looking aspects, the first set of empirical results consider the intercept term in the standard one-factor-model:

$$r_{it} = \alpha + \beta_i r_{Mt} + \epsilon_t \quad (20)$$

where $\beta_i \equiv \rho_i \sigma_i / \sigma_M$. Define the (annualized) empirical estimate $\hat{\alpha}$ as the abnormal realized return during the period of estimation t . In light of the result

$$\frac{1}{T} \ln \left[\frac{E(F_{iT})}{F_{i0}} \right] = \rho_{it} \sigma_{it} \lambda_t,$$

conditional on the value of ρ_{it} , consider the *normalized* (i.e., adjusted for risk) value $\hat{\alpha} / \sigma_{it}$ assumed stationary. Finally, consider empirical estimates of a 45-day moving average of $\hat{\alpha} / \sigma_{it}$ for ten deciles of $\rho_{it} \equiv \text{Corr}(r_{it}, r_{Mt})$:

Table 7 — Decile Range of $\text{Corr}(r_{it}, r_{Mt})$

Dates: 10/01/07 to 07/31/14				Dates: 10/01/07 to 06/30/17			
Decile	max	min	Ave. $\left(\frac{\hat{\alpha}}{\sigma_{it}}\right)$	Decile	max	min	Ave. $\left(\frac{\hat{\alpha}}{\sigma_{it}}\right)$
1	0.7746	0.6855	0.120	1	0.7746	0.6581	0.362
2	0.6850	0.6216	0.766	2	0.6579	0.5819	0.082
3	0.6211	0.5750	-0.312	3	0.5816	0.5289	-0.239
4	0.5748	0.5295	-0.321	4	0.5288	0.4746	0.102
5	0.5292	0.4897	0.203	5	0.4741	0.4050	0.456
6	0.4896	0.4042	0.498	6	0.4043	0.3129	0.325
7	0.4032	0.2977	0.976	7	0.3122	0.2212	0.199
8	0.2973	0.1065	1.325	8	0.2210	0.0980	0.250
9	0.1048	-0.0996	1.581	9	0.0976	-0.0271	0.548
10	-0.1002	-0.5479	1.782	10	-0.0273	-0.5479	1.927

There are two conclusions we can infer from Table 7:

1. Generally, the more negative the correlation between crude-oil futures and the stock-market — i.e., the more profound the geopolitical crisis which gave rise to a negative correlation — the greater the positive realized excess return $\hat{\alpha}$, normalized by the measure of ex-ante risk σ_{it} .

2. In contrast to the earlier time period 10/01/07 to 07/31/14, the results may not be as clearly monotonic for the full data period 10/01/07 to 06/30/17. This may be due to the consequence of a “negative” supply-side shock in the 2014 – 2015 time period, when prices declined precipitously due to an injection of the additional supply due to the increase of hydraulic fracking (“fracking”) activity in the continental U. S.

We now turn to the formal empirical tests of the different oil-price forecasts, including those offered by professional forecasters as well as those defined by different applications of the CAPM. There are several steps in the description of these empirical tests:

1. From the data and two of the models, we have forward-looking oil betas:
 - (a) The historical beta: $\hat{\beta}_{tT}$
 - (b) From the two-factor equity-and-oil model, $\beta_{tT} = \hat{\beta}_{tT} + \alpha_{7t}$
 - (c) From the one-factor oil model, $\beta_{tT} = \hat{\beta}_{tT} + \alpha_{9t}$

Combining these betas with also-forward looking σ_M and λ_t :

$$\begin{aligned}\mu_{tT} &= \beta_{tT} (\mu_M - r) \\ &= \beta_{tT} \sigma_M \lambda_t,\end{aligned}$$

which in turn results in the financial-economics price forecast:

$$E(F_{iT}) \equiv F_{i0} \exp\{\mu_{iT}T\} \tag{21}$$

In the empirical tests, $T = 1$ yr. Let BF be the Bloomberg price forecast. Designate the respective price forecasts $E(F_{iT})$ in eq. (21): HF, MF₂, MF₁.

2. Before turning to these results, it is instructive to compare and contrast the two sets of expectations — realizations and price forecasts — in terms of a measure of consistency. Under the CAPM, recall a positive Corr(CL1, SPX) implies the expected spot price should exceed the futures prices. What is of interest is to see whether that is indeed consistent with the Bloomberg price forecasts. To see this, consider the results of Table 8:

Table 8 — Measure of Sign Agreement of Correlations with
Bloomberg CPFC Forecasts

$$\frac{\text{sign} \left[\text{Corr} (\text{CL1}, \text{SPX}), E \left(\frac{\text{Forecast}}{\text{Futures}} \right) - 1 \right]}{\text{Correl Measured Over 45-Day Moving Window}}$$

Model	Measure of Sign Agreement
Historical Beta with Bloomberg Forecast	76.4%
One-Factor Forward Beta and Bloomberg Forecasts	75.9%
Two-Factor Forward Beta and Bloomberg Forecasts	42.5%

With respect to this one measure, the historical beta does slightly better than the one-factor forward beta, and substantially better than the two-factor beta. Esp. with respect to the two-factor beta, the model's highly non-linearity may adversely affect the empirical results.

