

Option-Implied Libor Rate Expectations across Currencies

Gebbia, Nick

Please cite paper as:
Gebbia, Nick (2016). Option-Implied Libor Rate Expectations
across Currencies.
International Finance Discussion Papers 1182.

<http://dx.doi.org/10.17016/IFDP.2016.1182>



International Finance Discussion Papers

Board of Governors of the Federal Reserve System

Number 1182
October 2016

Board of Governors of the Federal Reserve System

International Finance Discussion Papers

Number 1182

October 2016

Option-Implied Libor Rate Expectations across Currencies

Nick Gebbia

NOTE: International Finance Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References to International Finance Discussion Papers (other than an acknowledgment that the writer has had access to unpublished material) should be cleared with the author or authors. Recent IFDPs are available on the Web at www.federalreserve.gov/pubs/ifdp/. This paper can be downloaded without charge from the Social Science Research Network electronic library at www.ssrn.com.

Option-Implied Libor Rate Expectations across Currencies*

Nick Gebbia

Federal Reserve Board of Governors

October 13, 2016

Abstract

In this paper, I study risk-neutral probability densities regarding future Libor rates denominated in British pounds, euros, and US dollars as implied by option prices. I apply Breeden and Litzenberger's (1978) result regarding the relationship between option prices and implied probabilities for the underlying to estimate full probability density functions for future Libor rates. I use these estimates in case studies, detailing the evolution of probabilistic expectations for future Libor rates over the course of several important market events. Next, I compute distributional moments from density functions estimated for fixed horizons and test for Granger causality across the three Libor rate distributions considering their mean, standard deviation, skewness, and kurtosis. I further break these relationships down by various fixed horizon lengths, as well as the slope and curvature in the term structure of moments over different horizons. The results show a rich interconnectedness among these three Libor rates that extends well beyond levels of future mean expectations.

JEL Classifications: C14, E43, G13

Keywords: options, futures, Libor, pdf, distribution, moments, Granger causality

*The views presented in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System. I would especially like to thank Stephanie Curcuru, Michiel De Pooter, Rob Martin, Marius Rodriguez, and Charles Thomas, whose insights and guidance were necessary conditions to this paper's existence, and are very much appreciated. I also would like to acknowledge the Global Monetary and Sovereign Markets (GMSM) and Global Capital Markets (GCM) sections for thoughtful feedback on an earlier version of this project. All errors are my own. Email: nick.j.gebbia@gmail.com

1 Introduction

Option prices provide a unique insight into the probabilities assigned by markets to various future outcomes for a particular economic variable. An option contract allows its owner to buy or sell an underlying security or commodity at a pre-determined price. As such, options can be used to hedge against downside or upside risks, or alternatively to speculate on the future value of a particular security or commodity. In the hedging case, options provide an effective ceiling or floor to the price of the underlying. On the other hand, as the price of an option's underlying exceeds (falls below) the price at which a speculator agreed to buy (sell), the speculator earns a return by exercising the option to buy (buying at the market price) and selling at the market price (exercising the option to sell). In either case, the price an individual is willing to pay for a particular option is based intimately upon the likelihood assigned to various outcomes for the option's underlying.

Cox and Ross (1976) formalize this relationship by expressing the price of a European option¹ as the discounted expected payout of the option at maturity. Breeden and Litzenberger (1978) directly relate the second partial derivative of call option price with respect to strike to the probability density function (pdf) assigned to the underlying. Since these results were established, option prices have been applied in various cases to derive estimates of market-assigned probabilities for different underlying securities or commodities (for example, see Malz (1997) in the case of foreign exchange rates, or Melick and Thomas (1997) and Datta, Londono, and Ross (2014) in the case of oil prices).

Option prices have been applied to estimate market-assigned probabilities for future interest rates, as well. Clews, Panigirtzoglou, and Proudman (2000) estimate pdf's for future British pound-denominated Libor interest rates and use distributional characteristics to shed light on the impact of important events on market expectations. For example, they show that the advent of operational independence for the Bank of England in May, 1997 had little contemporaneous effect on market uncertainty around future GBP Libor rates. Ivanova and Gutiérrez (2014) use options to study probabilities assigned to future euro-denominated Libor (Euribor) rates, investigating the forecasting performance of such measures and using forecast biases to infer risk aversion.

Two general approaches to estimating option-implied probabilities have emerged in the literature as particularly useful. One involves assuming a functional form for the terminal distribution of the underlying (for example, a weighted mixture of two lognormal or two normal distributions), and determining the parameters of the distribution by minimizing the resulting option pricing errors (see Bahra (1997); Melick and Thomas (1997)). The other works by interpolating a fine set of option prices and using the

¹ A European option is one that can only be exercised on its maturity date. Contrastingly, American options allow the option holder to exercise at any time prior to expiration.

result of Breeden and Litzenberger (1978) to produce a pdf by taking the second derivative of option prices numerically. This paper does not take a stance on whether one method outperforms the other. However, previous research comparing these two methodologies in Monte Carlo simulation settings using similar data has favored the method of interpolating option prices and taking numerical derivatives (Cooper (1999); Bliss and Panigirtzoglou (2002)).

In this paper, I use options on Libor futures for interest rates denominated in euros (EUR), British pounds (GBP), and US dollars (USD) in order to study market-assigned probabilities for the three rates. This paper contributes to the existing literature on option-implied pdf's in a few ways. First, I detail the evolution of pdf estimates for Libor rates denominated in EUR, GBP, and USD over important case studies, including events critical to the financial crisis of 2007-08. These case studies serve as meticulous overviews of the ways in which market-assigned probabilities for future interest rates were affected by major events. I then move to a more thorough investigation of the time series dynamics among distributional characteristics for the three Libor rates. I perform tests of Granger causality across the three Libor rate distributions considering the option-implied mean, standard deviation, skewness, and kurtosis related to pdf's at fixed future horizons; these tests are carried out for the levels of the four distributional moments at different horizons, as well as their slope and curvature across horizons. By comparing estimates of distributional characteristics across the three separate Libor rates, this paper offers unique insights into their interrelationships extending beyond comparisons of spot or even expected future levels. Finally, as a methodological contribution, I take up the issue of the American option premium as it relates to option-implied pdf's using a Monte Carlo simulation approach; while previous research has focused on the level of the American option premium (see Tian (2011)), I tie the premium directly to its impact on probability function estimates. The remainder of the paper is structured as follows: Section 2 reviews the data and methodology; Section 3 presents empirical case studies of the evolution of future Libor rate pdf's over important market events; Section 4 details the results of Granger causality tests over distributional moments across the three Libor rates; and Section 5 concludes, along with suggestions for further research.

2 Estimating Probability Density Functions and Distributional Moments from Option Prices²

2.1 Data

Futures contracts for 3-month Libor rates denominated in GBP and EUR trade on the London International Financial Futures and Options Exchange (LIFFE), and similar contracts for Libor denominated in

² Appendix A contains further technical detail regarding option pricing models covered in this section.

USD trade on the Chicago Mercantile Exchange (CME). In order, these futures contracts are typically referred to as Short Sterling, Euribor, and Eurodollar futures. The futures contracts are cash settled at expiration based on the prevailing 3-month Libor rate in the relevant currency at that time, as well as the notional value of the contract. Prices are quoted as 100 minus the annualized 3-month rate – for example, an effective rate of 2.79% would imply a contract price of $100 - 2.79 = 97.21$.

Option contracts written on 3-month Libor futures are available on the same exchanges. For futures on GBP and EUR Libor, options exist with maturities in the nearest eight March-quarterly months (i.e. March, June, September, December); for options on USD Libor futures, available option maturities span the nearest sixteen March-quarterly months. All sets of options also offer serial maturities in the four nearest months that do not follow the March-quarterly cycle, so that there is always an available option maturity in each of the six upcoming months. However, when an option is exercised, it delivers the futures contract expiring in the upcoming March-quarterly month. Therefore, the value of a serial option at expiration will be based on a futures rate (for example, a February option's value at expiration is based on the concurrent price of the March futures contract). On the other hand, March-quarterly options expire on the same date as the underlying futures contract, so that the value of a March-quarterly option at expiration is based on a spot Libor rate.³ For this reason, I restrict analysis to option contracts maturing in the March-quarterly cycle. Futures and options data for EUR and GBP Libor comes from Thomson Reuters, and USD Libor data is from CME.

Figures B1 and B2 in Appendix B detail option liquidity trends for March-quarterly Eurodollar options from May 4th, 1998 through March 31st, 2016. Figure B1 shows open interest, and Figure B2 shows volume. Not surprisingly, option liquidity peaked in the early stages of the 2007-08 financial crisis as investors likely sought insurance against an uncertain monetary policy response. As interest rates settled to historically low levels following the crisis, trading of Eurodollar options dipped significantly through 2013. More recently, Eurodollar options have seen a secular rise in liquidity alongside the possibility and realization of monetary policy “liftoff” from near the Zero Lower Bound (ZLB) in the United States. Figures B3 and B4 show open interest and volume as they evolve with an option's time to maturity, aggregated over the same period. Trading is heaviest between 3 and 6 months prior to an option's expiration. Figures B1 - B4 also disaggregate trends in market liquidity to detail out-of-the-money option trading independently. It is apparent that out-of-the-money options capture the solid majority of trading action in the Eurodollar market. In total, out-of-the-money Eurodollar options account for 73% of open interest and 82% of volume over the full period. Finally, Figure B5 separates share of total open interest and volume by month of option maturity. Trading tends to be heaviest for December options, which account for 30% of both open interest and volume over the period.

³ Table B1 in Appendix B reports the maturity date conventions for Short Sterling, Euribor, and Eurodollar futures and options in the March-quarterly cycle.

The payment and exercise structures of the option contracts prove important to the following analysis. All three Libor-based options are American, meaning the buyer has the right to exercise at any time prior to the option's maturity. This is in contrast to European-style options, which only allow exercise at contract maturity. Eurodollar option contracts require payment of the option premium at the time of purchase; however, Short Sterling and Euribor options use a futures-style margining system. For Short Sterling and Euribor options, a premium is agreed at the time of purchase, but it is not paid at that time. Instead, the net value of the option is marked-to-market daily, and the payout of the contract at exercise becomes the value of the option net of the agreed premium. As a result, pricing of Short Sterling and Euribor options involves no discounting.⁴ Additionally, the margining structure results in zero opportunity cost of holding the option, meaning the American early exercise premium can be ignored and Short Sterling and Euribor options can be priced as though they were European options.⁵ The analysis in this paper assumes options are European. Appendix C takes up the issue of the American early exercise premium included in Eurodollar option prices, concluding that there is little accuracy lost in this paper's analysis by treating Eurodollar options as though they were European.

Finally, it is important to note that Libor rates may include a risk and/or term premium component in comparison to the interest rates targeted by central banks in policy decisions; therefore, estimated probability distributions for future Libor rates necessarily blur together market views of future central bank policy along with anticipated risk or term premia. These factors may reinforce or counteract one another: for example, if stresses on the financial sector increase, markets may expect monetary policy easing to lower rates, while simultaneously expecting a greater risk premium between Libor and the corresponding policy rate to push Libor upward. The following analysis does not formally attempt to distinguish between implied expectations for central bank policy rates and other factors contributing to a spread between Libor and the corresponding policy rate, but instead treats Libor as the primary focus.

2.2 Basic Methodology

In an efficient (and risk-neutral) market, the price of a European option is equal to the discounted expectation of its value at maturity.⁶ In the case of a put option on the futures price F_T struck at K and lasting from time t to a maturity date T ,

$$\begin{aligned} P(K, t, T) &= e^{-r\tau} E_t[\max(K - F_T, 0)] \\ &= e^{-r\tau} \int_{-\infty}^K (K - f) \pi_{F_T, t}(f) df, \end{aligned} \tag{1}$$

⁴ See Chen and Scott (1993).

⁵ See Lieu (1990) and Chen and Scott (1993).

