The LIX : a model-independent liquidity index

Reference:
Guillaume Florence.- The LIX : a model-independent liquidity index
Journal of banking and finance - ISSN 0378-4266 - 58(2015), p. 214-231
DOI: http://dx.doi.org/doi:10.1016/j.jbankfin.2015.04.015
Handle: http://hdl.handle.net/10067/1261260151162165141

The LIX: a model-independent liquidity index

Florence Guillaume*

April 27, 2015

Abstract

This paper provides a new model-free indicator of liquidity, the so-called LIX index. The computation of the LIX index combines the conic finance theory, which recognizes the two-price economy and is built upon the concept of indices of acceptability of Cherny & Madan (2010), with the option payoff spanning formula of Breeden & Litzenberger (1978). Matching the conic finance bid and ask prices of the stock with those observed in the market allows us to derive a model-free and unit-less indicator of spot liquidity. Just as the VIX and the SKEW index quantify the volatility and the tail risk perceived by today's investors, the resulting LIX index measures, in a similar market-implied fashion, the liquidity risk. The maximum likelihood estimation of popular mean-reverting processes applied to model-free liquidity time series indicates that spot liquidity tends to dry up during distress periods whereas a global drying-up of liquidity could not be detected during turmoil periods in the option market.

Keywords: *spot liquidity, model-free liquidity index, option liquidity surface, conic finance, pre- and post-crisis liquidity*

1 Introduction

The recent credit crunch has highlighted the need for strengthening capital and liquidity requirements. Quoting the Basel Committee on Banking Supervision (BCBS (2009)), "the crisis illustrated how quickly and severely liquidity risks can crystallize and certain sources of funding can evaporate, compounding concerns related to the valuation of assets and capital adequacy. A key characteristic of the financial crisis was the inaccurate and ineffective management of liquidity risk." Although the framework for liquidity risk management proposed by the Basel Committee is addressed to financial institutions, a sound liquidity risk management is crucial for each market player, and this for both funding and market liquidity[1]. Indeed, these two facets of liquidity are highly interconnected during market turmoil periods due to the massive fire sales which typically occur at those times (see, for instance, Brunnermeier (2009)). The aim of this paper is to develop a model independent framework for market liquidity which can be used to derive an index for stock (or other) market liquidity or to incorporate liquidity risk into risk management tools such as hedging or VaR models.

The wide variety of spot liquidity indicators (trading volume, number of trades, dollar bid-ask spread, proportional (i.e. relative) bid-ask spread, order imbalances (see Chordia et al. (2002)),

*University of Antwerp, Department of Mathematics and Computer Sciences, Middelheimlaan 1 , B-2020 Antwerpen, Belgium. E-mail: florence.guillaume@uantwerpen.be

[1]Funding liquidity is defined as the ease with which investors can obtain funding whereas market liquidity is the ability to find a counterparty for the other side of a trade.

Amihud's liquidity ratio (see Amihud (2002)), ...) demonstrates the difficulty of finding a reference proxy to assess liquidity in the stock market (see, for instance, Aitken & Comerton-Forde (2003) or Chou et al. (2011)). In this paper, adopting a similar methodology as Corcuera et al. (2012), we propose a new unit-less and model-free indicator of spot liquidity. Corcuera *et al.* introduced the concept of implied liquidity which originates from the conic finance theory where illiquid markets are modeled as a counterparty (see Cherny & Madan (2010)). The conic finance theory extends the classical finance theory by dropping the law of one price in favor of a two-price economy, making the price of any financial instrument dependent on the direction of the trade. In the classical theory, the liquid market acts as counterparty to investors and accepts any amount of financial assets traded at the going market price, whatever the direction of the trade. Conversely, in the conic finance theory, market participants buy from the market at the ask price but sell to the market at the lower bid price. Conic finance applications originate from the framework of indices of acceptability introduced by Cherny & Madan (2009) where risk measures are defined in terms of distorted expectations of zero cost cash-flows. More precisely, a cash-flow is said to be acceptable or marketed if its expectation under the different test measures belonging to a convex set of supporting measures is positive. Using different families of concave distortion functions defining different kinds of indices of acceptability, Cherny & Madan (2010) proposed parametric models for the cone of marketed cash-flows, and consequently for the bid and ask prices. Under these parameterizations, cones of acceptability depend only on the distribution function of the zero cost cash-flow. In particular, Cherny & Madan (2009) derived the expression of the call and put bid and ask prices in terms of the distribution function of the underlying asset. The conic finance theory, and in particular these closed-form expressions, led Corcuera *et al.* to propose an alternative tool to assess the liquidity of financial instruments, the so-called implied liquidity. This new proxy of liquidity is defined as the parameter of the concave distortion function (also called liquidity parameter), λ, that replicates as well as possible the bid and ask quotes of the financial instrument. The higher the λ, the higher the degree of distortion and the lower the liquidity of the instrument. In the limit, an implied liquidity of zero leads to a unique equilibrium price for the bid and the ask and hence to the classical finance theory. Corcuera *et al.* illustrated the computation of the implied liquidity for the particular case of call options, assuming that the stock price process is modeled by the famous Black-Scholes model. Moreover, they showed how the liquidity of at-the-money call options dried up during the recent crisis period. Albrecher et al. (2013) extended the concept of implied liquidity to more suited Lévy implied volatility models proposed by Corcuera et al. (2009) and investigated the model sensitivity of implied liquidity for call options. In particular, they tested the goodness of fit of CEV models for implied liquidity time series and showed that the choice of the option pricing model impacts significantly the maximum likelihood estimates of the long-run liquidity and volatility of liquidity parameters of the CEV model. The methodology proposed in this paper differs from the one adopted by Corcuera *et al.* and Albrecher *et al.* as the stock price distribution is directly inferred from the volatility surface in such a way that we do not have to assume any prevailing model for the asset log-return process. With such an adjustment for model-dependency, we can infer a market-implied liquidity index, the so-called LIX index, reflecting the degree of market liquidity in the stock market perceived by today's investors. Therefore, our approach addresses the criticism of *model-dependency of the derived implied liquidity parameter* raised by Corcuera et al. (2012). The LIX is the model-free counterpart of implied liquidity as the VIX is the model-independent counterpart of implied volatility. Hence, just as the VIX paved the way for volatility derivatives and volatility hedging, the LIX methodology could become the precursor of tradable liquidity derivatives which could be used to hedge against a drying up of

liquidity as detailed in Corcuera et al. (2012). Moreover, unlike existing measures of spot liquidity which are either historical (i.e. based on some time series) or today's proxies, the LIX index is a benchmark for the expected future liquidity in the market since it combines the current bid and ask prices of the asset with the distribution of its future cash flow as perceived by today's market investors. In other words, the LIX is a model free forward looking gauge of expected liquidity in the stock market over the reference period of 30 calendar days.

As numerical study, we first derive historical time series for the new forward looking measure of liquidity for the S&P 500 index and options. To this end, we use daily bid and ask prices of the S&P 500 index and quoted bid an ask closing option prices[2]. We will then propose and calibrate plausible models for liquidity and in particular for the LIX index. Given the empirical evidence that market liquidity presents the same characteristics as market volatility (see, for instance, Corcuera et al. (2012) or Albrecher et al. (2013)), viz. clustering and mean-reversion, we focus our interest on the class of mean-reverting processes which are commonly used to model the VIX volatility index. The maximum likelihood calibration of such models to the time series of the new model-free liquidity proxy we performed, constitutes the first step towards stochastic liquidity modeling. Armed with such models to simulate future liquidity scenarios, and consequently also future bid and ask scenarios, we could then incorporate market liquidity risk into popular risk management tools such as hedging or VaR models[3], leading to a more accurate and effective management of market liquidity risk.

The remaining of this paper is organized as follows. Section 2 describes how the distribution function of the stock price is extracted from current vanilla option prices. Section 3 details the computation of the new model-free liquidity index LIX. Section 4 provides a statistical study of the LIX index time series spanning from January 1998 to October 2009 and including therefore the recent credit crunch as well as the dot-com bubble. Section 5 introduces the concept of model-free implied liquidity surfaces, one for the call and one for the put options, summarizing the dependence of call and put liquidity on the strike price and the time to maturity of the option. In Section 6, we fit popular mean-reverting processes to the LIX time series but also to the model-free liquidity time series of different call and put options using maximum likelihood estimation. This allows for a comparison of the dynamics of liquidity in the underlying and derivative markets, but also during periods characterized by different market fear levels. Section 7 formulates our conclusions.

2 Extraction of the risk-neutral distribution function

The risk-neutral density of any asset can be extracted from option prices quoted for a continuum of strike prices spanning the range of possible asset price outcomes (see Breeden & Litzenberger (1978)). Indeed, the fundamental theorem of asset pricing states that the fair price of any financial

[2]Although the data used to derive the LIX index consist of daily market quotes, the methodology can be extended as it is to intraday high frequency data, leading to a LIX index based on the effective bid-ask spread rather than on the quoted daily spread.

[3]The importance of embedding market liquidity in the VaR model has been evidenced, among others, by Bangia et al. (2002) and by Bervas (2006). Bangia et al. (2002) integrated exogeneous liquidity risk by considering a parametric model where the mid price is modeled by a log-normal distribution with a volatility parameter adjusted for leptokurtosis and where the cost of liquidity is assumed to be proportional to half the current mid price and to the average relative spread plus a multiple of the relative spread volatility. Bervas (2006) proposed an historical simulation method to adjust the VaR for exogeneous liquidity risk and which is based on return and Roll coefficient time series. An approach based on Monte Carlo simulations of liquidity will arise as a natural alternative since the VaR will then be directly inferred from the distribution of the bid (liquidation) value.

derivative is equal to its discounted expected terminal payoff. In particular, for European call and put options with strike price K and time to maturity T, we have

$$C(K,T) = \exp(-rT) \int_K^{+\infty} (S_T - K) f(S_T) dS_T$$

and

$$P(K,T) = \exp(-rT) \int_0^K (K - S_T) f(S_T) dS_T,$$

respectively, where r is the risk-free interest rate and S_T the price of the underlying at the future time T. Taking the first derivative with respect to K and solving for the distribution function of S_T in K, $F(K)$, give

$$F(K) = \exp(rT) \frac{\partial C(K,T)}{\partial K} + 1 \quad \text{and} \quad F(K) = \exp(rT) \frac{\partial P(K,T)}{\partial K}.$$

Hence, for a continuum of quoted strikes, the distribution function of the underlying price is directly inferred from the first derivative of the option price with respect to the strike price K. In practice, due to the discrete nature of listed strikes, one can derive only a model-free approximation of the distribution function of S_T by using weighted finite differences. Denoting by $\mathbf{K} = \{K_1 < K_2 < \ldots < K_N\}$ the set of strikes of liquidly traded options with time to maturity T, we have

$$\frac{\partial C(K_i,T)}{\partial K_i} = \frac{C(K_{i+1},T) - C(K_i,T)}{K_{i+1} - K_i} \frac{K_i - K_{i-1}}{K_{i+1} - K_{i-1}} - \frac{C(K_{i-1},T) - C(K_i,T)}{K_i - K_{i-1}} \frac{K_{i+1} - K_i}{K_{i+1} - K_{i-1}}.$$

Similarly, we have for the put derivative

$$\frac{\partial P(K_i,T)}{\partial K_i} = \frac{P(K_{i+1},T) - P(K_i,T)}{K_{i+1} - K_i} \frac{K_i - K_{i-1}}{K_{i+1} - K_{i-1}} - \frac{P(K_{i-1},T) - P(K_i,T)}{K_i - K_{i-1}} \frac{K_{i+1} - K_i}{K_{i+1} - K_{i-1}}.$$

To extract the stock price distribution function, we consider the more liquid options only, i.e. out-of-the-money options. Following a selection procedure similar to the one adopted by the Chicago Board Options Exchange (CBOE) to infer the VIX volatility index (see CBOE (2003)), we select put options with a strike lower than the forward price $F_0 = S_0 \exp((r-q)T)$ and call options with a strike higher than F_0, where r and q denote the risk-free interest rate and the dividend yield, respectively, and where S_0 is the spot price. More precisely, sorting the call and put options by ascending and by descending order of strike, respectively, we select the options with a positive bid price until we encounter two successive options with zero bid. Using both call and put options allows us to broaden the set of strikes $\mathbf{K}$ for which we can compute an approximation of the distribution function F. For sake of consistency, the forward price and the interest rate are inferred from the put-call parity

$$\exp(-rT)F_0 - \exp(-rT)K = C(K,T) - P(K,T)$$

by performing, for each time to maturity T, a linear regression of the call-put spread, i.e. $C(K,T) - P(K,T)$, on the strike price K.

The distribution function of the stock price at time T evaluated at the listed strike K_i can thus be approximated by

$$F(K_i) = \begin{cases} \exp(rT)\left(\frac{P(K_{i+1},T)-P(K_i,T)}{K_{i+1}-K_i}\frac{K_i}{K_{i+1}} + \frac{P(K_i,T)}{K_i}\frac{K_{i+1}-K_i}{K_{i+1}}\right) & K_i = K_1 \\ \exp(rT)\left(\frac{P(K_{i+1},T)-P(K_i,T)}{K_{i+1}-K_i}\frac{K_i-K_{i-1}}{K_{i+1}-K_{i-1}} - \frac{P(K_{i-1},T)-P(K_i,T)}{K_i-K_{i-1}}\frac{K_{i+1}-K_i}{K_{i+1}-K_{i-1}}\right) & K_1 < K_i < F_0 \\ \exp(rT)\left(\frac{C(K_{i+1},T)-C(K_i,T)}{K_{i+1}-K_i}\frac{K_i-K_{i-1}}{K_{i+1}-K_{i-1}} - \frac{C(K_{i-1},T)-C(K_i,T)}{K_i-K_{i-1}}\frac{K_{i+1}-K_i}{K_{i+1}-K_{i-1}}\right)+1 & F_0 \leq K_i < K_N \\ \exp(rT)\frac{C(K_N,T)-C(K_{N-1},T)}{K_N-K_{N-1}}+1 & K_i = K_N, \end{cases} \tag{1}$$

the first equation coming from the fact that $K_0 = 0$ and $P(0,T) = 0$ by definition.

