The Implied Futures Financing Rate*

Nicholas L. Gunther† Robert M. Anderson‡ Lisa R. Goldberg§
Alex Papanicolau¶

February 20, 2021

*We thank our present and former colleagues at the Center for Risk Management Research and the Consortium for Data Analytics in Risk for many interesting and stimulating discussions that enhanced the quality of this article.

†Department of Economics and Consortium for Data Analytics in Risk, University of California at Berkeley, nlgunther@berkeley.edu.

‡Department of Economics and Consortium for Data Analytics in Risk, University of California at Berkeley, robert.anderson@berkeley.edu.

§Departments of Economics and Statistics and Consortium for Data Analytics in Risk

¶Consortium for Data Analytics in Risk, University of California at Berkeley,

Contents

The Implied Futures Financing Rate

Abstract

We explore the cost of implicit leverage associated with an S&P 500 Index futures contract and derive an implied financing rate (the *Futures-Implied Rate* or *FIR*), based on a simple model of stock and futures, without any explicit arbitrage or other relationship to market interest rates. We develop new estimation methods for the FIR, including point and interval estimates, with important advantages over existing methods, and extend our methods to bivariate estimates of the FIR and equity volatility based on Wishart distributions. While our resulting FIR estimates were often attractive relative to market rates on explicit financings, the relationship between the implicit and explicit financing rates was volatile and varied considerably based on legal and economic regimes. Among our significant findings was the effectiveness of regulatory reform in reducing substantially the spreads between this FIR and contemporaneous LIBOR and US Treasury rates. Other findings include estimates of the convexity adjustment associated with the FIR, futures-based estimates of stock volatility and stock-rate correlation, and new test statistics for the significance of these estimates.

1 Introduction.

What is the true financing cost of a levered equity investment strategy? Generally speaking, an investor who wishes more exposure than her free cash permits has two choices: she can borrow and add the debt proceeds to her cash investment, or, as discussed in this article, she can obtain implicit leverage through the derivatives market.[1] In this paper, we explore the cost of implicit leverage associated with a prominent equity index futures contract. In summary, our research shows that the related implicit financing rate has often been attractive relative to market rates on explicit financing; however, the relationship between the implicit and explicit financing rates has been volatile and varied considerably based on legal and economic regimes.

The purchase of an S&P 500 Index[2] futures contract[3] results in exposure to a notional amount of the underlying in exchange for a current payment much less than that notional amount.[4] Intuitively, this suggests financing; a levered purchase of the underlying equities provides similar exposure and cash flows. In this paper, we consider a stock and two futures contracts on the stock, from which we derive an implied forward financing rate (the *futures implied rate*, or *FIR*). This rate represents the implicit cost of financing associated with such an investment in a futures contract. The spread between this implied rate and market interest rates evolved dynamically over a series of four regimes that we identified during our total observation period from January 1996 to June 2019. The Commodity Futures Modernization Act of 2000 ("CFMA") reduced this spread; the 2008 financial crisis and the subsequent recovery each altered it further.

[1] For an exploration of explicit leverage, *see, e.g.*, Anderson et al. (2014).

[2] The term 'S&P 500 Index' is a registered copyright of Standard & Poors, Inc.

[3] Unless otherwise stated, the terms *futures contract* in this article refers to such an S&P 500 Index futures contract. The terms of such futures contracts are specified by the Chicago Mercantile Exchange (*CME*) for the big (or *floor*) and the small (or *E-mini*) futures contracts traded through the CME. Current terms for such futures contracts are available on the CME's website. The CME has other futures contracts related to the S&P 500 Index, including contracts that reflect dividends on the index components. These other contracts lack the volume and the history of the traditional contract we choose.

[4] The payment required to execute a futures contract consists principally of initial and variational margin; leverage of 10:1 or higher is often possible. *See* footnote 3.

We begin our investigation of the FIR with few restrictions on the stock and futures price processes, and then for more explicit results specialize to a simple model of stock and futures, without any explicit arbitrage or other relationship to market interest rates.

2 Literature

The futures literature is extensive[5] and the relationship between futures prices and interest rates is well-studied in the customary context of no-arbitrage pricing.[6] There is however, little on that relationship outside that no-arbitrage context, where market rates and rates implied by futures can be unequal.[7] Recently, recognition that simple static models of futures hedges are inadequate has appeared.[8] Our current article investigates the interest rate implied by S&P 500 futures prices determined independently from market interest rates, to our knowledge for the first time.[9]

Anderson et al. (2014) analyze the cost of explicit financing in a levered equity investment strategy; the current article provides a companion analysis of the cost of implicit financing.

[5] Dating back at least to, *e.g.*, Black (1976).

[6] See, *e.g.*, Cox et al. (1981) (showing, *inter alia*, that futures and forward prices are equal if rates are deterministic.)

[7] For example, because, unlike forward contracts, as a result of variational margin (see footnote 4), hedges between positions in the underlying physical stock and in futures contracts are dynamic and thus leave residual open risk.

[8] Contrast Dwyer et al. (1996) (using a simple static cost of carry model to explore the nonlinear relationship between the S&P 500 Futures and the underlying stock) with Hilliard et al. (2021) (using a 'short-lived arbitrage' model proposed in (Otto, 2000) to compute minimum variance hedge ratios.)

[9] The futures literature discusses many other topics not relevant to the current investigation. *See, e.g.,* Chen et al. (2016) (comparing price discovery in futures, index options and ETFs) and references cited therein; Chan et al. (1991) (comparing volatility in the cash and index futures markets); Bali et al. (2008)(testing mean reversion in S&P 500 futures) and Lien (2004)(investigating the effect of cointegration on hedging effectiveness.)

3 Interest Rates Implied by Futures

3.1 General Case without Dividends

The basic equations for futures and implied financing rates used in this paper come from the well-known no-arbitrage relationship that relates the futures price $F(t, T)$ at the current time t to the expected stock price $S(T)$ at the maturity date T of the relevant futures contract, conditioned on the information at time t:

Theorem 3.1. *Futures as Expectation*

$$F(t,T) = E_t[S(T)] \equiv E_Q[S(T)|t] \equiv E_Q[S(T)|\mathscr{F}_t] \tag{1}$$

This result is true quite generally, independent of any particular model for the stock or interest rates. The Computational Appendix contains a proof in continuous time.[10]

In Equation (1), $E_Q[S(T)|\mathscr{F}_t]$ denotes the expectation of S(T) conditional on the information available at an earlier time t under the risk neutral measure Q associated with riskless investment in a *money market account* we denote by $M(t_0, t_1), t_0 < t_1$. We will assume there is a short rate process $r(t)$ associated with this riskless investment. The related (stochastic) *discount factor* (or just *discount*) is also denoted by M but with time order reversed:

$$M(t_0, t_1) = e^{\int_{t_0}^{t_1} r(s)} ds, \; M(t_1, t_0) = \frac{1}{M(t_0, t_1)}$$

We assume that r_t is attainable[11] from the stock, the near futures contract and the next futures contract on the S&P 500 Index.[12] Such a short rate process need not be equivalent to

[10]For a proof in the discrete case, see, e.g., (Anderson and Kercheval, 2010, Chapter 4).

[11]A claim is *attainable* from other claims if there is a self-financing trading strategy in those other claims that has the same payoffs as the first claim. *See, e.g.* (Anderson and Kercheval, 2010, Chapter 2.3).

[12]These contracts are traded on the Chicago Mercantile Exchange. In market terminology, the *near* contract is the one with the earliest possible maturity after the observation time, and the *next* contract is the one with next earliest maturity. The difference in the maturity dates of two consecutive S&P 500 Index futures contracts, such as the near and the next, is typically approximately three months. In this research,

a market interest rate, where two rates are *equivalent* if they may be swapped into each other at no cost.[13] Instead, it is the rate of riskless investment attainable through a self-financing portfolio in the stock and the futures, and as such represents the FIR.

We assume there is no arbitrage and, at first, also that S pays no dividends. Accordingly S(t) equals its expected discounted future value under the standard Martingale Pricing Formula (*MPF*):[14]

$$S(t) = E_t[S(T)M(T,t)], T > t \tag{2}$$

Using this and the well-known relationship between the covariance and the expectation of a product of two random variables, the *inverse futures ratio (IFR)* is given by:

$$\begin{aligned}
\text{IFR} \equiv \frac{S(t)}{F(t,T)} &= \frac{E_t[S(T)M(T,t)]}{F(t,T)} \\
&= \frac{E_t[S(T)]E_t\left[M(T,t)\right]}{F(t,T)} + \frac{\text{cov}_t\left(S(T),M(T,t)\right)}{F(t,T)} && (3)\\
&= E_t\left[M(T,t)\right] + \frac{\text{cov}_t\left(S(T),M(T,t)\right)}{E_t[S(T)]} && (4)\\
&= E_t\left[M(T,t)\right] + \text{cov}_t\left(\frac{S(T)}{E_t[S(T)]}, M(T,t)\right) && (5)\\
&= E_t\left[M(T,t)\right]\left(1 + \text{cov}_t\left(\frac{S(T}{E_t[S(T)]}, \frac{M(T,t)}{E_t\left[M(T,t)\right]}\right)\right) && (6)
\end{aligned}$$

Equation (4) follows from Equation (3) by Equation (1). Equation (5) follows from Equation

we have used the *big* or *floor* contract because of its longer existence; in later years the volume of the *E-mini* is typically larger but the closing prices are almost always the same.

[13]Two equivalent rates are said in the market to have a *zero basis* between them. 'Equivalence' in this sense requires a static hedge, so that each rate is attainable from the other in a very simple way using market instruments. *Compare* Footnote 11 above. In the absence of arbitrage, two equivalent rates should differ by their credit risk, and as a result of variational margin and other credit mitigation features the risk associated with futures is extremely low. Were the FIR and a market rate equivalent, one would expect the FIR to track closely the risk-free rate. Figure (1) does not support such equivalence.

[14]The MPF states that the current price of a security or other self-financing trading strategy equals the conditional expectation of its discounted future cash flows, including its terminal future value, discounting by the riskless investment or other numeraire. (Anderson and Kercheval, 2010, Chapter 3). In the absence of dividends, the terminal future value is the only future cash flow, as in the equation in the text.

(4) since $E_t[S(T)]$ is known at time t. Finally, Equation (6) follows because $E_t\left[M(T,t)\right]$ is also known at time t.

The IFR is the ratio of the current stock price to the futures price and thus measures the discount implied by the futures. Equation (5) says simply that this measure equals the expected discount plus a covariance term. The first term on the right in Equation (5) is the expected term discount factor over $[t,T]$, corresponding to the FIR. The final term on the right in Equation (5) equals the covariance of the stock price and the FIR discount factor at time T around their expected values, viewed from time t.

It is instructive to compare Equation (5) with the analogous equation for a fair forward contract with forward price $\equiv F_o(t,T)$. Since the price of a fair forward contract should be zero, we have from the MPF:

$$S(t) = F_o(t,T)E_t[M(T,t)]$$

or equivalently,

$$\frac{S(t)}{F_o(t,T)} = E_t[M(T,t)] \tag{7}$$

Comparing Equations (5), (6) and (7) and recalling the definition of the IFR, we see that the covariance term may be thought of as a convexity adjustment.[15]

For empirical reasons,[16] it is easier to work with two futures contracts, the near and the

[15]Lacking a covariance term, the determination of an implied financing rate from a forward appears superficially more appealing than the more complex determination from a futures contract. More basically, an equity forward on a stock has two cash flows: a current fixed cash payment and a future equity-linked payment on the notional amount at maturity. Because there are only two cash flows, the forward-implied financing rate can be directly observed. By contrast, a futures contract has daily, often substantial cash flows in respect of variational margin, in effect marking the contract to market daily. Therefore determining the associated implied financing rate requires financial theory, as Equation (5) demonstrates. However, for credit and other reasons, forwards are OTC, not exchange traded, illiquid and bespoke, and so the contract terms, from which financing rates could be directly observed, are private and not fungible.

[16]These reasons include high price volatility, asynchronicity between the close of the cash and the futures markets, market segmentation and the substantial volume of daily transactions in both the spot and the

next,[17] and the inverse *forward* futures ratio than the inverse futures ratio itself. We will also need the forward discount $M(T_2, T_1)$. We define the *inverse forward futures ratio* or *IFFR* by:

$$IFFR \equiv F(t, T_2, T_1) \equiv \frac{F(t, T_1)}{F(t, T_2)},$$

where $T_2 > T_1$ are the maturities of the next and the near futures, respectively, and the final quotient is the near contract futures price divided by the next contract futures price, in each case observed at time t.

In this article, generally consistent with our data, we assume that

$$\Delta T \equiv T_2 - T_1 = 0.25 \quad \text{(A1)}$$

$$T_2 \leq 0.5 \quad \text{(A2)}$$

We measure time in years, so $\Delta T = 0.25$ of one year, or one-quarter.[18] By Assumption (A1), Assumption (A2) is equivalent to assuming $T_1 \leq 0.25$.

futures markets. *See generally* Anderson et al. (2013). A more detailed description of our methodology appears in the Computational Appendix.

[17] *See* Footnote 12 above.

[18] *Id.* These assumptions come from the approximately quarterly calendar of futures contracts.

3.2 The Expected Discount and the IFFR

We can now begin the derivation of our basic equation. We start with the IFFR definition:

$$
\begin{aligned}
IFFR &= \frac{F(t,T_1)}{F(t,T_2)} = \frac{E_t[S(T_1)]}{F(t,T_2)} && (8)\\
&= \frac{E_t[E_{T_1}[S(T_2)M(T_2,T_1)]]}{F(t,T_2)} = \frac{E_t[S(T_2)M(T_2,T_1)]}{F(t,T_2)} && (9)\\
&= \frac{E_t[S(T_2)]E_t[M(T_2,T_1)]}{F(t,T_2)} + \frac{\mathrm{cov}_t(S(T_2),M(T_2,T_1))}{F(t,T_2)} && (10)\\
&= E_t[M(T_2,T_1)] + \mathrm{cov}_t\left(\frac{S(T_2)}{E_t[S(T_2)]}, M(T_2,T_1)\right) && (11)\\
&= E_t[M(T_2,T_1)]\left(1 + \mathrm{cov}_t\left(\frac{S(T_2)}{E_t[S(T_2)]}, \frac{M(T_2,T_1)}{E_t[M(T_2,T_1)]}\right)\right) && (12)
\end{aligned}
$$

where Equation (9) follows from the Tower Law and Equations (11) and (12) follow from Equation (10) since $F(t,T_2) = E_t[S(T_2)]$ which is known at time t, as is $E_t[M(T_2,T_1)]$. By analogy to Equations (5) and (7), we see that the covariance term in Equation (12) is substantially the same as a *forward* forward-futures adjustment.[19]

Thus the inverse forward futures ratio $F(t,T_2,T_1)$ is the expected forward discount factor $E_t[M(T_2,T_1)]$ multiplied by a *convexity adjustment* based on a *covariance term*

$$CX_t \equiv 1 + CT_t \equiv 1 + \mathrm{cov}_t\left(\frac{S_{T_2}}{E_t[S_{T_2}]}, \frac{M(T_2,T_1)}{E_t[M(T_2,T_1)]}\right) = \frac{E_t\left[S_{T_2}M(T_2,T_1)\right]}{E_t[S_{T_2}]E_t\left[M(T_2,T_1)\right]} \qquad (13)$$

The last quotient will prove useful in computing the SDE for CX_t in our model.

