

Simulation of Implied Volatility Surfaces via Tangent Lévy Models

R. Carmona, Y. Ma and S. Nadtochiy*

Current version: Sep 30, 2016

Abstract

In this paper, we implement and test a market-based model for European-type options, based on the tangent Lévy models proposed in [4] and [3]. As a result, we obtain a method for generating Monte Carlo samples of future paths of implied volatility surfaces. These paths and the surfaces themselves are free of arbitrage, and are constructed in a way that is consistent with the past and present values of implied volatility. We use market data to estimate the parameters of this model and conduct an empirical study, to compare the performance of the chosen market-based model with the classical SABR model and with the method based on direct simulation of implied volatility, described in [7]. We choose the problem of minimal-variance portfolio choice as the main measure of model performance and compare the three models. Our study demonstrates that the tangent Lévy model does a better job at finding a portfolio with the smallest variance than the SABR model. In addition, the prediction of return variance, provided by the tangent Lévy model, is more reliable and the portfolio weights are more stable. We also find that the performance of the direct simulation method, at the portfolio choice problem, is not much worse than that of the tangent Lévy model. However, the direct simulation method of [7] is not arbitrage-free. We illustrate this shortcoming by comparing the direct simulation method and the tangent Lévy model at a different problem – estimation of Value at Risk of an options’ portfolio. To the best of our knowledge, this paper is the first example of empirical analysis, based on real market data, which provides a convincing evidence of the superior performance of market-based models for European options, as compared to the classical spot models.

1 Introduction

The existence of liquid markets for equity and volatility derivatives, as well as a well-developed over-the-counter market for exotic derivatives, generates a need for a modeling framework that is consistent across time and across financial instruments. Within this framework, once a model is chosen so that it matches both the present prices of liquid instruments and their past dynamics, it is expected to produce more realistic results for the problems of pricing and hedging of exotic instruments. In addition, such models can be used to quantify the risk embedded in portfolios of derivative contracts. Needless to say, evaluating and managing the risk of such portfolios is crucial for proper functioning of the financial markets: recall, for example, that VIX index, itself, is a portfolio of European options written on S&P 500.

In this paper we investigate an arbitrage-free modeling framework for multiple European-type options written on the same underlying, which is consistent across time and products. In particular, this framework allows to resolve one of the nagging challenges of *quant* groups supporting equity trading: i.e. how to generate realistic Monte Carlo scenarios of implied volatility surfaces which are consistent with present and historical observations? As mentioned above, such models can be used to address the problems of pricing, hedging and risk management. Herein, we implement several such models using real market data and *conduct a numerical experiment which demonstrates clearly the advantages of this modeling approach*.

The attempts to model the dynamics of implied volatility surface directly can be dated back as early as the “sticky smile model” and the “sticky delta model” (also known as “floating smile model”) (see Section 6.4 of [23] for the definitions). As an improvement of the two models, Cont et al. later proposed a multi-factor model of implied volatility surface in [7] and [8], where they applied a Karhunen-Loève decomposition on the daily variations of implied volatilities. It turns out that the first three eigenvectors could explain most of the daily variance, and a mean-reverting factor model based on the three eigenvectors is then constructed for future implied

*We thank the anonymous referees and the Associate Editor for the constructive comments that helped us improve the paper significantly.