⁶ This result is originally due to Cox and Ross (1976).

where r is the risk-free rate used for discounting,⁷ $\tau \equiv T - t$, and $\pi_{F_T,t}(f)$ is the pdf for F_T from time t evaluated at a terminal futures contract price of f . It is important to note that F_T is considered as a random variable.

Using the insights of Breeden and Litzenberger (1978), it is possible to extract the market-assigned cumulative distribution function (cdf) and pdf for F_T directly from Equation (1). Taking the first derivative of Equation (1) with respect to K yields

$$\frac{dP(K, t, T)}{dK} = e^{-r\tau} \int_{-\infty}^K \pi_{F_T,t}(f) df .$$

Where the cdf of F_T evaluated at K is defined by $\Pi_{F_T,t}(K) \equiv \int_{-\infty}^K \pi_{F_T,t}(f) df$, it is then true that

$$\Pi_{F_T,t}(K) = e^{r\tau} \frac{dP(K, t, T)}{dK} . \quad (2)$$

Equation (2) states that the cdf of F_T is equal to the (future value of the) first derivative of the put pricing function with respect to strike. Taking the derivative of Equation (2) with respect to K yields

$$\pi_{F_T,t}(K) = e^{r\tau} \frac{d^2P(K, t, T)}{dK^2} , \quad (3)$$

showing that the pdf for F_T is equal to the (future value of the) second derivative of the put pricing function with respect to strike.

After estimating $\pi_{F_T,t}(\cdot)$, it is possible to evaluate the overall methodology using Equation (1). Observed options can be re-priced using the estimated pdf, and these prices can then be compared against actual market prices. Smaller magnitudes of re-pricing errors generally indicate that the estimated pdf is increasingly reflective of market expectations. Pricing errors must also be balanced with the plausibility of the shape of the resulting pdf, a point further elaborated below.

2.3 Interpolation of Option Prices

In order to apply Equations (2) and (3), the instantaneous first and second derivatives of the put pricing function with respect to strike are required. Of course, we do not observe the latent option pricing

⁷ I set $r = 0$ for Short Sterling and Euribor options. Eurodollar discounting is based on a curve constructed using an overnight rate and Treasury yields at various tenors.

function, but only prices at discrete intervals. As a result, derivatives are computed numerically after interpolating put option prices to a sufficiently high granularity.

The question arises of how best to produce a stable and accurate interpolation function. The choice of interpolation specification is not trivial as we are directly interested in the slope and curvature of the estimated pricing function, which of course depend on the interpolation approach. Further, it is clear that the actual pricing function will exhibit a high degree of curvature, especially so near the mode of the pdf for the underlying (Equation (3) directly links the pdf with curvature in the put pricing function).

In order to replicate this curvature in a robust way I first transform prices to implied volatilities and strikes to deltas (specifically, put deltas), both of which are derived via the Black (1976) model for pricing options on futures.⁸ Only out-of-the-money put and call options are used as inputs, as these tend to be more liquid than in-the-money options and therefore embed more reliable information.⁹ Transforming option prices to implied volatilities was introduced by Shimko (1993). The implied volatility function tends to entail a lower degree of curvature than the price function, yielding a simpler space over which to interpolate. An option's delta expresses the rate of change in the option price relative to a change in the price of the underlying security. As a result, converting strikes to deltas provides two main benefits.¹⁰ First, deltas compress the domain under consideration: far out-of-the-money options have deltas approaching zero, while far in-the-money put options have deltas approaching $-e^{-r\tau}$. This eases extrapolation beyond the range of observed strikes as well as interpolation across time (both discussed further below). Second, because deltas change more rapidly near-the-money but group together for far in-the-money or out-of-the-money options, interpolating over deltas increases interpolation granularity near-the-money to better capture the inherent curvature of the pricing function.

A cubic smoothing spline is used to perform interpolation of implied volatilities over deltas. The cubic smoothing spline is a piecewise cubic polynomial function that has continuous first and second derivatives at observed data points, where the function is segmented (so-called "knots" of the spline). In addition to providing a very flexible functional form, the cubic smoothing spline offers an explicit balancing between smoothness of the interpolated function and fit errors that result from the smoothing process. Specifically, the cubic smoothing spline $g(\cdot)$ minimizes the following objective function:

$$(1 - \lambda) \sum_{i=1}^n w(i) [y(i) - g(x(i))]^2 + \lambda \int_x [g''(t)]^2 dt ,$$

⁸ Note that the Black (1976) assumption of a lognormally distributed underlying futures price at option expiration does not carry through to estimated pdf's. Rather, the Black (1976) model simply provides a mapping between prices and implied volatilities, as well as strikes and deltas.

⁹ Ivanova and Gutiérrez (2014) verify that out-of-the-money liquidity is greater for Euribor futures options, while Clews, Panigirtzoglou, and Proudman (2000) argue the same holds for Short Sterling. Figures B1 - B4 in Appendix B demonstrate the relatively high liquidity of out-of-the-money options on Eurodollar futures.

¹⁰ Interpolating implied volatilities over deltas was first performed by Malz (1997).

where x represents n observed deltas and y represents n implied volatilities, w is a set of weights, and λ is known as the “smoothing parameter”. The first term in the above objective function is simply the weighted sum of squared residuals at the spline’s knots, while the second term is meant to reflect the overall degree of curvature in the spline. The smoothing parameter fully determines the balance between fit residuals and spline smoothness. In fact, as $\lambda \rightarrow 1$, the spline approaches a weighted least squares (WLS) regression estimate. Weights w are imputed as each option’s vega, equal to the rate of change in price relative to implied volatility. As pointed out by Bliss and Panigirtzoglou (2002), because vega is the change in price relative to implied volatility, weighting by vega over implied volatilities is equivalent to constant weighting over prices.

The smoothing parameter λ presents a challenge, as it is not readily apparent that a particular value will prove ideal in any sense. One approach to determining λ automatically based on input data involves cross validation. This approach selects λ so as to minimize the average leave-one-out residual (i.e. for each i of the n observations, $g(\cdot)$ is estimated while ignoring data point i , after which the residual for data point i is calculated). In the case of options on Libor futures, testing of the cross validation approach yields clearly undesirable results. In a number of trials, cross validation selects a smoothing parameter at or near zero, which in turn produces implausibly choppy implied pdf’s. Examples of such pdf’s can be seen in Figures B6, B7, and B8 in Appendix B. These figures show estimates based on options maturing in December, 2010,¹¹ which were priced on January 4th, 2010, for each of Short Sterling, Euribor, and Eurodollar in turn. Respectively, cross validation selected smoothing parameters of 0.00000, 0.00000, and 0.00002 on a granularity of 0.00001 between 0 and 1.

Alternatively, one might think to implement a similar approach which would instead be based on pricing errors. As mentioned previously, Equation (1) can be used to re-price a set of options based on an estimated pdf for the underlying security. From these estimated prices, pricing errors can be computed as the difference between estimated and actual option prices. Instead of selecting λ to minimize leave-one-out implied volatility residuals, one could alternatively attempt to minimize leave-one-out pricing errors. However, this turns out to be prohibitively expensive computationally.¹²

Due to the undesirable performance of cross validation in selecting the smoothing parameter, I instead follow the alternative approach in the literature, which is to select a smoothing parameter based on subjective balancing between the plausibility of resulting pdf estimates and the magnitude of pricing

¹¹ Maturity dates are December 13, 2010 for EUR and USD Libor, and December 15, 2010 for GBP Libor.

¹² As elaborated below, the spline is evaluated at a large number of points, which quickly adds to computing time required for a cross validation approach based on pricing errors.

errors.¹³ I find $\lambda = .03$ tends to produce reasonable pdf's while keeping option pricing errors generally within one or two ticks.¹⁴ Figures B9 through B14 in Appendix B show example pdf's alongside the pricing errors that result after re-pricing options using the estimated pdf. Pricing errors are shown for the set of options used as inputs to each estimation, which includes out-of-the-money options meeting basic pricing assumption restrictions (as detailed further below). In Figures B9 through B14, the vast majority of pricing errors are within two ticks, while many are within one tick.

In order to ensure the estimated pdf integrates to 1, the fitted spline function is extrapolated so that the full delta range between $-e^{-r\tau}$ and 0 is covered.¹⁵ Extrapolation is performed with a quadratic polynomial function on either end of the observed delta range, as in Ivanova and Gutiérrez (2014). The full resulting function of implied volatilities over deltas is then evaluated at 50,000 evenly spaced delta values between the minimum and maximum deltas.¹⁶ Minimum and maximum deltas used for evaluation are offset by a small value from the actual bounds of $-e^{-r\tau}$ and 0, because these deltas imply strikes approaching ∞ and 0, respectively (see Appendix A). In the case that any implied volatilities are evaluated as non-positive, the estimation procedure is quit. Finally, the evaluated implied volatilities and deltas are converted back to prices and strikes via an inverse of the Black (1976) model.

Because the estimated put option prices remain discrete – as opposed to a continuous function – derivatives and integrals must be taken numerically. The formal definition of a derivative or an integral provides the basis for approximating continuous derivatives and integrals with discrete (but highly granular) data points. The following equations summarize these relationships for any function $f(x)$:

$$\frac{df(x)}{dx} \equiv \lim_{h \rightarrow 0} \frac{f(x) - f(x-h)}{h}, \quad (4)$$

and

$$\int_a^b f(x) dx \equiv \lim_{\max_{2 \leq i \leq n} (\Delta x_i) \rightarrow 0} \sum_{i=2}^n f \left(a + \sum_{j=2}^i (\Delta x_j) \right) (\Delta x_i), \quad (5)$$

¹³ For example, Datta, Londono, and Ross (2014) demonstrate graphically the tradeoffs involved in increasing or decreasing the smoothing parameter when estimating option-implied pdf's for future oil prices. They settle on a constant value for λ , which keeps pricing errors low while producing reasonable pdf estimates.

¹⁴ The tick size is the minimum price increment. For all three Libor-based contracts, the minimum increment is .005. An exception is made for Eurodollar option contracts that deliver the nearest expiring futures contract, as well as for low-priced options expiring in the two upcoming March-quarterly months, in which cases the minimum tick size is .0025. For simplicity and consistency, I will refer to the minimum tick size as .005, though observed prices with ticks of .0025 are not altered in any way.

¹⁵ Resulting pdf's are only reported if their area integrates to between 0.99 and 1.01

¹⁶ I use a very large number of points for evaluation in order to closely approach a continuous setting, while accounting for some balancing of the computation time required. As we are interested in approximating instantaneous first and second derivatives, it is important that the option price granularity is very fine in order to avoid bias. Further, applications that require integration over the estimated pdf - for example, re-pricing options or computing moments of the distribution - gain accuracy as the pdf approaches a near-continuous setting.

where in the second definition $n - 1$ is the number of discrete sub-intervals on (a, b) , and $\Delta x_i \equiv x_i - x_{i-1}$. Note that, while the second equation simplifies in the case where Δx_i is a constant value for all i , this does not hold for our purposes. Discrete observations are evaluated in implied volatility - delta space at a constant delta interval. However, as previously mentioned, an option's delta does not change at a constant rate with respect to strike – in fact, the gamma of an option specifies this rate of change across strikes. Therefore, the interval between strikes is not constant, so the above formula is applied when integrating.

The process of smoothing implied volatilities over put option deltas does introduce the possibility of negative estimated probabilities, even after cleaning input option prices as described further below. In rare cases, the estimated set of implied volatilities over deltas results in a subset of implied volatilities exhibiting large swings when placed over strikes, a pattern carried through to the relationship between prices and strikes. It is apparent in these cases that the set of implied volatilities represents a situation in which implied volatility – and, thereby, price – is not a proper function of strike (i.e. at least one strike exists that corresponds to more than one implied volatility), which contradicts Black's (1976) option pricing model. This is analogous to considering the points (0,0), (1,-1), (4,2), and (9,-3): at some degree of granularity, it becomes apparent that the function being traced out is neither piecewise linear nor sinusoidal, but rather displays $y = \pm\sqrt{x}$. Unfortunately, there is no clear restriction to impose for the discrete set of implied volatilities over strikes (for example, monotonicity). However, because such a set of implied volatilities tends to produce swings in prices over strikes and thereby negative estimated probabilities (and because negative probabilities are not sensible), any estimated pdf's are only used in further analysis if they produce no negative probabilities.