In practice, option price imperfections might lead to unacceptable features of the estimated distribution function as proposed, and in particular to ranges in which it is decreasing. To tackle this problem inherent to market data, many parametric and non-parametric procedures have been proposed in the financial literature in order to obtain a distribution function which is free of static arbitrage (see, for instance, Figlewski (2010), Gatheral (2006) or Shimko (1993)). In this paper, we focus on two particular methodologies that turn out to work out generally rather well for the time period under consideration. The first one, referred to as the spline methodology, consists of fitting a smooth curve to the implied volatility smile by least-square minimization, as first proposed by Shimko (1993). More particularly, we follow a methodology similar to the one adopted by Figlewski (2010), where the smooth function is chosen to be a fourth degree spline with one knot, i.e. two segments of fourth order polynomials joined at the at-the-money strike. We also consider the SVI (stochastic volatility inspired) parameterization of the total implied variance as proposed by Gatheral (2006)

$$\mathrm{IV}^2(k,T)\cdot T = a + b\left(\rho(k-m) + \sqrt{(k-m)^2+\sigma^2}\right),$$

where $k = \log\left(\frac{K}{F_0}\right)$ and where $a \in \mathbb{R}$, $b \geq 0$, $|\rho| < 1$, $m \in \mathbb{R}$, $\sigma > 0$ and $a + b\sigma\sqrt{1-\rho^2} \geq 0$.

The left panel of Figure 1 shows the "arbitrage-free" implied volatility smiles obtained by the spline methodology and the SVI parameterization for October, 1, 2008 and for a maturity $T = 0.0466$. We clearly see that the two arbitrage-free smiles differ only within the spread since a penalty term compels each arbitrage-free volatility to fall between the quoted bid and ask implied volatilities. Once we have obtained a volatility curve free of static arbitrage, we can extract the adjusted option prices by inverting the Black-Scholes formula as shown on the center panel of Figure 1. Finally, the distribution function of the stock price at time T is inferred from the adjusted option prices by using the second order difference approximation (1) (see right panel of Figure 1). We observe that the piecewise linear distribution function reconstructed from liquid out-of-the-money option prices under the two methodologies differs only slightly and spans essentially the whole distribution function co-domain $[0,1]$. This is easily explained by the fact that the S&P 500 option market has deep and active trading across a broad range of strike prices. The two methodologies will consequently generally lead to a similar LIX index and hence to a quite similar trend in the historical LIX time series. In the numerical study, we will report the results corresponding to the spline methodology only; similar conclusions could be drawn from the SVI parameterization approach[4]. Note that we focus on the spline procedure since it is computationally less demanding and does not require any initial parameter guess, contrary to the SVI procedure. Indeed, the spline parameters can be initialized by fitting a unique fourth degree polynomial to the implied

[4]For a comparison of empirical values of the LIX index under the two methodologies, we refer the reader to A.

volatility curve, whereas the SVI parameterization requires an initial guess for the parameter set $\{a, b, \rho, m, \sigma\}$. In particular, it turns out that for a few quoting days under investigation, the optimal SVI parameters obtained as output of the search algorithm actually correspond to a "bad" local minimum of the least-square objective function.

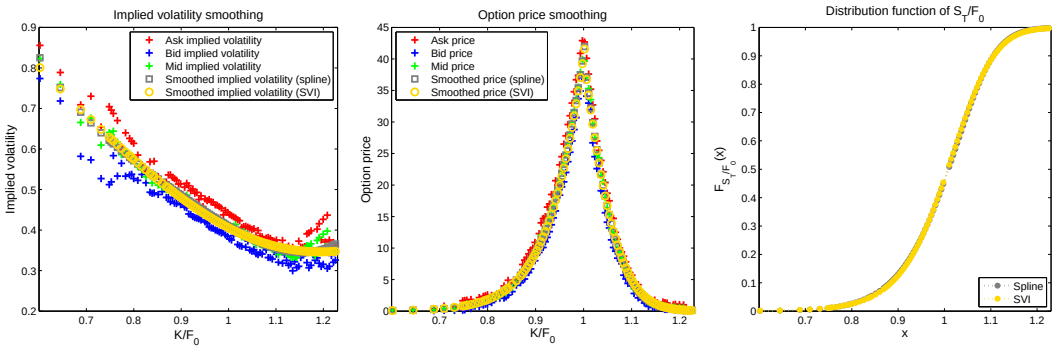

Figure 1: Construction of the risk-neutral distribution function of S_T by using either the spline procedure or the SVI parameterization (October, 1, 2008, $T = 0.0466$): smoothing of implied volatilities (left), adjusted vanilla option prices (center) and risk-neutral stock price distribution function (right).

3 The LIX index

Under the framework of indices of acceptability, a risk X is said to be acceptable or marketed, i.e. $X \in \mathcal{A}$ if

$$E_Q[X] \geq 0, \ \forall Q \text{ in a convex set } \mathcal{M}.$$

The convex set $\mathcal{M}$ consists of test measures under which the expected cash-flow has to be positive in order for X to be acceptable. The larger the set $\mathcal{M}$, the more tests have to be passed and thus the smaller the set of marketed cash-flows X. Under this framework, parametric models for cones of acceptability which contain the non-negative random variables are obtained by parameterizing the test measures in terms of some family of concave distortion functions. More particularly, $X \in \mathcal{A}$ if its distorted expectation with respect to some distortion function Ψ_λ is positive:

$$X \in \mathcal{A} \Leftrightarrow \int_{-\infty}^{+\infty} x d\Psi_\lambda(F_X(x)) \geq 0. \tag{2}$$

By construction, the such obtained operational cone of acceptability depends only on the distribution function $F_X(x)$ of X and on the parametric family of distortion functions $\{\Psi_\lambda, \lambda \geq 0\}$, where the parameter λ quantifies the degree of the distortion. The higher the λ, the higher the degree of distortion, i.e. the heavier the weight allocated to the downside. The most commonly used family of distortion functions is the so-called MINMAXVAR family, which is defined as

$$\Psi_\lambda(y) = 1 - \left(1 - y^{\frac{1}{1+\lambda}}\right)^{1+\lambda}, \ \lambda \geq 0,$$

and upon which we will focus in this paper. For other families of concave distortion functions used in the context of indices of acceptability, we refer the reader to Cherny & Madan (2009). Using such a parametric family of distortion functions allows for parametric models for the cone of marketed cash-flows, and hence also for the bid and ask prices, in terms of the single parameter λ. In conic finance theory, the ask price of a claim is defined as the smallest price for which the cash-flow of selling the claim at its ask price is acceptable for the market. Denoting by X the cash-flow generated by the claim at the future maturity date T, we have

$$a(X) = \min\{a : a - \exp(-rT)X \in \mathcal{A}\};$$

or, using Definition (2) of acceptability

$$a(X) = -\exp(-rT)\int_{-\infty}^{+\infty} x d\Psi_\lambda(1 - F_X(-x)). \tag{3}$$

Similarly, the bid price of a claim paying X at time T is defined as the highest price for which the cash-flow of buying the claim at its bid price is acceptable for the market:

$$b(X) = \max\{b : \exp(-rT)X - b \in \mathcal{A}\};$$

or, using (2)

$$b(X) = \exp(-rT)\int_{-\infty}^{+\infty} x d\Psi_\lambda(F_X(x)). \tag{4}$$

From Equation (3) and Equation (4), it is clear that the bid and ask prices of a contingent claim with a cash-flow X delivered at time T depend only on the distribution function of X and on one unobservable parameter λ. Hence, as long as the distribution function of X is available, we can derive λ such that the conic finance bid and ask prices match the bid and ask prices observed in the market. This methodology was first proposed by Corcuera et al. (2012) to define the so-called implied liquidity of vanilla call options. The terminology *implied liquidity* comes from the assumption made in Corcuera et al. (2012) that the price of the underlying asset is modeled by the famous Black-Scholes model, making an obvious reference to the widespread concept of implied volatility. With such a convention, the implied liquidity inevitably depends on the implied volatility of the option which is chosen as estimate for the Black-Scholes volatility parameter and which typically differs from strike price to strike price given the intrinsic smiling nature of the implied volatility curve. The implied liquidity parameter of options with different moneyness[5] is thus computed by considering different stock price distributions. Hence, it does not seem so natural to extend the procedure as it is in order to define the implied liquidity of the underlying. Indeed, which volatility benchmark should we then consider? The at-the-money implied volatility, some out-of-the-money implied volatility or some in-the-money implied volatility? The different specifications will lead to a different implied liquidity for the underlying. In particular, the deeper in-the-money the volatility benchmark, the more spread the bulk behavior of the underlying distribution under the Black-Scholes model and hence the wider the conic finance bid-ask spread for a fixed distortion

[5]The moneyness is defined as the difference between the spot price and the strike price.

parameter λ. As an illustrative example of this dependence, we consider an artificial implied volatility curve parameterized by

$$\sigma_{\mathrm{BS}}(K) = 0.2/\sqrt{K/S_0}. \tag{5}$$

This particular volatility skew parameterization is chosen to replicate the typical Black-Scholes volatility skew adjustment observed in equity markets and represented in Figure 2.

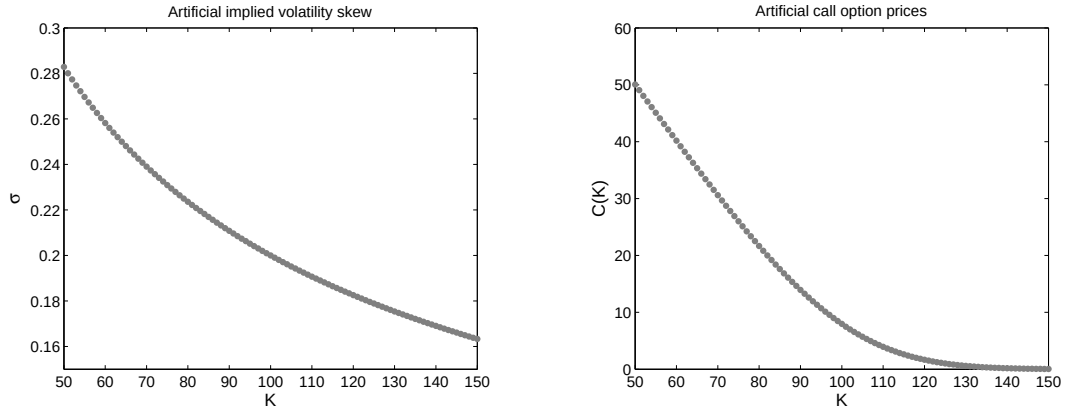

Figure 2: Artificial Black-Scholes implied volatility skew $\sigma_{\mathrm{BS}}(K) = 0.2/\sqrt{K/S_0}, (S_0 = 100)$ (left) and corresponding call prices (right).

By way of example, we focus on the following setting[6]:

$$T = 1,\ r = 0,\ q = 0,\ S_0 = 100,\ K = 50,\ 51,\ \ldots,\ 149,\ 150.$$

We consider both the model-free and the Black-Scholes approximation of the distribution function of S_T. Under the Black-Scholes model, the stock price at maturity follows a log-normal distribution:

$$S_T \sim \mathrm{logN}\left(\log(S_0) + (r - q - \frac{\sigma^2}{2})T\,, \sigma^2 T\right)$$

where the cumulative distribution function of a log-normal random variable $X \sim \mathrm{logN}(\mu, v)$ is equal to $F_X(x; \mu, v) = \Phi\left(\frac{\log(x)-\mu}{\sqrt{v}}\right)$, where Φ denotes the distribution function of a standard normal random variable. Increasing the volatility parameter of the Black-Scholes model (i.e. decreasing the strike K) leads to an increase of the variance, skewness and kurtosis of the distribution of S_T, as can be seen from the left above panel of Figure 3. Hence, as shown in the right above panel of

[6]The set of strikes is chosen such that the strike range is wide and refined enough to avoid substantial numerical errors in the model-free bid and ask prices for the range of distortion parameters λ considered. Such numerical errors can arise from the truncation and the discretization of the integrals included in the expression of the bid and ask prices when the strike range does not span the whole distribution co-domain or when the strike grid is too coarse. Besides, the option maturity is chosen in such a way that the one-year at-the-money implied volatility equates to the common value of 20 percent in a standard market regime.

Figure 3, the higher the volatility, the bigger the conic finance bid-ask spread for a given distortion λ, or, equivalently, the lower the degree of distortion for a given conic finance spread. Hence, for a given market spread to match, the Black-Scholes implied liquidity parameter λ decreases with the Black-Scholes implied volatility and so the spot liquidity increases with the Black-Scholes volatility as shown in the below panels of Figure 3. Besides, we also note the huge difference between the Black-Scholes implied liquidity parameter and the model-free liquidity, even for relatively low bid-ask spreads and whatever the Black-Scoles volatility parameter. This is due to the inadequacy of the Black-Scholes model to replicate the artificial market situation, and in particular to its underestimation of the tails of the stock price distribution. We note that, by comparison, the impact of the choice of the methodology aimed at obtaining arbitrage-free data is far from being significant. This is evidenced by the similitude of the market implied distribution function extracted by using either the spline or the SVI methodology, as shown in Figure 1, and by the scatter plot of the spline and SVI spot implied liquidities shown in Figure 13.

A more appropriate spot liquidity proxy could be defined by extending the approach to more suited option pricing models that allow for a flat volatility curve (i.e. an implied volatility independent of the strike price), such as, for instance, the Lévy space implied volatility models proposed by Corcuera et al. (2009). Such an extension of the concept of implied liquidity to Variance Gamma implied volatility models has been proposed by Albrecher et al. (2013). Nevertheless, the implied liquidity still remains, by nature, a model-dependent proxy of liquidity. This is why we propose to drop the assumption of some prevailing model for the stock price process and to infer its distribution at time T from liquid out-of-the-money vanilla options with time to maturity T, as described in Section 2, allowing for an extension of the concept of implied liquidity to the underlying asset itself in a straightforward way. This model-free approach answers the criticism of *model-dependency of the derived implied liquidity parameter* λ raised by Corcuera et al. (2012).

Taking as claim the underlying asset, the conic finance bid and ask prices at level λ can be approximated by using the risk-neutral model-free distribution function (1):

$$a(\lambda;T) = \exp(-rT) \sum_{i=0}^{N-1} \frac{K_i + K_{i+1}}{2} \left(\Psi_\lambda(1 - F(K_i)) - \Psi_\lambda(1 - F(K_{i+1}))\right);$$

$$b(\lambda;T) = \exp(-rT) \sum_{i=0}^{N-1} \frac{K_i + K_{i+1}}{2} \left(\Psi_\lambda(F(K_{i+1})) - \Psi_\lambda(F(K_i))\right).$$

The model-free conic finance bid and ask prices are thus computed by approximating the area below the bid curve $(\Psi_\lambda(F_X(K)), \exp(-rT)K)$ and the ask curve $(\Psi_\lambda(1-F_X(K)), -\exp(-rT)K)$, respectively, using a trapezoidal integration rule (see Figure 4).