volatility surface. The major issue with these early attempts is that the proposed models for the dynamics of implied volatility are either too restrictive, not allowing to match the historical evolution of implied volatility, or too loose, so that they may contain *arbitrage opportunities*. While the importance of the first issue for any time-series analysis is very clear, the second one deserves a separate discussion. Indeed, what do we mean by arbitrage opportunities in a model for implied volatility and why do we need to avoid it? There are two types of arbitrage opportunities we refer to: *static* and *dynamic*. A given implied volatility surface contains static arbitrage if it is impossible to obtain such a surface in any arbitrage-free model for the underlying. The fact that not every surface can be an arbitrage-free implied volatility simply follows from the well-known static no-arbitrage restrictions on the shape of a call price surface: e.g. monotonicity and convexity in strikes, etc (cf. [10] and [11]). Notice that a violation of any of these conditions leads to an obvious arbitrage opportunity which is very easy to implement, hence, it is natural to assume that every implied volatility surface is free of static arbitrage. This, in turn, implies that any realistic simulation algorithm for future implied volatility surfaces has to produce surfaces that are arbitrage-free: otherwise, the algorithm generates outcomes that are simply impossible (see Subsection 3.4.2). The static no-arbitrage conditions are rather difficult to state explicitly, in terms of the implied volatility surface itself (without mapping it to a call or put price surface first). Nevertheless, it is not hard to deduce from the existing necessary conditions (cf. [20]) that the set of arbitrage-free implied volatility surfaces forms a “thin” set in the space of all (regular enough) functions of two variables. Hence, it is a non-trivial task to construct a modeling framework that excludes static arbitrage in the implied volatility surface. The dynamic arbitrage adds to this problem, and it refers to a restriction on the evolution (i.e. the time increments) of implied volatility surface, rather than its values at a fixed moment in time. This restriction follows from the same arbitrage considerations for option prices. However, its associated arbitrage strategies are not as straightforward as in the case of static arbitrage. In addition, the simulated implied volatility surfaces that contain only dynamic arbitrage are, typically, very close to the ones that are arbitrage-free, when the time horizon is small (it is related to the fact that dynamic arbitrage only changes the drift term of the implied volatility, which is much smaller than the diffusion term, for small times). This is why, eliminating the dynamic arbitrage in a model for implied volatility surface is often viewed as a “second priority” for risk management. Nevertheless, we believe that a good model should exclude both types of arbitrage, in order to produce realistic dynamics of implied volatility surface (for risk management) and eliminate the possible arbitrage opportunities (for pricing).

We have already mentioned that it is not a trivial task to construct a model of implied volatility that excludes arbitrage opportunities. In fact, when trying to model the surface directly, the first challenge that one faces is: how to describe the space of possible implied volatility surfaces? Note that, as discussed above, the existing characterizations of arbitrage-free implied volatility surfaces are rather implicit. In addition, if the resulting space is not an open subset of any linear space (which it is not), what kind of mathematical tools can be used to describe evolution in space? Recall, for example, that all statistical models of time-series are defined on linear spaces (or those that can be easily mapped in to a linear space). Hence, it appears natural to map the space of possible implied volatility surfaces to an open set in a linear space, and then proceed with the construction of arbitrage free models. Such mapping became known as a *code-book* mapping, and it turns out that it can be constructed by means of the so-called *tangent models* (cf. [2], [4], [3]). The concept of a tangent model is very close to the method of *calibrating* a model for underlying to the target derivatives’ prices (in the present case, European options calls). Consider a family of arbitrage-free models for the underlying, $\mathcal{M}(\theta)$, parameterized by θ , taking values in a “convenient” set Θ (an open set of a linear space). For any given surface of option prices (or, equivalently, any given implied volatility surface), we can try to calibrate a model from this family to a given surface of option prices (or, equivalently, to a given implied volatility surface). In other words, we attempt to find $\theta \in \Theta$ such that: $C^\theta(T, K) = C(T, K)$, for all given maturities T and strikes K , where $C(T, K)$ is the given call price, and $C^\theta(T, K)$ is the call price produced by the model $\mathcal{M}(\theta)$. If the above calibration problem has a unique solution, we obtain a one-to-one correspondence between the call price surfaces and the models in a chosen family: $\theta \leftrightarrow C^\theta$. For every call price surface $C = C^\theta$, the associated (calibrated) model $\mathcal{M}(\theta)$ is called a tangent model.¹ Notice that C^θ is always arbitrage-free, hence, we obtain the desired code-book mapping $C = C^\theta \mapsto \theta$. Now, the problem of static arbitrage has been resolved, and one simply needs to prescribe the distribution of a stochastic process (θ_t) , taking values in a convenient set Θ , in order to obtain a model for the dynamics of call prices ($C_t = C^{\theta_t}$), and, in turn, the dynamics of implied volatility surface. Finally, one needs to characterize all possible dynamics of (θ_t) that produce no dynamic arbitrage in the associated call prices (C^{θ_t}). An interested reader is referred to [3], for a more detailed description of this general algorithm, and, for example, to [2], [4], [15], [30], [19], [24], for the analysis of specific choices of the families of models $\{\mathcal{M}(\theta)\}$.