2.4 Converting Futures Prices to Interest Rates

The preceding methodology is used to obtain a highly granular set of put option prices over strikes. These prices apply to options written on the relevant futures contract price, which is based on a simple transformation of the Libor rate being referenced. Specifically, where F_T is the random variable denoting the futures contract price at time T , then $R_T = 100 - F_T$ is the corresponding annualized 3-month Libor rate in percentage points. Of course, we are directly interested in probability assignments for future Libor interest rate realizations, rather than for futures contract prices.

Because prices are interpolated and extrapolated in the setting of options on the futures price, Equations (2) and (3) are applied initially in the same setting. Probabilities are transformed to the interest rate setting after being estimated in a futures price setting, as

$$\begin{aligned}
& \text{Prob}_t(R_T \leq a) \\
&= \text{Prob}_t(100 - R_T \geq 100 - a) \\
&= 1 - \text{Prob}_t(100 - R_T \leq 100 - a) \\
&= 1 - \text{Prob}_t(F_T \leq 100 - a),
\end{aligned} \tag{6}$$

which gives the desired cdf over interest rates.¹⁷ Similarly, when considering pdf's,

$$\pi_{R_T,t}(a) = \pi_{100-F_T,t}(a) = \pi_{F_T,t}(100 - a). \tag{7}$$

Interpolating and extrapolating option prices in the setting of the futures contract price (F_T) provides a useful tool for circumventing limitations embedded in the Black (1976) option pricing model. The Black (1976) model assumes a lognormal probability distribution for the underlying security, which in turn implies no probability mass assigned below zero (see Appendix A). As noted previously, this assumption does not carry through to probability functions estimated as in this paper; the Black (1976) model is employed only to provide a mapping between prices and implied volatilities, as well as between strikes and deltas. However, these mappings break down for strikes below zero. As a result, while the above methodology could be applied directly to the interest rate setting by subtracting option strikes from 100 and reversing call and put options (e.g. 97.25 put \rightarrow 2.75 call), this approach would break down for options with converted strikes below zero (or, original strikes above 100). Deriving probabilities for the futures contract price before converting to the interest rate setting then provides the advantage of retaining information embedded in options with strikes above 100. In this way, one can actually view the futures price as a useful transformation of the Libor rate as relates to estimating pdf's.

2.5 Cleaning Input Prices

In addition to restricting input prices to out-of-the-money options, inputs are screened so they meet basic option pricing assumptions. These assumptions can be thought of in terms of the price level, slope, and curvature over strikes. First, all option prices of zero or lower are immediately removed.

The next two restrictions regarding slope and curvature of option prices are simple results of Equations (2) and (3), combined with the fact that probabilities are defined as non-negative. Equations (2)

¹⁷ Note the implicit assumption that the probability of a random variable Z being equal to any particular value is assumed to be zero. Or, $\lim_{\epsilon \rightarrow 0} \text{Prob}(k - \epsilon < Z < k + \epsilon) = 0$ for any constant k . This allows equalities and inequalities to be used interchangeably.

and (3) reveal that the first and second derivatives of the put pricing function must be non-negative, as the corresponding cdf and pdf are non-negative functions. Setting up an equation for the call pricing function $C(K, t, T)$ analogous to Equation (1) and taking derivatives yields similarly that¹⁸

$$\Pi_{F_T, t}(K) = e^{r\tau} \frac{dC(K, t, T)}{dK} + 1 \quad (8)$$

and

$$\pi_{F_T, t}(K) = e^{r\tau} \frac{d^2C(K, t, T)}{dK^2}. \quad (9)$$

Equation (8) implies that

$$-e^{-r\tau} \leq \frac{dC(K, t, T)}{dK} \leq 0.$$

Therefore, the call pricing function must be non-increasing over strikes, and from Equation (9) it must have non-negative curvature.

When option prices violate either the slope or curvature assumption, it is not obvious which option is “in the wrong” and should be removed. In any subset of prices that violates the slope or curvature assumption, I remove the option furthest from the money. There are two reasons for this selection. First, deeper out-of-the-money options are likely to be less liquid, so their prices are more likely to be stale. Second, the minimum tick size imposed on prices can be seen as introducing an error to the prices that would be observed in a continuous setting, which bears relatively more importance to deeper out-of-the-money options with already low prices. In practice, I allow for slight deviations from the slope and curvature assumptions. This accommodates the possibility that prices are not perfectly arbitrated, and accounts for a small degree of price inaccuracy due to discrete pricing ticks.

After option prices are cleaned, the estimation procedure is only carried out if there remain at least three option prices comprising at least one call and one put.

2.6 Fixed Horizon Estimates

The procedure outlined above lends itself naturally to estimating the put option pricing function for fixed horizons, which in turn can be used to generate fixed horizon probabilities. Option contracts allow

¹⁸ A full derivation can be found in Appendix A.

directly for estimation of pdf's corresponding to a specific option maturity date. However, we might also be interested in estimates for a fixed horizon, for example 180 days ahead from the option pricing date. While the former type of estimate can show how expectations for interest rates at a particular point in the future evolve over time, fixed horizon estimates control for the impact of time to maturity. Pdf's for a particular maturity date will tend to narrow with the passing of time, as more information is garnered on the likely price of the underlying futures contract at expiration. On the other hand, fixed horizon estimates demonstrate the evolution of market expectations unaffected by time to maturity.

Fixed horizon estimates are computed by interpolating implied volatility - put delta functions across time, similar to Datta, Londono, and Ross (2014). For a given hypothetical fixed horizon maturity date, data is collected for the March-quarterly contracts with expirations that most closely precede and follow that date. Each set of options is then used to estimate a function of implied volatilities over put option deltas. For each of 25,000 evenly spaced deltas over the delta range corresponding to the fixed maturity date, the corresponding implied volatility is the linear interpolation between the two functions, weighted based on the fixed maturity date's proximity to each contract's actual maturity date. Implied volatilities and deltas are again transformed back to prices and strikes, from which probabilities can be estimated.

Fitting implied volatilities over deltas instead of strikes makes for more consistent interpolation of implied volatilities across time, as the domain covered is always $(-e^{-r\tau}, 0)$. For a hypothetical fixed maturity date, the range of strikes that should be used for interpolation would not be clear, as the support of the pdf $\pi_{F_T,t}(\cdot)$ over strikes varies with t and T – even for a constant τ .

2.7 Calculating Distributional Moments

Once a pdf has been estimated from option prices, computation of distributional moments provides useful summary statistics, particularly when these moments are viewed over a span of time. Central moments of the distribution are computed using the traditional formulas, which are applied to the discrete case using the definition found in Equation (5) to compute integrals. Appendix A contains the continuous-case formulas referenced for calculating moments, with skewness and kurtosis being normalized by the standard deviation (square root of variance). In the remainder of this paper, references to skewness and kurtosis are in terms of the normalized computations. Additionally, I focus on the standard deviation rather than variance, as it has a more natural interpretation with units being in percentage points.

While distributional moments are discussed in terms of implied pdf's on future interest rates rather than futures contract prices, moments are initially computed in regard to the futures price and are then transformed to the interest rate setting. It is straightforward to show that the following relationships hold:

$$\text{Mean}(R_T) = 100 - \text{Mean}(F_T)$$

$$\text{Variance}(R_T) = \text{Variance}(F_T)$$

$$\text{Skewness}(R_T) = -1 \times \text{Skewness}(F_T)$$

$$\text{Kurtosis}(R_T) = \text{Kurtosis}(F_T)$$

It is worth briefly noting the potential influence that the risk-neutral assumption of this paper's methodology might have on estimated distributional moments. If investors are more risk averse to higher (lower) relative to lower (higher) interest rates, then option payouts in the case of higher (lower) rates will be priced above what the actual market-assigned probabilities would imply, resulting in a higher (lower) estimated pdf mean and skewness. Similarly, if investors are relatively more risk averse to a more (less) extreme interest rate environment, estimates of kurtosis might be overstated (understated). Variance (and thereby standard deviation) are likely to be overstated to the extent that investors are risk averse generally, as this would increase the price of out-of-the-money options used as inputs, implying a wider spread of probability mass. Ivanova and Gutiérrez (2014) provide evidence that risk-neutral pdf estimates for future Euribor rates likely overstate the tails of the distribution in general, and particularly the right-hand tail as relates to interest rate levels (i.e. investors are more risk averse to larger deviations from the expected future rate, and are particularly risk averse to higher interest rates).

3 Empirical Case Studies

The methodology laid out above is now applied to four sample case studies. These case studies demonstrate the efficacy of such an approach in illuminating the evolution of market sentiment regarding future interest rates. Particularly volatile historical cases with respect to financial market conditions are selected for study, as these cases entail immediate and apparent changes in interest rate expectations.

3.1 Case 1: BNP Paribas Freezes Funds¹⁹

Paramount in triggering the 2007-08 financial crisis was heavy exposure throughout financial markets – especially the exposure carried by large and systemically important banks – to the performance of subprime home mortgages. As derivative products including asset-backed securities (ABS) (in particular, mortgage-backed securities (MBS)) and collateralized debt obligations (CDO) spread, the potential for

¹⁹ News sources: “BNP Paribas suspends funds” (2007, August 9); Kingsley (2012, August 6).

a significant market crash on the back of a spike in mortgage defaults grew. So markets were on edge when French bank BNP Paribas announced on August 9th, 2007 that it would suspend valuation for three of its ABS-based funds, claiming it was “impossible to value certain assets fairly” in those funds alongside disappearing liquidity. Indeed, the BNP Paribas announcement was a watershed moment in the development of the financial crisis; as Adam Applegarth – former CEO of the nationalized British bank Northern Rock – put it, August 9th was “the day the world changed.”

Figure 1: PDF Evolution, Case 1

EUR Libor, Maturity = Dec. 17, 2007





Figure 1 shows estimated option-implied pdf's for the 3-month Euribor rate on December 17th, 2007, using options priced on August 8th and 9th, 2007. While actual and expected interest rates would continue to slide downward as the financial crisis unfolded, it is clear that the BNP Paribas announcement made an early contribution to investor expectations of monetary policy easing. Figure 1 reveals that much of the change in interest rate expectations following the BNP Paribas announcement occurred through a movement of probability mass toward the left tail of the distribution and largely out of the center, rather than a simple shift of the entire distribution. In fact, though the spot 3-month Euribor rate had been steadily trending upward since late 2005 to reach 4.35% by August 8th, 2007, the assigned probability of falling back below 4.25% by mid-December nearly doubled from 6.84% to 12.53% in the course of one day. Uncertainty for where interest rates would lie in December increased as well: the standard deviation widened from 16 basis points to 18 basis points.

A similar trend can be seen in expectations for the December, 2007 USD Libor rate in Figure 2 below. Such similarity is likely reflective of the extent of financial market exposure to asset-backed securities whose returns were highly correlated. Normalized skewness and kurtosis decreased as the standard deviation jumped from 42 to 58 basis points. The mean of the distribution fell by 16 basis points from 5.12% to 4.96%

Figure 2: PDF Evolution, Case 1

USD Libor, Maturity = Dec. 17, 2007



3.2 Case 2: Lehman Brothers Files for Bankruptcy²⁰

Of course, the bankruptcy of Lehman Brothers was a further watershed moment in the unfolding of the financial crisis. On September 15th, 2008, Lehman Brothers filed for bankruptcy following the decision by Federal Reserve and United States Treasury officials not to bail out the bank. As broader economic conditions continued to sour, the Federal Reserve would hold an unscheduled meeting weeks later and decide to resume interest rate cuts, which had seen a respite during the preceding six months. The collapse of Lehman Brothers marks an important turning point in the depth of the financial crisis.