The spot implied liquidity for a time horizon T can then be obtained by equating the conic finance bid-ask spread with the bid-ask spread observed in the market, as illustrated in Figure 5,

$$\mathrm{IL}(T) = \lambda : a(\lambda;T) - b(\lambda;T) = a^M - b^M,$$

where a^M and b^M denote the market ask and bid prices of the underlying asset, respectively. Figure 5 clearly indicates that the lower the implied liquidity, the more liquid the underlying asset. In the limit, as $\lambda \to 0$, the bid and ask prices coincide and the market is not illiquid anymore.

The model-free spot implied liquidity for a maturity of 30 days is then obtained by interpolating linearly the maturity components:

$$\mathrm{IL}(T_0) = \frac{1}{T_0}\left(\omega\, T_1\, \mathrm{IL}(T_1) + (1-\omega)\, T_2\, \mathrm{IL}(T_2)\right),$$

where

$$\omega = \frac{T_2 - T_0}{T_2 - T_1},$$

and where T_0 denotes the reference maturity of 30 calendar days and (T_1, T_2) the two nearest-term expiration months, with $T_1 \leq T_0 < T_2$. A similar methodology can be used for any other reference maturity T_0. Note that, as for the VIX index, if the shortest maturity is less than 8 days, the first month maturity contracts are replaced by the third month maturity contracts in order to reduce pricing anomalies which might occur close to expiration. Moreover, we exclude any maturities for which less than 21 liquid options are quoted to ensure that the market-implied stock price distribution is captured precisely enough by the limited range of available option prices.

The LIX index is obtained by considering the actively traded S&P 500 stock index and can thus be seen as the model-free equivalent of the implied liquidity introduced in Corcuera et al. (2012) just as the VIX index is the model-independent counterpart of the implied volatility. The LIX index thus conveys and summarizes the level of liquidity risk in the market as perceived by today's investors, as the VIX and the SKEW index issued by the CBOE quantify the volatility and the tail risk. Hence, we are a hair's breadth away from seeing options on the LIX traded and therefore from being able to hedge against liquidity risk.

Although the focus of this paper is on stock market liquidity, the proposed methodology can be used to derive a liquidity index related to other markets. In particular, considering FX options or options on money market futures could lead to a currency market liquidity index and to a money market liquidity index, respectively.

4 The LIX historical time series: a statistical analysis

This section features a statistical analysis of the time series of the LIX index for a period ranging from 2 January 1998 to 30 October 2009. The LIX time series is derived by following the methodology described in Section 3. More precisely, for a given trading day, we first determine the liquidity parameter λ which replicates as well as possible the spot bid-ask spread observed in the market, and this for a time horizon equal to any of the quoted (liquid) maturities. We then infer the LIX index for that day by interpolating linearly the spot liquidity term structure to obtain the spot implied liquidity corresponding to the reference time horizon of 30 days. This procedure is repeated for each quoting day included in the sample period which ranges from January 1998 to October 2009. The resulting time series is hereafter referred to as the LIX time series. A similar methodology is repeated to obtain additional spot liquidity time series, corresponding to a time horizon of 3 and 6 months.

In the numerical study, we will typically divide the sample period into four distinct periods: the 1998-2000 pre-dot-com crisis period, the 2000-2002 dot-com crisis period, the 2002-2007 pre-credit crisis period and the 2007-2009 credit crisis period[7]. The dot-com crash is taken as the period

[7]We are grateful to the referee for suggesting comparing the overall level of stock market liquidity during the 2007-2009 liquidity crisis (also referred to as credit crisis or global financial crisis) and during the 2000-2002 dot-com crisis.

characterised by a gradual decrease of the "barometer of the dot-com mania", viz. the Nasdaq index, i.e. from 11 March 2000 to 9 October 2002 (see Authers & Mackenzie (2010)) while the ending day of the period preceding the global financial crisis is chosen as the day when the S&P 500 reached its maximum value, viz. 9 October 2007[8].

Figure 6 shows the time series of the S&P 500 index price, of the VIX index and of the LIX index, whereas Figure 7 shows the time series of different spot liquidity proxies, namely the trading volume, the relative bid-ask spread (i.e. the dollar bid-ask spread expressed as a percentage of the spot price) and the model-free spot implied liquidity for three time horizons, i.e. one, three and six months. Figure 6 seems to indicate that the VIX and the LIX are positively correlated and that both are negatively correlated with the spot price[9]. We also clearly see that the VIX index is much more persistent than the LIX index which reverts much faster to its mean level of 0.0592. Moreover, the results indicate a huge increase of the volatility during the global financial crisis: the average VIX went up by 48.13 percent during the credit crisis whereas the increase of the VIX average value barely reached 1.80 percent during the dot-com crisis.

Figure 7 indicates that the trading volume has gradually increased as from the start of the sample period, although a downward trend is observed from March 2009 onwards. The steady increase in trading volume over the sample period can be explained by the emergence of high frequency trading resulting from technological breakthroughs such as algorithmic trading, but also by the decrease in trading commission and in the tick size (see, for instance, Chordia et al. (2011)). Moreover, the trading volume is significantly higher after 9 October 2007; the average trading volume increasing by 164.28 percent after that day, compared to the 2002-2007 pre-crisis period. This trend can be further explained by the fact that the collapse of the structured-credit market in mid-2007 led to a drying up of the overall funding liquidity which in turn caused fire sales from hedge funds and financial institutions triggered by the necessity of meeting margin calls, reserve requirements and redemption demands (see, for instance, Barrell & Davis (2008)). A similar observation of increased trading volume during turmoil periods was also made by Aitken & Comerton-Forde (2003) for the Jakarta Stock Exchange when investigating the ability of different liquidity measures to reflect the Asian liquidity crisis of the late 1990s, and by Chordia et al. (2001) when analyzing the aggregate correlation between spread changes and the number of transactions for U.S. equities between 1988 and 1998. On the other hand, both the relative bid-ask spread and the spot implied liquidity tend to be higher during the two turmoil periods considered. Indeed, the average relative spread went up by 16.47 percent and by 118.46 percent whereas the average LIX went up by 21.62 and by 24.22 percent only, during the dot-com crash and the credit crisis, respectively. Widening the time horizon for the spot implied liquidity computation inevitably leads to a rise in the overall level of liquidity[10], and, generally, to a more marked average increase of spot implied liquidity during turmoil periods. Indeed, the average relative increase of the model-free spot implied liquidity amounts to 27.37 and to 26.06 percent after the dot-com crash and to 30.07 and to 39.23 percent after October 2007, and this for a time horizon of 3 and 6 months, respectively. Overall, turmoil periods tend to be characterized by a decrease of market liquidity, which is more pronounced during funding liquidity driven crises. Hence, although the credit crisis mainly affected liquidity in the money market, the results indicate that liquidity also dried up in the stock market. This is in line with the empirical

[8]This date (or more precisely the next day) turns out to be close to the month of August 2007, which is seen by many as the actual start of the credit crunch.

[9]For more evidence about the correlation structure, see the cross-correlation function reported in Table 2.

[10]This trend can be easily explained by the fact that the longer the time horizon, the more spread the bulk behavior of the underlying distribution and hence the wider the conic finance bid-ask spread for a fixed distortion parameter λ.

evidence of liquidity co-movement across markets, and in particular across the bond and stock markets (see, for instance, Chordia et al. (2005) or Goyenko & Ukhov (2009)).

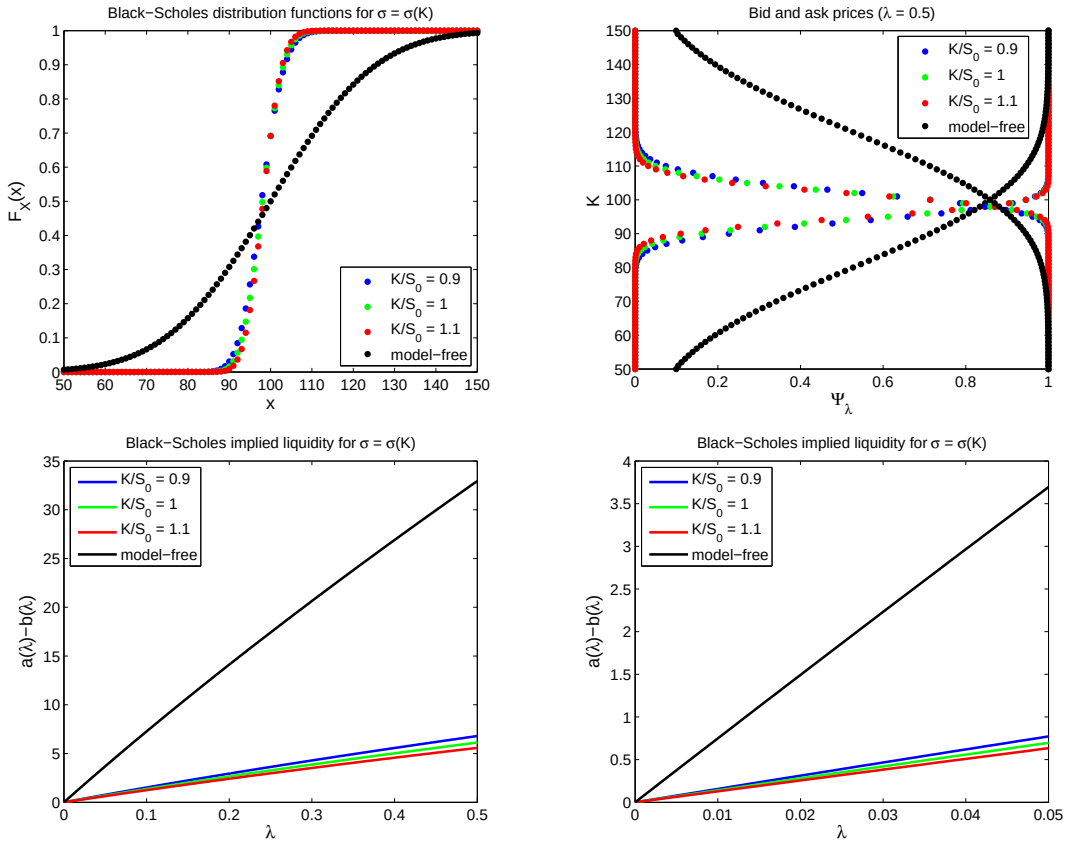

Figure 3: Model-free distribution function and Black-Scholes distribution function for different volatility parameters (upper left), corresponding bid and ask curves (upper right) and conic finance bid-ask spread as a function of the liquidity parameter λ (lower) for the artificial market situation shown in Figure 2.

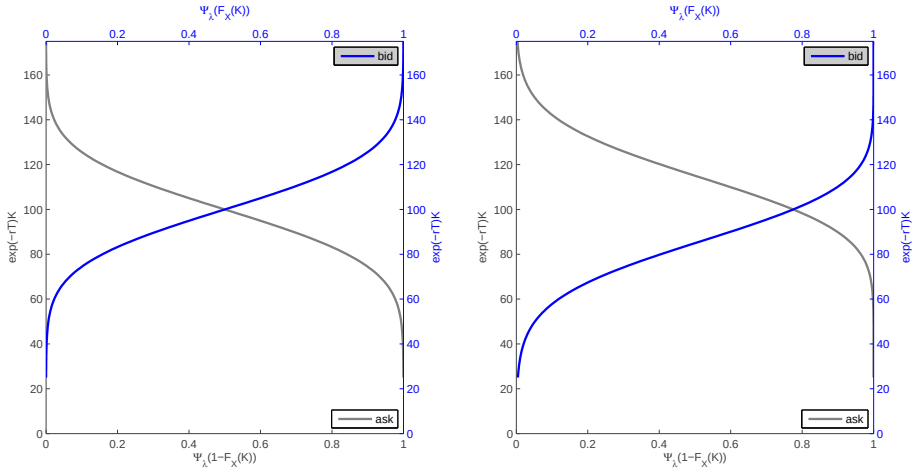

Figure 4: Computation of the model-free bid and ask prices for the artificial market situation shown in Figure 2 in a liquid market ($\lambda = 0$, left) and an illiquid market ($\lambda = 0.5$, right). The ask and the bid prices are computed as the area below the ask and the bid price curve, respectively.

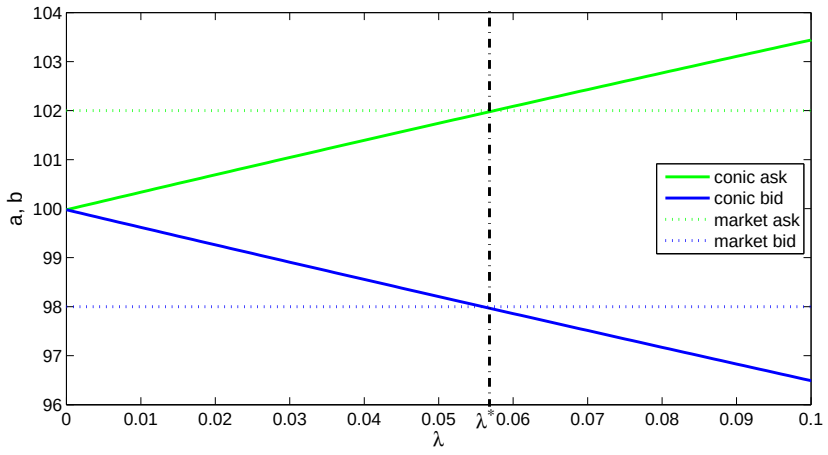

Figure 5: Computation of the model-free implied liquidity. The model-free implied liquidity (λ^*) is computed as the value of the distortion parameter λ for which the conic finance bid and ask prices coincide with the market bid and ask prices.