Condensing, we have our basic equation:

$$IFFR = E_t[M(T_2,T_1)] \times CX_t \implies E_t[M(T_2,T_1)] = \frac{IFFR}{CX_t} \qquad (14)$$

Equation (14) expresses the unobservable expected forward FIR term discount as the observable IFFR scaled by CX_t. Because they are conditioned on $\mathscr{F}_t$, both CT_t and CX_t

[19] *Compare* Equations (5), (7) and Footnote 15 and accompanying text.

are deterministic functions of t determined by the unknown parameters of the processes M and S.

While S and IFRR are both observable,[20] M is not, and therefore CX_t is also unobservable. Equation (14) thus expresses the unobservable (expected) FIR on the left-hand side in terms of the unobservable CX_t on the right-hand side, which also depends on the unobservable $M(T_2, T_1)$. The unobservable $M(T_2, T_1)$ therefore appears on both sides of Equation (14), creating an apparent circularity in the estimation expression we wish to use - to use Equation (14) to estimate the FIR that appears by itself on the left-hand side, we must already know it to determine the denominator ($= CX_t$) on Equation (14)'s right-hand side.

Were these quantities simply unknown numbers, standard algebra might provide a solution, but instead they are time-dependent and determined from an unobservable stochastic quantity ($M(T_2, T_1)$) with unknown parameters. In Section 5.5 below, we develop a new calculus to resolve this circularity and provide an explicit estimate of CX_t based on our model.

$\mathbf{CX_t}$'s Effect Should be Insignificant We sketch a rough argument that on our data, the effect of CX_t is likely to be small and thus $E_t[M(T_2, T_1)]$ should equal the IFFR with modest error. We then preview our later result that this error is in fact insignificant outside a small minority of periods, making the FIR directly observable as a practical matter, with negligible error. *See* Figures (7) and (8).

[20]But only imperfectly. The S&P 500 is the hypothetical portfolio of stocks underlying an index, observing which raises asynchronicity and other issues. For example, the index at each instant in time is a weighted average of the most recent trade prices of the constituent stocks. Since the S&P 500 stocks are relatively large and liquid, on most days there will be at least one trade of most or all of the stocks, but there may be days on which some stocks do not trade at all, in which case the price information would be stale by at least 24 hours. The more common situation in these relatively large stocks is that the last trade occurs several minutes or even several hours before the close; for example, there are occasional trading halts in specific stocks. For different stocks, the time of the last trade may therefore differ. For another example, trading does not actually stop at the nominal close (4 pm). At the close, there will be some limit orders and buy/sell on close orders that remain to be executed. These orders are crossed, with trading times up to ten minutes after the close. To that extent the closing price of each of the individual constituent stocks may not be known at the close. *See, e.g.*, Anderson et al. (2013)

S&P 500 volatility typically averages less than 20%, interest rate volatility is approximately 1% and correlation never exceeds 1 in absolute value. CT_t measures covariance over at most six months for stock volatility and three months (one quarter) for rate volatility; over such short periods drift may generally be ignored and variation is almost entirely from volatility, which scales with the square root of time. Therefore on average and approximately

$$|CT_t| \leq 20\% \times 0.7 \times 1\% \times 0.5 \approx 7 \textit{ basis points}. \tag{15}$$

Because this is a multiplicative factor, CX_t would be less than eight additive basis points even were the FIR as high 10%.

This approximation relies on several assumptions, including that FIR volatility is similar to market interest rate volatility. Because the FIR is unobservable, this not obvious. The following sections develop a novel method for estimating FIR volatility based on a bivariate model of stock and the FIR, and use it to show the convexity adjustment materially exceeds a basis point in only a small number of periods, based on related estimates from the futures of the S&P 500's volatility and correlation with the FIR. Details are deferred to the Computational Appendix.

The practical conclusion is that the FIR can generally be determined with negligible error from the quotient of two observable futures prices, after adjustment for dividends as described in the following subsection. Investigating those few periods in which the convexity adjustment is significant, although potentially interesting, is outside the scope of this article.

3.3 General Case with Dividends

The previous equations assumed that the stock $S(t)$ pays no dividends. However, many stocks in the S&P 500 Index pay dividends that have been significant in many periods. We must therefore modify Equation (14) to reflect dividends.

We obtained implied dividend rates $r_{div}(t)$ on the S&P 500 from the option markets, based on put-call parity.[21] As derived in the Computational Appendix,[22] the required modification to Equation (14) is:

$$E_t[M(T_2, T_1)] = \frac{k_{div}(t) \times IFFR}{CX_t}. \tag{16}$$

where[23]

$$k_{div}(t) = 1 - r_{div}(t)\Delta T.$$

The value $k_{div}(t)$ lies between 0 and 1.[24] Thus in the presence of dividends the left-hand side of Equation (14) decreases by $k_{div}(t)$. As a result of this formula, we can include the effect of dividends by applying to the IFFR a discount of $k_{div}(t)$ (or equivalently increasing the forward futures ratio by $\frac{1}{k_{div}(t)}$) and then computing as though there were no dividends.

Equation (16) can be written in identical form to Equation (14) if we replace IFFR with $k_{div}(t) \times IFFR$; except where explicitly otherwise indicated, we will do so hereafter to simplify notation.

Figure (1) shows the FIR determined from the IFFR as described at the end of Section 3.2, adjusted for dividends but not convexity, expressed as a rate and appropriately annualized. This graph, with US Treasury and LIBOR rates included for comparison, provides informal, visual confirmation that the FIR behaves similarly to market interest rates, although it is not equivalent to any of them. A more precise statement requires the model in the following section.

[21] Wharton Research Data Services' OptionMetrics contains a time series of estimated implied dividend rates $r_{div}(t)$ obtained from a regression with the property that the present value at time t of the implied dividends over $(t, T]$ equals $r_{div}(t) \times S(t) \times (T - t)$, where $S(t)$ is the price of the associated stock at time t. *See* Ivy DB File and Data Reference Manual 30, Version 3.0, revised May 19, 2011.

[22] Computational Appendix [**?? ??**].

[23] Time is measured in full years, so a six-month interval measures 0.5.

[24] Since (1) ΔT is on the order of a quarter and (2) dividend rates on the S&P 500 have always been positive but less than 400% (an extremely high rate).

4 The Model of Stock and Interest Rates

From Equation (16), the FIR cannot be observed, but must be estimated. The necessary quantity to estimate is CX_t, because the IFFR is observable. We will propose a simple bivariate model for rates and the stock underlying the futures contracts that will allow this estimation. To preview, we will find the following form for $CX_t \left(= \frac{E_t\left[S_{T_2} M(T_2,T_1)\right]}{E_t[S_{T_2}]E_t[M(T_2,T_1)]}\right)$:

$$CX_t = e^{a(t)(\sigma^r)^2 + b(t)\sigma^r\sigma^S\rho} \implies \text{IFFR} = E_t\left[M(T_2,T_1)\right] \times e^{a(t)(\sigma^r)^2 + b(t)\sigma^r\sigma^S\rho} \tag{17}$$

where a(t), b(t) depend on the underlying (unobserved) vector of parameters $\theta \equiv (\sigma^r, \sigma^S, \rho)$ of our model but are otherwise deterministic. We will estimate θ along with confidence regions, and will then use these confidence regions for the vector θ to derive 95% confidence intervals for the unknown scalar CX_t. Observing that the convexity adjustments CX_t are negligible in the large majority of these confidence intervals, we will conclude that with 95% confidence, the convexity adjustments CX_t are negligible in the corresponding large majority of our observation periods, from which we will finally deduce that the FIR discount $E_t\left[M(T_2,T_1)\right]$ may generally be taken to equal the observable IFFR with negligible error, and the FIR estimated accordingly.

From Equation (17), the estimate $\widehat{CX_t}$ follows from estimates for $\widehat{(\sigma^r)^2}, \widehat{\sigma^S}, \widehat{\rho}$. The balance of this section will described our model and the estimation of $\widehat{(\sigma^r)^2}$; estimation of $\widehat{\sigma^S}, \widehat{\rho}$ will be summarized and the details deferred to the Computational Appendix.

4.1 Specification

We rely on a Ho-Lee interest rate model and a standard stock lognormal model. In matrix form, it is

$$\begin{bmatrix} dr_t \\ \frac{dS_t}{S_t} \end{bmatrix} = \begin{bmatrix} \mu^r dt + \sigma^r dB_t^r \\ r_t dt + \sigma^S dB_t^S \end{bmatrix} \tag{18}$$

with $\mathbf{dB}_t \equiv \left(dB_t^S, dB_t^r\right)$ a two-dimensional Brownian motion with instantaneous correlation ρ and $r_t \equiv \mathrm{FIR}(t)$. The parameters $\mu^r, \sigma^r, \sigma^S, \rho$ are all constant.

4.2 Preliminary Results for Later Use

Restricting to our model, we have the following useful technical propositions. The proofs are deferred to the Computational Appendix.

Let F(t,T) denote the futures price at time t under a futures contract on the stock S_t with expiration at time T.

Proposition: 4.1. *Stock Price Evolution*

$$S(T) = S_t e^{\int_t^T \left(r_s - \frac{\sigma^{S2}}{2}\right) ds + \int_t^T \sigma^S B_s^S} \tag{19}$$

The next proposition, which follows from the first and Theorem (3.1), describes the futures price evolution in our model.

Proposition: 4.2. *Futures SDE:*

$$\frac{dF(t,T)}{F(t,T)} = (T-t)\sigma^r dB_t^r + \sigma^S dB_t^S \tag{20}$$

The third proposition describes the FIR discount evolution in our model.

Proposition: 4.3. *FIR Discount Evolution*

$$\frac{dE_t\left[M(T_2,T_1)\right]}{E_t\left[M(T_2,T_1)\right])} = (t \vee T_1 - T_2)\sigma^r dB_t^r, T_2 \geq T_1, t \tag{21}$$

From Proposition 4.3, consistent with one's intuition, $E_t\left[M(T_2,T_1)\right]$ continues to evolve for $T_1 < t < T_2$, but with diminishing volatility.

The final proposition provides the explicit formula for CX_t discussed at the beginning of this section.

Proposition: 4.4. *Convexity Evolution*

$$\begin{aligned}
&E_t[M(T_2,T_1)S_{T_2}] = E_t\left[M(T_2,T_1)\right] E_t[S_{T_2}]e^{U_1(T_2,\theta)-U_2(t,\theta)} \\
&\textit{where } U_1(T_2,\theta) - U_2(t,\theta) \\
&= \frac{(\sigma^r)^2}{4}\frac{1}{2}\left(\frac{1}{48} - (T_2-t)^2\right) - \frac{\sigma^r}{4}\sigma^S\rho\left(T_1 - t + \frac{1}{8}\right) \\
&\equiv a(t)\left(\sigma^r\right)^2 + b(t)\sigma^r\sigma^S\rho
\end{aligned}$$

Proposition 4.4 follows from Propositions (4.1) and (4.3). Details are deferred to the Computational Appendix.

It will be helpful to express Proposition 4.4 as a quadratic form:

$$U_1(T_2,\theta) - U_2(t,\theta) = \begin{bmatrix}\sigma^r & \rho\sigma^S\end{bmatrix} \mathbf{Q}(t) \begin{bmatrix}\sigma^r \\ \rho\sigma^S\end{bmatrix} \tag{22}$$

$$\mathbf{Q}(t) \equiv \begin{bmatrix} \frac{1}{8}\left(\frac{1}{48} - (T_2-t)^2\right) & \frac{1}{8}\left(T_1 - t + \frac{1}{8}\right) \\ \frac{1}{8}\left(T_1 - t + \frac{1}{8}\right) & 0 \end{bmatrix}$$

The form $\mathbf{Q}$ is deterministic and can be computed once, independent of the particular (next, near) pair of futures for which the convexity term is estimated.

4.3 Canonical Coordinates for Brownian Diagonalization

Let $F_i, i = 1, 2$ be two futures (a *futures pair*) on the stock S_t with sequential maturities (expiration dates) $T_1 < T_2$. From Proposition 4.2:

$$\frac{dF_i}{F_i} = (T_i - t)\sigma^r dB^r + \sigma^S dB^s, i = 1, 2$$

We wish to diagonalize these two SDEs in terms of the bivariate Brownian motion $\mathbf{dB}_t$.

Using a change of variables to isolate the Brownian motion in terms of stochastically-integrable quantities

$$X = \frac{\log F_2 - \log F_1}{T_2 - T_1},$$
$$Y = \frac{T_2 - t}{T_2 - T_1} \log F_1 - \frac{T_1 - t}{T_2 - T_1} \log F_2$$

Note that $X = \frac{1}{T_2 - T_1} \times log\left(\frac{F_2}{F_1}\right)$ and so functions as a proxy for the IFFR. X, Y are both observable and have the following SDEs:

$$\begin{bmatrix} dX \\ dY \end{bmatrix} = \begin{bmatrix} \sigma^r dB^r + dg(t;\theta) \\ \sigma^S dB^s + dh(t;\theta) + X dt \end{bmatrix} \tag{23}$$

where $\theta = (\sigma^r, \sigma^S, \rho), dB^r dB^s = \rho$ and

$$g(t;\theta) = -\frac{1}{2}\int_0^t \left\{ \frac{(T_2 - s)^2 - (T_1 - s)^2}{T_2 - T_1} (\sigma^r)^2 + 2\sigma^r\sigma^S\rho \right\} ds,$$
$$h(t;\theta) = \frac{1}{2}\int_0^t \left\{ (T_2 - s)(T_1 - s)\,(\sigma^r)^2\, ds - (\sigma^S)^2 \right\} ds$$

g, h are deterministic but depend on θ.

Integrating Equation (23), rearranging and using the approximation[25]

$$\int X dt \approx \frac{1}{2}(X_{t+\Delta t} + X_t)\Delta t \equiv X_{Avg}\Delta t$$

we may discretize Equation (23):

$$\begin{bmatrix} \Delta X \\ \Delta Y \end{bmatrix} = \begin{bmatrix} \sigma^r \Delta B_t^r - (\sigma^r)^2 (T_1 - t)\Delta t \\ \sigma^S \Delta B_t^S + X_{Avg}\Delta t - \frac{1}{2}(\sigma^r)^2 \Delta t\,(t - T_1)(t - T_2)) \end{bmatrix} + \begin{bmatrix} c \end{bmatrix} \tag{24}$$

where c is a (vector) constant depending on θ.[26]

We rewrite Equation (24) in terms of the Brownian motion increments:

$$\begin{bmatrix} \sigma^r \Delta B_t^r \\ \sigma^S \Delta B_t^S \end{bmatrix} = \begin{bmatrix} \Delta X + (\sigma^r)^2 (T_1 - t)\Delta t \\ \Delta Y - X_{Avg}\Delta t + \frac{1}{2}(\sigma^r)^2 \Delta t\,(t - T_1)(t - T_2)) \end{bmatrix} + \begin{bmatrix} c \end{bmatrix} \tag{25}$$

Equation (25) illustrates, in the context of our model, the circularity referred to in Section 3.2 above. If we knew the unknown parameter $(\sigma^r)^2$ we could determine the right-hand side, up to the unknown constant, detrending the observations of ΔX and ΔY and leaving a series of fully-observed iid samples of the bivariate Brownian increments on the left-hand side. We could then estimate the covariance matrix of the left-hand side with the usual methods. However, the parameter $(\sigma^r)^2$ we need to estimate the covariance matrix in this manner is itself one of the entries the covariance matrix we wish to estimate. The next section describes the modification required to the usual methods to resolve this circularity.

[25] Since we observe both $X_{t+\Delta t}, X_t$, X is based on a Brownian Bridge. [The Computational Appendix provides further justification for this approximation.]

[26] Herein 'c' denotes a constant, generally differing based on the context.