Figure 3: PDF Evolution, Case 2

USD Libor, Maturity = Dec. 15, 2008




Figure 3 shows estimated pdf's for USD Libor on December 15th, 2008, using options priced on September 12th and 15th of the same year. Though the Federal Reserve had left its target interest rate

²⁰ News source: Sorkin (2008, September 14).

unchanged at the most recent three Federal Open Market Committee (FOMC) meetings, Lehman's failure appears to have immediately set off expectations for further monetary policy easing. Similar to the BNP Paribas case, the change in interest rate expectations resulting from Lehman was not a straightforward shift of the entire distribution, but rather movement of probability from the center of the distribution toward the left tail. The standard deviation of the distribution concomitantly increased from 43 to 57 basis points.

3.3 Case 3: European Markets Deteriorate; ECB Responds²¹

Late summer 2011 was a tumultuous period for financial markets, particularly so in Europe. Broadly, global growth concerns were on the rise as composite purchasing managers indexes (PMI's) dipped for both the United States and the Euro area. Especially salient in Europe was burgeoning skepticism over the ability of so-called European periphery countries – including Greece, Ireland, Italy, Portugal, and Spain – to meet creeping debt obligations alongside rising bond yields. On August 4th tensions boiled over: an ECB intervention in bond markets ignoring Spanish and Italian bonds was interpreted as signalling economic conditions in those countries so dire as to be unworthy of attempted rescue, while a letter written by European Commission president José Manuel Barroso indicated that risks were not confined to the periphery; stock markets tumbled, with the Dow Jones Euro Stoxx index down 3.75% on the day and the S&P 500 down 4.78%.

On August 4th, the ECB announced that it would provide loans to banks at fixed rates through early 2012, while simultaneously announcing a longer-term refinancing operation. Days later, on August 7th, the ECB decided to resume bond purchases under the Securities Markets Programme (SMP), which had not been active since March 2011. The reinstatement of SMP purchases was intended to relieve the pressures of high interest rates on government borrowing.


Figure 4 details the evolution of pdf's for the 3-month Euribor rate on December 19th, 2011 over the most volatile few days of this episode. The evolution was not a linear one. As market pessimism came to a head on August 4th, skewness dropped from near-zero into more negative territory, reflecting greater tail probability for downward interest rate moves. Interestingly, uncertainty decreased slightly as the standard deviation dropped from 42 to 39 basis points, possibly reflective of a stabilizing effect from the ECB's August 4th policy measures. Much of the remainder of the distribution caught up with the left tail on August 5th as skewness returned to near-zero, while the standard deviation bounced back from 39 to 45 basis points. Following the revival of SMP, the mean of the distribution pushed further downward on August 8th (with a net decline of 37 basis points from August 3rd). Skewness turned positive as the

²¹ News sources: Bowley (2011, August 4); European Central Bank (2011, September); Rettman (2011, August 4).

left tail of the distribution became increasingly constrained by the ZLB while the right tail remained significant.

Figure 4: PDF Evolution, Case 3

EUR Libor, Maturity = Dec. 19, 2011




3.4 Case 4: BOE Governor Carney Says Rates May Remain Low “for some time”²²

A few months into 2014, market participants had come to expect that the Bank of England (BOE) might begin increasing its target interest rate in the coming year – which would mark the first change to the BOE’s policy rate since it had bottomed out at 0.50% in March, 2009. However, when Governor Mark Carney delivered the Bank’s regular Inflation Report on May 14th, 2014, these expectations were subdued with language considerably more dovish than anticipated. Figure 5 below shows pdf’s for the 3-month GBP Libor rate on December 15th, 2014, based on option prices from May 13th and 14th. Following Governor Carney’s remarks, the standard deviation fell as the distribution compressed back toward zero, while normalized skewness and kurtosis increased. In the process, the mean rate expectation dropped from 80 to 72 basis points.

²² News source: “Bank of England dampens interest rate rise speculation” (2014, May 14).

Figure 5: PDF Evolution, Case 4

GBP Libor, Maturity = Dec. 15, 2014



4 Cross-Rate Relationships in the Term Structure of Distributional Moments²³

Using fixed horizon pdf estimates at varying horizons, it is possible to view elements of the term structure of distributional moments for future interest rates across time. Much attention has been paid to the term structure of interest rates, especially concerning Treasury yield curves. Examining the term structure of distributional moments around future interest rate expectations complements traditional yield curve analysis by painting a fuller picture of market views. Further, with 3-month Libor rates as the underlying, emphasis here is given to short rates viewed over the near term. This section compares simple measures of the level, slope, and curvature of distributional moments on future 3-month Libor rates denominated in US dollars, euros, and UK pounds. Analysis focuses on more recent history, which is due in part to greater consistency of data availability, as well as an effort to focus attention on the period of historically low interest rates following the financial crisis of 2007-08.

Figures B15 through B18 in Appendix B show the estimated mean, standard deviation, (normalized) skewness, and (normalized) kurtosis for the three Libor rates at fixed horizons of 180, 360, and 540 days ahead. The series span from the beginning of 2009 through March 31st, 2016. A cursory view of these figures yields some insight into market expectations over this period, which is made more rich by considering the term structure of multiple moments in tandem. For example, an important trend in financial markets over recent years has been the divergence in interest rates between the Euro area and other advanced economies including the United Kingdom and the United States. Figure B15 reflects that markets have consistently expected an increasing near-future path for GBP and USD Libor rates

²³ The “term structure” here is considered across varying forward horizons with the same forward terms (3 months), as opposed to being across spot rates with varying terms.

since mid-2013. This timing coincides with the “taper tantrum”, when in May, 2013 Federal Reserve Board Chairman Ben Bernanke stated the Fed might begin slowing the pace of bond purchases under its quantitative easing program. The “taper tantrum” period did see an increased slope for projected Euribor rates, though this quickly faded, and the difference between 540-day-ahead and 180-day-ahead Euribor expectations has remained near zero since early 2014. As the slopes of mean rate expectations for GBP and USD Libor have become elevated, so too has uncertainty increased for these rates: Figure B16 shows that the standard deviation for each rate 540 days ahead has remained higher since mid-2013. As Euribor mean expectations approached – and eventually breached – zero over the same period, the standard deviation generally fell as the distribution compressed, with a slight trend upward since mid-2014. Figures B17 and B18 show that skewness and kurtosis both jumped for GBP Libor in mid-2013, especially for rates 180 days ahead, while both moments have broadly decreased across the term structure since then for both GBP and USD Libor. Overall, as mean expectations for GBP and USD Libor rates have increased, the bulk of each distribution has generally moved upward, while the right-hand tail has not seen a comparable shift. Skewness and kurtosis for Euribor rates have not seen a clear secular trend over the past few years.

4.1 Granger Causality across Rates

It is useful to understand the time series dynamics of distributional moments across Libor rates for various components of the term structure. I examine these cross-rate dynamics by testing for Granger causality in a vector autoregression (VAR). The general VAR is expressed as follows:

$$Y_t = C + Z_1 Y_{t-1} + Z_2 Y_{t-2} + \dots + Z_k Y_{t-k} + U_t, \quad (10)$$

where each Y_i is a 3-by-1 vector of the given moment and term structure component pairing (for example, standard deviation level 360 days ahead) for each Libor rate at time i , C is a 3-by-1 vector of constants, each Z_i is a 3-by-3 matrix of coefficients, U_t is a 3-by-1 vector of error terms at time t , and k is the number of lags included in the VAR.²⁴ In total, 24 VAR’s of this form are estimated, corresponding to the number of combinations between distributional moments and term structure components being considered. The moments include mean, standard deviation, skewness, and kurtosis. The term structure components are levels at 180, 360, and 540 days ahead; the difference between levels at 360 and 180 days ahead (referred to as “Slope 1”); the difference between levels at 540 and 360 days ahead (referred to as “Slope 2”); and the difference between Slope 2 and Slope 1 (referred to as “Curvature”). Slope and curvature measures capture information regarding the trajectory of moments abstracted from their levels.

²⁴ Lag length is determined by minimization of Schwarz’s Bayesian Information Criterion.

After a VAR is estimated, Granger causality is determined by Wald tests. Let each Y_i retain the same ordering with respect to the three Libor rates. For a particular VAR, consider the test that the relevant pairing of moment and term structure component for the second-positioned rate Granger causes that for the first-positioned rate (e.g. 540-day-ahead kurtosis for Euribor Granger causes that for GBP Libor). This is equivalent to testing whether the entry in the first row, second column of each Z_i matrix of coefficients is equal to zero, with the null hypothesis being that all such entries are equal to zero. A rejection of the null hypothesis implies Granger causality, meaning previous values of the relevant pairing for the second-positioned rate contain explanatory power for the current value of the pairing for the first-positioned rate, controlling for previous values of the pairing for the first- and third-positioned rates (e.g. past values of 540-day-ahead kurtosis for Euribor contain explanatory power for the current value of 540-day-ahead kurtosis for GBP Libor, controlling for previous values of 540-day-ahead kurtosis for GBP and USD Libor). Such tests are revealing of the directions in which expectations for different Libor rates may affect one another across various distributional and term structure characteristics. It is worth noting that Granger causality is not the same as formal causality, but should be interpreted simply as past values of one variable containing residual explanatory power for the current value of another variable.

When testing for Granger causality in a VAR with possibly integrated time series, Toda and Yamamoto (1995) demonstrate that the distribution under the null hypothesis is non-standard. However, they introduce a simple solution: after determining the lag length k of the VAR by use of an information criterion, one estimates a VAR using lag length $k + d$ with d being the maximum order of integration, then performs Wald tests for Granger causality using k lags of explanatory variables. This approach resolves asymptotic inference. Relevant to this paper, augmented Dickey-Fuller (ADF) tests fail to reject the null hypothesis of nonstationarity in many of the time series considered. As a result, tests for Granger causality are performed both in the traditional manner (the VAR still being estimated with k lags) as well as using Toda and Yamamoto's (1995) method where applicable.

Tables B2 and B3 in Appendix B synthesize the results of all Granger causality tests. The sample period begins on May 7th, 2009 (the date on which the European Central Bank cut its main policy rate target to 1%, rounding out the series of rate cuts in the United States, United Kingdom, and the Euro area following the financial crisis onset) and continues through March 31st, 2016. The results in Table B2 use the traditional approach to testing for Granger causality. In Table B3, Toda and Yamamoto's (1995) method (with $d = 1$) is used for testing Granger causality in all VAR's in which an ADF test fails to reject the null hypothesis of nonstationarity in at least one of the time series at the 5% confidence level, and uses the traditional approach in remaining VAR's. The two tables serve as a robustness check on one another, as well as a basic indication of the extent to which nonstationarity might be affecting the results of traditional Granger causality tests in this case. VAR's are set up as in Equation (10), and

resulting bivariate Granger causality results are reorganized in both tables to ease reading of relationships between pairs of Libor rates. An “x” denotes Granger causality significant at the 5% confidence level. As a concrete example, the “x” under the GBP-180 column and the EUR-Mean row in either table denotes that in a VAR of 180-day-ahead mean rates among all three Libor rates, the Euribor mean Granger causes the GBP Libor mean at a 5% significance level under the relevant testing procedure. The sections of each table corresponding to tests of autocorrelation are left blank, as all series are autocorrelated. The exception is the EUR-EUR section in Table B3, which is used as a marker to denote which results come from VAR’s using Toda and Yamamoto’s (1995) method; a blue “O” signifies that Granger causality tests for the corresponding pairing of distributional moment and term structure component are performed using Toda and Yamamoto’s (1995) approach.

I note here some broad takeaways from Tables B2 and B3 in an attempt to summarize the results; this is not meant to be comprehensive, but rather to draw attention to some important patterns that emerge. The most basic trend to note in Tables B2 and B3 is the striking overall degree of interconnectedness among the three Libor rates, present across varying distributional moments and term structure components. It is immediately clear that – in addition to levels of mean expectations – combinations of higher distributional moments and simple measures of the broader term structure bear relationships across rates as well. In other words, the cross-rate dynamics are not thoroughly captured by simply examining the levels of future mean expectations; there is a rich interplay among rates involving standard deviation, skewness, and kurtosis, as well as measures of the slope and curvature of future paths.