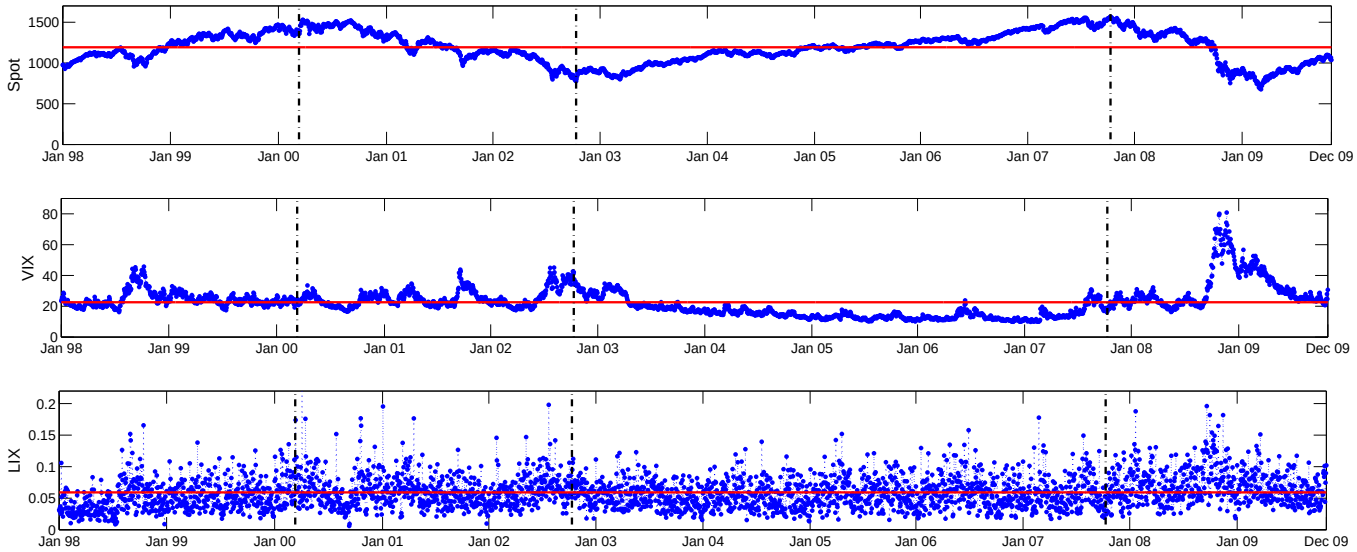

Figure 6: Time series plot of the S&P 500 stock price (upper), of the VIX volatility index (center) and of the LIX index (lower). The vertical dotted lines separate the dot-com crash period (March 2000 - October 2002) and the global financial crisis period (October 2007 - October 2009) from the pre-crisis periods.

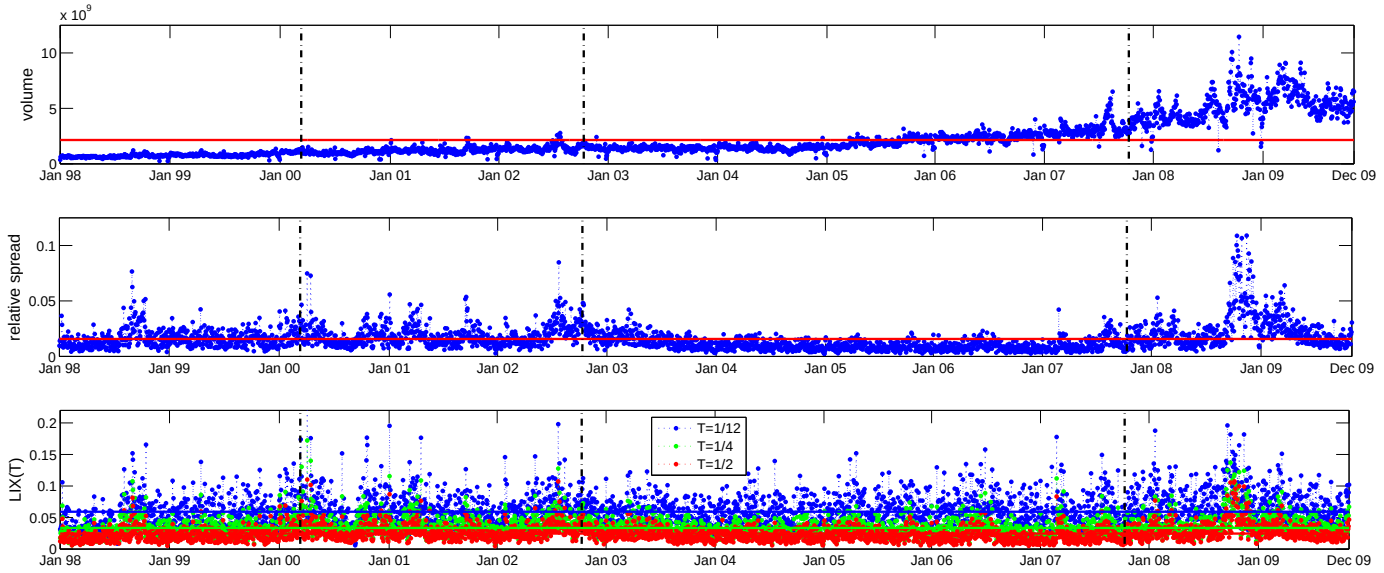

Figure 7: Time series plot of different liquidity measures: trading volume (upper), relative bid-ask spread (center) and model-free spot implied liquidity for a time horizon of one, three and six months (lower). The vertical dotted lines separate the dot-com crash period (March 2000 - October 2002) and the global financial crisis period (October 2007 - October 2009) from the pre-crisis periods.

In order to gain more insight into the historical behavior of the different risk factors, Table 1 shows the autocorrelation function (ACF) of volatility (VIX) and liquidity indicator (LIX and relative bid-ask spread) levels and daily changes for the first five lags. The results indicate a significant positive autocorrelation for the different volatility and liquidity proxies up to five lags, which is a sign that the different original time series exhibit some dependence. To further test for liquidity and volatility clustering, Figure 8 shows the ACF of the different time series up to 100 lags. The ACF of the LIX remains significant (and positive) up to 66 lags, suggesting a long-run dependence in the liquidity index: periods of high (respectively low) liquidity are more likely to be followed by periods of high (respectively low) liquidity. However, the long memory characteristic is less marked (both in amplitude and duration) than for the relative spread and the volatility index VIX. Nevertheless, we clearly see that widening the time horizon of the spot implied liquidity leads to a more persistent spot liquidity time series since the ACF of the 3 months, and to a larger extent 6 months model-free spot implied liquidity remains significant and positive for a larger number of lags, namely 75 and more than 100, respectively.

Moreover, the significant negative ACF of the LIX first differenced time series at the first lag reported in Table 1 (-0.5325) suggests a mean-reverting trend in the LIX time series. This observation corroborates the evidence of negative serial correlation in differenced time series of liquidity for NYSE listed stocks, as highlighted by Chordia et al. (2001). A similar conclusion can be drawn for the relative spread and for the VIX index, even so to a smaller extent. Indeed, the first lag ACF of the daily change time series is higher, in absolute value, for the LIX and, to a smaller extent, for the relative bid-ask spread than for the VIX. The ACF of the original and differenced series suggests that, although both historical liquidity and volatility exhibit a mean-reverting trend, liquidity indicators revert faster towards their long-run mean than volatility which is more persistent through time.

Table 1: Autocorrelation function of liquidity and volatility indicator time series and differenced time series (referenced as Δ) for the first five lags. Figures with * are significant at a level $\alpha = 0.05$. A correlation is significant if it lies outside the test bounds which are approximated by $\pm 2/\sqrt{N}$, where N is the number of observations. The results are based on the whole sample period (January 1998 - October 2009).

Time series	ACF(1)	ACF(2)	ACF(3)	ACF(4)	ACF(5)
LIX	0.2045*	0.2557*	0.2166*	0.1868*	0.2060*
Δ LIX	-0.5325*	0.0571*	-0.0054	-0.0317	0.0382*
VIX	0.9832*	0.9709*	0.9613*	0.9514*	0.9433*
Δ VIX	-0.1297*	-0.0837*	0.0115	-0.0544*	-8.9E-06
relative spread	0.6206*	0.6262*	0.6041*	0.5889*	0.5806*
Δ relative spread	-0.5073*	0.0366	-0.0087	-0.0099	0.0157*

Table 2 shows the cross-correlation function between the differenced series of the liquidity proxies on one hand and the log-asset return series or the volatility daily change series on the other. From order-based liquidity proxies, an increase in the volatility of the S&P 500 tends to be associated

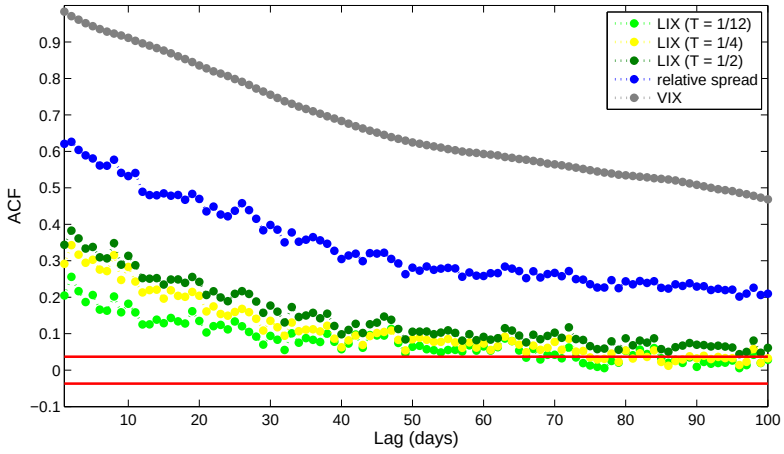

Figure 8: Sample autocorrelation function of liquidity and volatility proxies with test bounds (approximated by $\pm 2\sqrt{N}$, N being the number of observations) for the first 100 lags. The results are based on the whole sample period (January 1998 - October 2009).

with a decrease in the S&P 500 liquidity. However, looking at the trading volume leads to the opposite conclusion, i.e. an increase in the volatility tends to be concurrent with an increase in the trading volume of the S&P 500 stock index, and so with an increase in trade-based liquidity. Those observations confirm the conclusions drawn from the time series plots and the empirical results of Aitken & Comerton-Forde (2003) but also of Chou et al. (2011) which have shown that an increase in spot liquidity, measured as relative bid-ask spreads, is associated with a decline in the level of the implied volatility curve. Moreover, an increase in the asset price tends to be concurrent with a decrease in the LIX index and in the relative bid-ask spread, but also with a decrease in the trading volume, the latter being not significant at a level $\alpha = 0.05$. Hence, from the order-based liquidity proxies (i.e. the relative bid-ask spread and the LIX index), spot liquidity tends to dry up when the returns are below and more volatile than average.

To conclude the statistical analysis, we compare the LIX index with another popular measure of illiquidity, viz. the Amihud's (il)liquidity ratio (see Amihud (2002)):

$$\text{ILLIQ} = \frac{|R|}{\text{Vol}},$$

where R is the daily return and Vol the daily volume expressed in dollars. Table 3 shows the cross-correlation function between the differenced series of the LIX index (with a time horizon of 1, 3 or 6 months) and the differenced series of the Amihud's liquidity ratio. An increase in the LIX index tends to be concurrent with an increase in the Amihud's (il)liquidity ratio, whatever the time horizon of the LIX index. Hence, the two measures of (il)liquidity are significantly positively correlated (see also Figure 9) so that the LIX index could be used as an alternative (forward looking) measure of illiquidity.

Table 2: Cross-correlation function between the differenced series (referenced as Δ) of liquidity proxies on one hand and the log-asset return series or the volatility daily change series on the other. Figures with * are significant at a level $\alpha = 0.05$. A correlation is significant if it lies outside the test bounds which are approximated by $\pm 2/\sqrt{N}$, where N is the number of observations. The results are based on the whole sample period (January 1998 - October 2009).

Time series	CF(-3)	CF(-2)	CF(-1)	CF(0)	CF(1)	CF(2)	CF(3)
Δ LIX/Log return	0.0136	0.0224	-0.0932*	-0.0509*	0.0355	-0.0186	0.0225
Δ LIX/Δ VIX	0.0029	-0.0077	0.0071	0.0937*	-0.0256	-0.0027	-0.0081
Δ volume/Log return	-0.0018	0.0289	-0.1313*	-0.0206	0.0416	-0.0405	0.0501
Δ volume/Δ VIX	-0.0257	-0.0157	0.0824*	0.1062*	-0.0685*	0.0582*	-0.0549*
Δ relative spread/Log return	0.0104	0.0459*	-0.1184*	-0.1404*	0.0498*	-0.0044	0.0086
Δ relative spread/Δ VIX	-0.0059	-0.0277	0.0075	0.2270*	-0.0406*	-0.0155	-0.0066

Table 3: Cross-correlation function between the differenced series (referenced as Δ) of the LIX index (with a time horizon of 1, 3 or 6 months) and the differenced series of the Amihud's (il)liquidity ratio. Figures with * are significant at a level $\alpha = 0.05$. A correlation is significant if it lies outside the test bounds which are approximated by $\pm 2/\sqrt{N}$, where N is the number of observations. The results are based on the whole sample period (January 1998 - October 2009).

Time series	CF(-3)	CF(-2)	CF(-1)	CF(0)	CF(1)	CF(2)	CF(3)
Δ LIX(1m)/Δ ILLIQ	0.0135	0.0237	-0.2877*	0.5423*	-0.3084*	0.0426*	-0.0170
Δ LIX(3m)/Δ ILLIQ	0.0119	0.0276	-0.2958*	0.5587*	-0.3153*	0.0365	-0.0060
Δ LIX(6m)/Δ ILLIQ	0.0134	0.0221	-0.3070*	0.5881*	-0.3321*	0.0402*	-0.0066

5 The implied liquidity surfaces

In this section, we extend the theoretical framework of Section 3 in order to introduce the model-free option liquidity, and hence the so-called implied liquidity surfaces. Taking as claim the vanilla call option with strike K and maturity T, the conic finance ask and bid prices (3) and (4) can be derived from the distribution function F_X of the underlying asset at time T by

$$a^{\text{call}}(\lambda; K, T) = \exp(-rT) \int_K^\infty (K - x) d\Psi_\lambda(1 - F_X(x)) \tag{6}$$

and

$$b^{\text{call}}(\lambda; K, T) = \exp(-rT) \int_K^\infty (x - K) d\Psi_\lambda(F_X(x)). \tag{7}$$

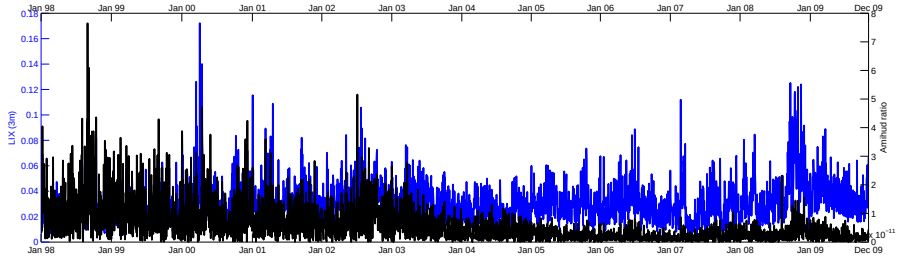

Figure 9: Amihud's (il)liquidity ratio (black) and 3 months LIX index (blue).