¹It is important to remember that any such model serves *only as a static description* of option prices, and it does not describe their dynamics!

The idea of modeling prices of derivative contracts directly dates back to the work of Heath, Jarrow and Morton [14], who analyzed the dynamic of bond prices along with the short interest rate. Such models have become known as the *market-based models* (or simply *market models*), as opposed to the classical *spot models*, since the former are designed to capture the evolution of the entire market, including the liquid derivatives. This approach has been extended to more general mathematical settings, as well as to other derivatives' markets. The list of relevant works includes [11], [25], [26], [28], [27], [12], in addition to those mentioned in the previous paragraph. Even though the notions of code-book and tangent models never appear in these papers, almost all of them follow the algorithm outlined in the previous paragraph (and described in more detail in [3]), in order to construct a market-based model.

Even though various code-books for implied volatility surface (or, equivalently, for call price surface) have been proposed and the corresponding arbitrage-free dynamics have been characterized, it was not until very recently that some of these models were implemented numerically. As is shown in the rest of the paper, the lack of such results is not a surprise given the complexity of the models. So far, the numerical implementations are mostly based on *tangent Lévy models* proposed in [4] and [3]: as the name suggests, this corresponds to a code-book which is constructed using non-homogeneous Lévy (or, additive) models as the tangent models. Karlsson [16] implements a class of tangent Lévy models with absolutely continuous Lévy densities and no continuous martingale component. Zhao [30] and Leclercq [19], on contrary, implemented the tangent Lévy models whose Lévy measure is purely atomic in the space variable. As opposed to [30], the work of Leclercq [19] allows for tangent models with continuous martingale component and includes options with multiple maturities, but it does require that the Lévy density possess certain symmetry, which may limit the ability of the model to capture the skew of the implied smile. All of the works [16], [30], [19] estimate the parameters of the model from real market data. In addition, [19] conducts a numerical experiment comparing the performance of a market-based model to a classical spot model. The actual results of this experiment, however, do not provide a convincing evidence in favor of the market-based approach. We believe that the latter is simply due the choice of experiment and not to the deficiency of the theory, and we intend to demonstrate it in the present work.

The purpose of this paper is to propose an implementation method for a class of tangent Lévy models and to test its performance using market data. This method provides an algorithm for simulating future arbitrage-free implied volatility surfaces, which are consistent with both present and past observations. Our method is similar to the one used in [16], but with a different “dynamic fitting” part. However, the most important original contribution of this paper is the *numerical experiment which uses real market data to demonstrate clearly the advantages of market-based models for implied volatility (or, option prices)*, as compared to the classical spot models. In addition, we conduct a smaller-scale experiment in which we illustrate the advantages of using an arbitrage-free model for simulating future options' prices. To the best of our knowledge, this is the first convincing empirical analysis that justifies the use of market-based approach for modeling option prices (or, equivalently, for modeling the implied volatility surface).

The rest of the paper is organized as follows. Section 2 starts by reviewing the work on tangent Lévy models with continuous Lévy density, developed in [4]. We, then, proceed to describe the parametric estimation of the parameters of this model, which is partially based on the double exponential jump process. In Section 3, the estimated model is tested against a popular stochastic volatility model and a model based on direct simulation of implied volatility – in a portfolio optimization problem and for estimating Value at Risk. Section 4 concludes the paper by highlighting the main contributions and the future work. Appendix A contains technical proofs and derivations, Appendix B contains all tables and graphs.