- Turning now to generate forward-looking forecasts, we seek to use a forward-looking Sharpe Ratio in eq. (19). With that in mind, consider the Doran, Ronn and Goldberg (2009) model for an equity-market expected rate of return:

$$\begin{aligned} \mu_{Mt} &= r_{St} + \left(0.46 - 0.162 \frac{\text{S\&P } 500_t}{\text{S\&P } 500_{t-5,t-6}} \right) \text{VIX}_t \\ \implies \lambda_t &\equiv \frac{\mu_{Mt} - r_{St}}{\text{VIX}_t} = 0.46 - 0.162 \frac{\text{S\&P } 500_t}{\text{S\&P } 500_{t-5,t-6}} \end{aligned} \quad (22)$$

where

μ_{Mt} = the expected rate of return on the Market portfolio at time t

r_{St} = the one-month short-term rate of interest

$\text{S\&P } 500_{t-5,t-6}$ = average value of the S&P 500 Index for a one-year period centered 5.5 yrs. ago

VIX_t = contemporaneous value of the VIX implied-vol index

The model's parameters 0.46 and 0.162 were obtained from a proxy for the market's expected risk premium (*not* realized returns), inserted into a linear regression on a constant plus the ratio $\text{S\&P } 500_t / \text{S\&P } 500_{t-5,t-6}$.

Fig. 10 below depicts the time series of λ as given by eq. (22) over the relevant time period. Immediately following Fig. 10 is Fig. 11, which depicts the four price forecasts denoted BF, HF, MF₂ and MF₁.

While the other price forecasts have varied errors to whose numerical examination we soon turn, what is clear is that the Bloomberg price forecasts as presented on the online system suffer from a lag effect: In addition understandably to appearing “sluggish” relative to daily price movements, they appear to lag the price increases and then lag the price declines.

Finally, we turn to two sets of regression-based analyses of these price forecasts. The first set considers day-ahead prediction, whereas the second uses Newey-West regressions to perform one-year ahead regressions. In Tables 9a and 9b, we test the log levels of the forecasts as well as whether the forecast *relative to* the prevailing one-year or one-month futures contracts are meaningful:

Regression Tests of Forward-Looking Oil Betas and their Oil-Price Forecasts

Table 9a — OLS Regressions — One-Day Ahead Regressions

- Notation:

CL13_t = Futures contract with one-year maturity on day *t*

CL2_t = Futures contract with two-months to maturity on day *t*

MF_{it} = One-Year price forecast from Model *i* on day *t*, *i* = 1, ..., 4

i = 1 : Historical-Beta Forecast

i = 2 : One-Factor Model

i = 3 : Two-Factor Model

i = 4 : Bloomberg Forecast

-

Regression #	Regression Equation			
1	$\ln \text{CL13}_{t+1} - \ln \text{CL13}_t = a_1 + a_2 (\ln \text{MF}_{it} - \ln \text{CL13}_t)$			
2	$\ln \text{CL13}_{t+1} - \ln \text{CL2}_{t+1} = a_1 + a_2 (\ln \text{MF}_{it} - \ln \text{CL2}_t)$			
3	$\ln \text{CL13}_{t+1} = a_1 + a_2 \ln \text{MF}_{it}$			
Regression #	Price Forecast by	<i>a</i> ₂ Coefficient	<i>t</i> -stat	<i>R</i> ²
1	Historical Beta	−0.019	−1.93	0.1%
	MF1	−0.02	−2.99	0.4%
	MF2	0.0004	0.071	0.0%
	Bloomberg Forecast	−0.025	−7.81	2.5%
2	Historical Beta	0.61	92.7	78%
	MF1	0.5	72.1	68%
	MF2	0.45	50.3	51%
	Bloomberg Forecast	0.26	37.4	37%
3	Historical Beta	1.03	331	97%
	MF1	1.02	229	95%
	MF2	0.98	202	94%
	Bloomberg Forecast	1.07	112	84%

Table 9b — Newey-West One-Year-Ahead Regressions

• Notation:

CL13_t = Futures contract with one-year maturity on day *t*

CL2_t = Futures contract with two-months to maturity on day *t*

MF_{it} = One-Year price forecast from Model *i* on day *t*, *i* = 1, ..., 4

i = 1 : Historical-Beta Forecast

i = 2 : One-Factor Model

i = 3 : Two-Factor Model

i = 4 : Bloomberg Forecast

•

Regression #	Regression Equation				
1	$\ln \text{CL2}_{t+1\text{yr}} - \ln \text{CL13}_t$	=	$a_1 + a_2 (\ln \text{MF}_{it} - \ln \text{CL13}_t)$		
2	$\ln \text{CL2}_{t+1\text{yr}} - \ln \text{CL2}_t$	=	$a_1 + a_2 (\ln \text{MF}_{it} - \ln \text{CL2}_t)$		
3	$\ln \text{CL2}_{t+1\text{yr}}$	=	$a_1 + a_2 \ln \text{MF}_{it}$		
Regression #	Price Forecast by	a_2 Coefficient	<i>t</i> -stat	R^2	
1	Historical Beta	8.5	2.33	7%	
	MF1	3.09	4.77	37%	
	MF2	1.25	3.16	8%	
	Bloomberg Forecast	-0.7	-0.61	0.4%	
2	Historical Beta	2.624	5.81	55%	
	MF1	2.237	6.52	52%	
	MF2	1.887	4.26	34%	
	Bloomberg Forecast	1.026	2.6	22%	
3	Historical Beta	0.719	3.35	32%	
	MF1	0.745	3.67	33%	
	MF2	0.642	2.95	27%	
	Bloomberg Forecast	0.635	2.48	20%	

Other than the difference in the econometrics between OLS and Newey-West, there is an important distinction between Tables 9a and 9b: Whereas 9a is one-day-ahead regressions, 9b provides for a one-year-ahead time frame.

A careful review of Tables 9a and 9b gives rise to several conclusions:

- (a) Regression 3 in Table 9a, $\ln \text{CL13}_{t+1} = a_1 + a_2 \ln \text{MF}_{it}$, unsurprisingly gives rise to a_2 values quite close to unity. While comforting, it does not confirmation fails to constitute a true test of a falsifiable hypothesis.

Regression 3 in Table 9b, $\ln \text{CL2}_{t+1\text{YR}} = a_1 + a_2 \ln \text{MF}_{it}$, gives rise to coefficients which are less than 1.0, with the one-factor beta forecast delivering the coefficient closest to unity.

- (b) Regression 2 in Table 9a, $\ln \text{CL13}_{t+1} - \ln \text{CL2}_{t+1} = a_1 + a_2 (\ln \text{MF}_{it} - \ln \text{CL2}_t)$, constitutes a test of whether the price forecast deviation from the second futures contract informs the change of the one-year ahead futures contract deviation from the same second futures contract. Since it is a deviation, that is a more meaningful test. While all coefficients are positive, the historical beta has the coefficient closest to unity.

- (c) Regression 1 in Table 9b, $\ln \text{CL2}_{t+1\text{YR}} - \ln \text{CL13}_t = a_1 + a_2 (\ln \text{MF}_{it} - \ln \text{CL13}_t)$, tests whether the one-year-ahead perturbation of the realized second futures contract deviation from its one-year-prior value is positively related to the difference between the price forecast and the one-year-prior value. This is one test where the forward-beta forecasts outperform the historical beta and Bloomberg Forecast values. Although its coefficient of 3.09 is in excess of unity, the R^2 for the one-factor beta is an impressive 37%.

- (d) Regression 2 in Table 9b, $\ln \text{CL2}_{t+1\text{YR}} - \ln \text{CL2}_t = a_1 + a_2 (\ln \text{MF}_{it} - \ln \text{CL2}_t)$, tests the one-year-ahead deviation of the second month contract from its one-year-prior value, as a function of the deviation of the forecast from that one-year-prior value. All three forward-looking forecasts outperform the historical beta in terms of slope coefficients being closer to unity, with the one-factor beta having an R^2 essentially equal to the historical beta's.

Although not in all cases, the forward-looking betas outperform the historical beta in terms of its forecasting ability in several important features, in both

one-day-ahead as well as one-year-ahead time frames. These results indicate the CAPM merits serious attention in terms of applying an intuitive asset pricing model to the forecasting of oil prices.

6 Conclusions

Harkening back to the long tradition of viewing market prices as conveyors of important forward-looking information, this paper has documented the potential for using options on equities, oil futures contracts and a market index to infer forward-looking statistics of relevance to investors and portfolio managers. We summarize those conclusions here:

1. Applying a single-factor model to equity and index options, we were able to use historical data parsimoniously to obtain meaningful forward-looking equity betas and idiosyncratic variances. In comparing and contrasting these signals with other forward-looking measures, such as VIX and ICJ, we are able to extract what might be termed the “Message from Markets.”

In terms of relating the one-factor model’s forward-looking idiosyncratic variances, the critical dates are Sep. 12, 2008 for the date anticipating a crisis, and Feb. 9, 2009 as the date for when evidence of recovery is shown to begin. While the former date is already well within the “crisis period,” the latter date is one month ahead of the equity market’s trough.