Similarly, for a put option with strike K and maturity T, we have

$$a^{\text{put}}(\lambda; K, T) = -\exp(-rT) \int_0^K (x - K) d\Psi_\lambda(F_X(x)) \tag{8}$$

and

$$b^{\text{put}}(\lambda; K, T) = -\exp(-rT) \int_0^K (K - x) d\Psi_\lambda(1 - F_X(x)). \tag{9}$$

The model-free bid and ask prices of a call option with strike $K_j \in \mathbf{K}$ and maturity T are approximated by considering the market-implied expression of the distribution function (1):

$$a^{\text{call}}(\lambda; K_j, T) = \exp(-rT) \sum_{i=j}^{N-1} \frac{(K_j - K_i) + (K_j - K_{i+1})}{2} \left(\Psi_\lambda(1 - F(K_{i+1})) - \Psi_\lambda(1 - F(K_i))\right); \tag{10}$$

$$b^{\text{call}}(\lambda; K_j, T) = \exp(-rT) \sum_{i=j}^{N-1} \frac{(K_i - K_j) + (K_{i+1} - K_j)}{2} \left(\Psi_\lambda(F(K_{i+1})) - \Psi_\lambda(F(K_i))\right). \tag{11}$$

For a put option with strike $K_j \in \mathbf{K}$ and maturity T, we have

$$a^{\text{put}}(\lambda; K_j, T) = -\exp(-rT) \sum_{i=0}^{j-1} \frac{(K_i - K_j) + (K_{i+1} - K_j)}{2} \left(\Psi_\lambda(F(K_{i+1})) - \Psi_\lambda(F(K_i))\right); \tag{12}$$

$$b^{\text{put}}(\lambda; K_j, T) = -\exp(-rT) \sum_{i=0}^{j-1} \frac{(K_j - K_i) + (K_j - K_{i+1})}{2} \left(\Psi_\lambda(1 - F(K_{i+1})) - \Psi_\lambda(1 - F(K_i))\right). \tag{13}$$

Equations (10) to (13) indicate that, for a given distortion parameter λ, the deeper the option is in-the-money (i.e. the lower and the higher the strike for call and put options, respectively), the wider the conic finance bid-ask spread. This is illustrated in Figure 10 for the artificial implied volatility skew governed by Equation (5) and depicted in Figure 2. Moreover, Figure 10 seems to

indicate that, for a given distortion parameter λ, the conic finance option spread is symmetric for call and put options: a call and a put option with strikes (approximately) equidistant from the forward price and the same time to maturity have the same conic finance bid-ask spread. This can be explained by the symmetry of the underlying stock price distribution around the forward level. Indeed, using the fact that $d\Psi_\lambda(y) = -d\left(1-\Psi_\lambda(y)\right)$ to express the bid price and integrating by parts Equations (6) to (9) give the following equivalent expressions for bid and ask option prices (see also Cherny & Madan (2010)):

$$a^{\text{call}}(\lambda;K,T) = \exp(-rT)\int_K^\infty \Psi_\lambda(1-F_X(x))dx, \quad b^{\text{call}}(\lambda;K,T) = \exp(-rT)\int_K^\infty \left(1-\Psi_\lambda(F_X(x))\right)dx;$$

$$a^{\text{put}}(\lambda;K,T) = \exp(-rT)\int_0^K \Psi_\lambda(F_X(x))dx, \quad b^{\text{put}}(\lambda;K,T) = \exp(-rT)\int_0^K \left(1-\Psi_\lambda(1-F_X(x))\right)dx.$$

The bid-ask spread is thus equal to

$$\text{spread}^{\text{call}}(\lambda;K,T) = \exp(-rT)\int_K^\infty \left(\Psi_\lambda(1-F_X(x)) - (1-\Psi_\lambda(F_X(x)))\right)dx;$$

$$\text{spread}^{\text{put}}(\lambda;K,T) = \exp(-rT)\int_0^K \left(\Psi_\lambda(F_X(x)) - (1-\Psi_\lambda(1-F_X(x)))\right)dx.$$

Hence, if $F_X(x)$ is symmetric around the forward price F_0, as it is almost the case for the artificial market situation (see upper left panel of Figure 3), then we have that $F_X(F_0-\Delta K) = 1-F_X(F_0+\Delta K)$ such that the call and put spreads coincide as long as $K^{\text{call}} = F_0 \pm \kappa$ and $K^{\text{put}} = F_0 \mp \kappa$, i.e. as long as the call and put strike prices are equidistant from the forward level. This explains why the conic finance option spread related to the artificial volatility skew appears to be symmetric for call and put options, as shown in Figure 10. Nevertheless, this symmetry will not hold anymore under more realistic market situations for which the stock price distribution has a non-zero skewness.

Matching the model-free approximation of the conic finance bid-ask spread with the bid-ask spread observed in the market for each quoted option allows us to construct an implied liquidity surface for the call and the put options separately. These two liquidity surfaces can be seen in Figure 11 (upper panel) for 1 October 2009. For the sake of comparison, the Black-Scholes model-dependent liquidity surfaces, for which the volatility parameter is set equal to the option implied volatility, are shown in the center panel of Figure 11 and the relative market spreads (i.e. the dollar spreads expressed as a percentage of the mid price) in the lower panel of Figure 11. Figure 11 indicates that the deeper the option is in-the-money, the more liquid the option since the model-free implied liquidity λ increases and decreases with the strike price for call and put options, respectively. Moreover, comparing the option model-free implied liquidity with the option relative bid-ask spread highlights a great similarity between the two liquidity proxies, both across strikes and times to maturity. Hence, both the relative bid-ask spread and the conic finance liquidity proxy allow for a scaling of the dollar bid-ask spread by taking into account the option price, either explicitly or implicitly. On the other hand, comparing the model-free and the Black-Scholes implied

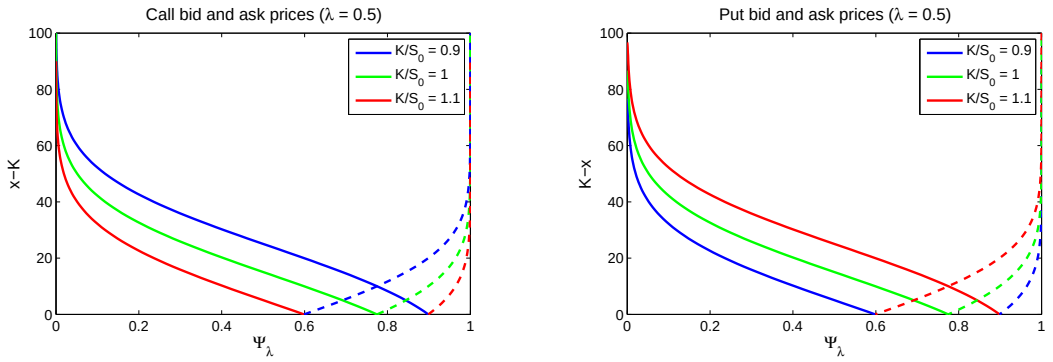

Figure 10: Model-free bid and ask curves for one-year call (left) and put options (right) with a strike price of 90, 100 and 110 percent of the forward price under the artificial market situation shown in Figure 2. The bid-ask spread is computed as the difference between the area below the ask curve (solid line) and the area below the bid curve (dashed line).

liquidity surfaces underpins the model-dependency of the implied liquidity as proposed in Corcuera et al. (2012), both in the shape and in the scale of the liquidity skew.

Figure 12 shows the term structure of the at-the-money model-free implied liquidity for call and put options. It clearly indicates that the at-the-money implied liquidity decreases with the time to maturity. We note that the changes of concavity in the implied liquidity term structure correspond to the less liquid maturities (i.e. those with the fewest liquidly traded options) and can be partially explained by the poor precision of the integration scheme used to approximate the option bid and ask prices (see Table 4 for the number of liquidly traded out-of-the-money options as a function of maturity for 1 October 2008). Figure 11 indicates that the decreasing trend of the model-free option implied liquidity as a function of the time to maturity holds whatever the option moneyness. Besides, a similar term structure can be observed for the relative bid-ask spread of call and put options, as shown in the lower panel of Figure 11. Hence, the longer the option term, the more liquid the option. This dependence is not new and has already been observed in different option markets. In particular, Wei & Zheng (2010) showed that options with a shorter maturity tend to have a wider relative bid-ask spread when considering actively traded U.S. options between 1996 and 2007 whereas the empirical study of George & Longstaff (1993) relative to the S&P 100 index option market highlighted a wider dollar bid-ask spread for near-maturity options than for long-maturity options, on average. Moreover, Chong et al. (2003) gave evidence that the relative volatility spread of at-the-money currency call options is negatively related to their time to maturity.

Table 4: Number of liquidly traded out-of-the-money options for each option maturity T (October, 1, 2008).

T	0.0438	0.1397	0.2164	0.2493	0.2932	0.4658	0.4959	0.7151	0.7452	0.9644	0.9973	1.2137	1.7123	2.2110
#	97	147	114	27	48	43	24	38	23	41	22	53	48	51

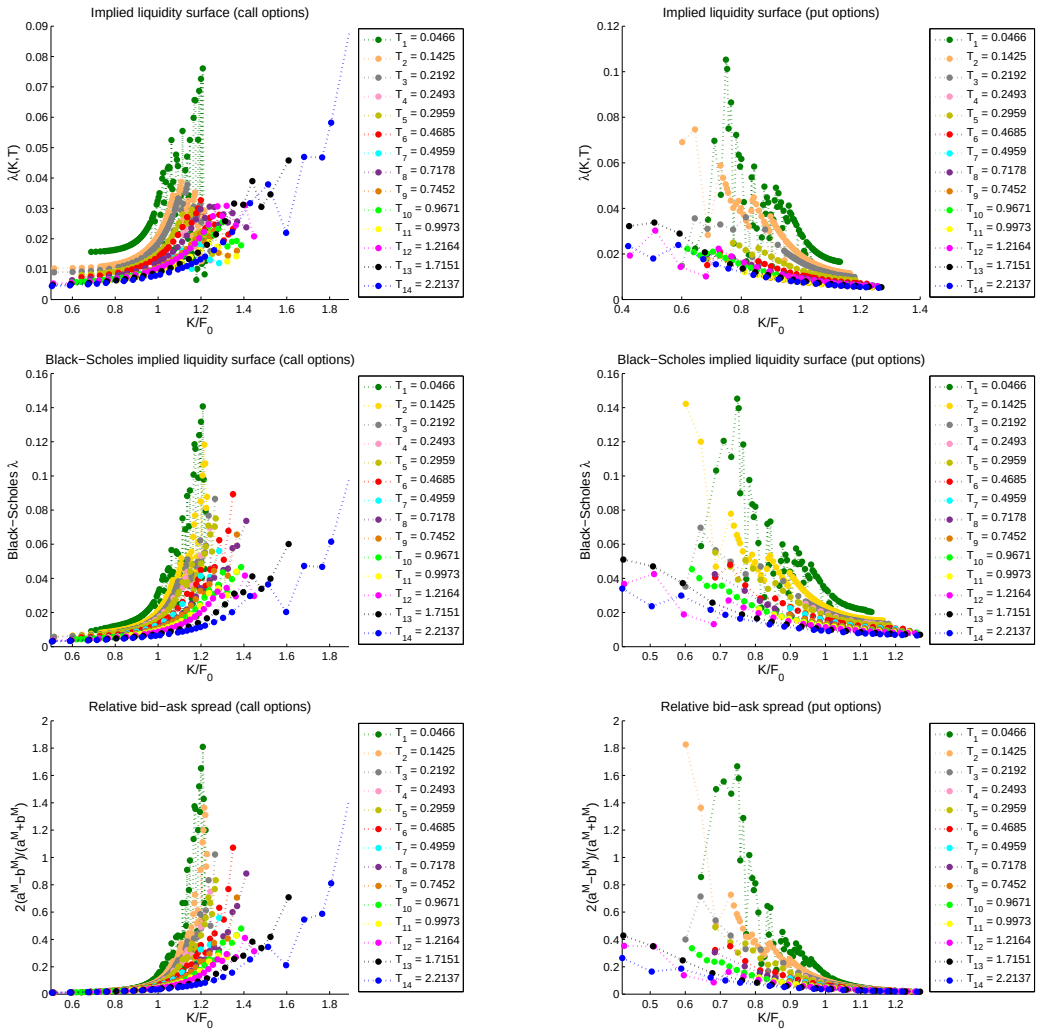

Figure 11: Model-free implied liquidity surface (upper), Black-Scholes implied liquidity surface (center) and relative bid-ask spreads (lower) for call options (left) and put options (right) (October, 1, 2008).

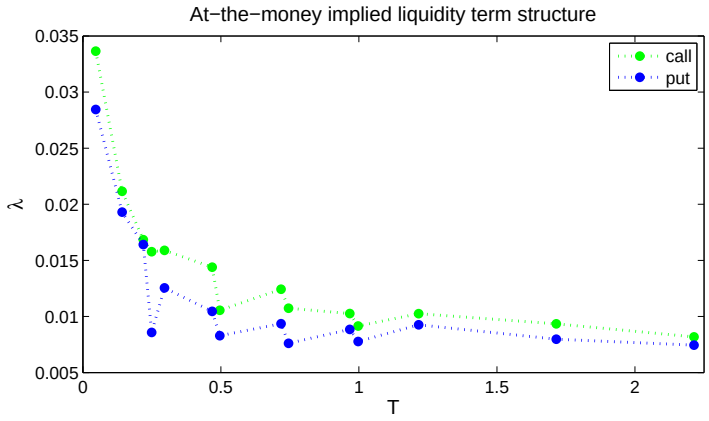

Figure 12: At-the-money model-free implied liquidity term structure for call and put options (October, 1, 2008).