5 Estimating the Model

5.1 Objective and Methodology for Model Estimation

Objective. Our basic equation for the FIR is Equation (17). In that equation, the IFFR is observed directly, but from Proposition 4.4 and Equation (22) the convexity adjustment that multiplies the IFFR to determine the FIR discount depends on the two unknown quantities, a variance and a covariance $(\sigma^r)^2, \sigma^r \sigma^S \rho$, with the observed deterministic coefficients given in Proposition 4.4.

Equation (1) relies on the absence of arbitrage, as does our model at certain other points, so the relevant variance and covariance derive from the risk-neutral measure. Instead, we estimate these quantities in the physical (real-world) measure, for two reasons. First, the FIR is not directly observable and so the traditional Breeden-Litzenberger framework does not extend naturally to provide from a schedule of prices for options on S&P 500 futures contracts an estimate of the underlying bivariate stock-rate distribution.[27] Second, more practically, options on the S&P 500 futures are illiquid at many strikes[28] rendering the Breeden-Litzenberger estimates unreliable.

By contrast, the S&P 500 futures trade liquidly in substantial volume and so we have based our estimates on their closing prices.[29]

To estimate confidence intervals for the FIR thus requires confidence regions for these two quantities, over which we may take the maximum and minimum of Equation (22) to determine a confidence interval for the convexity adjustment. We will see empirically that

[27] *Cf.* Equation (18). The Breeden-Litzenberger framework was first described in Breeden-Litzenberger (1978).

[28] For example, the CME Volume and Open Interest report on December 5, 2020 showed the average closing volume over available strikes of the Dec 20 puts was less than 2% of the volume for the underlying futures contract. The average volume for the Dec 20 calls was significantly less than that of the puts.

[29] Although implied and realized volatility of course differ from each other, the former has been viewed as predicting the latter. *See, e.g.,* Poon and Granger (2003). In this article, we are concerned primarily with average levels of volatility and covariance over substantial periods, over which the two may reasonably be viewed as approximate surrogates for each other.

outside a small minority of periods the convexity adjustment is negligible with 95% confidence across all futures pairs and periods; therefore the FIR may be estimated directly from the observed IFFR with a high degree of confidence.

In the final sections of this article, we will use these estimates over our entire observation period, more than two decades long, to identify the four regimes mentioned in the introduction.

Methodology. We proceed by (i) estimating, for each period and corresponding futures pair, an associated confidence region for the three parameters in θ, and then (ii) finding the corresponding confidence interval for the convexity adjustment by maximizing and minimizing that adjustment over the confidence region for the period.

In the following subsection, we start by deriving, again for each period and corresponding futures pair, a point and confidence interval estimate for the associated FIR variance $\equiv (\sigma^r)^2$ in that period. In the ensuing subsections, we use that point estimate $\widehat{(\sigma^r)^2}$ to detrend Equation (25) and estimate a two-dimensional confidence region for $\sigma^S, \rho \mid \widehat{(\sigma^r)^2}$ conditional on that point estimate, as the dependence of the drift term in Equation (25) on the unknown parameter $(\sigma^r)^2$ suggests. We then use Equation (22) and the previously-computed[30] matrices $\mathbf{Q}(t)$ at each time to maximize and minimize the convexity adjustment over that conditional confidence region, deriving a conditional confidence interval for the convexity adjustment supporting the conclusion that, conditional on the estimate $\widehat{(\sigma^r)^2}$, the convexity adjustment is typically negligible over our entire observation period with similar confidence.

In the Computational Appendix, we vary the estimates $\widehat{(\sigma^r)^2}$ over a $(\sigma^r)^2$ confidence interval to estimate unconditional maxima and minima for the convexity adjustment and show that, outside a few periods, it remains negligible with a high degree of confidence.

[30] *See* the final paragraph in Section 4.2 *supra*.

5.2 Estimating $(\sigma^r)^2$

Throughout this subsection N_a, N_q denote the observations per year and per quarter, generally around 248 and 62, respectively.

The first row of Equation (24) implies:

$$\Delta X \in N\left((\sigma^r)^2 (t - T_1)\Delta t + c_1, \frac{(\sigma^r)^2}{N_a}\right)$$

To derive iid random variables for estimation, we rewrite this as

$$\Delta X + (\sigma^r)^2 (T_1 - t)\Delta t \in N\left(c_1, \frac{(\sigma^r)^2}{N_a}\right) \tag{26}$$

Taking the unbiased sample variance of the left-hand side of Equation (26),

$$N_a \times s^2 \left(\Delta X + (\sigma^r)^2 (T_1 - t)\Delta t\right) \in (\sigma^r)^2 \chi^2_f \tag{27}$$

where we annualize based on N_a observations per year. We estimate quarterly using N_q observations $\implies f \equiv N_q - 1$ degrees of freedom for the estimated variance of these iid normal random variables with unknown means.

Equation (27) exhibits the circularity problem described previously: to estimate the sample variance of the iid samples in Equation (26), we first need to detrend the observable series ΔX, to do which requires the variance we're trying to estimate.

To express this circularity in more conventional statistical terms, define *the conditional sample variance estimate* by

$$s^2_{a,f}|\sigma^2 \equiv \frac{N_a}{f} \times s^2 \left(\Delta X + \sigma^2(T_1 - t)\Delta t\right) \tag{28}$$

$s^2_{a,f}|\sigma^2$ is thus the ordinary annualized sample variance estimate[31] after detrending the time series ΔX using (conditional on) σ^2, the true variance. $s^2_{a,f}|\sigma^2$ is not a statistic because of its dependence on the unknown σ^2; however, a simple modification of the usual methods nonetheless provides a solution.

To simplify the notation in the remainder of this section, we will dispense with the superscript denoting the FIR variance and set $\sigma^2 \equiv (\sigma^r)^2$. We divide both sides of Equation (27) by σ^2 to get a draw d from a standard sampling distribution:

$$d(\sigma^2) \equiv \frac{f \times s^2_{a,f}|\sigma^2}{\sigma^2} \in \chi^2_f \tag{29}$$

Here, σ^2 is the true variance and Equation (29) is a conventional frequentist statement about repeated samples ΔX. Letting $\sigma^2 = \sigma_0^2$ be a particular null hypothesis about the value of σ^2, the right-hand side of Equation (29) becomes a *test statistic* in the usual way:

$$\mathcal{T}(\sigma_0^2) \equiv d(\sigma_0^2) = \frac{f \times s^2_{a,f}|\sigma_0^2}{\sigma_0^2} \in \chi^2_f \tag{30}$$

Using the sampling distribution χ^2_f we can as usual determine critical regions for any significance level α and the p-value of the observed test static $\mathcal{T}(\sigma_0^2)$ but we do not have a conventional point estimate $\widehat{\sigma^2}$ and so cannot determine confidence intervals in the conventional manner. We can only construct such a point estimate conditional on a null value σ_0^2 to use for detrending, another manifestation of the circularity referred to in Section 3.2 above.

We now derive point estimates in a novel manner, using a method related to fiducial inference.[32] Note that $s^2_{a,f}|\sigma_0^2$ is the ordinary annualized sample variance estimate after detrending the time series ΔX using (conditional on) a particular hypothesized null value

[31]The 'a' subscript denotes annualization and reflects multiplication by N_a, and the 'f' subscript denotes the number of observations $N_q - 1$. Although we know the true mean is 0, there is an unknown constant in Equation (25) to be estimated, implying $N_q - 1$ degrees of freedom, by which we must divide to determine the sample estimate.

[32]Fisher (1935).

σ_0^2. As previously remarked, we do not in our context have a single sample variance estimate $\frac{N_a}{f} \times s^2()$; rather we have a continuum of estimates $s^2_{a,f}|\sigma_0^2$, each conditional on a different null value $\sigma^2 = \sigma_0^2$ used to detrend the observations.

Letting the variable u_{σ^2} represent the possible null values σ_0^2, we may view $\mathcal{T} : \mathbb{R}^{++} \to \mathbb{R}^{++}$ as a function $\mathcal{T}(u_{\sigma^2})$ of a positive real variable u_{σ^2}. In the Computational Appendix, we show that this function $\mathcal{T}(u_{\sigma^2})$ has a restricted inverse[33] $u_{\sigma^2}(d) \equiv \mathcal{T}^{-1}(d) : [\mathcal{T}_{min}, \infty) \subset \mathbb{R}^{++} \to (0, u_{\sigma^2,argmin}] \subset \mathbb{R}^{++}$ (these domain and range intervals are defined in the Computational Appendix) mapping values $d \in [\mathcal{T}_{min}, \infty)$ of the test statistic bijectively to (possible) null values in its range.[34]

We will assume, and confirm empirically, that the intervals $[\mathcal{T}_{min}, \infty), (0, u_{\sigma^2,argmin}]$ are large enough, meaning that the latter interval is likely to contain any reasonable hypotheses $(\sigma_0^r)^2$ consistent with our data. For example, $f > \mathcal{T}_{min}$ in all cases on our data, as will be important for the point estimate below.

The standard χ^2_f distribution for the test statistic has distinguished values (such as the mean, the mode and quantiles) by virtue of being a distribution, and each such value d in the domain of $\mathcal{T}^{-1}$ then corresponds to a particular null value $u_{\sigma^2}(d)$ that in turn gives rise as described below to a particular *conditional sample variance estimate given the test statistic value d*, which we will shorten to *conditional estimate*, $\widehat{\sigma^2}|d \equiv u_{\sigma^2}(d)$.[35] In general, $\widehat{\sigma^2}|d \neq s^2_{a,f}|u_{\sigma^2}(d)$; the sample estimate typically differs from the true value in conventional

[33] As used herein : means "such that".

[34] One could find an approximate inverse for a given d from Equation (??) by simple trial-and-error. In the Computational Appendix we derive a closed-form solution. To see that some restriction on $\mathcal{T}$ is needed, note that $\mathcal{T} \to \infty$ as $\sigma_0^2 \to 0, \infty$. Thus, $\mathcal{T}()$ is not univalent, but since $\mathcal{T} > 0, \mathcal{T}$ has a global minimum $\mathcal{T}_{min}$ attained by a $(\sigma^r)^2_{0,argmin}$. $\mathcal{T}$ is univalent above and below $(\sigma^r)^2_{0,argmin}$. There is no *a priori* measure on null values $\sigma_0^2 \in R^+$, but we can induce one (a *pull-back measure*) in a standard way from the χ^2_f distribution via $\mathcal{T}$, once restricted to be univalent, under which $\mathcal{T}$ will preserve measure by construction: the pull-back measure is the push-forward measure associated with $\mathcal{T}^{-1}$. Details are left to the Computational Appendix.

[35] The notation is intended to emphasize that, rather than starting with a null hypothesis σ_0^2, we start with a draw $d \in \chi^2_f$ and then treat the associated null hypothesis $u_{\sigma^2}(d)$ as an estimate corresponding to d. The pull-back measure on null values from Footnote 34 provides further motivation for constructing these estimates; under this measure $\mathcal{T}$ is measure-preserving by construction, and the significance levels of conventional confidence intervals are determined by their pull-back measure in a natural way.

statistics, and the same is true in our context.

A Point Estimate. The mean of a χ^2_f variable is f, which treated as a draw $d = f$ from χ^2_f corresponds to the estimate $\widehat{\sigma^2}|f \equiv u_{\sigma^2}(f) \equiv \mathcal{T}^{-1}(f)$. From Equation (**??**) with $\sigma_0^2 = u_{\sigma^2}(f), \mathcal{T}(u_{\sigma^2}(f)) = f \iff \widehat{\sigma^2}|f \equiv u_{\sigma^2}(f) = s^2_{a,f}|u_{\sigma^2}(f)$. Thus, the estimate $\widehat{\sigma^2}|f$ corresponds to a conditional sample variance estimate $s^2_{a,f}|u_{\sigma^2}(f)$ that itself equals that $\widehat{\sigma^2}|f$, which we can therefore think of as a fixed point estimate. An equivalent description of this fixed point estimate that generalizes well to the multivariate case is to think of it as the *zero error* estimate, explained as follows. We may think of the quantity $\frac{s^2_{a,f}|u_{\sigma^2}}{u_{\sigma^2}} - 1$ as a measure of the error (the *null error*)[36] in our sample variance estimate $s^2_{a,f}|u_{\sigma^2}$ relative to a hypothesized null value u_{σ^2}.

In the conventional case where the sample variance estimate does not depend on the null value and we have simply

$$\Delta X \in N\left(c_1, \frac{(\sigma^r)^2}{N_a}\right) \implies s^2_{a,f}|u_{\sigma^2} = \frac{N_a}{f} \times s^2(\Delta X) \equiv \widehat{\sigma^2}$$

independent of any null hypothesis (equivalently, independent of any choice for the value for u_{σ^2}). The zero error estimate is then just the (conventional) sample variance estimate itself: $\frac{s^2_{a,f}|u_{\sigma^2}(f)}{u_{\sigma^2}(f)} - 1 \equiv \frac{\widehat{\sigma^2}}{u_{\sigma^2}(f)} - 1 = 0 \iff u_{\sigma^2}(f) = \widehat{\sigma^2}$.

$\widehat{\sigma^2}|f$, which can be defined quite generally, thus reduces to the conventional sample variance estimate in the conventional case. This observation is the intuition behind viewing $\widehat{\sigma^2}|f$ as a generalization of the conventional sample variance estimate to more general contexts, such as the one in this article.[37]

[36] The null error varies generally with the p-value from statistics, and thus adds intuitive content to the probability measure on null values constructed abstractly in Footnote 34.

[37] This observation also applies to $\widehat{\sigma^2}|d, \forall d \in \chi^2_f$. For another example of a conditional estimate, the mode of a χ^2_f distribution is the draw d = f - 2, which from Equation (**??**) corresponds to the estimate $s^2_{a,f}|u_{\sigma^2}(f-2) = \frac{f-2}{f} \times u_{\sigma^2}(f-2)$, an estimate with (nonzero) error $\frac{s^2_{a,f}|u_{\sigma^2}(f-2)}{u_{\sigma^2}(f-2)} - 1 = -\frac{2}{f}$. In this manner, every draw from the sampling distribution χ^2_f corresponds to a different estimate of σ^2 with a different error.

Remark. To reduce the possibility of confusion, $\widehat{\sigma^2}|d$ means in words "given $d \in \chi^2_f$, $\widehat{\sigma^2}|d$ is the (unique) value for σ^2 such that, when it is used both to detrend and to divide in Equation (29), the resulting test statistic (i.e., the quotient) equals d." The "hat" means that it is an estimate, as discussed above; however, as mentioned above, it generally will not equal the conditional sample variance estimate $s^2_{a,f}|\left(\widehat{\sigma^2}|d\right)$. Indeed, $\widehat{\sigma^2}|d = s^2_{a,f}|\left(\widehat{\sigma^2}|d\right)$ if and only if $d = f$; that is to say, only for the zero error estimate. It is an estimate, not a null value, but, when taken as a null value, the resulting test statistic from Equation (**??**) equals d, by construction.

Interval Estimates. The zero error estimate $\widehat{\sigma^2}|f$ above provides a distinguished point estimate. For interval estimates, given a (null) hypothesis $\sigma^2 = \sigma_0^2$ and further a significance level α, the test statistic in Equation (30) determines, through a modification of the usual methods, an interval in which the true value of sigma is likely to lie with the corresponding significance. The Computational Appendix describes the needed modification.