USD Libor expectations contain the most consistent explanatory power across moments and term structure components, particularly so in relation to GBP Libor rates. USD Libor expectations are found to Granger cause GBP Libor expectations across both Tables B2 and B3 in 20 of 24 combinations of distributional moment and term structure component, with the only exceptions being levels of skewness and kurtosis 540 days ahead, Slope 2 for kurtosis, and Curvature for skewness. Similarly, the mean and standard deviation of USD Libor expectations Granger cause those of Euribor expectations across all term structure components except the Curvature of standard deviation; USD Libor contains little explanatory power for Euribor in terms of skewness and kurtosis, however. There are a couple potential explanations for the relatively great importance of USD Libor rate expectations to those for EUR and GBP Libor. It is certainly possible that USD-denominated interest rates hold greater causal sway relative to alternative interest rates; for example, Brusa, Savor, and Wilson (2016) find that Federal Reserve policy decision announcements bear an outsized influence across global stock market risk premia, which is not true for any other central bank they examine. It may also be true, however, that USD Libor expectations simply serve as a stronger proxy for omitted macro variables, which in turn are linked causally to EUR and GBP Libor expectations. Distinguishing between these explanations is beyond the scope of this paper, though

the evidence of Brusa, Savor, and Wilson (2016) seems to imply the former explanation holds water.

While USD Libor rate expectations hold much explanatory power for EUR and GBP Libor expectations, the converse is not true. In fact, Euribor expectations are statistically important precursors of USD Libor expectations in only 2 out of 24 total pairings of moment and term structure component when applying Toda and Yamamoto's (1995) testing procedure. Similarly, while the levels of GBP Libor mean expectations consistently Granger cause their USD Libor counterparts, higher distributional moments of GBP Libor expectations bear little importance to USD Libor when following Toda and Yamamoto (1995). On the whole, higher moments than the mean of USD Libor expectations are consistently not found to be Granger caused by those of either EUR or GBP Libor rates, while it is of course important to note that levels of mean expectations for GBP Libor do Granger cause USD Libor mean expectations.

EUR and GBP Libor expectations share clear relationships running in both directions, especially so for means and standard deviations. The levels of mean expectations for GBP Libor Granger cause those for Euribor across fixed horizons of 180, 360, and 540 days ahead. Additionally, while levels of GBP Libor standard deviations are not very consistently related to those for Euribor, the slopes and curvature of standard deviations over time are; so, while the actual levels of uncertainty in GBP Libor rates are less reliably linked to levels of uncertainty in Euribor, the path at which GBP Libor uncertainty is evolving does help predict the path of Euribor uncertainty. In the opposite direction, the standard deviation of Euribor expectations Granger causes that of GBP Libor expectations across term structure components, with the exception of the 180-day-ahead level. Mean expectations for Euribor have little predictive power for GBP Libor expectations across term structure components.

5 Conclusion

This paper uses options on Libor futures for rates denominated in EUR, GBP, and USD in order to estimate risk-neutral market-assigned pdf's for future interest rates. The methodology is reviewed, and estimates are applied to produce case studies of probability estimates as they evolve over important market events. Distributional moments are compared across rates to study their interrelationships in a Granger causality framework.

The case studies carried out in this paper elucidate the significant impact of specific market events on market expectations for future interest rates by detailing precisely how expectations changed along the full probability distribution. For example, while the mean expectation for USD-denominated Libor fell upon news of the collapse of Lehman Brothers in September, 2008, the probability distribution did not undergo a simple level shift; rather, probability moved toward the left tail and primarily out of the center of the distribution, resulting not only in a lower mean expectation but also a higher standard deviation.

The evolution of expectations for GBP-denominated Libor rates as BOE Governor Mark Carney stated interest rates might remain low “for some time” provides evidence on the critical role of statements by central bank authorities in shaping market beliefs.

Granger causality tests considering the future path of distributional moments across the three Libor rates demonstrate the large extent to which expectations for the three rates are interrelated. The dynamics among mean expectations for future interest rates do not fully capture the rich interconnectedness among broader probabilistic expectations for these rates. I find generally that expectations for USD-denominated Libor rates bear the most consistent predictive power across combinations of distributional moment and term structure component for EUR and GBP Libor, while EUR and GBP Libor expectations do not provide much consistent explanatory power for USD Libor. USD Libor holds consistent explanatory power for GBP Libor across mean, standard deviation, skewness, and kurtosis, while being linked to Euribor primarily in mean and standard deviation. GBP Libor holds predictive power for mean levels as well as the slope and curvature of standard deviation in Euribor expectations, while Euribor standard deviation Granger causes that of GBP Libor across levels (excluding the shortest horizon of 180 days), slope, and curvature.

Further research might build upon the Granger causality results established here with estimates of the magnitude of various relationships, for example in a multivariate vector error correction model among cointegrated series. Additionally, it would be useful to study the relationships of the distributional moment time series used in this paper with other economic variables of interest, including equity indexes or exchange rates. Time series of distributional moments might be derived from calibrated real-world (i.e. not risk-neutral) probability densities as well, as detailed in Ivanova and Gutiérrez (2014); it would be instructive to test whether results in this paper are sensitive to the assumption of risk-neutrality. Finally, further analysis might view the relationships among Libor rate expectations (and possibly other economic variables) as they have evolved over time.

References

- Bahra, Bhupinder (1997). Implied risk-neutral probability density functions from option prices: Theory and application. Working paper no. 66, Bank of England.
- Bank of England dampens interest rate rise speculation. (2014, May 14). BBC. Retrieved from <http://www.bbc.com>
- Barone-Adesi, Giovanni, and Robert E. Whaley (1987). Efficient analytic approximation of American option values. *The Journal of Finance*, 42(2), 301-320.
- Black, Fischer (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3, 167-179.
- Bliss, Robert R., and Nikolaos Panigirtzoglou (2002). Testing the stability of implied probability density functions. *Journal of Banking and Finance*, 26, 381-422.
- BNP Paribas suspends funds because of subprime problems. (2007, August 9). *The New York Times*. Retrieved from <http://www.nytimes.com>
- Bowley, Graham (2011, August 4). Stocks plunge on fears of global turmoil. *The New York Times*. Retrieved from <http://www.nytimes.com>
- Breeden, Douglas T., and Robert H. Litzenberger (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business* 51(4), 621-651.
- Brusa, Francesca, Pavel Savor, and Mungo Wilson (2016). One central bank to rule them all. Working paper.
- Chen, Ren-raw, and Louis Scott (1993). Pricing interest rate futures options with futures-style margining.
- Clews, Roger, Nikolaos Panigirtzoglou, and James Proudman (2000). Recent developments in extracting information from options markets. *Bank of England Quarterly Bulletin*, February 2000, 50-60.
- Cooper, Neil (1999). Testing techniques for estimating implied RNDs from the prices of European-style options. Working paper.
- Cox, John C., and Stephen A. Ross (1976). The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, 3, 145-166.

- Datta, Deepa D., Juan M. Londono, and Landon J. Ross (2014). Generating options-implied probability densities to understand oil market events. International Finance Discussion Paper no. 1122, Board of Governors of the Federal Reserve System.
- European Central Bank (2011, September). Monthly bulletin: September 2011. 47-53.
- Ivanova, Vesela, and Josep Maria Puigvert Gutiérrez (2014). Interest rate forecasts, state price densities and risk premium from Euribor options. *Journal of Banking and Finance*, 48, 210-223.
- Kingsley, Patrick (2012, August 6). Financial crisis: Timeline. *The Guardian*. Retrieved from <https://www.theguardian.com>
- Lieu, Derming (1990). Option pricing with futures-style margining. *Journal of Futures Markets*, 10(4), 327-338.
- Malz, Allan M. (1997). Option-implied probability distributions and currency excess returns. Staff Reports no. 32, Federal Reserve Bank of New York.
- Melick, William R. and Charles P. Thomas (1997). Recovering an asset's implied PDF from option prices: An application to crude oil during the Gulf Crisis. *The Journal of Financial and Quantitative Analysis*, 32(1), 91-115.
- Rettman, Andrew (2011, August 4). Barroso raises alarm about severity of euro crisis. *EUobserver*. Retrieved from <https://euobserver.com>
- Shimko, David (1993). Bounds of probability. *Risk*, 6, 33-37.
- Sorkin, Andrew Ross (2008, September 14). Lehman files for bankruptcy; Merrill is sold. *The New York Times*. Retrieved from <http://www.nytimes.com>
- Tian, Yisong S. (2011). Extracting risk-neutral density and its moments from American option prices. *Journal of Derivatives*, 18(3), 17-34.
- Toda, Hiro Y., and Taku Yamamoto (1995). Statistical inference in vector autoregressions with possibly integrated processes. *Journal of Econometrics*, 66, 225-250.

Appendix A: Technical Appendix

Black (1976) Model

Conversion between prices and implied volatilities occurs through the Black (1976) model for pricing options on futures contracts. For a put option, the Black (1976) price is equal to

$$P = e^{-r\tau} (K\Phi(-d_2) - F\Phi(-d_1)) ,$$

where

$$d_1 = \frac{\ln\left(\frac{F}{K}\right) + \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}}$$

and

$$d_2 = \frac{\ln\left(\frac{F}{K}\right) - \frac{\sigma^2}{2}\tau}{\sigma\sqrt{\tau}} = d_1 - \sigma\sqrt{\tau} .$$

Adding to prior notation, $\Phi(\cdot)$ is the cumulative distribution of the standard normal, F is the futures contract price, and σ is the implied volatility.

Relating Deltas and Strikes

A put option's delta, Δ_p , denotes the sensitivity of the option price to a change in the underlying futures rate, or

$$\Delta_p = \frac{\partial P}{\partial F} = e^{-r\tau} (\Phi(d_1) - 1) .$$

From this, one can solve for the strike of a put option given its delta as the following:

$$K = Fe^\alpha ,$$

where

$$\alpha = \frac{\sigma^2}{2}\tau - (\sigma\sqrt{\tau})\Phi^{-1}(e^{r\tau}\Delta_p + 1) .$$

The above formula makes clear the bounds on Δ_p . The inverse of the standard normal cdf, $\Phi(\cdot)$, can only be evaluated on $(0, 1)$. Therefore, it must be that $0 < (e^{r\tau}\Delta_p + 1) < 1$, or $-e^{-r\tau} < \Delta_p < 0$. Additionally, it is the case that

$$\lim_{\Delta_p \rightarrow -e^{-r\tau}+} K = \infty,$$

and

$$\lim_{\Delta_p \rightarrow 0^-} K = 0.$$

Because of this, when evaluating implied volatilities as a function of deltas I place the end points for evaluation a small value above $-e^{-r\tau}$ and a small value below 0.

Vega

An option's vega, v , is the sensitivity of the option price to a change in implied volatility. Specifically, in the case of a put option,

$$v_P = \frac{\partial P}{\partial \sigma} = Fe^{-r\tau} \phi(d_1) \sqrt{\tau},$$

where $\phi(\cdot)$ is the standard normal density function.

Proof of Equations (8) and (9)

As in the case of put options, we begin by describing the call option pricing function as

$$\begin{aligned} C(K, t, T) &= e^{-r\tau} E_t[\max(F_T - K, 0)] \\ &= e^{-r\tau} \int_K^{\infty} (f - K) \pi_{F_T, t}(f) df. \end{aligned}$$

Then, taking the derivative with respect to K using Leibniz's rule for differentiating under an integral yields

$$\begin{aligned} \frac{dC(K, t, T)}{dK} &= -e^{-r\tau} \int_K^{\infty} \pi_{F_T, t}(f) df \\ &= -e^{-r\tau} (1 - \Pi_{F_T, t}(K)) \\ &= e^{-r\tau} (\Pi_{F_T, t}(K) - 1). \end{aligned}$$

Solving for $\Pi_{F_T, t}(K)$,

$$\Pi_{F_T,t}(K) = e^{r\tau} \frac{dC(K,t,T)}{dK} + 1,$$

which is Equation (8). It is then straightforward that taking the derivative of Equation (8) with respect to strike K yields Equation (9).