To compare the liquidity dynamics in the spot and the option market, we can compute rank correlations between the spot implied liquidity and the option implied liquidity for different option terms and moneynesses. Table 5 shows the Kendall's tau rank correlation coefficient for different pairs of model-free spot and option implied liquidity time series for a sample period spanning from January 1998 to October 2009. For each trading day included in the sample period, we consider the options with the closest quoted strike price to 0.95, 1 and 1.05 times the forward index level F_0. The lowest strike leads to the selection of in-the-money call and out-of-the-money put options whereas the highest strike results in the selection of out-of-the-money call and in-the-money put options. For the spot implied liquidity, we follow the methodology of Section 3 aimed at building the LIX index. More precisely, we interpolate linearly the spot implied liquidity term structure in order to obtain a model-free spot implied liquidity for a time horizon of one, three and six months. On the other hand, for the model-independent option implied liquidity, we consider the quoted option which has the closest strike and time to maturity to the corresponding reference levels, rather than interpolating the implied liquidity surface across these two dimensions. The results indicate that the spot implied liquidity for different time horizons tends to move in a comonotonic way, as evidenced by the high value of the Kendall's τ ($\tau(S_{im}, S_{jm})$ greater than 70 percent $\forall i, j \in \{1, 3, 6\}$). This can be easily explained by the fact that the spot implied liquidity depends, by construction, on the market spread and on the asset price distribution, where only the asset price distribution changes with the time horizon. Moreover, the rank correlation between the spot implied liquidity and the call option implied liquidity is higher than the one between the spot and the put option implied liquidity; the latter being either rather small or not significant at a level α of 0.05. Overall, the rank correlation between the spot and the option implied liquidity is not high (at most equal to 25 percent), which indicates that the liquidity dynamics is different in the two markets. This result is in line with previous empirical studies (see for instance Cao & Wei (2010) or Wei & Zheng (2010)). The positive correlation between the spot and the option implied liquidity is consistent with the *derivative hedge* model proposed by Cho & Engle (1999), stating that the spreads in the derivative market are (partially) determined by the spreads in the underlying market since

market makers in derivative markets can hedge (imperfectly) their positions by trading in the underlying market. Moreover, the option implied liquidities are positively correlated, whatever the option type, moneyness and time to maturity. This is again supported by the *derivative hedge* theory. Nevertheless, the option implied liquidities tend to move more strongly together when they correspond to options of the same type or with the same time to maturity. Furthermore, despite being quite high, the concordance between the option implied liquidities is less marked than the one between the spot implied liquidities, which can be explained by the fact that both the option cash-flow distribution and the option market bid-ask spread change from option to option.

Table 5: Pairwise Kendall's τ rank correlations between spot, call and put model-free implied liquidities. S stands for spot, C for call, P for put, A for at-the-money, I for in-the-money and O for out-of-the-money. Moreover the symbols *1m*, *3m* and *6m* represent the time horizon of 1, 3 and 6 months, respectively. Figures with $^{(*)}$ are significant at a level $\alpha = 0.05$. The results are based on the whole sample period (January 1998 - October 2009).

	IC1m	AC1m	OC1m	IP1m	AP1m	OP1m	S3m	IC3m	AC3m	OC3m	IP3m	AP3m	OP3m	S6m	IC6m	AC6m	OC6m	IP6m	AP6m	OP6m
S1m	0.1671 (*)	0.1964 (*)	0.2428 (*)	0.0113	0.0380 (*)	0.0725 (*)	0.7303 (*)	0.1234 (*)	0.1109 (*)	0.1578 (*)	0.0275 (*)	0.0063	0.0122	0.7243 (*)	0.1319 (*)	0.1220 (*)	0.1114 (*)	0.0537 (*)	0.0479 (*)	0.0197
IC1m		0.5460 (*)	0.4279 (*)	0.5785 (*)	0.3329 (*)	0.3044 (*)	0.0619 (*)	0.4383 (*)	0.4142 (*)	0.2440 (*)	0.4455 (*)	0.3804 (*)	0.3306 (*)	0.0421 (*)	0.4184 (*)	0.4095 (*)	0.3363 (*)	0.4847 (*)	0.4584 (*)	0.4315 (*)
AC1m			0.4395 (*)	0.3065 (*)	0.2634 (*)	0.2268 (*)	0.1014 (*)	0.2797 (*)	0.2901 (*)	0.2434 (*)	0.2304 (*)	0.2186 (*)	0.1987 (*)	0.0914 (*)	0.2671 (*)	0.2622 (*)	0.2287 (*)	0.2589 (*)	0.2556 (*)	0.2315 (*)
OC1m				0.1379 (*)	0.0994 (*)	0.1468 (*)	0.1438 (*)	0.2356 (*)	0.2069 (*)	0.2321 (*)	0.1382 (*)	0.1003 (*)	0.0864 (*)	0.1354 (*)	0.2316 (*)	0.2170 (*)	0.1831 (*)	0.1806 (*)	0.1694 (*)	0.1269 (*)
IP1m					0.4585 (*)	0.3556 (*)	-0.0547 (*)	0.3609 (*)	0.3699 (*)	0.1325 (*)	0.5222 (*)	0.4816 (*)	0.3966 (*)	-0.0750 (*)	0.3448 (*)	0.3538 (*)	0.2959 (*)	0.5297 (*)	0.5006 (*)	0.4983 (*)
AP1m						0.3319 (*)	-0.0045	0.2111 (*)	0.2322 (*)	0.1325 (*)	0.2814 (*)	0.2926 (*)	0.2674 (*)	-0.0074	0.2205 (*)	0.2272 (*)	0.2191 (*)	0.2876 (*)	0.2922 (*)	0.2895 (*)
OP1m							0.0023	0.1857 (*)	0.1940 (*)	0.1083 (*)	0.2421 (*)	0.2273 (*)	0.2272 (*)	-0.0007	0.1856 (*)	0.1797 (*)	0.1530 (*)	0.2521 (*)	0.2479 (*)	0.2458 (*)
S3m								0.1745 (*)	0.1603 (*)	0.2529 (*)	0.0027	-0.0063	0.0149	0.8271 (*)	0.1444 (*)	0.1336 (*)	0.1308 (*)	0.0133	0.0092	-0.0209
IC3m									0.7748 (*)	0.5047 (*)	0.5646 (*)	0.4760 (*)	0.4244 (*)	0.0982 (*)	0.5321 (*)	0.5066 (*)	0.4149 (*)	0.5266 (*)	0.4932 (*)	0.4368 (*)
AC3m										0.5193 (*)	0.5196 (*)	0.4895 (*)	0.4389 (*)	0.0821 (*)	0.4409 (*)	0.4471 (*)	0.3937 (*)	0.4639 (*)	0.4447 (*)	0.4291 (*)
OC3m											0.2077 (*)	0.1892 (*)	0.2109 (*)	0.1886 (*)	0.3200 (*)	0.3157 (*)	0.3231 (*)	0.2023 (*)	0.1935 (*)	0.1737 (*)
IP3m												0.7291 (*)	0.5719 (*)	-0.0481 (*)	0.4332 (*)	0.4287 (*)	0.3452 (*)	0.6861 (*)	0.6290 (*)	0.5851 (*)
AP3m													0.6131 (*)	-0.0538 (*)	0.3531 (*)	0.3629 (*)	0.3212 (*)	0.5584 (*)	0.5385 (*)	0.5516 (*)
OP3m														-0.0343 (*)	0.3159 (*)	0.3170 (*)	0.3025 (*)	0.4448 (*)	0.4493 (*)	0.4649 (*)
S6m															0.1675 (*)	0.1577 (*)	0.1643 (*)	-0.0094	-0.0087	-0.0318 (*)
IC6m																0.8399 (*)	0.6987 (*)	0.5658 (*)	0.5401 (*)	0.4753 (*)
AC6m																	0.7378 (*)	0.5514 (*)	0.5296 (*)	0.4846 (*)
OC6m																		0.4272 (*)	0.4253 (*)	0.4226 (*)
IP6m																			0.8140 (*)	0.7079 (*)
AP6m																				0.7289 (*)

6 Maximum likelihood estimation of Vasicek and CIR model parameters

In this section, we adopt a similar methodology as the one followed by Albrecher et al. (2013) to evidence the model sensitivity of implied liquidity, but in order to highlight the fundamental differences in the liquidity dynamics during periods characterized by different levels of overall market fear, and this for both the underlying and option markets. For that purpose, we consider the five time series windows defined in Section 4: the whole sample period (January 1998 - October 2009), the internet boom (January 1998 - March 2000), the dot-com crash (March 2000 - October 2002), the pre-credit crisis period (October 2002 - October 2007) and the global financial crisis (October 2007 - October 2009). The aim of this section is twofold: to evidence the differences in liquidity dynamics in the stock and option markets during the different periods and to highlight the ability of mean-reverting models to fit historical implied liquidity time series. Allowing for both clustering and mean-reversion, these models are adequate candidates to model both volatility and liquidity. Indeed, the sample liquidity ACF has evidenced the presence of liquidity clusters: periods characterized by a low (high) market liquidity are more likely to be followed by periods of low (high) market liquidity whereas the first lag differenced liquidity ACF has highlighted a significant mean-reverting trend in the implied liquidity time series (see Section 4). Hence, we assume that the model-free implied liquidity and the VIX volatility index follow either a Vasicek process

$$dX_t = \kappa(\eta - X_t)dt + \zeta dW_t \tag{14}$$

or a CIR process (see Cox et al. (1985)):

$$dX_t = \kappa(\eta - X_t)dt + \zeta\sqrt{X_t}dW_t \tag{15}$$

where $\boldsymbol{\theta} = \{\kappa, \eta, \zeta\}$ is the set of model parameters and where $X = \{X_t, t \geq 0\}$ denotes the volatility or liquidity process and where $W = \{W_t, t \geq 0\}$ is a standard Brownian motion.

The model parameters $\boldsymbol{\theta}$ can be directly inferred from the time series of the model-free implied liquidity or volatility since the transition density[11] is known in closed-form under the CIR and Vasicek models (see for instance Aït-Sahalia (1999)). We have, under the CIR model,

$$p_X(x|x_0;\boldsymbol{\theta},\Delta t) = c\exp(-u-v)\left(\frac{v}{u}\right)^{q/2} I_q\left(2\sqrt{uv}\right),$$

where $q \equiv \frac{2\kappa\eta}{\zeta^2} - 1 \geq 0$, $c \equiv \frac{2\kappa}{\zeta^2(1-\exp(-\kappa\Delta t))}$, $u = cx_0\exp(-\kappa\Delta t)$ and $v = cx$ and where I_q is the modified Bessel function of the first kind of order q.

Under the Vasicek model, the transition density is given by

$$p_X(x|x_0;\boldsymbol{\theta},\Delta t) = \left(\frac{\pi\xi^2}{\kappa}\right)^{-1/2} \exp\left(-\left(x-\eta-(x_0-\eta)\exp\left(-\kappa\Delta t\right)\right)^2\frac{\kappa}{\xi^2}\right),$$

where $\xi^2 = \zeta^2(1-\exp(-2\kappa\Delta t))$.

The parameter set $\boldsymbol{\theta}$ can thus be efficiently estimated by the maximum likelihood estimation technique. Given the Markovian nature of the Vasicek process (14) and the CIR process (15), the

[11] The transition density is defined as the conditional density that X is at level x at time $t+\Delta t$ given that it is at level x_0 at time t.

maximum likelihood estimate $\widehat{\boldsymbol{\theta}}$ is obtained by maximizing

$$\log L_N(\boldsymbol{\theta}) = \frac{1}{N-1}\sum_{i=1}^{N-1} \log(p_X(x_{t_{i+1}}|x_{t_i};\boldsymbol{\theta},\Delta t_i)),$$

where N denotes the number of observations in the time series, $\{x_{t_1}, x_{t_2}, \ldots, x_{t_N}\}$ denotes the collection of historical implied liquidity or volatility values and where $\Delta t_i = t_{i+1} - t_i$ is the time between two successive observations.

Table 6 to Table 8 show the maximum likelihood (ML) estimate of the Vasicek and CIR model parameters for the VIX time series and the implied liquidity time series of the S&P 500 index and options for the whole sample period as well as for the different pre- and post-crisis sub-periods. The results show that the VIX and the LIX index (as well as its three and six months variants) are better modeled by a CIR process, whatever the sample period, whereas the best fitting model for the model-free option implied liquidity time series depends on the sample period and the option characteristics (viz. type, moneyness and time to maturity). In particular, option implied liquidity is better modeled by a CIR process for the credit crisis period, whatever the option type, moneyness or time to maturity and for the call options when considering the whole sample period. On the other hand, during the internet boom and the period preceding the global financial crisis, put option implied liquidity is better captured by the Vasicek model, except for the one-month out-of-the-money option. We also observe that the maximum likelihood is significantly higher for the option implied liquidity time series than for the VIX or the LIX time series. Moreover, except for two rare cases, the maximum likelihood of mean-reverting processes is higher for the put option implied liquidity time series than for the call option ones. Besides, the deeper the option is in-the-money and the longer its maturity, the better the fit of mean-reverting processes for modeling its implied liquidity. Indeed, $L(\hat{\boldsymbol{\theta}})$ decreases and increases with the strike price K for call and put options, respectively whereas it increases with the time to maturity of the option[12]. Furthermore, the goodness of fit of mean-reverting processes for modeling the spot implied liquidity time series is of the same magnitude order as the one corresponding to the VIX time series and it improves with the time horizon of the spot implied liquidity. We also note that the higher the degree of market confidence, the better the fit of mean-reverting models for the LIX (as well as for its three and six months variants) time series. Indeed, $L(\hat{\boldsymbol{\theta}})$ is significantly lower for the dot-com crash and the global financial crisis period than for the internet boom and the period preceding the credit crisis. For the VIX index, $L(\hat{\boldsymbol{\theta}})$ is also significantly lower for the global financial crisis than for the period preceding it but it is higher for the dot-com burst than for the dot-com boom. The ML results corroborate the empirical evidence shown in Albrecher et al. (2013) that mean-reverting processes are adequate candidates to model spot implied liquidity, and, to a larger extent, option implied liquidity, in a similar fashion as the volatility is modeled in stochastic volatility models, such as the Heston model.