In broad overview, the standard $(1-\alpha)$ χ^2_f confidence interval for the quotient $\mathcal{T}(\sigma_0^2)$ on the right-hand side of Equation (30) is $I^d_\alpha \equiv [Q_{\chi^2_f}(\frac{\alpha}{2}), Q_{\chi^2_f}(1-\frac{\alpha}{2})]$, where $Q \equiv$ the *quantile function.* (For notational convenience, references to these intervals omit the multiplier $f\times$ where the context makes it clear.) The preimage $\mathcal{T}^{-1}(I^d_\alpha)$ of I^d_α under $\mathcal{T}()$

$$I^{\sigma^2}_\alpha \equiv \{\sigma_0^2 : \mathcal{T}(\sigma_0^2) \in I^d_\alpha\}$$

defines the corresponding confidence interval[38] for σ^2. Because the median of χ^2_f is approximately its mean, $\widehat{\sigma^2}(f) \in I^{\sigma^2}_\alpha$ for reasonably small significance levels α. Details are left to the Computational Appendix.

Standard Confidence Interval. In Section 5.4 below, we will be interested in a three-

[38]If an inverse exists, the preimage is the image under that inverse. By construction, $\mathcal{T}$ preserves measure with the pull-back measure on its domain $\subset \mathbf{R}^{++}$, so the measure of a confidence interval is equivalent to its significance. See Footnote 34.

dimensional confidence region for all of $\theta \equiv \sigma^r, \sigma^S, \rho$ with measure $95\% < (98.4\%)^3$. A *standard confidence interval* ${}^0I^d \equiv I^d_{1.6\%}$ will therefore be of particular interest to us. We standardize[39] f $\equiv$ number of observations - 1 = 60 $\implies$ ${}^0I^d \equiv [0.84, 1.41]$, following the convention that the interval's halves around its center at 1 should have equal weight.[40] We then have the associated standard confidence interval for σ^2 : ${}^0I^{\sigma^2} \equiv \mathcal{T}^{-1}\left({}^0I^d\right)$.

5.3 Is the Estimate $\widehat{\sigma^2}|d$ Reasonable?

We have provided a theoretical construction of new estimates $\{\widehat{\sigma^2}|d, d \in \chi^2_f\}$. Does this construction produce reasonable estimates in practice?

The answer in the context of our futures data is a qualified "yes". As discussed below, the new zero error estimate agrees closely with the conventional MLE estimate from our data, as well as a naive estimate that assumes the drift is small enough to be ignored. The latter observation raises the qualification, since when applied to our data the new estimate cannot therefore be shown to be definitively superior to that naive estimate.

MLE. Maximum likelihood estimation (*MLE*) offers a different, more conventional method for estimating the three unknown parameters σ^r, σ^S, ρ jointly.[41] The Computational Appendix has some details of our MLE of these parameters. As shown in Figure (2), our zero error estimate $\widehat{(\sigma^r)^2}|f$ is close to the bivariate maximum likelihood estimate $\widehat{\sigma^r}_{MLE}$, but typically exceeds $\widehat{\sigma^r}_{MLE}$ when they differ. Following convention we graph volatilities not variances. Using estimates of σ^S, ρ from the next subsection, Figure (3) extends this comparison to all of θ. The similarity of these two different estimates provides some corroboration of our estimate $\widehat{(\sigma^r)^2}|f$.

[39] An idealized quarter with three months of 20 trading days each, plus one additional trading day.

[40] As discussed above, we choose this confidence interval because there are three parameters to estimate and $98.4\%^3 > 95\%$.

[41] MLE is significantly more computationally expensive than our method $\widehat{(\sigma^r)^2}|f$ described in the text. In our context, MLE is computationally-tractable only because (X,Y) in our model are jointly Markov, as is X, but not Y, alone.

Figure (4) shows the zero error and MLE estimates inside a 98.4% confidence interval for the FIR vol σ^r computed using our methods.[42] The width of the confidence interval averaged 36 basis points over our observation period.

Driftless Approximation. Another confirmation of the reasonableness of $\widehat{(\sigma^r)^2}|f$ is provided by near-equality of $\widehat{(\sigma^r)^2}|f$ and the conventional driftless estimate $\frac{N_a}{f} \times s^2(\Delta X)$, as discussed in the Computational Appendix. As mentioned above, this confirms the variation in ΔX exceeds the effect of the associated trend by a sufficient order of magnitude that the trend may be ignored.

5.4 Estimating σ^S, ρ

We estimate σ^S, ρ in a similar way to estimating σ^r. Specifically, we will extend the estimation methods of Section 5.2 to all of θ, and the resulting vector estimation will return the $(\sigma^r)^2$ estimates of Section 5.2 as one entry in an estimated covariance matrix.

Extending Equation (26) via Equation (25), we have

$$\begin{bmatrix} \Delta X + (\sigma^r)^2 (T_1 - t)\Delta t \\ \Delta Y - X_{Avg}\Delta t + \frac{1}{2} (\sigma^r)^2 \Delta t (t - T_1)(t - T_2) \end{bmatrix} \in N\left(\begin{bmatrix} c_1 \\ c_2 \end{bmatrix}, \frac{1}{N_a}\mathbf{\Sigma} \equiv \begin{bmatrix} (\sigma^r)^2 & \rho\sigma^r\sigma^S \\ \rho\sigma^r\sigma^S & (\sigma^S)^2 \end{bmatrix} \right) \tag{31}$$

with iid observations on the left-hand side.

Taking the unbiased sample vector covariance of the left-hand side we extend Equation (27):

$$N_a \times S^2 \left(\begin{bmatrix} \Delta X + (\sigma^r)^2 (T_1 - t)\Delta t \\ \Delta Y - X_{Avg}\Delta t + \frac{1}{2} (\sigma^r)^2 \Delta t (t - T_1)(t - T_2) \end{bmatrix} \right) \in \mathcal{W}_2(\mathbf{\Sigma}, f) \tag{32}$$

[42] Such confidence intervals are computed exactly in constant time. By contrast, MLE confidence intervals must be bootstrapped or determined asymptotically, neither of which are sufficiently accurate in this context.

where the right-hand side is a draw from a *Wishart* distribution with (scale) parameter $\mathbf{\Sigma}$, dimension 2 and f degrees of freedom. We have used an upper-case S to indicate a standard vector sample covariance estimate on the left-hand side.

Given a bivariate null hypothesis on the covariance matrix $\mathbf{\Sigma} = \mathbf{\Sigma}_0$, the left-hand side of Equation (32) depends on $\mathbf{\Sigma}_0$ only through $(\sigma_0^r)^2 = \mathbf{\Sigma}_0[0,0]$. We generalize the definition of $s_{a,f}^2|\sigma_0^2$ in Equation (28) by defining

$$S_{a,f}^2 \mid \mathbf{\Sigma}_0 \equiv \frac{N_a}{f} \times S^2(\Delta X, \Delta Y, X_{Avg}, (\sigma_0^r)^2),$$

so we may write Equation (32) subject to the null as

$$\left(f \times S_{a,f}^2 \mid \mathbf{\Sigma}_0\right) \in \mathcal{W}_2\left(\mathbf{\Sigma}_0, f\right). \tag{33}$$

To standardize the sampling distribution we pre- and post-multiply by the factors of the Cholesky decomposition[43] of $\mathbf{\Sigma}_0 \equiv \mathbf{L}_{\mathbf{\Sigma}_0}$ to get

$$\mathbf{D}(\mathbf{\Sigma}_0) \equiv f \times \mathbf{L}_{\mathbf{\Sigma}_0}^{-1}\left(S_{a,f}^2 \mid \mathbf{\Sigma}_0\right)\mathbf{L}_{\mathbf{\Sigma}_0}^{-T} \in \mathcal{W}_2\left(\mathbb{1}(2), f\right) \tag{34}$$
$$\mathbf{L}^{-T} \equiv \left(\mathbf{L}^{-1}\right)^T = \left(\mathbf{L}^T\right)^{-1}$$

In Equation (34), ${}^2\mathcal{T}(\mathbf{\Sigma}_0) \equiv \mathbf{D}(\mathbf{\Sigma}_0)$ is a bivariate test statistic generalizing the univariate test statistic in Equation (30). We follow our practice of naming the particular draw from the standard sampling distribution that our data represents, in this case with an upper-case $\mathbf{D}$. Because of our choice of the Cholesky decomposition as matrix square root, this

[43]The *Cholesky decomposition* of a real, symmetric and positive definite matrix $\equiv \mathbf{M} \in \mathbf{S}_+^2$ is a lower-triangular matrix $\mathbf{L_M}$ such that $\mathbf{L_M}\mathbf{L_M}^T = \mathbf{M}$, i.e. $\mathbf{L_M}$ is a lower-triangular square root of $\mathbf{M}$. Any square root of $\mathbf{\Sigma}_0$ will standardize the sample covariance in the same manner; the Cholesky decomposition fits particularly well with our prior estimation of $(\sigma^r)^2$, as shown below.

generalization can be made more specific. We have[44]

$$\mathbf{D}[0,0] = \left({}^2\mathcal{T}(\boldsymbol{\Sigma})\right)[0,0] = \mathcal{T}\left(\boldsymbol{\Sigma}[0,0]\right) = \mathcal{T}\left(\sigma^2\right) \tag{35}$$

Otherwise put, the following diagram commutes, where the vertical arrows are projection $\pi_{0,0}$ of a matrix onto its first entry $[0,0]$ and the horizontal arrows are the test statistics ${}^2\mathcal{T}(), \mathcal{T}()$, respectively:

$$\begin{array}{ccc} \boldsymbol{\Sigma}_0 & \xrightarrow{{}^2\mathcal{T}()} & \mathbf{D} \\ \downarrow{\scriptstyle \pi_{0,0}} & & \downarrow{\scriptstyle \pi_{0,0}} \\ \sigma_0^2 & \xrightarrow{\mathcal{T}()} & d \end{array} \tag{36}$$

Restricting Diagram (36) to the domains of $\mathcal{T}, {}^2\mathcal{T}$ and their chosen inverses, we have, by elementary diagram-chasing,[45] that the following diagram also commutes.

$$\begin{array}{ccc} \boldsymbol{\Sigma}_0 & \xleftarrow{{}^2\mathcal{T}^{-1}()} & \mathbf{D} \\ \downarrow{\scriptstyle \pi_{0,0}} & & \downarrow{\scriptstyle \pi_{0,0}} \\ \sigma_0^2 & \xleftarrow{\mathcal{T}^{-1}()} & d \end{array}$$

From these two diagrams we see that conditioning on a value $\sigma_0^2 = \boldsymbol{\Sigma}_0[0,0]$ conditions on $\mathbf{D}[0,0] = d$ and conversely.

In the Computational Appendix, we compute an explicit inverse $\boldsymbol{\Sigma}_0(\mathbf{D}) \equiv {}^2\mathcal{T}^{-1}(\mathbf{D}) : \mathbf{S}^2_+ \to \mathbf{S}^2_+$ for this test statistic mapping symmetric positive semi-definite matrix values $\mathbf{D} \in \mathcal{W}_2\left(\mathbb{1}(2), f\right)$ of the test statistic to null values $\boldsymbol{\Sigma}_0 = \boldsymbol{\Sigma}_0(\mathbf{D})$.[46] Similarly to the univariate case

[44] As discussed in the Computational Appendix, this follows because the entry in the first row and column of the product of a matrix and an upper-triangular matrix depends only on the entries in the first rows and columns of the two factor matrices, as is also true of the product of a lower-triangular matrix and another matrix.

[45] Meaning starting at the top right with $\mathbf{D}$, following the top arrow to $\boldsymbol{\Sigma}_0$ and then using that Diagram (36) commutes.

[46] After suitable restriction of the domain and range $\subset \mathbf{S}^2_+$ by analogy to Section 5.2. As before, this inverse induces a measure on null values $\boldsymbol{\Sigma}_0$ in the domain and we assume that the domain and range subsets of $\mathbf{S}^2_+$ are sufficiently large for our purposes and confirm this assumption empirically. *See* Footnotes 34, 35 and accompanying text.

discussed above, the standard $\mathcal{W}_2(\mathbb{1}(2), f)$ distribution for the test statistic has distinguished values and each such value $\mathbf{D}$ then corresponds to a particular null value $\mathbf{\Sigma}_0(\mathbf{D})$, providing a *conditional sample covariance estimate given the test statistic value* $\mathbf{D}$, which we will shorten to *conditional estimate*, $\widehat{\mathbf{\Sigma}} \mid \mathbf{D}$.

A Point Estimate. The mean of the standard Wishart $\mathcal{W}_2(\mathbb{1}(2), f)$ equals $f \times \mathbb{1}(2)$, and therefore choosing $\mathbf{D}_f \equiv f \times \mathbb{1}(2)$ yields the *bivariate zero error estimate* $\widehat{\mathbf{\Sigma}} \mid \mathbf{D}_f$ by analogy to $\widehat{\sigma^2}|f$ in the univariate case. From Equation (35) above, $\left(\widehat{\mathbf{\Sigma}} \mid \mathbf{D}_f\right)[0,0] = \widehat{(\sigma^r)^2}|f$.

Region Estimates. From Equation (34), if $\mathbf{R}^{\mathbf{D}}_\alpha \ni \mathbf{D}_f$ is a confidence region for $\mathcal{W}_2(\mathbb{1}(2), f)$ with significance α, then $\mathbf{R}^{\Sigma}_\alpha \equiv \mathbf{\Sigma}_0\left(\mathbf{R}^{\mathbf{D}}_\alpha\right) = {}^2\mathcal{T}^{-1}\left(\mathbf{R}^{\mathbf{D}}_\alpha\right) \ni \mathbf{\Sigma}_0(\mathbf{D}_f)$ is a confidence region for $\widehat{\mathbf{\Sigma}} \mid \mathbf{D}_f$ with the same significance.[47]

As shown in the Computational Appendix, the density for a $\mathcal{W}_2(\mathbb{1}(2), f)$ is a product of three densities: two χ^2_f densities corresponding to the two variances of the associated bivariate Gaussian, and a density $\equiv \mathcal{D}_\rho$, up to a normalization factor, of the form $p(\rho) = (1-\rho^2)^{\frac{f-3}{2}}$ corresponding to their correlation.[48] We may therefore choose for our standard $\mathcal{W}_2(\mathbb{1}(2), f)$ confidence region $\equiv {}^0\mathbf{R}^{\mathbf{D}}_{5\%}$ a rectangular solid equal to a product of three $I_{1.6\%}$ confidence intervals corresponding to each of (i) the FIR variance, (ii) the stock variance and (iii) their correlation.[49]

As we did with the interval estimate for $(\sigma^r)^2$, we seek the three-dimensional confidence region $\mathbf{R}^{\Sigma}_{5\%} \equiv {}^2\mathcal{T}^{-1}\left({}^0\mathbf{R}^{\mathbf{D}}_{5\%}\right)$. Because de-trending the bivariate series in Equation (32) requires only a univariate hypothesis $(\sigma^r_0)^2$, for mechanical ease and efficiency, and simple 2D visualization, we start with our standard $(\sigma^r)^2$ confidence interval ${}^0I^{(\sigma^r)^2}$ as before and

[47]By construction of the pullback measure described in Footnote 34.

[48]In our context the first two of these densities correspond to the FIR and stock variances.