2. In recognizing the negative correlation between the model’s intercept α_1 and its slope α_2 , the one-factor equity model provided insights into what might be called a mean-reversion of equity betas.
3. The one-factor oil model provides interesting indicia of market changes on two dates. The first is in mid-Sep. 2008, as oil prices continued their crash from the July 3rd peak. The second is the third week of Jan. 2011, as the “Arab Spring” was continuing to be recognized as an important political and economic driver.
4. Although not directly a test of the forward-looking model, Table 7 constituted a test of the distinction between the supply- and demand-side shocks: The

results demonstrated the *ex-post* risk-adjusted return increases with the onset of a supply-side (i.e., negative oil-price beta) effect.

5. Finally, the exhaustive price-forecasts tests demonstrated both the relevance of the CAPM in generating oil-price forecasts. This applied both to the historical-beta, as well as the outperformance by forward-looking (one- and two-factor) betas in several important categories reported by the results of Tables 9a and 9b.

While the work here can admit extensions in several directions, perhaps an interesting one would be to explore the efficacy of the forward-looking methodologies utilized here — that is, the use of a confluence of historical statistics and options' forward-looking risk-neutral expectations — in other areas of investment analysis in finance.

References

1. Bakshi, G., and N. Kapadia, 2003a, "Delta-hedged Gains and the Negative Market Volatility Risk Premium." *Review of Financial Studies* 16, 527-566.
2. Bakshi, G., and N. Kapadia, "Volatility Risk Premiums Embedded in Individual Equity Options Some New Insights," *Journal of Derivatives*, Fall 2003b, Vol. 11, No. 1, pp. 45 – 54.
3. Basak, Suleyman and Anna Pavlova, "A Model of Financialization of Commodities," London Business School CEPR Working Paper, May 2013.
4. Buss, Adrian and Grigory Vilkov, "Measuring Equity Risk with Option-Implied Correlations," *Review of Financial Studies*, 2012, Vol. 25, Issue 10, pp. 3113 – 3140.
5. Bollerslev, Tim, Michael Gibson and Hao Zhou, "Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities," *Journal of Econometrics*, Volume 160, Issue 1, January 2011, pp. 235 – 245.
6. Chang, Bo-Young, Peter Christoffersen, Kris Jacobs and Gregory Vainberg, "Option-Implied Measures of Equity Risk," *Review of Finance*, 2012, Volume 16, Issue 2, pp. 385 – 428.
7. Chicago Board Options Exchange, "CBOE S&P 500 Implied Correlation Index," White Paper, 2009.
8. D'Ecclesia, Rita, "Correlations among energy commodities: New evidences," Working Paper, 2017.
9. Doran, James, Ehud I. Ronn and Robert S. Goldberg, "A Simple Model for Time-Varying Expected Returns on the S&P 500 Index," *Journal of Investment Management*, Vol. 7, Number 2, 2nd Quarter 2009.
10. Fouque, Jean-Pierre and Eli Kollman, "Calibration of Stock Betas from Skews of Implied Volatilities," *Applied Mathematical Finance*, Vol. 18 , Iss. 2, 2011.

11. Gutierrez, Roberto G., Jean Marie Linhart and Jeffrey S. Pitblado, “From the help desk: Local polynomial regression and Stata plugins,” *The Stata Journal*, 2003, 3, Number 4, pp. 412 – 419.
12. Hamilton, J.D. and J.C. Wu, 2012, “Effects of index-fund investing on commodity futures prices,” Working Paper, UCSD.
13. Henderson, B.J., N.D. Pearson and L. Wang, 2012, “New evidence on the financialization of commodity markets,” working paper, George Washington University.
14. Jackwerth, J. C. and M. Rubinstein, 1996. “Recovering probability distributions from option prices.” *Journal of Finance* 51, 1611-1631.
15. Latané, Henry and Richard Rendleman, “Standard Deviation of Stock Price Ratios Implied in Option Prices,” *Journal of Finance*, 31 (May 1976), pp. 369 – 381.
16. Low, Buen Sin and Shaojun Zhang, “The Volatility Risk Premium Embedded in Currency Options,” *Journal of Financial and Quantitative Analysis*, Vol. 40, No. 4, Dec. 2005, pp. 803 – 832.
17. Margrabe, William, “The Value of an Option to Exchange One Asset for Another,” *Journal of Finance*, 33 (March 1978), pp. 177 – 186.
18. Murphy, Finbarr and Ehud I. Ronn, “The Valuation and Informational Content of Options on Crude-Oil Futures Contracts,” *Review of Derivatives Research*, 2014.
19. 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.
20. Singleton, K., 2012, “Investor flows and the 2008 boom/bust in oil prices,” working paper, Stanford University.