Comparing the parameters calibrated on the different liquidity time series allows us to highlight the difference between the market liquidity dynamics in the stock and option markets. The ML estimate of the mean reversion rate κ is clearly much higher for the option and spot implied liquidity

[12]Nevertheless, we note two exceptions: the first one for the in-the-money call for which the likelihood of the Vasicek model is higher for the one month than for the three months call option implied liquidity during the dot-com crash and the second one for the 3 months call for which the likelihood of the Vasicek model is lower for the in-the-money than for the at-the-money call option implied liquidity during the dot-com crash.

time series than for the VIX. This confirms the results of Section 4 indicating that the spot implied liquidity reverts much faster to its long-run mean level than the volatility does. The ML estimate $\hat{\kappa}$ is lower for the 3 months spot implied liquidity than for the 1 month spot implied liquidity and decreases with the time horizon during the credit crisis period and the period preceding it. Besides, the ML estimate $\hat{\zeta}$ decreases with the time horizon for the spot implied liquidity. Hence, the longer the time horizon, the less volatile the spot implied liquidity and, in general, the higher its tendency to cluster; which is in line with the conclusions drawn from the time series plots and ACFs in Section 4. Moreover, options with a longer maturity are characterized by a lower volatility of liquidity $\hat{\zeta}$ and, generally, by a lower mean reversion rate $\hat{\kappa}$, indicating that option liquidity is less volatile and tends to cluster more for longer term options. Besides, the lower and the higher the strike price for call and put options, respectively, the lower $\hat{\zeta}$ and, in general, $\hat{\kappa}$, indicating that the deeper the option is in-the-money, the less volatile and the more persistent its liquidity through time. Hence, liquidity is less volatile and tends to cluster more for long term in-the-money call and put options. Comparing implied liquidity for call and put options with the same time to maturity and moneyness indicates that put option implied liquidity tends to revert faster to its long-run mean since $\hat{\kappa}$ is generally higher for put than for call options, with some clear exceptions for the dot-com crash period. Moreover, the estimate of the volatility of option implied liquidity $\hat{\zeta}$ is higher for the call than for the put options with the same time to maturity and moneyness during the dot-com crash and the credit crisis period[13], evidencing that call option implied liquidity is more volatile, at least during turmoil periods. Furthermore, for a fixed term, the ML estimates $\hat{\kappa}$ and $\hat{\zeta}$ are almost always higher for the spot implied liquidity than for option implied liquidities, indicating that liquidity is more volatile and tends to cluster less in the stock market than in the derivative market.

Comparing the ML estimates of η with the time series average levels (not reported) shows that the long-run mean parameter can be adequately approximated by the sample mean. Indeed, the relative discrepancy amounts to less than 2 percent, whatever the liquidity time series and the sample period. With few exceptions, the deeper the option is in-the-money or the longer its time to maturity, the lower the long-run liquidity parameter $\hat{\eta}$, i.e. the more liquid the option, on average. These findings are in line with the stylized facts of option liquidity proxies recalled in Section 5. Comparing the results of Table 7 and Table 8, we observe that put options tend to be more liquid than call options with the same characteristics (i.e. time to maturity and (absolute value of) moneyness). Indeed, the long-run liquidity ML estimate $\hat{\eta}$ is generally higher for call options. Besides, this trend intensifies during turmoil periods since the (relative) difference of $\hat{\eta}$ between call and put options doubled or even more during the dot-com crash and the subprime crisis period, except for the shortest term out-of-the money options during the global financial crisis.

We now compare the fitted Vasicek and CIR models for the five different time series windows under investigation. For the VIX time series, the estimate $\hat{\kappa}$ is larger for the four sub-periods than for the whole sample period, the biggest difference corresponding to the internet boom period. For the LIX index and its three and six months variants, the estimate $\hat{\kappa}$ is significantly lower during the credit crisis period than before, whereas it actually increases from the dot-com boom to the dot-com crash sub-period. Hence, volatility, measured as the VIX index, was more persistent once the internet bubble had burst whereas spot liquidity, measured as the LIX index, tended to persist more during the global financial crisis, i.e. during the liquidity-driven turmoil period. For option implied liquidity time series, the estimate $\hat{\kappa}$ is higher during the dot-com crash than during the

[13]Nevertheless, we note two exceptions for which $\hat{\zeta}$ is higher for the put than for the call option during the global financial crisis under the CIR model, namely the 3-months in-the-money and 6-months at-the-money options.

internet boom period, except for the longest-term at-the-money put option whereas it decreases from the pre-credit crisis to the credit crisis period. Hence, option liquidity tends to cluster more during liquidity-driven crisis periods, which is in line with the dynamics of the model-dependent implied liquidity reported in Albrecher et al. (2013). Such a trend is nevertheless not intrinsic to market turmoil periods since a completely opposite behavior is observed for the dot-com bubble. Regarding the VIX time series, the calibrated long-term volatility $\hat{\eta}$ and the calibrated volatility of volatility $\hat{\zeta}$ are lower and higher, with respect to the whole period setting, during the pre-credit crisis and the credit crisis period, respectively. Hence, the market overall volatility, measured as the VIX, tends to be higher and more volatile during liquidity-driven distress periods. These findings confirm the conclusions drawn from model-dependent volatility measures such as the Black-Scholes or Lévy implied at-the-money volatilities (see Albrecher et al. (2013)). On the other hand, the calibrated long-term volatility $\hat{\eta}$ only slightly increases from the boom to the burst phase of the dot-com bubble whereas the volatility of volatility estimate $\hat{\zeta}$ decreases. Besides, the calibrated long-run spot implied liquidity $\hat{\eta}$ is lower and higher, with respect to the whole sample period, during the two periods preceding a crisis and the two turmoil periods, respectively. The volatility estimate $\hat{\zeta}$ of the spot implied liquidity is higher during the burst phase of the two bubbles, except for the housing bubble under the CIR model. For the option implied liquidity time series, we usually observe a completely different phenomenon: the calibrated volatility of liquidity $\hat{\zeta}$ is higher during the burst phase than the boom phase of the internet bubble whereas it is generally (with few exceptions, mainly for call options under the Vasicek model) higher during the boom phase than the burst phase of the housing bubble.

The calibrated long-run option implied liquidity $\hat{\eta}$ is higher during the burst phase than during the boom phase of the dot-com bubble, except for the shortest term in-the-money put option. Comparing the estimate $\hat{\eta}$ for the four sub-periods highlights the overall drying-up of liquidity during the housing bubble and the subsequent global financial crisis. For put options, the long-run option implied liquidity estimate $\hat{\eta}$ is generally higher for the housing boom than for the credit crisis period. Some exceptions prevail for the shortest term at-the-money and out-of-the-money options as well as for the longest term in-the-money option. Regarding call options, the ML estimate $\hat{\eta}$ of long-run implied liquidity is higher during the global financial crisis than during the period preceding it, except for 1 month in-the-money and out-of-the-money options as well as for the 3-months at-the-money option. Hence, spot liquidity tends to dry up, on average, during distress periods, being liquidity-driven or not, whereas, put and call options generally turn out to be, on average, more and less liquid, respectively during the credit crunch than during the period preceding it. These findings have to be distinguished from the ones drawn from the model-dependent at-the-money call option implied liquidity time series considered in Albrecher et al. (2013).

7 Conclusion

This paper features a new model-free proxy of spot liquidity obtained by combining the conic finance theory, allowing for bid and ask parametric expressions, with the option payoff spanning formula, allowing for the extraction of the stock price distribution from the volatility surface. In particular, we construct a market liquidity index, the LIX, in a similar market-implied fashion as the one used by the CBOE to extract the volatility index VIX and the skewness index SKEW. Hence, the LIX can be seen as the model-free equivalent of the implied liquidity concept, just as the VIX is the model-independent counterpart of the implied volatility. The LIX index thus conveys the level of liquidity risk in the market as perceived by today's investors, as the VIX and the SKEW index quantify

Table 6: Maximum likelihood estimate of mean-reverting model parameters for volatility and spot implied liquidity time series.

		Vasicek				CIR			
	period	$\hat{\kappa}$	$\hat{\eta}$	$\hat{\zeta}$	$L(\hat{\boldsymbol{\theta}})$	$\hat{\kappa}$	$\hat{\eta}$	$\hat{\zeta}$	$L(\hat{\boldsymbol{\theta}})$
VIX	1998-2009	4.257	0.22123	0.26640	14.541	4.000	0.22133	0.48565	17.648
	dot-com boom	15.431	0.23874	0.27881	14.201	16.436	0.23878	0.52514	15.589
	dot-com crash	8.137	0.25276	0.24292	16.068	8.639	0.25223	0.46248	17.314
	pre-crisis	5.856	0.15315	0.16261	23.897	6.499	0.15394	0.38002	26.158
	global crisis	6.293	0.32047	0.43499	8.941	6.555	0.32004	0.67366	10.747
LIX (1m)	1998-2009	382.217	0.05880	0.73776	9.295	381.557	0.0588	2.94954	10.247
	dot-com boom	407.950	0.05097	0.70328	10.026	376.083	0.05097	2.91102	11.190
	dot-com crash	467.145	0.06118	0.80368	9.318	461.685	0.06118	3.02921	10.516
	pre-crisis	484.247	0.05623	0.75698	10.056	499.727	0.05623	3.20230	10.837
	global crisis	378.333	0.07027	0.83356	8.191	398.082	0.07027	3.11509	8.978
LIX (3m)	1998-2009	271.316	0.03160	0.37165	16.132	276.253	0.03160	1.97514	18.569
	dot-com boom	307.154	0.02684	0.33751	18.600	327.078	0.02684	1.99288	21.344
	dot-com crash	374.947	0.0364	0.47009	14.469	380.293	0.03645	2.23389	17.083
	pre-crisis	408.666	0.02852	0.37308	18.915	366.001	0.02852	2.07241	20.641
	global crisis	210.850	0.03799	0.39092	14.102	241.013	0.03797	1.97234	16.146
LIX (6m)	1998-2009	254.353	0.02378	0.27175	21.566	282.121	0.02378	1.69712	25.047
	dot-com boom	372.102	0.02212	0.28974	23.401	419.041	0.02212	1.91174	26.910
	dot-com crash	382.906	0.02806	0.35287	19.448	398.382	0.02806	1.95824	22.591
	pre-crisis	390.868	0.02081	0.26132	26.492	355.771	0.02081	1.71248	28.819
	global crisis	157.289	0.02735	0.27289	18.625	204.999	0.02732	1.62073	22.223

the volatility and the tail risk. Besides, we extend the methodology to build the so-called implied liquidity surfaces, one for the call and one for the put options, summarizing the dependency of the model-free option implied liquidity across moneynesses and times to maturity. The numerical study consists of an empirical analysis as well as a maximum likelihood analysis of popular mean-reverting processes to test the dynamics of volatility and liquidity time series. Although the numerical part of this paper focuses on the S&P 500 spot and option implied liquidity, the LIX methodology is of broad scope since it can be implemented as long as the option market has deep and active trading across a broad range of strike prices. Our findings confirm the stylized fact that the deeper the option is in-the-money and the longer its time to maturity, the more liquid the option, but also the more persistent its liquidity through time. Moreover, the results indicate that liquidity is more volatile and tends to cluster less in the underlying market than in the derivative market whereas volatility, measured as the VIX index, tends to cluster more than liquidity. Furthermore, the subdivision of the sample period into the dot-com and the housing bubble has highlighted that both spot and option implied liquidities are more persistent during liquidity-driven crisis periods. Our findings also indicate that spot liquidity tends to dry up, on average, during distress periods, being liquidity-driven or not, whereas a global drying-up of liquidity could not be detected during turmoil periods in the option market. Indeed, although option liquidity dried up, on average, during the dot-com crash whatever the option type, such a general trend was not observed for the global financial crisis during which put and call options became generally more and less liquid, respectively. Finally, the results confirm the evidence that mean-reverting processes could be considered to model

liquidity in both the underlying and derivative markets in a similar way as volatility is modeled under stochastic volatility models, paving the way for stochastic liquidity-volatility models. Such models could be used, in future research, to forecast liquidity and therefore to incorporate market liquidity risk into popular risk management tools, which will eventually facilitate the development of a more accurate and effective management of market liquidity risk.

A Influence of the density extraction method on the LIX index

In this section, we compare the LIX index (as well as its 3 and 6 months variants) obtained by using the spline methodology and the SVI parameterization. For that purpose, we perform a linear regression of the LIX series, where the spline LIX is the response variable and the SVI LIX the explanatory variable. The predicted regression line differs only slightly from the identity-line, as shown in Figure 13. Besides, the determination coefficient is almost equal to one: $R^2 = 0.9636$, $R^2 = 0.9930$ and $R^2 = 0.9954$ for a time horizon of 1, 3 and 6 months, respectively. The fact that the fraction of variability in the spline implied liquidity that can be explained by the variability in the SVI implied liquidity increases with the time horizon is in line with the fact that the option spread is wider for shorter term options (see, for instance, George & Longstaff (1993) or Wei & Zheng (2010)). Indeed, larger option spreads increase the flexibility in the fit of the implied volatility, and hence in the shape of the arbitrage-free density.

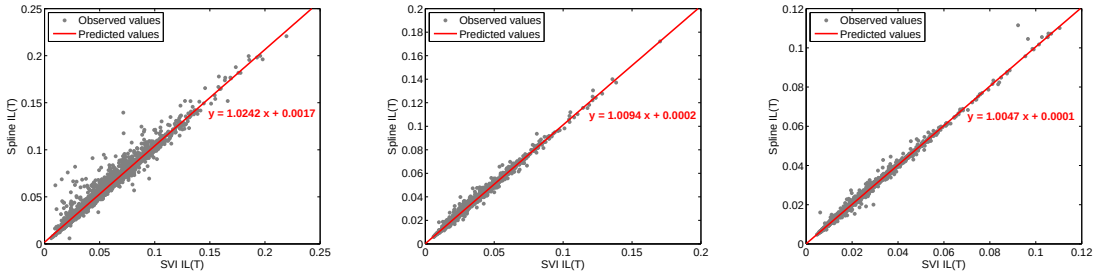

Figure 13: Linear regression of the model-free implied liquidity series for three time horizons: 1 month (left), 3 months (center) and 6 months (right). The response variable is the spline implied liquidity and the explanatory variable the SVI implied liquidity.

References

Aït-Sahalia, Y. (1999). Transition densities for interest rate and other diffusions. *Journal of Finance*, *54*, 1361–1395.

Aitken, M., & Comerton-Forde, C. (2003). How should liquidity be measured? *Pacific-Basin Finance Journal*, *11*, 45–59.

Albrecher, H., Guillaume, F., & Schoutens, W. (2013). Implied liquidity: Model sensitivity. *Journal of Empirical Finance*, *23*, 48–67.

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, *5*, 31–56.

Authers, J., & Mackenzie, M. (2010). Techs reflect on decade since dotcom boom. *The Financial Times*.

Bangia, A., Diebold, F., Schuermann, T., & Stroughair, J. (2002). Modeling liquidity risk: With implications for traditional market risk measurement and management. In S. Figlewski, & R. Levich (Eds.) *Risk Management: The State of the Art*, (pp. 3–13). Kluwer Academic Publishers.

Barrell, R., & Davis, E. (2008). The evolution of the financial market crisis in 2007-8. *National Institute Economic Review*, *206*, 5–14.