[49]As previously stated, we choose these confidence intervals because there are three parameters to estimate and $(100\% - 1.6\%)^3 = 98.4\%^3 = 95.3\% > 95\%$. A significance level of 1.6% is thus a conservative cube root for a significance level of 5%: ${}^0\mathbf{R}^{\mathbf{D}}_{5\%} \equiv {}^0I^{(\sigma^r)^2} \times {}^0I^{(\sigma^S)^2} \times {}^0I^{\rho} \approx [0.84, 1.41] \times [0.84, 1.41] \times [-0.3, 0.3]$. We have followed the convention that the intervals should be symmetric (after weighting) about their center, 1 in the case of χ^2_f and 0 in the case of $\mathcal{D}_\rho$. *See Standard Confidence Interval* at the end of Section 5.2.

then, conditional on a particular $\left(\sigma_0^r\right)^2$ in this interval, estimate a two-dimensional cross-sectional confidence region for $\rho, \left(\sigma^S\right)^2 : \equiv \left(\mathbf{R}_{3.2\%}^{\rho,\left(\sigma^S\right)^2} \mid \left(\sigma_0^r\right)^2\right) \subset \mathbf{R}_{5\%}^{\Sigma}$. We recover the full confidence region $\mathbf{R}_{5\%}^{\Sigma}$ disjointly from the cross-sections:

$$\mathbf{R}_{5\%}^{\Sigma} = \bigcup_{\left(\sigma_0^r\right)^2 \in {}^0 I^{(\sigma^r)^2}} \mathbf{R}_{3.2\%}^{\rho,\left(\sigma^S\right)^2} \mid \left(\sigma_0^r\right)^2 .$$

Details are left to the Computational Appendix. Conditioned on FIR volatility zero error estimates $\widehat{\left(\sigma^r\right)^2}_0 \equiv \widehat{\left(\sigma^r\right)^2}|f$, Figure (5) shows for ten selected futures pairs the associated ${}^0\mathbf{R}_{3.2\%}^{\rho,\left(\sigma^S\right)^2}$, defined as the indicated two-dimensional confidence regions for stock volatility and stock-FIR correlation at the 3.2% significance level, conditional on the associated zero error estimate $\widehat{\left(\sigma^r\right)^2}_0$ for $\left(\sigma^r\right)^2$. In Figure (5), each region's color indicates the associated level of $\widehat{\left(\sigma^r\right)^2}_0$. Figure (6) shows all confidence regions associated with the ninety-five futures pairs. The Computational Appendix provides details, including similar confidence regions conditioned on other estimates $\widehat{\left(\sigma^r\right)^2}|d \in {}^0 I^{(\sigma^r)^2}$.

5.5 Estimating the Convexity Adjustment

We restrict the discussion to the confidence regions ${}^0\mathbf{R}_{3.2\%}^{\rho,\left(\sigma^S\right)^2} \equiv \mathbf{R}_{3.2\%}^{\rho,\left(\sigma^S\right)^2} \mid \widehat{\left(\sigma^r\right)^2}_0$ described in the prior subsection and leave the general case to the Computational Appendix.

The confidence regions from the prior section are interesting in their own right. In addition, they imply, via Proposition 4.4, confidence intervals for the FIR convexity adjustment $\equiv c(t) = e^{U_1(T_2,\theta) - U_2(t,\theta)}$, in which the exponent can be written as a quadratic form defined on these regions via Equation (22).

Figure (7) contains a graph of the maximum, minimum and average of the convexity exponent $U_1(T_2,\theta) - U_2(t,\theta)$ over a 95% confidence interval in each quarterly futures period, together with a graph of the unadjusted FIR (without any convexity adjustment, based solely on the IFFR) and contemporaneous market (forward) rates. A quick look at this convexity graph reveals that the average adjustment rarely exceeds a basis point and never exceeds six basis points. The average convexity adjustment histogram in Figure (8) confirms this conclusion.

On this basis we conclude that the convexity adjustment is typically negligible at the data frequencies and time horizons we consider, outside of a small number of periods we will treat as atypical.

Thus, we take as our estimate of the FIR the unadjusted FIR in the top graph of Figure (7) and Figure (1).

6 Spread Analysis

Our primary objective in this paper has been to understand the implicit financing cost of futures investment, which we have addressed above, and its relationship to market rates on explicit financing, to which we now turn, and which can be expressed through *spreads*. Our FIR estimates from the last section permit an estimate of the spreads[50] between the FIR and market rates. To match the FIR terms we take contemporaneous forward Treasury rates and LIBOR, respectively, as representative of risk-free and financial institution borrowing rates.[51]

Figure (9) shows a graph of the estimated FIR-LIBOR and FIR-Treasury spreads with four *regimes* superimposed. These regimes are based on a visual review suggesting three shifts in the behavior of the two spreads over our total observation period: a tightening of the spreads near the start of 2001, an increase in spread volatility and divergence between the two spreads midway through 2007, and a reduction in volatility and convergence of the spreads in 2009.

[50] The *spread* equals the FIR minus the indicated contemporaneous market rate. Section [•] of the Computational Appendix describes the methodology we used the determine the forward market rates and associated FIR spreads. *See also* the text accompanying Footnote 16 for a general indication of our methodology.

[51] *Id.*

Motivated by this visual review, we defined the following four separate regimes, generally determined by reference to the passage of the Commodity Futures Modernization Act of 2000 ("CFMA") and the 2008 financial crisis:

1. From 1996 until the passage of the CFMA, which we set at the end of 2000.
2. After the CFMA but before the financial crisis, which we treated as beginning in July of 2007, the month in which "Bear Stearns disclosed that ... two [of its] subprime hedge funds had lost nearly all of their value amid a rapid decline in the market for subprime mortgages."[52]
3. The financial crisis, which we treated as ending in March 2009, and
4. Recovery from the financial crisis.

The spreads series remained in a range between positive and negative three percent, suggesting mean-reversion. An AR(1) fit (shown in dotted red in Figure (10))) supports mean-reversion as well.

We therefore treat the spreads series in each regime as Ornstein-Uhlenbeck (OU) processes, which we estimate as AR(1) processes,[53] viewed as a discrete OU processes, allowing us, for example, to test hypotheses that the AR(1) reversion mean[54] of a spread in a particular regime was greater, or less, than zero at a specified confidence interval.[55]

Figure (11) shows the block bootstrapped distribution of the AR(1) means of the two spread series in each regime and over the entire observation period (last boxplot).[56] The colored regions represent block bootstrapped 95% confidence regions for those means, while the

[52] "Bear Stearns." *Wikipedia: The Free Encyclopedia.* Wikimedia Foundation, Inc. September 4, 2016. Web. October 17, 2016.

[53] *First order autoregressive processes.*

[54] For an AR(1) of the form $x_{t+1} = c + bx_t + a\epsilon_t$, the *reversion mean* is defined by $\mu = \frac{c}{1-b}$, equal to the long-term equilibrium determined by setting $\epsilon_t \equiv 0$.

[55] Equivalently, was the difference between the FIR and the LIBOR or Treasury rate statistically significant in that regime?

[56] We used a procedure with 10,000 (block) bootstrapped (re)samples and a block length of 20 for each subseries. The bootstrap procedure addressed the remaining serial correlation in the AR(1) residuals, but did not address the possibility that the volatility of the AR(1) residuals varied. Further research could include fitting AR(1) + GARCH(1,1) models to reflect this possibility.

	Average FIR Spread (bps)	
Period	LIBOR Rate	Treasury Rate
Total (1996 to 2019/6)	(6)	38
BCI	(10)-(1)	34-43
1996 to 2000 (pre-CFMA)	26	75
BCI	18-32	68-81
2001 to 2007/6 (post-CFMA)	(7)	15
BCI	(15)-2	6-25
2007/7 to 2009/3 (financial crisis)	(73)	56
BCI	(97)-(37)	29-77
2009/4 to 2019/6 (recovery)	(11)	30
BCI	(15)-(7)	27-33

Table 1. Forward FIR Spreads to Two Market Rates in Four Regimes with Confidence Intervals Based on Block Bootstraps

notches represent those confidence regions determined by conventional bootstraps, showing the importance of serial correlation to estimation. With the sole exception of the LIBOR spread in the second regime, each mean is nonzero at the 95% confidence level. The Treasury spread mean is consistently positive, while the LIBOR spread mean's sign is positive before the CFMA but negative thereafter. The associated numeric data appears in Table (1) above, from which we can see that, excluding the final Treasury spread, the means in consecutive regimes differ at the 90% significance level.

In this regard, our analysis provides three main findings. First, the mean spreads decreased after the CFMA's passage, and this decrease was significant. Second, the mean spreads in consecutive regimes differed, and this difference was generally significant. Third, over the entire observation period the mean spread to LIBOR (respectively, the Treasury rate) was generally negative (resp., consistently positive), but the spread, and in the case of LIBOR its sign, varied considerably over time and regime.

We can therefore conclude in broad terms that over most of our twenty-four year observation periods the FIR was above the Treasury rate but below LIBOR, though with many exceptions, some large. Because LIBOR is a market rate associated with unsecured financial financing, we further conclude in turn that the FIR was generally attractive relative to unsecured explicit financing available to large financial institutions.

In margin lending, however, the financing is collateralized by purchased securities. The associated rates vary considerably, but can be close to, or even below, the risk free rate, particularly if the securities are on "special".[57]

Putting all this together, we conclude that, since the FIR was generally significantly above the Treasury rate, the evidence from our observation period showed no meaningful systematic advantage or disadvantage for the FIR relative to the rates on explicit financings, including margin. On any given date, the FIR could be, and during the observation period was, above or below the explicit financing rate.

7 Summary

We explored the cost of implicit leverage associated with the S&P 500 equity index futures contract. We showed that the related implicit financing rate has often been attractive relative to market rates on explicit financing; however, the relationship between the implicit and explicit financing rates has been volatile and varied considerably based on legal and economic regimes. Only rarely was the FIR below the Treasury rate.

These findings depended on accurate estimates of the FIR, which depended in turn on estimates of the FIR convexity term. We developed new estimation methods for these purposes with strong similarities to fiducial inference, including point and interval estimates.

[57]Securities that are hard to borrow are said to be "on special" and can be lent on favorable terms. Since margin agreements often permit the lender to lend the hypothecated securities, the associated lending revenue can reduce the margin rate. Geczy et al. (2002). Each margin lender sets its own margin rate, and historical data on a market margin rate level is accordingly difficult to obtain. The *broker call money rate* is the rate lenders charge to brokers on loans to finance the brokers' margin loans to customer. This rate has been observed to be around 1.5 points above the Federal Funds rate. Fortune (2000). From this we see that margin lending rates vary considerably and exceed Treasury rates absent unusual circumstances. By way of comparison, in our observation period LIBOR exceeded the Treasury rate by 44 basis points on average.

In order to estimate the FIR convexity term, we extended our methods to bivariate estimates of the FIR and equity volatility, Our estimates showed the convexity term could be treated as insignificant in all but a small number of cases, allowing us to estimate the FIR simply from the observable futures ratio, adjusted for dividends.

References

Anderson, G. and Kercheval, A. N. (2010). *Lectures on Financial Mathematics: Discrete Asset Pricing*. Synthesis Lectures on Mathematics and Statistics. Morgan and Claypool Publishers.

Anderson, R. M., Bianchi, S. W., and Goldberg, L. R. (2012). Will my risk parity strategy outperform? *Financial Analysts Journal*, 68(6):75–93.

Anderson, R. M., Bianchi, S. W., and Goldberg, L. R. (2014). Determinants of levered portfolio performance *Financial Analysts Journal*, 70(5):53–72.

Anderson, R. M., Eom, K. S., Hahn, S. B. and Park, J. H. (2013). Autocorrelation and partial price adjustment *Journal of Empirical Finance*, 24:78–93.

Bali, T. G. and Demirtas,K. 0. (2008) Testing Mean Reversion In Financial Market Volatility: Evidence From S&P 500 Index Futures *The Journal of Futures Markets*, 28(1): 1–33

Black, F. (1976) The Pricing of Commodity Contracts *Journal of Financial Economics*, 3: 167-175

Breeden, D. T. and Litzenberger, R. H. (1978) Prices of State-contingent Claims Implicit in Option Prices *Journal of Business*, 51(4): 621–51

Chan, K., Chan, K. C. and Karolyi,G. A. (1991) Intraday Volatility in the Stock Index and Stock Index Futures Markets *The Review of Financial Studies*, 4(4): 657-684

Chen, W. P., Chung, H. and Lien, D. (2016) Price discovery in the S&P 500 index derivatives markets, *International Review of Economics and Finance*, 45: 438–452

Cox, J., Ingersoll, J. and Ross,S. (1981) The Relation Between Forward Prices and Futures Prices *Journal of Financial Economics*, 9: 321-347

Dwyer, G., Locke, P. and Yu, W. (1996). Index Arbitrage and Nonlinear Dynamics Between the S&P 500 Futures and Cash *The Review of Financial Studies*, 9(1):301-332

Fisher, R. A. (1935). The Fiducial Argument In Statistical Inference *Annals of Eugenics*, 6:391-98

Fortune, P. (2000). Margin Requirements, Margin Loans, and Margin Rates: Principles and Practice *New England Economic Review*, 20 (September/October):31

Geczy, C. et al., (2002). Stocks are Special Too: An Analysis of the Equity Lending Market *Journal of Financial Economics, 66, 241-269*, (10.1016/S0304-405X(02)00225-8)

Hilliard, J. E., Hilliard, J. and Ni, Y. (2021). Using the short-lived arbitrage model to compute minimum variance hedge ratios: application to indices, stocks and commodities *Quantitative Finance* 21:1, 125-142, available at:https://www.tandfonline.com/loi/rquf20

Jegadeesh, N. and Pennacch, b. (1996). The Behavior of Interest Rates Implied by the Term Structure of Eurodollar Futures Journal of Money, Credit, and Banking 28(3/2):426-46

Kenourgios, K. Samitas, A. and Droso, P. (2000). Hedge ratio estimation and hedging effectiveness: the case of the S&P 500 stock index futures contract *International Journal of Risk Assessment and Management* 9(1/2):121-34

Lien, D. (2004). Cointegration and the optimal hedge ratio: the general case *The Quarterly Review of Economics and Finance* 44: 654–658

Otto, M. (2000) Stochastic relaxational dynamics applied to finance: Towards non-equilibrium option pricing theory. Eur. Phys. J. B. 14:383–394

Poon, S. and Granger, C. (2003) Forecasting Volatility in Financial Markets: A Review Journal of Economic Literature 41(2): 478-539.