Distributional Moment Formulas

Continuous-case references used for calculating distributional moments are listed below. In practice, integrals are computed using the discrete-case equivalent via Equation (5) applied over all option strikes evaluated in the course of interpolating and extrapolating option strikes and prices. Below, μ_X is equivalent to Mean(X), and σ_X is the standard deviation of X, or the square root of Variance(X).

$$\begin{aligned} \text{Mean(X):} \quad E[X] &= \int_X t \pi_X(t) dt \\ \text{Variance(X):} \quad E[(X - \mu_X)^2] &= \int_X (t - \mu_X)^2 \pi_X(t) dt \\ \text{Skewness(X):} \quad E\left[\left(\frac{X - \mu_X}{\sigma_X}\right)^3\right] &= \frac{1}{\sigma_X^3} \int_X (t - \mu_X)^3 \pi_X(t) dt \\ \text{Kurtosis(X):} \quad E\left[\left(\frac{X - \mu_X}{\sigma_X}\right)^4\right] &= \frac{1}{\sigma_X^4} \int_X (t - \mu_X)^4 \pi_X(t) dt \end{aligned}$$

Appendix B: Figures and Tables

Figure B1: Eurodollar Options Aggregate Open Interest




Figure B2: Eurodollar Options Aggregate Volume

30-day lagged moving average




Figure B3: Eurodollar Options Aggregate Open Interest

By time to maturity




Figure B4: Eurodollar Options Aggregate Volume

By time to maturity




Figure B5: Eurodollar Options Trading

By month of maturity




Figure B6: PDF using Cross Validation for Smoothing Parameter - Short Sterling

Maturity = December 15, 2010




Figure B7: PDF using Cross Validation for Smoothing Parameter - Euribor

Maturity = December 13, 2010