BCBS (2009). International framework for liquidity risk measurement, standards and monitoring. Tech. rep., BCBS Consultative Document.

Bervas, A. (2006). Market liquidity and its incorporation into risk management. Tech. Rep. 8, Financial Stability Review, Banque de France.

Breeden, D., & Litzenberger, R. (1978). Prices of state contingent claims implicit in option prices. *Journal of Business*, *51*, 621–651.

Brunnermeier, M. (2009). Deciphering the liquidity and credit crunch 2007-2008. *Journal of Economic Perspectives*, *23*, 77–100.

Cao, M., & Wei, J. (2010). Option market liquidity: Commonality and other characteristics. *Journal of Financial Markets*, *13*, 20–48.

CBOE (2003). VIX: CBOE volatility index. Tech. rep., Chicago.

Cherny, A., & Madan, D. (2009). New measures for performance evaluation. *The Review of Financial Studies*, *22*, 2571–2606.

Cherny, A., & Madan, D. (2010). Markets as a counterparty: An introduction to conic finance. *International Journal of Theoretical and Applied Finance*, *13*, 1149–1177.

Cho, Y., & Engle, R. (1999). Modeling the impacts of market activity on bid-ask spreads in the option market. Tech. rep., National Bureau of Economic Research.

Chong, B.-S., Ding, D., & Tan, K.-H. (2003). Maturity effect on bid-ask spreads of OTC currency options. *Review of Quantitative Finance and Accounting*, *21*, 5–15.

Chordia, T., Roll, R., & Subrahmanyam, A. (2001). Market liquidity and trading activity. *The Journal of Finance*, *56*, 501–530.

Chordia, T., Roll, R., & Subrahmanyam, A. (2002). Order imbalance, liquidity, and market returns. *Journal of Financial Economics*, *65*, 111–130.

Chordia, T., Roll, R., & Subrahmanyam, A. (2011). Recent trends in trading activity and market quality. *Journal of Financial Economics*, *101*, 243–263.

Chordia, T., Sarkar, A., & Subrahmanyam, A. (2005). An empirical analysis of stock and bond market liquidity. *The Review of Financial Studies*, *18*, 85–129.

Chou, R., Chung, S.-L., Hsiao, Y.-J., & Wang, Y.-H. (2011). The impacts of liquidity risk on option prices. *The Journal of Futures Markets*, *31*, 1116–1141.

Corcuera, J., Guillaume, F., Leoni, P., & Schoutens, W. (2009). Implied lévy volatility. *Quantitative Finance*, *9*, 383–393.

Corcuera, J., Guillaume, F., Madan, D., & Schoutens, W. (2012). Implied liquidity: Towards liquidity modeling and liquidity trading. *International Journal of Portfolio Analysis and Management*, *1*, 80–91.

Cox, J., Ingersoll, J., & Ross, S. (1985). A theory of the term structure of interest rates. *Econometrica*, *53*, 385–407.

Figlewski, S. (2010). Estimating the implied risk neutral density for the U.S. market portfolio. In T. Bollerslev, J. Russell, & M. Watson (Eds.) *Volatility and Time Series Econometrics: Essays in Honor of Robert Engle*, (pp. 323–353). Oxford University Press.

Gatheral, J. (2006). *The Volatility surface - A practiotionner's guide.*. New York: Wiley.

George, T., & Longstaff, F. (1993). Bid-ask spreads and trading activity in the S&P 100 index options market. *Journal of Financial and Quantitative Analysis*, *28*, 381–397.

Goyenko, R., & Ukhov, A. (2009). Stock and bond market liquidity: a long-run empirical analysis. *Journal of Financial and Quantitative Analysis*, *44*, 189–212.

Shimko, D. (1993). The bounds of probability. *Risk*, *6*, 33–37.

Wei, J., & Zheng, J. (2010). Trading activity and bid-ask spreads of individual equity options. *Journal of Banking and Finance*, *34*, 2897–2916.

Table 7: Maximum likelihood estimates of mean-reverting model parameters for call option implied liquidity time series.

		Vasicek				CIR			
	period	$\hat{\kappa}$	$\hat{\eta}$	$\hat{\zeta}$	$L(\hat{\boldsymbol{\theta}})$	$\hat{\kappa}$	$\hat{\eta}$	$\hat{\zeta}$	$L(\hat{\boldsymbol{\theta}})$
IC (1m)	1998-2009	176.489	0.01105	0.08847	59.194	186.377	0.01105	0.85573	62.676
	dot-com boom	233.216	0.00798	0.05495	103.586	209.784	0.00798	0.60696	104.961
	dot-com crash	376.200	0.00901	0.09256	73.588	355.360	0.00901	0.96283	76.413
	pre-crisis	325.771	0.01285	0.11530	55.709	337.361	0.01285	1.13960	53.024
	global crisis	133.707	0.01257	0.09213	53.111	118.903	0.01258	0.67881	66.105
IC (3m)	1998-2009	177.217	0.00730	0.06122	85.635	142.112	0.00730	0.59594	102.826
	dot-com boom	150.298	0.00534	0.02909	172.763	154.929	0.00534	0.43982	163.126
	dot-com crash	433.536	0.00690	0.11965	60.528	320.958	0.00690	0.87513	93.184
	pre-crisis	185.837	0.00787	0.04848	109.569	188.242	0.00787	0.63688	98.405
	global crisis	83.266	0.00851	0.05062	88.714	64.946	0.00853	0.45358	110.000
IC (6m)	1998-2009	104.972	0.00572	0.03306	141.040	126.105	0.00571	0.47148	142.285
	dot-com boom	171.547	0.00429	0.02031	255.897	148.894	0.00429	0.35052	224.270
	dot-com crash	305.942	0.00531	0.04272	146.730	294.542	0.00531	0.62164	142.017
	pre-crisis	157.782	0.00590	0.03268	155.627	173.344	0.00590	0.51823	136.230
	global crisis	75.160	0.00733	0.04249	104.155	65.198	0.00734	0.42039	128.438
AC (1m)	1998-2009	302.092	0.01280	0.16813	37.102	284.304	0.01280	1.41058	41.974
	dot-com boom	206.961	0.00962	0.09087	60.322	179.987	0.00961	0.99970	58.704
	dot-com crash	489.562	0.01154	0.15230	50.233	440.847	0.01154	1.46453	49.866
	pre-crisis	438.013	0.01371	0.19713	36.905	452.533	0.01371	1.84068	37.729
	global crisis	305.849	0.01561	0.23207	27.008	248.882	0.01562	1.26667	39.584
AC (3m)	1998-2009	132.814	0.00928	0.06469	75.533	135.441	0.00928	0.69713	77.605
	dot-com boom	159.310	0.00667	0.04657	109.478	164.861	0.00667	0.58150	113.671
	dot-com crash	254.877	0.00879	0.07290	80.454	222.800	0.00879	0.79263	79.387
	pre-crisis	183.582	0.01045	0.07455	71.013	183.479	0.01045	0.81777	66.827
	global crisis	84.470	0.00988	0.06154	73.123	73.607	0.00988	0.51594	91.284
AC (6m)	1998-2009	110.763	0.00698	0.04275	110.160	126.076	0.00698	0.54890	111.132
	dot-com boom	167.337	0.00507	0.02629	196.412	157.385	0.00507	0.42040	175.403
	dot-com crash	313.334	0.00646	0.05350	118.237	311.832	0.00646	0.71668	114.310
	pre-crisis	145.922	0.00746	0.04535	110.050	155.046	0.00746	0.61587	100.078
	global crisis	79.867	0.00856	0.04970	89.810	67.345	0.00858	0.44644	111.778
OC (1m)	1998-2009	274.054	0.01818	0.97823	6.151	344.401	0.01818	3.44144	21.680
	dot-com boom	228.367	0.00799	0.13268	42.608	201.525	0.00799	1.30485	56.460
	dot-com crash	322.995	0.01083	0.17153	37.320	285.565	0.01083	1.53788	43.678
	pre-crisis	291.718	0.02597	1.51180	4.072	373.892	0.02597	5.09654	15.676
	global crisis	276.183	0.01951	0.27778	21.721	230.901	0.01952	1.55627	28.681
OC (3m)	1998-2009	135.777	0.01038	0.09068	54.142	121.601	0.01038	0.88380	58.633
	dot-com boom	136.217	0.00775	0.06246	78.668	130.363	0.00775	0.69788	85.600
	dot-com crash	228.567	0.01132	0.09445	59.875	187.083	0.01132	0.90754	58.462
	pre-crisis	152.242	0.01021	0.10410	48.435	140.130	0.01021	1.05167	52.729
	global crisis	109.045	0.01246	0.08796	53.386	97.463	0.01247	0.68377	64.233
OC (6m)	1998-2009	124.812	0.00873	0.06007	80.275	126.525	0.00873	0.68198	80.941
	dot-com boom	147.251	0.00649	0.03696	135.315	161.739	0.00649	0.55812	119.046
	dot-com crash	344.181	0.00851	0.07516	87.358	340.918	0.00851	0.86522	85.525
	pre-crisis	126.071	0.00931	0.06408	75.409	115.400	0.00931	0.74745	71.038
	global crisis	105.343	0.01001	0.06604	70.651	100.520	0.01001	0.56666	86.046

Table 8: Maximum likelihood estimates of mean-reverting model parameters for put option implied liquidity time series.

		Vasicek				CIR			
	period	$\hat{\kappa}$	$\hat{\eta}$	$\hat{\zeta}$	$L(\hat{\boldsymbol{\theta}})$	$\hat{\kappa}$	$\hat{\eta}$	$\hat{\zeta}$	$L(\hat{\boldsymbol{\theta}})$
IP (1m)	1998-2009	187.177	0.01087	0.08127	65.499	200.850	0.01087	0.81626	67.120
	dot-com boom	273.602	0.00886	0.06714	89.567	289.879	0.00886	0.76423	88.611
	dot-com crash	386.294	0.00871	0.07815	88.135	383.972	0.00871	0.90561	84.715
	pre-crisis	331.626	0.01284	0.10632	60.840	347.877	0.01284	1.05105	57.710
	global crisis	141.153	0.01100	0.07848	63.109	129.417	0.01100	0.62923	77.162
IP (3m)	1998-2009	135.708	0.00692	0.04285	114.558	138.094	0.00692	0.54692	113.684
	dot-com boom	128.793	0.00534	0.02637	184.086	149.958	0.00534	0.41354	171.651
	dot-com crash	306.541	0.00569	0.05140	122.036	269.621	0.00569	0.65499	126.709
	pre-crisis	244.605	0.00808	0.04997	115.723	246.737	0.00808	0.66942	99.859
	global crisis	107.775	0.00734	0.04695	99.794	86.736	0.00734	0.45594	121.659
IP (6m)	1998-2009	114.662	0.00519	0.03005	157.737	126.096	0.00519	0.44689	157.295
	dot-com boom	144.430	0.00388	0.01676	297.174	148.701	0.00388	0.32054	257.562
	dot-com crash	290.635	0.00411	0.03111	197.642	283.187	0.00411	0.52788	187.022
	pre-crisis	266.888	0.00598	0.03647	163.426	277.724	0.00598	0.57417	140.629
	global crisis	92.593	0.00606	0.03840	118.859	72.434	0.00607	0.39716	150.654
AP (1m)	1998-2009	352.109	0.01192	0.12198	54.330	354.588	0.01192	1.23756	52.300
	dot-com boom	347.008	0.01060	0.09141	72.068	308.401	0.01060	1.01618	63.221
	dot-com crash	443.180	0.01101	0.12020	60.841	442.547	0.01101	1.24217	59.190
	pre-crisis	707.955	0.01256	0.17542	51.998	630.912	0.01256	1.76770	46.616
	global crisis	175.358	0.01296	0.10367	50.427	194.299	0.01296	0.88355	56.461
AP (3m)	1998-2009	152.273	0.00847	0.05899	85.474	149.222	0.00847	0.67972	84.899
	dot-com boom	150.770	0.00645	0.03821	131.653	159.623	0.00645	0.55847	119.021
	dot-com crash	237.137	0.00722	0.06289	91.000	215.664	0.00722	0.74525	92.752
	pre-crisis	290.158	0.01006	0.07531	81.587	275.253	0.01006	0.86154	72.832
	global crisis	109.410	0.00834	0.05521	85.099	98.388	0.00834	0.50771	104.593
AP (6m)	1998-2009	139.128	0.00611	0.03934	125.478	149.327	0.00611	0.54527	123.916
	dot-com boom	272.130	0.00451	0.03001	199.993	258.154	0.00450	0.49790	182.193
	dot-com crash	258.476	0.00486	0.03668	160.658	254.732	0.00486	0.58261	150.837
	pre-crisis	323.524	0.00712	0.05033	127.278	331.808	0.00712	0.71502	111.162
	global crisis	112.198	0.00695	0.04639	101.768	97.230	0.00695	0.47131	123.340
OP (1m)	1998-2009	305.777	0.01388	0.17249	36.332	305.262	0.01388	1.40472	41.030
	dot-com boom	308.541	0.01112	0.11036	56.981	271.097	0.01112	1.05301	57.302
	dot-com crash	413.614	0.01166	0.13052	54.351	394.871	0.01166	1.29617	52.860
	pre-crisis	479.620	0.01507	0.23924	31.679	478.720	0.01507	1.92506	35.099
	global crisis	237.911	0.01677	0.17190	33.332	217.495	0.01677	1.08836	41.758
OP (3m)	1998-2009	184.843	0.01027	0.08176	64.875	187.313	0.01027	0.87307	64.146
	dot-com boom	154.029	0.00804	0.05279	95.783	167.086	0.00804	0.69367	87.796
	dot-com crash	282.734	0.00900	0.08387	72.559	248.847	0.00900	0.86847	74.523
	pre-crisis	327.907	0.01197	0.10828	59.474	307.743	0.01197	1.15375	52.812
	global crisis	112.420	0.01009	0.06819	69.256	112.075	0.01009	0.59404	83.177
OP (6m)	1998-2009	133.797	0.00750	0.05092	96.099	142.801	0.00750	0.63613	95.702
	dot-com boom	183.745	0.00540	0.03257	162.558	195.777	0.00540	0.52244	145.983
	dot-com crash	258.834	0.00578	0.04293	137.331	246.872	0.00578	0.63225	126.213
	pre-crisis	287.571	0.00922	0.06967	87.900	282.440	0.00922	0.84279	78.587
	global crisis	120.709	0.00771	0.05392	88.822	115.677	0.00771	0.54093	105.400