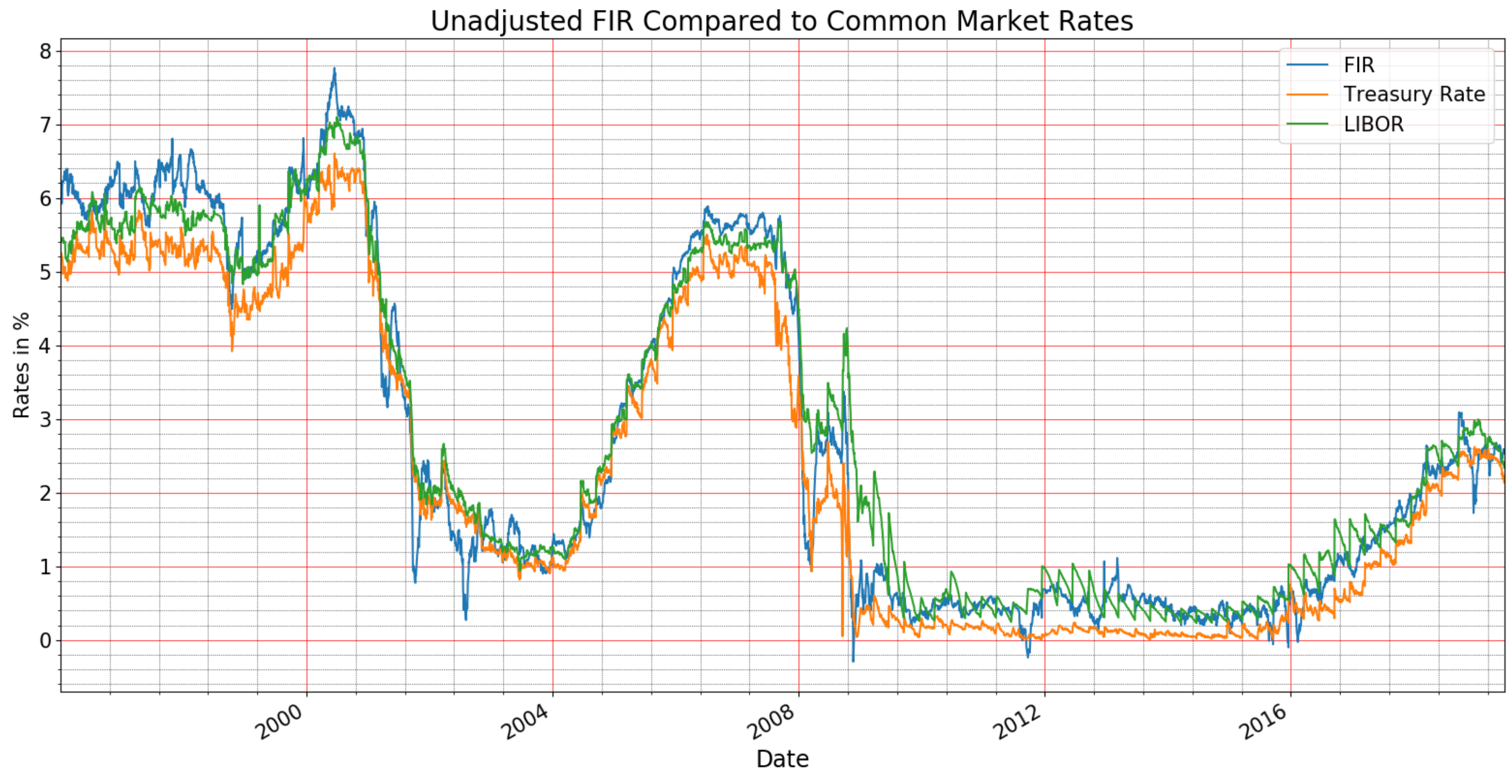

Figure 1. Unadjusted FIR Estimates vs Market Rates

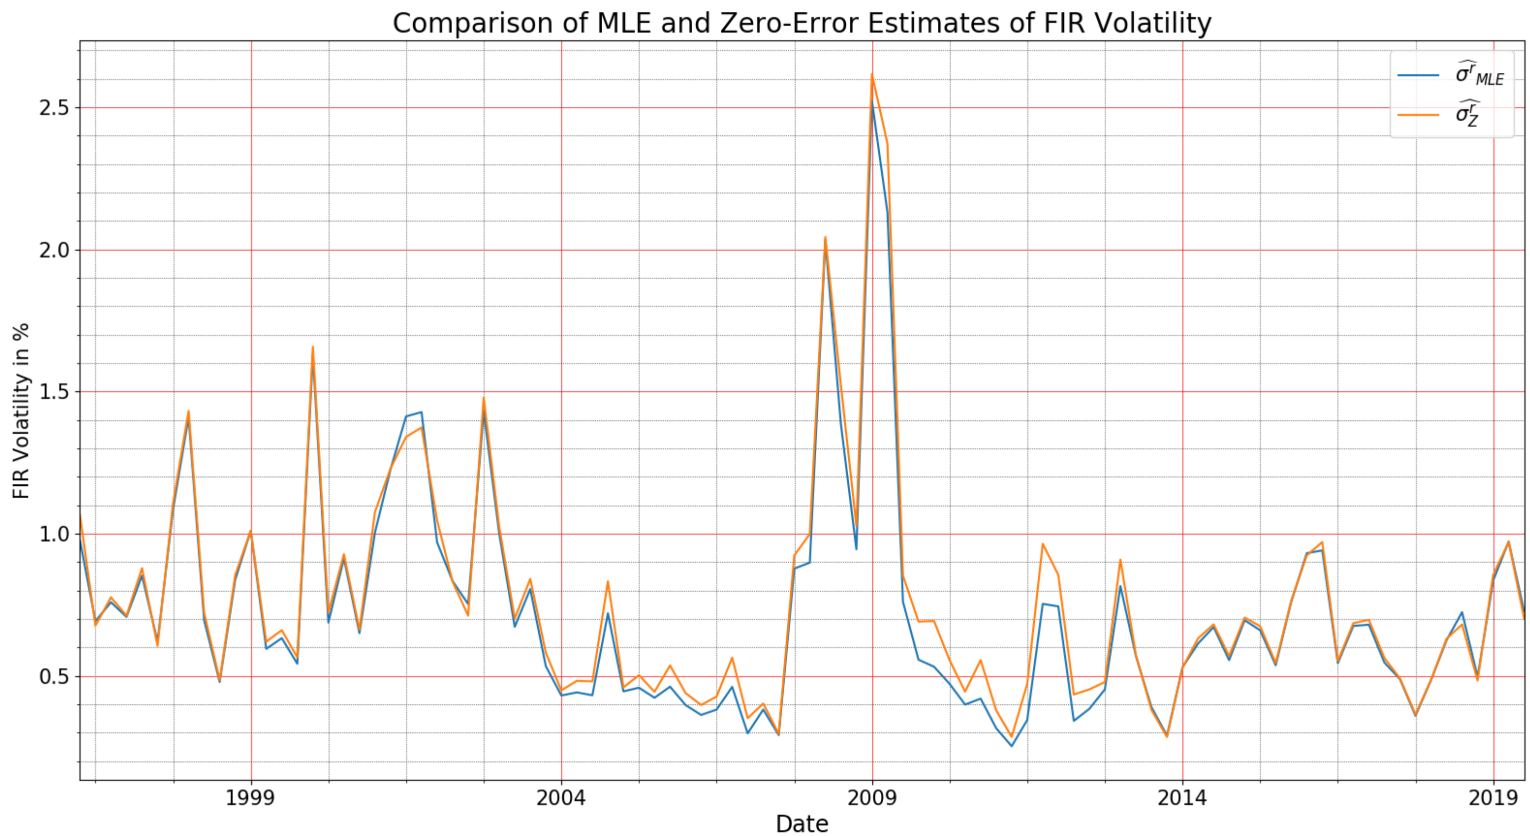

Figure 2. $\widehat{(\sigma^r)^2}|f$ vs MLE Estimates

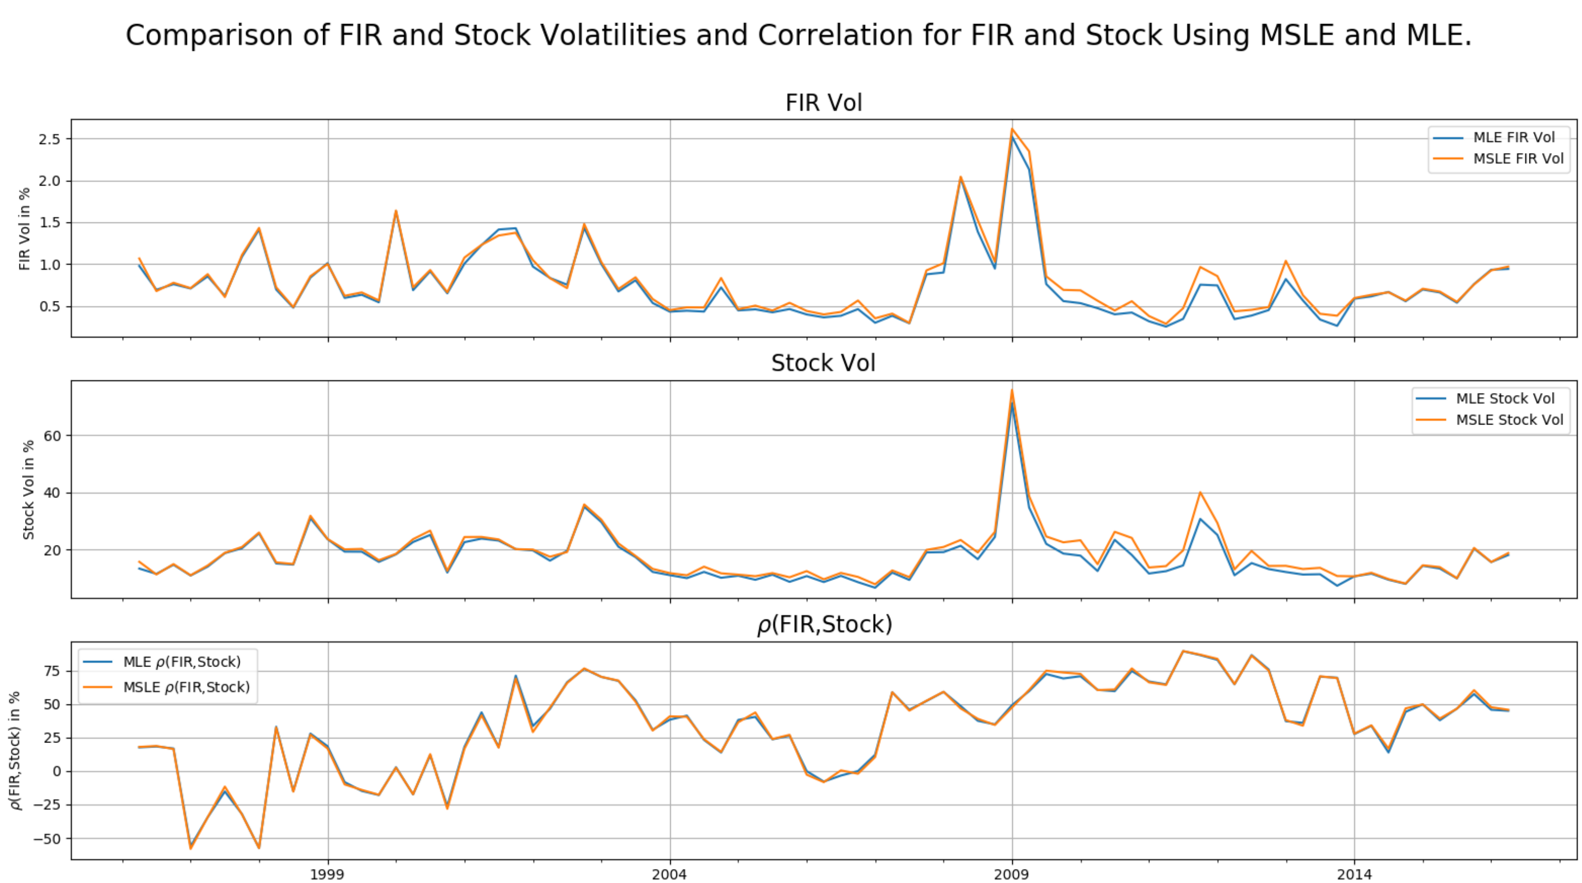

Figure 3. $\widehat{\theta} \mid f$ vs MLE Estimates

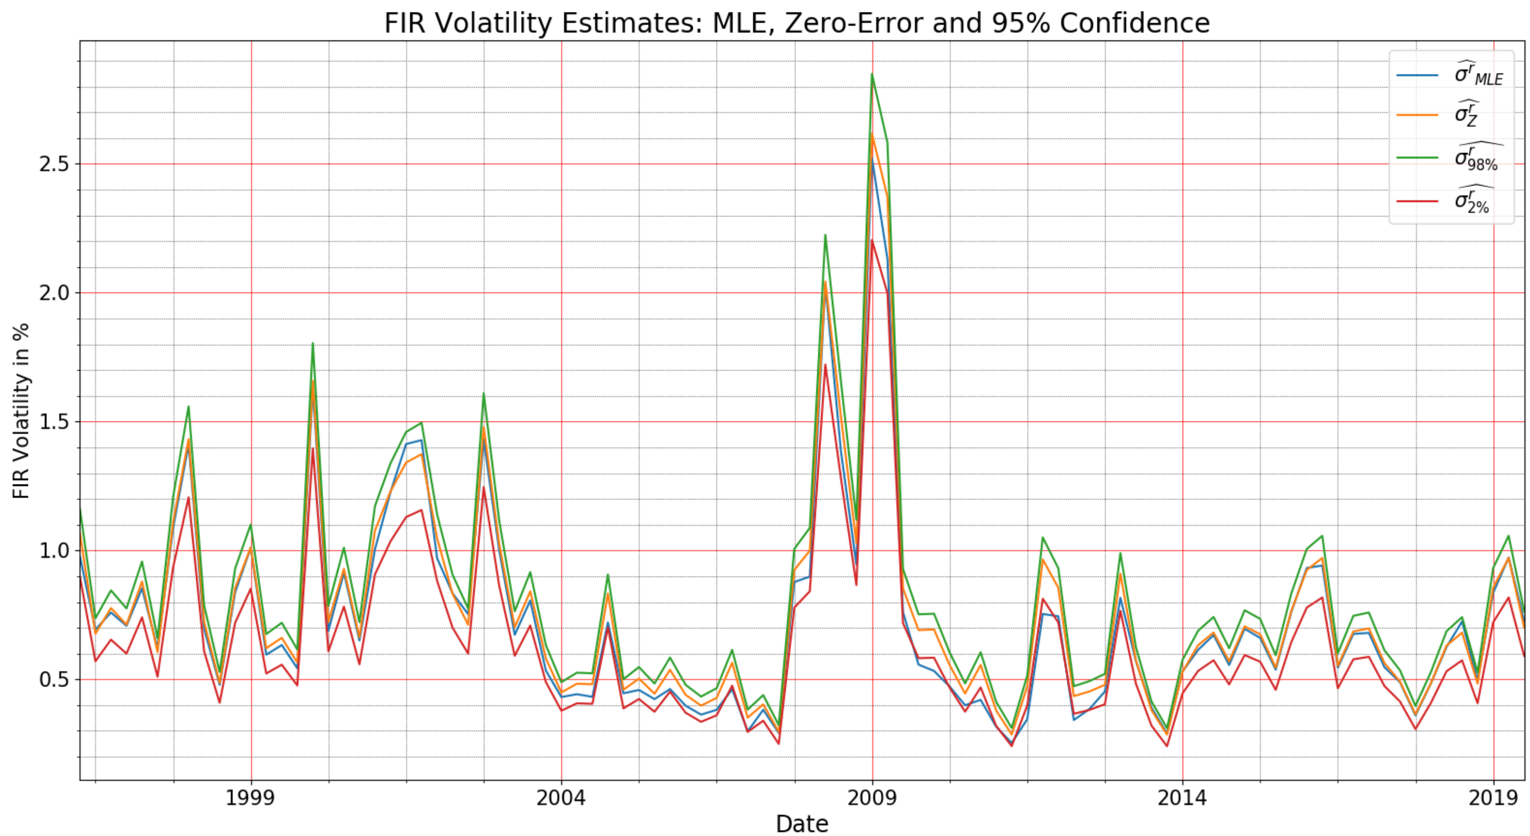