Figure B8: PDF using Cross Validation for Smoothing Parameter - Eurodollar

Maturity = December 13, 2010




Figure B9: PDF and Pricing Errors - Short Sterling

Maturity = December 15, 2010




Figure B10: PDF and Pricing Errors - Euribor

Maturity = December 13, 2010




Figure B11: PDF and Pricing Errors - Eurodollar

Maturity = December 13, 2010




Figure B12: PDF and Pricing Errors - Short Sterling

Maturity = December 17, 2014




Figure B13: PDF and Pricing Errors - Euribor

Maturity = December 15, 2014




Figure B14: PDF and Pricing Errors - Eurodollar

Maturity = December 15, 2014




Figure B15: Option-Implied Mean, by Days-to-Maturity

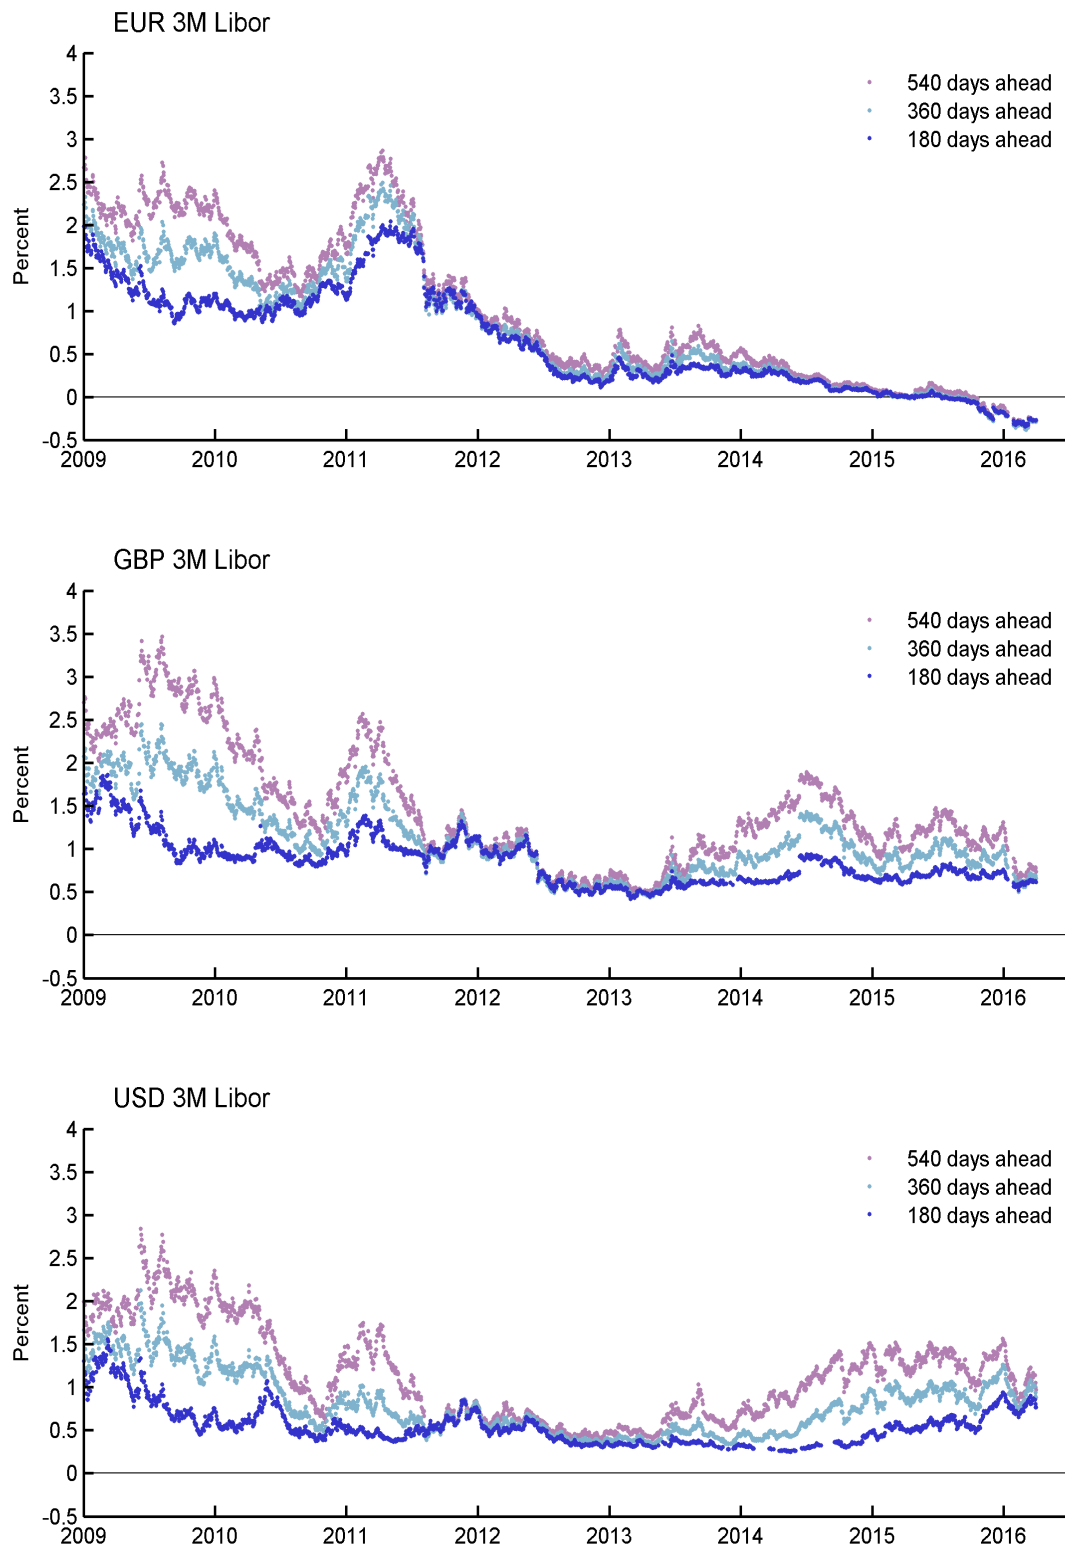


Figure B16: Option-Implied Standard Deviation, by Days-to-Maturity

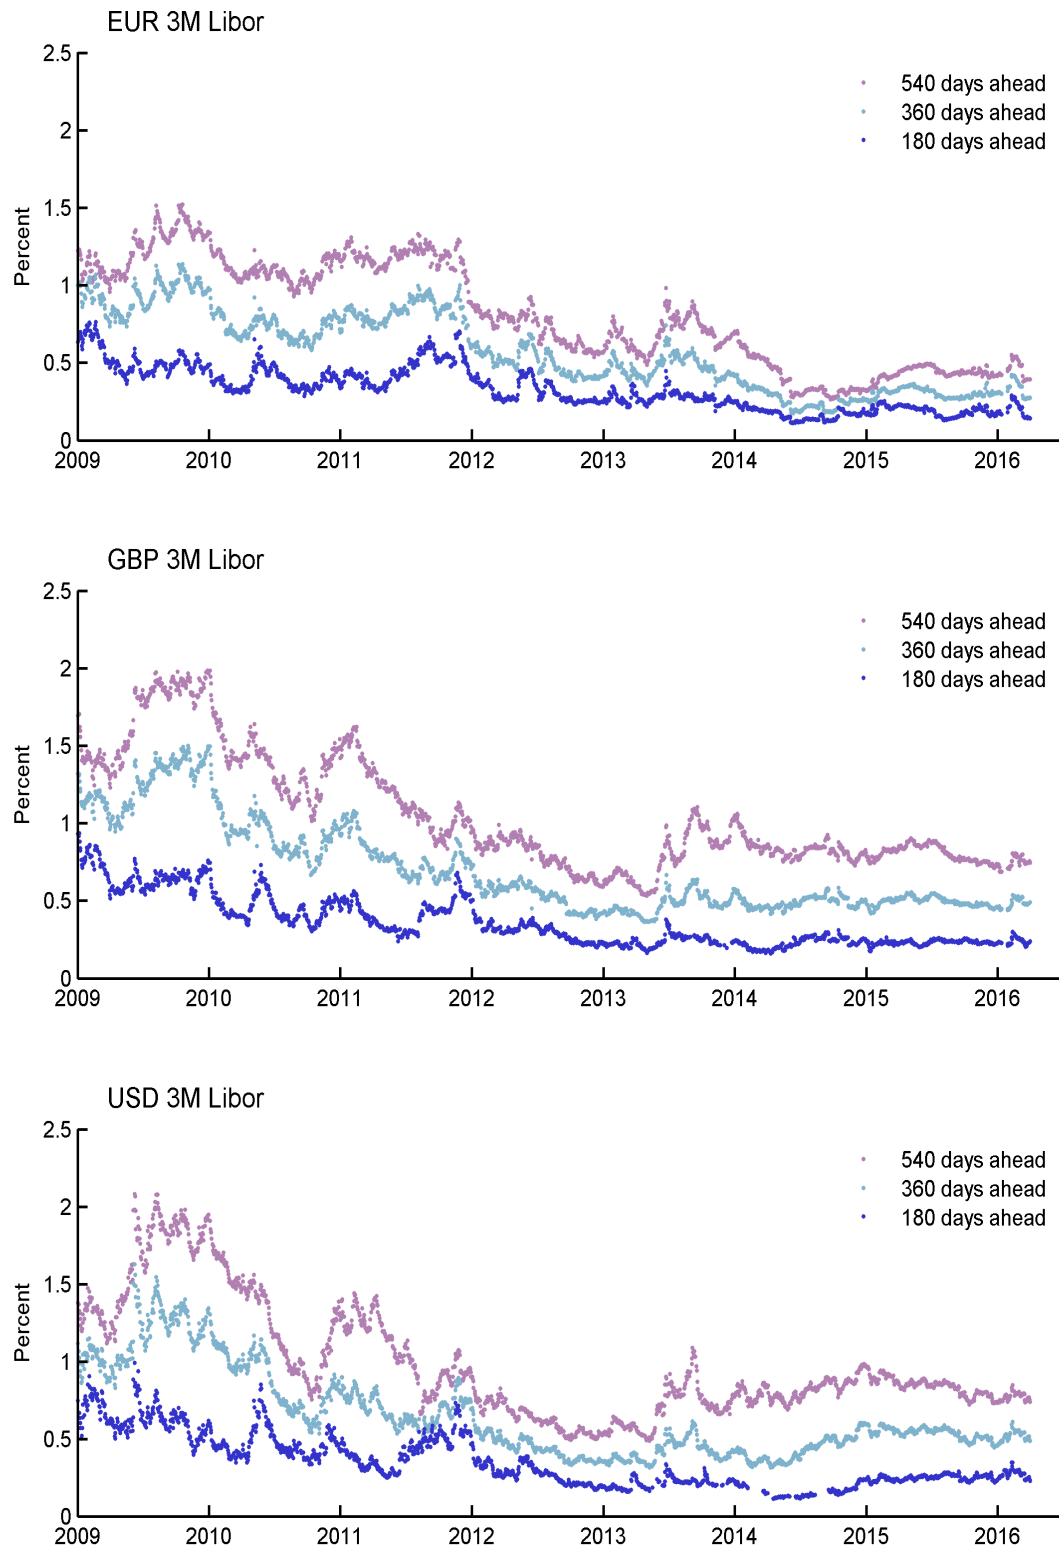


Figure B17: Option-Implied Skewness, by Days-to-Maturity

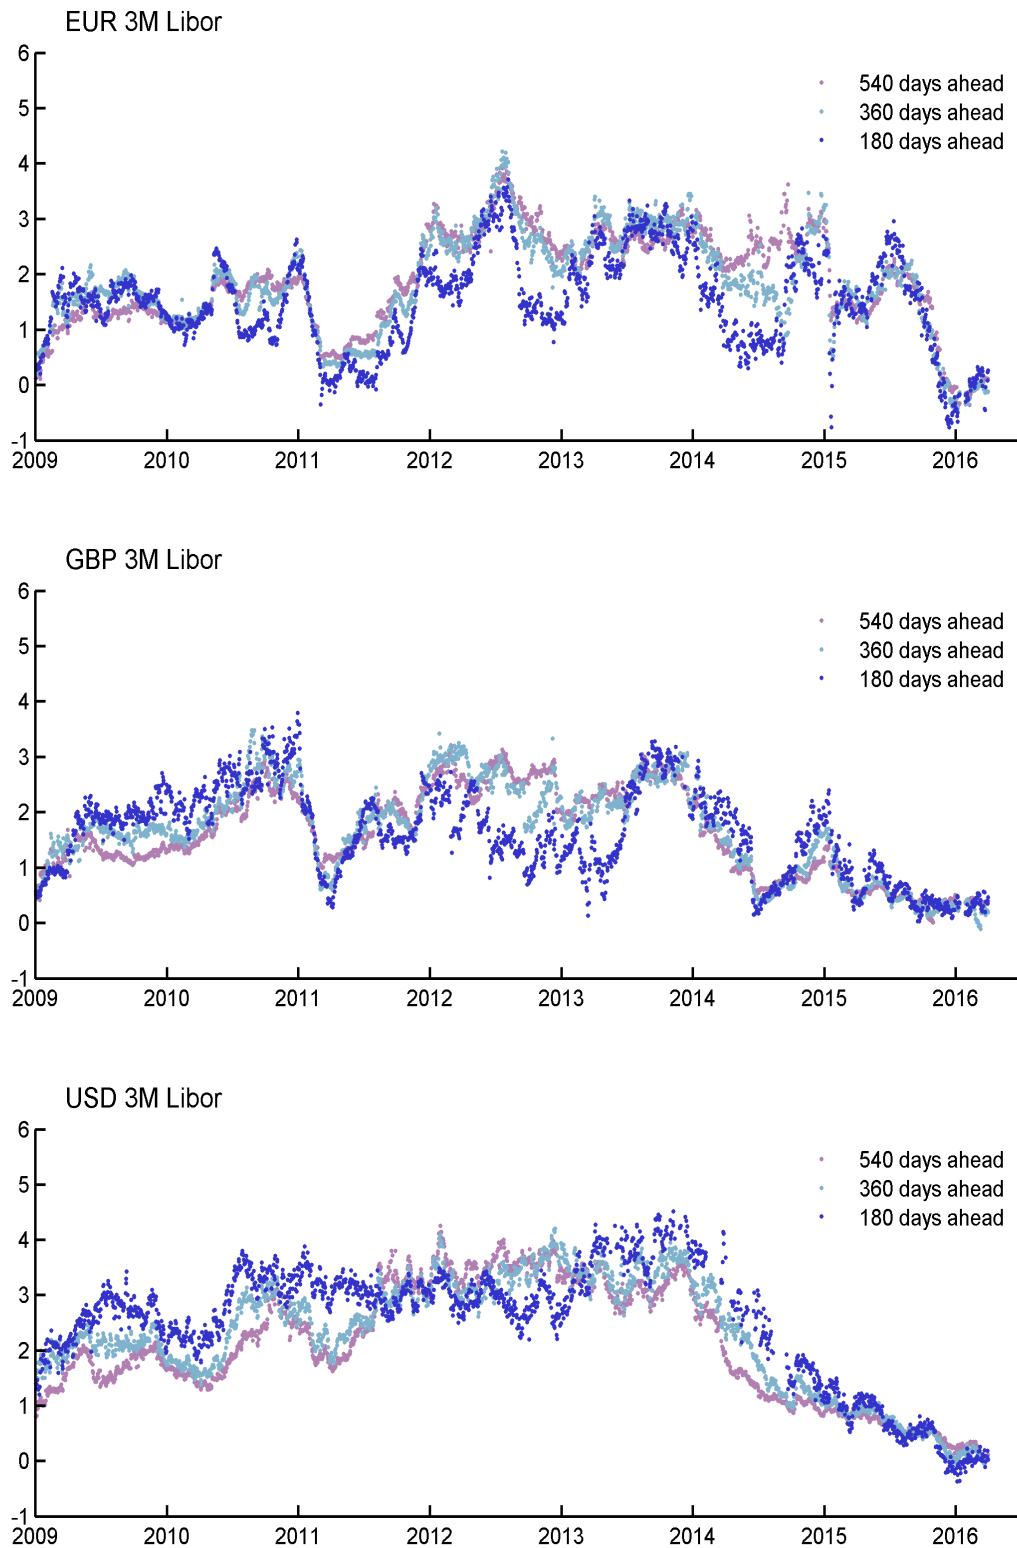


Figure B18: Option-Implied Kurtosis, by Days-to-Maturity

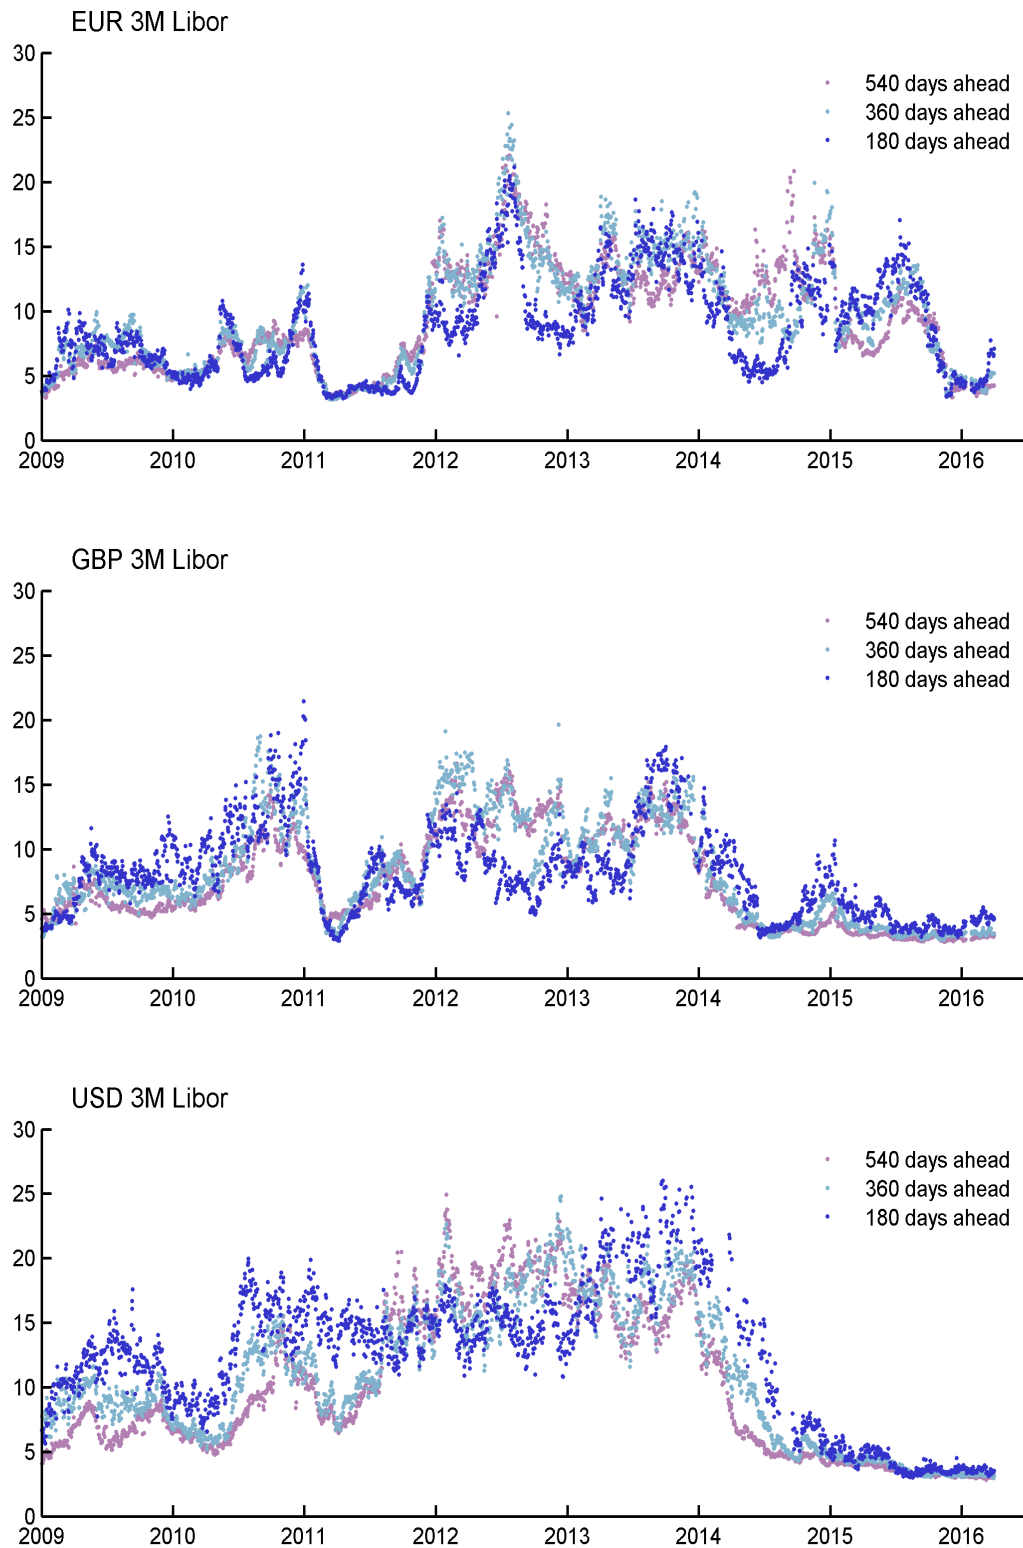


Table B1: Maturity Date Conventions for March-Quarterly Futures and Options

| | |
|----------------|---|
| Short Sterling | Third Wednesday of the month |
| Euribor | Two business days before the third Wednesday of the month |
| Eurodollar | Two London bank business days before the third Wednesday of the month |

Sources: Intercontinental Exchange and CME Group. Contract specifications for Short Sterling and Euribor futures and options can be found on www.theice.com, and for Eurodollar futures and options on www.cmegroup.com.

Table B2: 3M Libor Cross-Currency Granger Causality Significance

| Dependent Rate: | | EUR | | | | | GBP | | | | | USD | | | | |
|-----------------|----------|-----|-----|-----|---------|---------|-------|-----|-----|-----|---------|---------|-------|-----|-----|-----|
| | | 180 | 360 | 540 | Slope 1 | Slope 2 | Curve | 180 | 360 | 540 | Slope 1 | Slope 2 | Curve | 180 | 360 | 540 |
| EUR | Mean | | | | | | | x | . | . | . | . | x | . | x | x |
| | Std Dev | | | | | | | . | x | x | x | x | x | . | . | . |
| | Skewness | | | | | | | . | . | . | x | . | x | . | . | . |
| | Kurtosis | | | | | | | x | . | . | x | . | . | . | . | . |
| GBP | Mean | x | x | x | . | x | x | | | | | | | x | x | x |
| | Std Dev | x | x | . | x | x | x | | | | | | | x | x | x |
| | Skewness | . | . | . | x | . | x | | | | | | | . | x | . |
| | Kurtosis | . | x | . | x | . | x | | | | | | | . | x | . |
| USD | Mean | x | x | x | x | x | x | x | x | x | x | x | x | | | |
| | Std Dev | x | x | x | x | x | . | x | x | x | x | x | x | | | |
| | Skewness | . | . | . | x | . | x | x | x | x | x | x | . | | | |
| | Kurtosis | . | . | . | . | . | . | x | x | x | x | . | x | | | |

Notes: Estimated from VARs, with lag lengths selected by Schwarz's Bayesian information criterion. Sample period begins May 7, 2009 and ends March 31, 2016. Within each dependent rate grouping, the first three columns correspond to levels at 180, 360, and 540 days ahead; "Slope 1" and "Slope 2" denote the differences in levels at 360 vs. 180 days ahead, and 540 vs. 360 days ahead, respectively; and "Curve" is Slope 2 minus Slope 1. An "x" marks statistical significance at the 5% confidence level. VARs are run for each pairing of distributional moment and term structure component (for example, skewness levels 180 days ahead), with resulting bivariate relationships reorganized in this table for presentation. All series are autocorrelated; the corresponding sections are left blank to ease reading.

Table B3: 3M Libor Cross-Currency Granger Causality Significance, using Toda and Yamamoto's (1995) method

| Dependent Rate: | EUR | | | | | | GBP | | | | | | USD | | | | | |
|-----------------|----------|---|-----|---|-----|---|---------|---|---------|---|-------|---|-----|---|-----|---|-----|---|
| | 180 | | 360 | | 540 | | Slope 1 | | Slope 2 | | Curve | | 180 | | 360 | | 540 | |
| | | | | | | | | | | | | | | | | | | |
| EUR | Mean | O | O | O | O | O | O | . | . | . | . | X | . | . | . | . | . | X |
| | Std Dev | O | O | O | O | O | O | O | O | X | X | X | . | X | . | . | . | . |
| | Skewness | O | O | O | O | O | O | O | O | X | X | X | X | X | . | . | . | . |
| | Kurtosis | | O | O | O | O | O | O | O | X | X | X | X | X | . | . | . | . |
| GBP | Mean | X | X | X | X | X | X | . | X | X | X | X | X | X | X | X | X | . |
| | Std Dev | X | . | . | . | . | X | X | X | X | X | X | X | X | . | . | . | X |
| | Skewness | . | . | . | X | X | X | X | . | . | X | X | . | . | . | . | . | . |
| | Kurtosis | . | . | . | . | . | X | X | . | . | X | X | . | . | . | . | . | . |
| USD | Mean | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| | Std Dev | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X | X |
| | Skewness | . | . | . | . | . | X | X | . | X | X | X | X | X | X | X | X | X |
| | Kurtosis | . | . | . | . | . | . | . | . | . | . | . | X | X | X | X | X | X |

Notes: Estimated from VARs, with lag lengths selected by Schwarz's Bayesian information criterion. Sample period begins May 7, 2009 and ends March 31, 2016. Within each dependent rate grouping, the first three columns correspond to levels at 180, 360, and 540 days ahead; "Slope 1" and "Slope 2" denote the differences in levels at 360 vs. 180 days ahead, and 540 vs. 360 days ahead, respectively; and "Curve" is Slope 2 minus Slope 1. An "x" marks statistical significance at the 5% confidence level. VARs are run for each pairing of distributional moment and term structure component (for example, skewness levels 180 days ahead), with resulting bivariate relationships reorganized in this table for presentation. All series are autocorrelated; the corresponding sections are left blank to ease reading. A blue "O" in the EUR-EUR section is used to signify a VAR in which an Augmented Dickey-Fuller test failed to reject the null hypothesis of a unit root at the 5% confidence level for at least one currency; in each of these cases, the Granger causality significance results reported in this table employ the approach of Toda and Yamamoto (1995), assuming a maximum order of integration of one.

Appendix C: The American Early Exercise Premium in Eurodollar Option Prices

The American option early exercise premium is the premium paid for an American option compared against an equivalent European option. The only difference in value between the American and the European option in this case is the fact that the American option can be exercised at any time prior to the expiration date, while the European option can only be exercised upon expiration. Barone-Adesi and Whaley (1987) derive a method which can be used to estimate the early exercise premium embedded in the price of an American option. This method, however, relies on the assumptions of the Black-Scholes option pricing model, which include the assumption of geometric Brownian motion in the underlying security and thereby a lognormal terminal distribution. The approach outlined in this paper does not make such an assumption.

Melick and Thomas (1997) demonstrate that the upper bound on any American option early exercise premium (considered multiplicatively) is equal to $e^{r\tau}$. The intuition is that the greatest early exercise premium comes for an option which will be exercised as soon as possible with near certainty; and, for such an option, the price is simply the undiscounted expected payout (compare this against Equation (1), in which the price of a European option is the discounted expected payout). Considering the analysis in Section 4, the maximum USD-denominated risk-free rate used for discounting at 540 days ahead over the sample period from May 7th, 2009 to March 31st, 2016 is just over 1%. Then, the maximum upper bound on the difference between a European and American option price is about 1.5%. The multiplicative early exercise premium will be smaller as options are less in-the-money, or in fact are out-of-the-money. In this paper, I use only out-of-the-money option prices as model inputs. As the value of an out-of-the-money option if exercised today is zero, the American price will certainly involve some degree of discounting based on likely time to exercise. Therefore, the maximum daily upper bound on the early exercise premium for Eurodollar options used in the analysis of Section 4 is certainly less than 1.5%.

While the relatively small bound on the early exercise premium points to a small impact on pdf's estimated from Eurodollar option prices, it ignores the interaction of option prices with the actual estimation procedure. I therefore use Monte Carlo simulation to better understand the potential effect of the American early exercise premium in Eurodollar option prices on estimated pdf's. A set of European option prices for a given pricing and maturity date can be estimated by assuming a basic functional form for the early exercise premium and shocking observed American option prices based on the premium. Characteristics of the pdf's estimated from each set of option prices can then be compared in order to infer the degree to which the early exercise premium might be affecting implied pdf's.

In estimating European option prices, I work with the generous assumption that the daily upper

bound for the early exercise premium of $e^{r\tau}$ is effective at-the-money. From here, the multiplicative early exercise premium is linearly phased out over option strikes to reach 1 where the estimated cdf reaches 2.5% for puts or 97.5% for calls (i.e. the early exercise premium is assumed to be zero for out-of-the-money options where less than 2.5% probability is assigned to a positive option value at expiration). European option prices are then estimated by multiplying observed American option prices by the inverse of the corresponding premium.

I select all Wednesdays over the sample period from May 7th, 2009 to March 31st, 2016 for the simulation option pricing dates, and I estimate fixed horizon pdf's 540 days ahead as described in Section 2.6 using both observed option prices and simulated European option prices. I then calculate the mean, standard deviation, (normalized) skewness, and (normalized) kurtosis of each pdf. Figure C1 shows histograms of the percent differences in pdf moments, where positive (negative) values imply a higher (lower) value using estimated European option prices.

Figure C1: Percent Differences in PDF Moments after Option Price Shocks

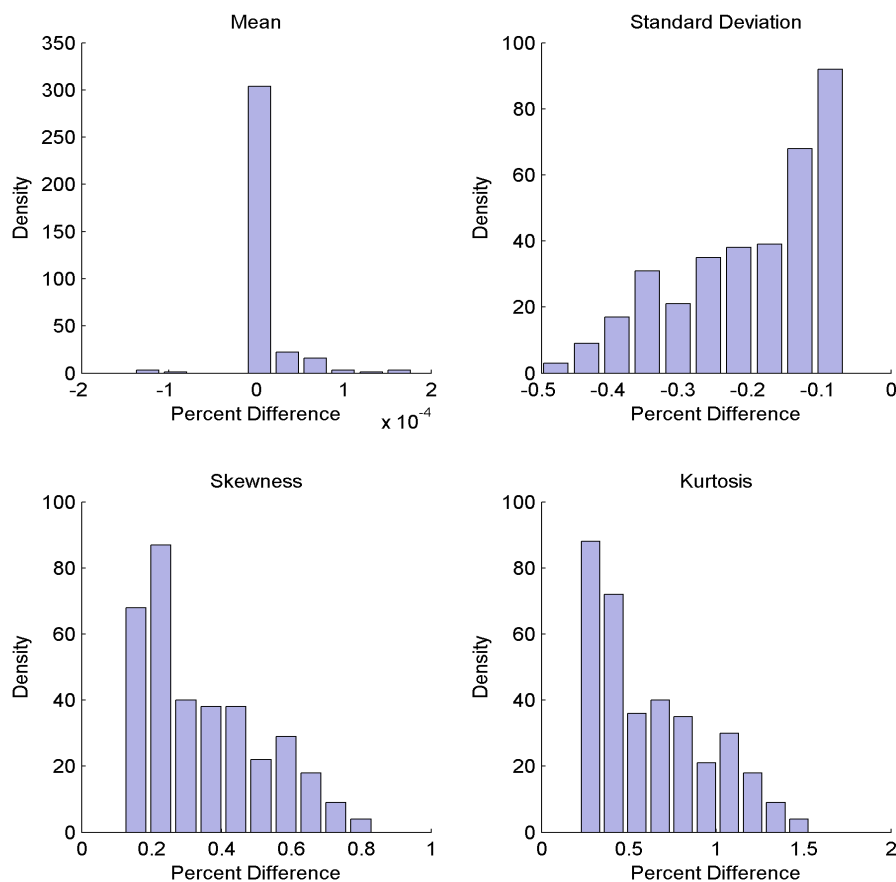


Figure C1 shows that the overall effects of even a generous specification for the early exercise premium are minor. Differences in the pdf mean are negligible, while differences in higher moments are all

within small margins. It is not surprising that the standard deviation is consistently lower when using estimated European option prices, or that the normalized skewness and kurtosis are consistently higher. Because European option prices will be lower than American prices, and because out-of-the-money options are used as inputs, estimated probabilities will be less dispersed in order to match lower out-of-the-money option prices. Further, as the early exercise premium falls for increasingly out-of-the-money options, simulated price shocks will push prices increasingly downward nearer the money, implying less probability mass near the center of the distribution relative to the tails in the case of hypothetical European options. Finally, as skewness is normalized, a lower standard deviation will result in higher skewness even if there is little to no change in non-normalized skewness. Table C1 below shows the average and median percent differences between moments estimated from hypothetical European prices and observed American prices, complementing Figure C1.

Table C1: Percent Differences in PDF Moments after Option Price Shocks

| | Mean | Std Dev | Skewness | Kurtosis |
|--------------------|-----------|---------|----------|----------|
| Average Difference | 7.24 e-06 | -0.20 | 0.35 | 0.63 |
| Median Difference | 0.00 | -0.17 | 0.29 | 0.53 |

Note: "Difference" is the percent difference between a given distributional moment as estimated from hypothetical European option prices or observed American option prices.