Figure 4. $\widehat{(\sigma^r)^2}|f$ Estimates with 98.4% Confidence Intervals

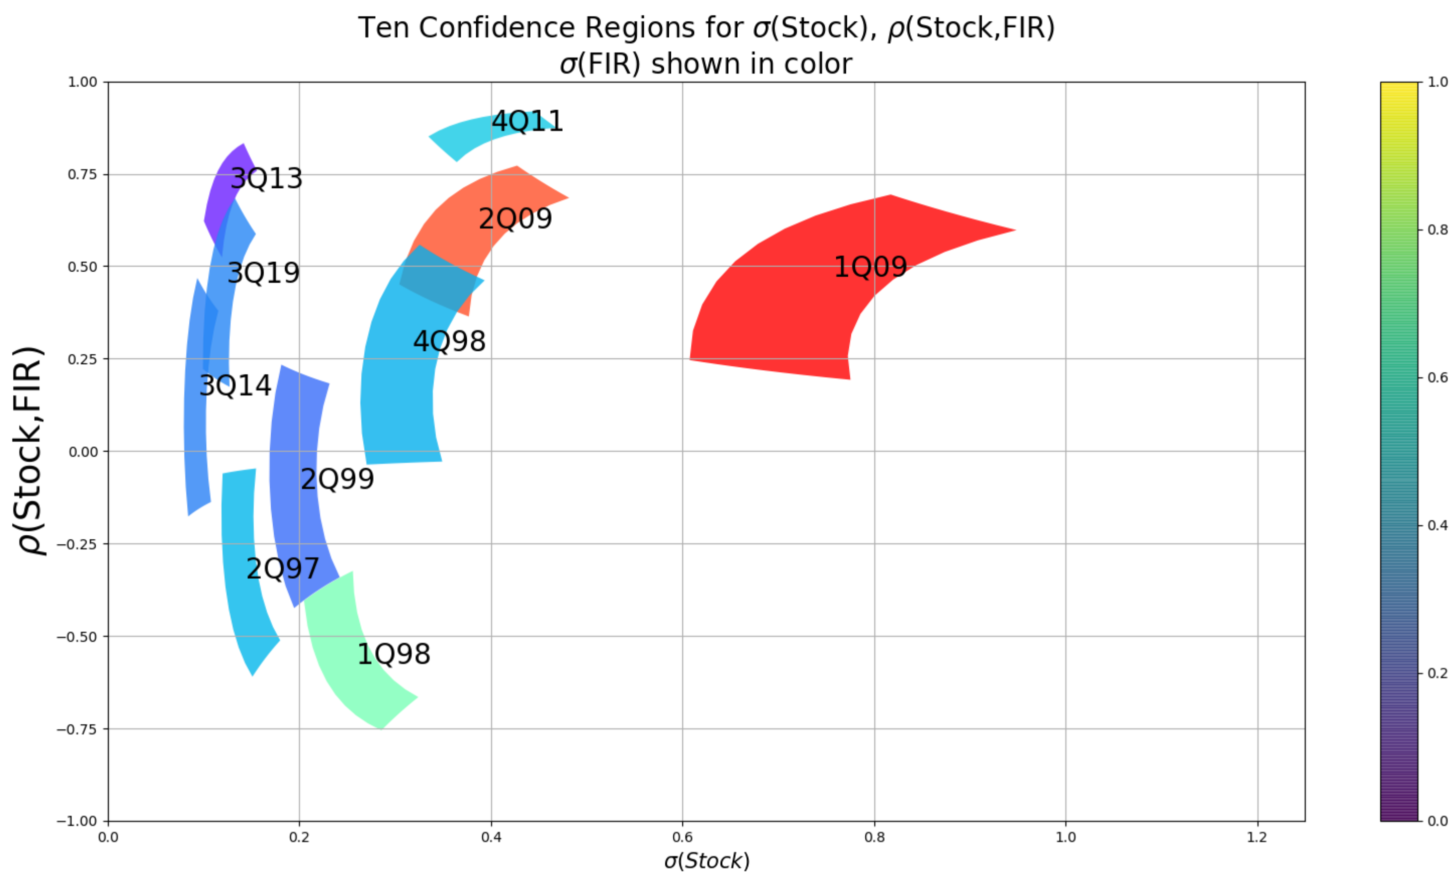

Figure 5. Select Confidence Regions

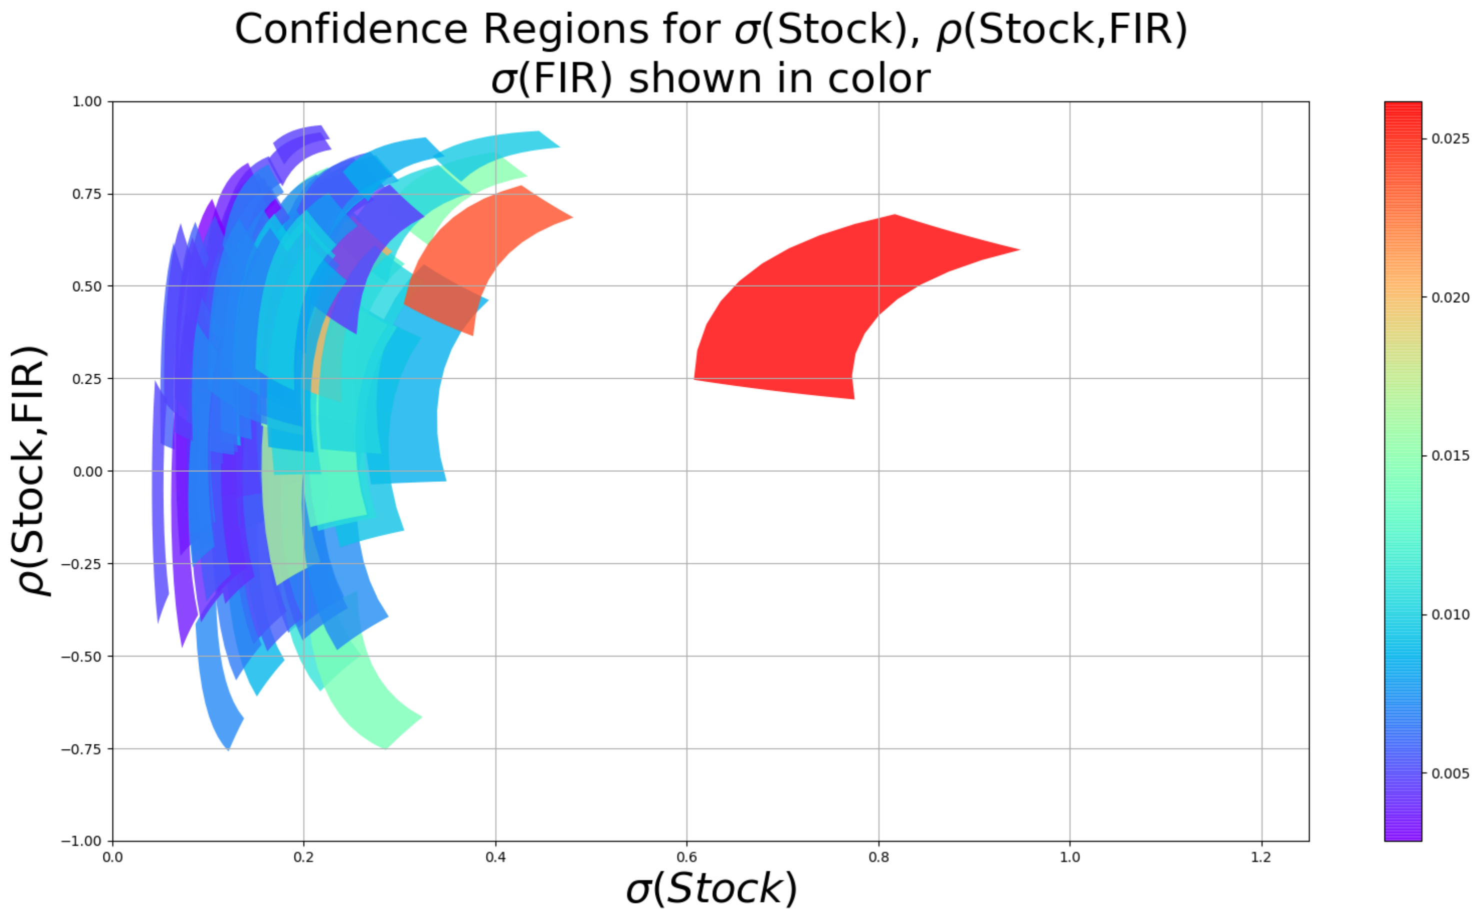

Figure 6. All Confidence Regions

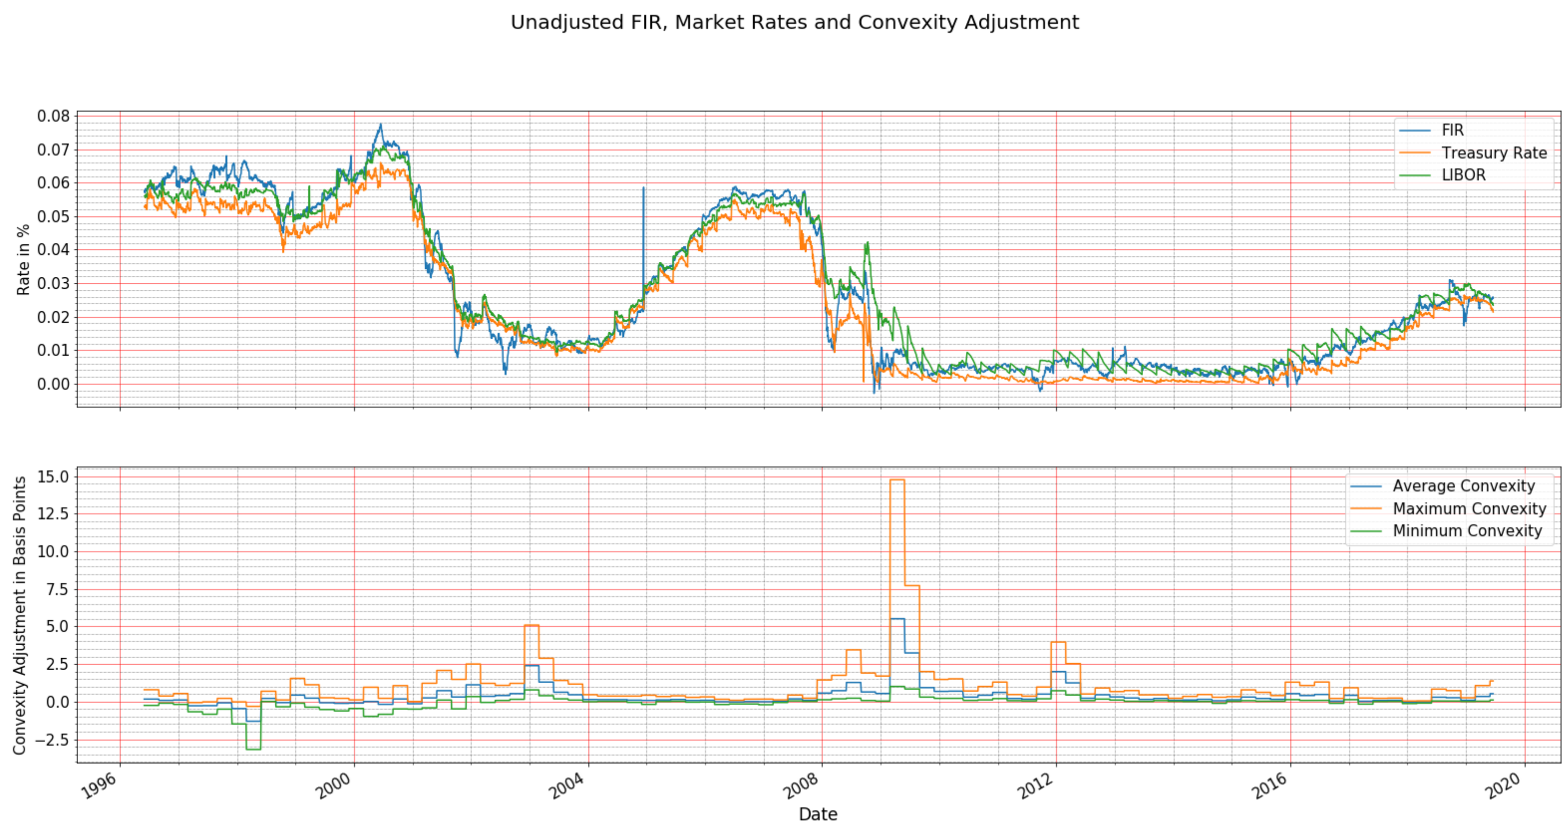

Figure 7. Unadjusted FIR, Market Rates and Convexity Adjustment 95% CIs

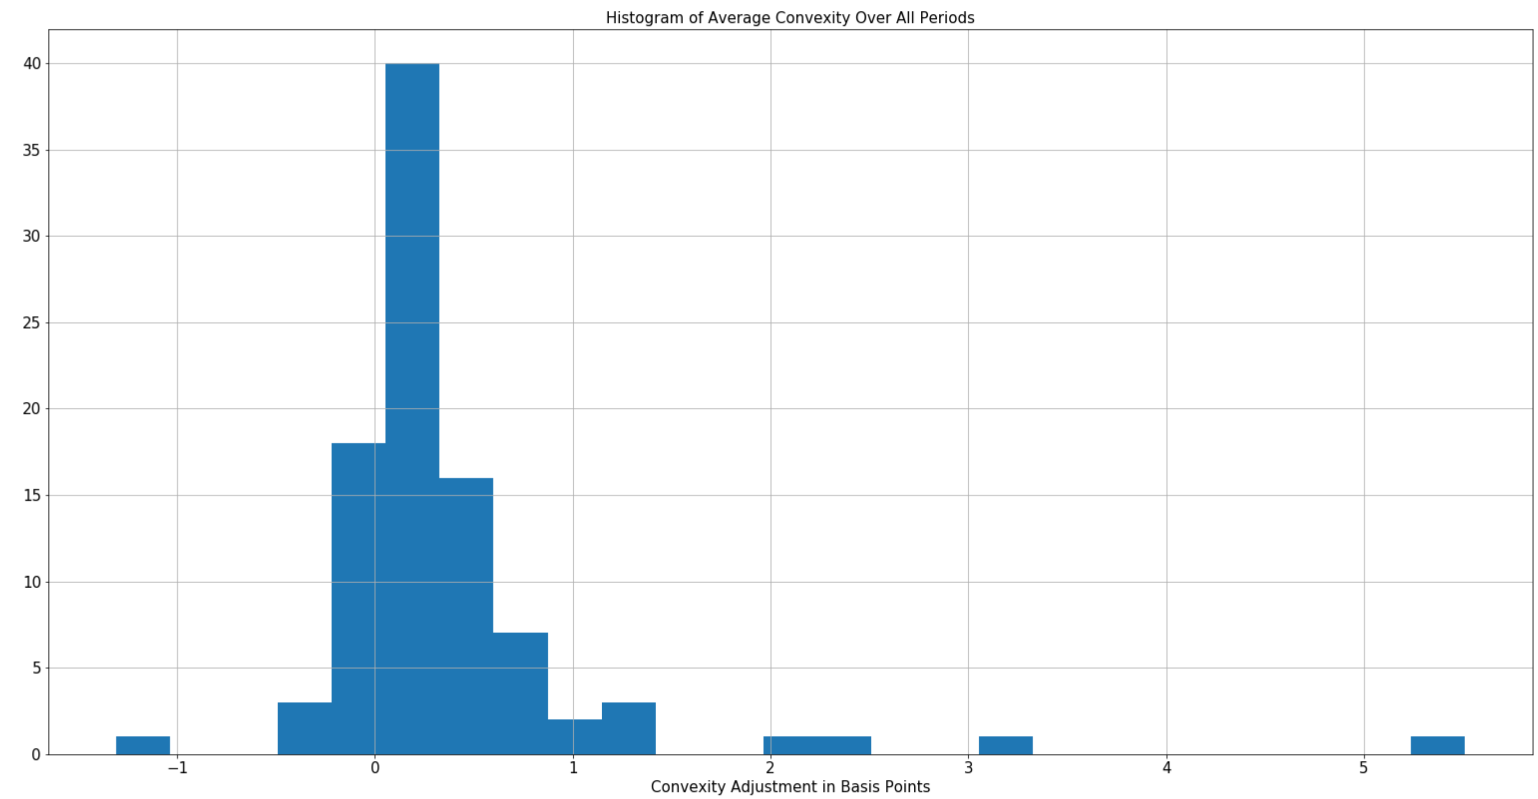

Figure 8. Histogram of Average Convexity Adjustment over 95% CIs

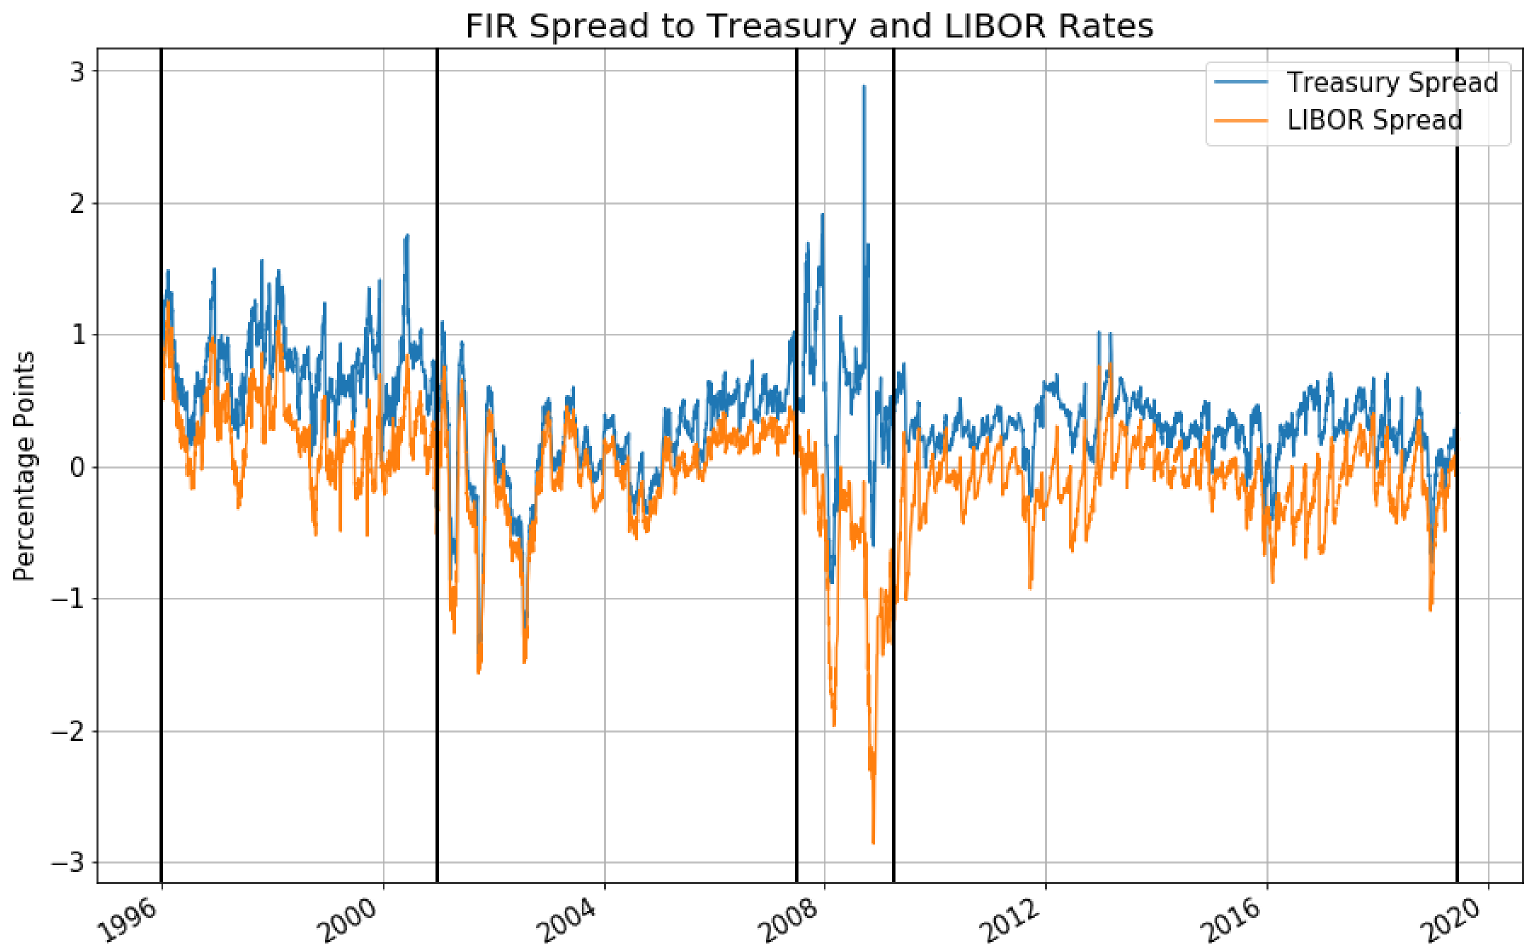

Figure 9. FIR Spread Mean Reversion in Four Regimes

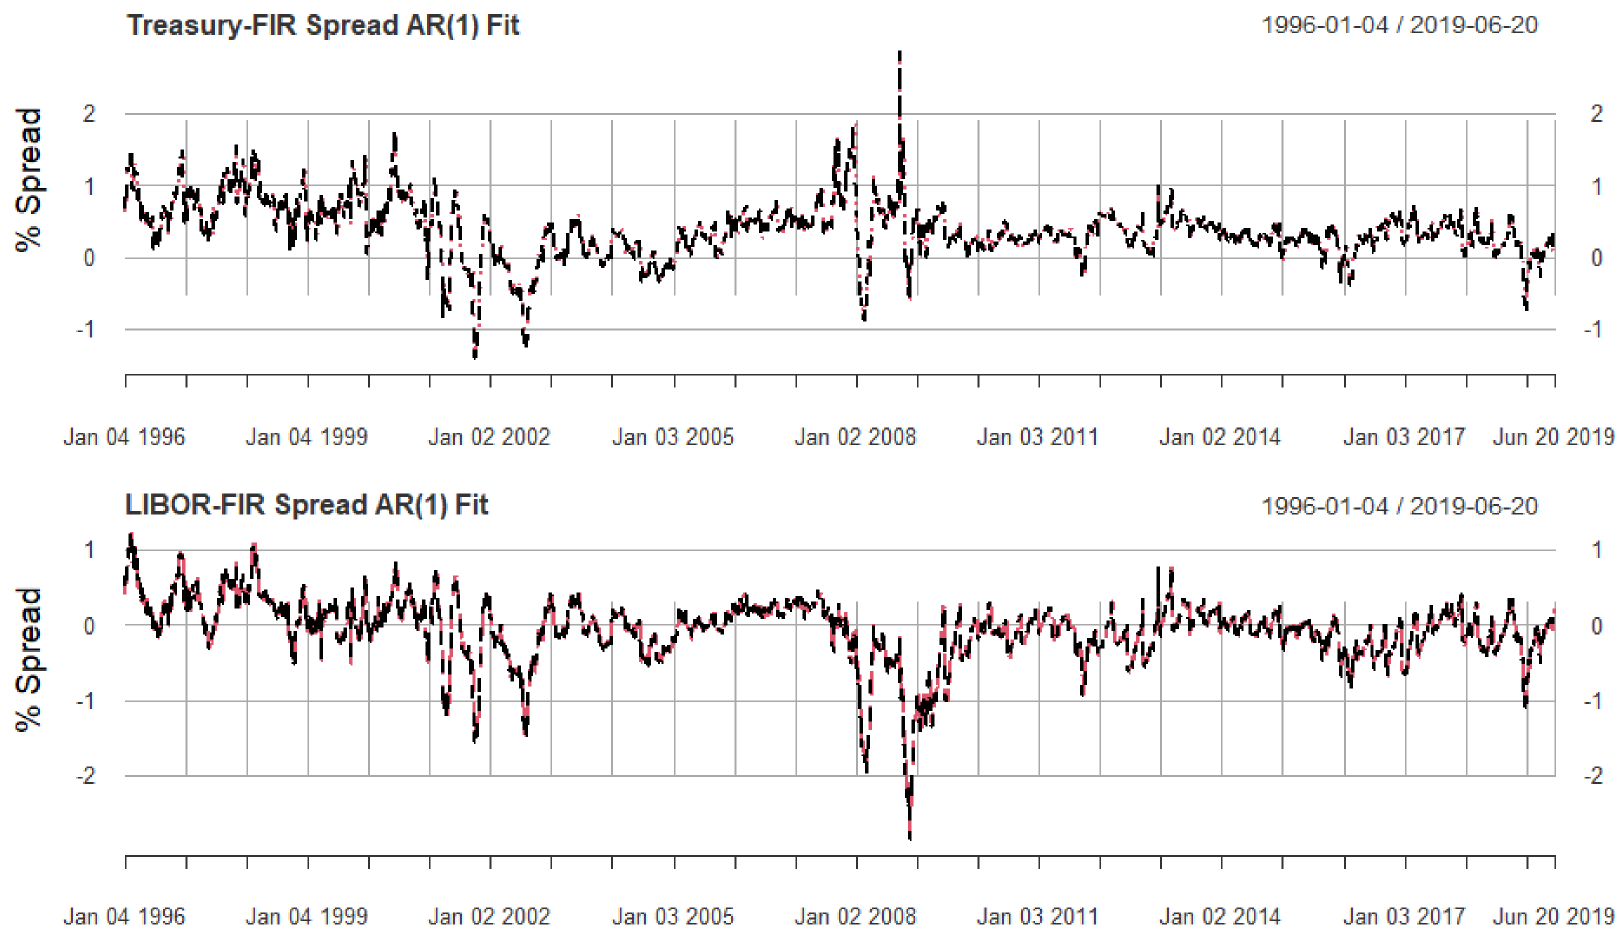

Figure 10. FIR Spread AR(1) Fits

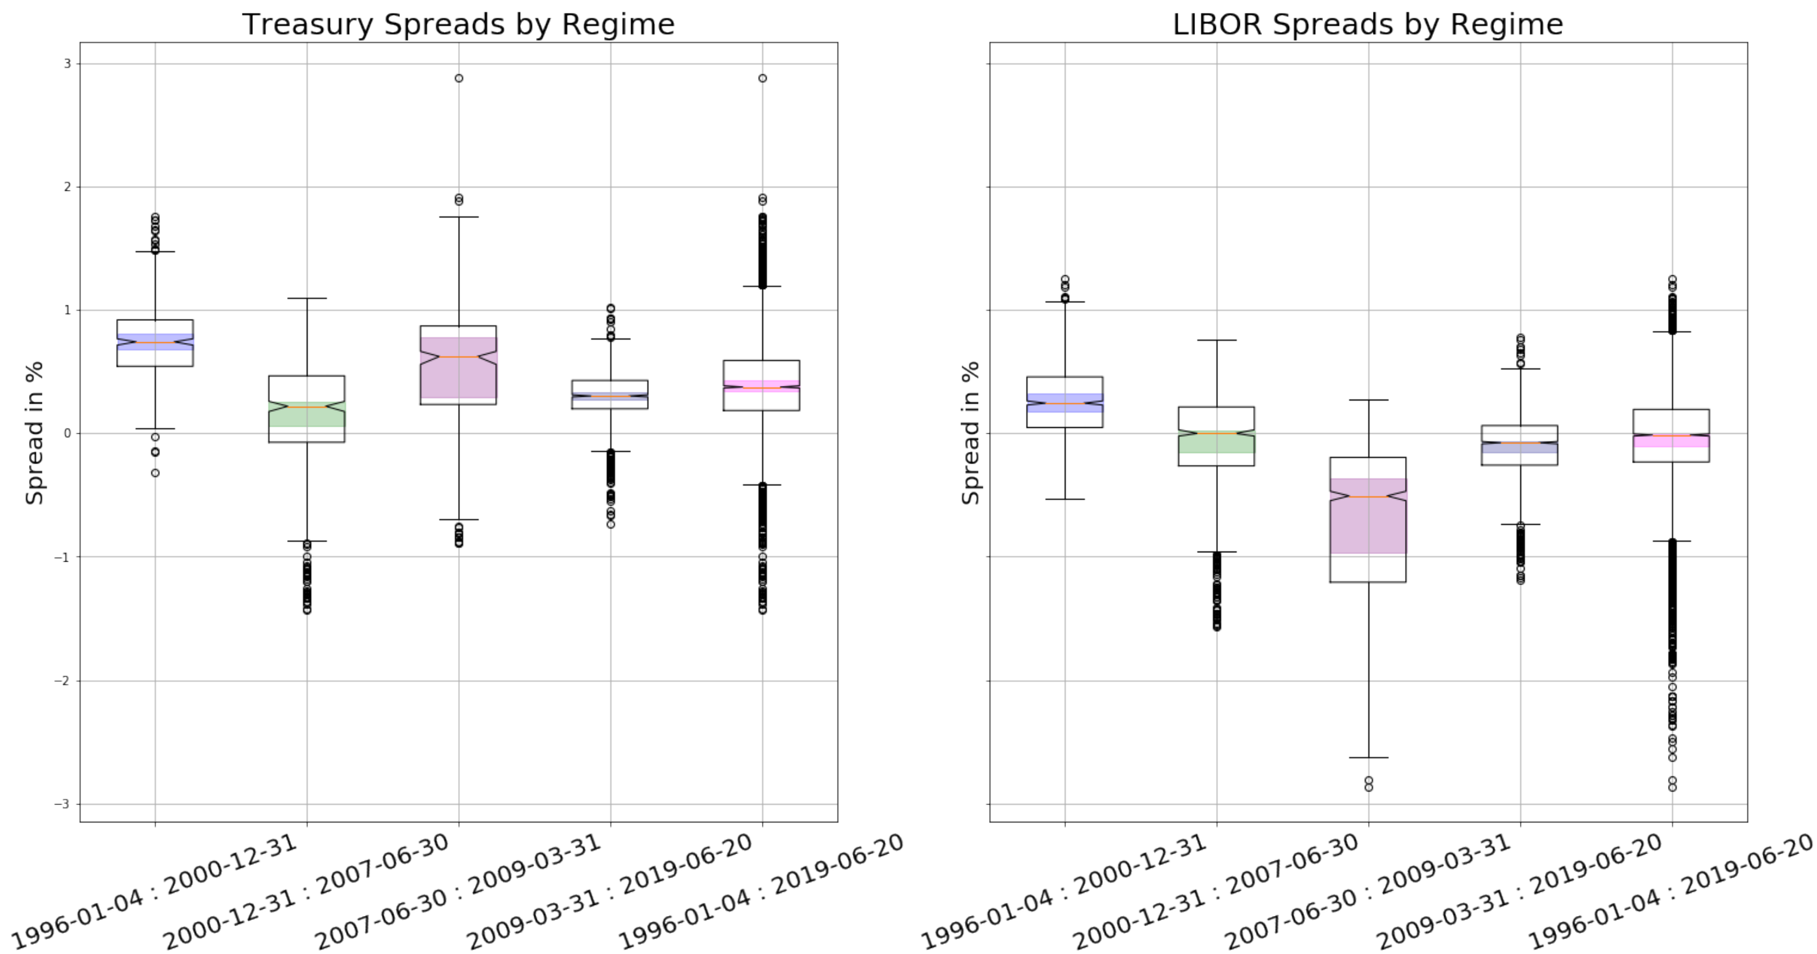

Figure 11. FIR Spread AR(1) Means