2 Double exponential tangent Lévy models

2.1 Model setup and consistency conditions

In this subsection, we review and update the results of [4], which serve as a foundation for the analysis in subsequent sections. Herein, we assume that the interest and dividend rates for the underlying asset are zero. In the implementation that follows, we discount the market data accordingly, to comply with this assumption. As in [3], we denote by $(S_t)_{t \geq 0}$ an adapted stochastic process defined on a stochastic basis $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{Q})$, with the filtration \mathbb{F} satisfying the usual hypotheses. The process S represents the dynamics of the underlying under the pricing measure \mathbb{Q} . We assume that it is of the form

$$S_t = S_0 + \int_0^t \int_{\mathbb{R}} S_{u-} (e^x - 1) [M(dx, du) - K_u(x) dx du]. \quad (2.1)$$

Here, M is a general integer-valued random measure (not necessarily a Poisson measure!), whose compensator is $K_{u,\omega}(x)dxdu$, where $(K_t)_{t \geq 0}$ is a predictable stochastic process taking values in the function space \mathcal{B}_0 , defined in (5.1).

For any fixed time $t \geq 0$ and a given value of S_t , a stochastic process $(\tilde{S}_T)_{T \geq t}$ is said to be *tangent* to the true model $(S_t)_{t \geq 0}$ if the time- t prices of all European call options written on S can be obtained by pretending the future risk-neutral evolution of the index value is instead given by $(\tilde{S}_T)_{T \geq t}$ from t on. Throughout this section, for any fixed $t \geq 0$, we assume that the tangent processes \tilde{S} is in the form

$$\tilde{S}_T = S_t + \int_t^T \int_{\mathbb{R}} \tilde{S}_{u-} (e^x - 1) [N_t(dx, du) - \kappa_t(u, x) dx du], \quad (2.2)$$

for $T \in [0, \bar{T}]$, where \bar{T} is a fixed terminal time horizon and $N_t(dx, du)$ is a Poisson random measure associated with the jumps of $\log \tilde{S}$, whose compensator is given by a deterministic measure $\kappa_t(u, x) dx du$. Notice that the law of \tilde{S} is uniquely determined by (S_t, κ_t) . Let $C_t^{S_t, \kappa_t}(T, x)$ denote the option prices generated by $(\tilde{S}_u)_{u \geq t}$, i.e

$$C_t^{S_t, \kappa_t}(T, x) := \mathbb{E} \left[(\tilde{S}_T - e^x)^+ | \tilde{S}_t = S_t \right], \quad \forall T \geq t, x \in \mathbb{R}. \quad (2.3)$$

The concept of a *tangent model*, then, requires that, for each fixed $t \in [0, \bar{T})$,

$$C_t^{S_t, \kappa_t}(T, x) = \mathbb{E} [(S_T - e^x)^+ | \mathcal{F}_t], \quad \forall T \geq t, \forall x \in \mathbb{R}. \quad (2.4)$$

Thus, at each time t , we obtain the *code-book* for call prices, given by (S_t, κ_t) . Of course, the value of the code-book may be different at a different time t . Hence, we consider the dynamic tangent Lévy models characterized by a pair of stochastic processes $(S_t, \kappa_t)_{t \in [0, \bar{T}]}$ that satisfies (2.4). Here, S is a positive martingale with dynamics given by (2.1); κ is progressively measurable positive stochastic process taking values in \mathcal{B} (cf. (5.2)). The dynamics of S_t and κ_t are given by

$$\begin{cases} S_t = S_0 + \int_0^t \int_{\mathbb{R}} S_{u-} (e^x - 1) [M(dx, du) - K_u(x) dx du], \\ \kappa_t(T, x) = \kappa_0(T, x) + \int_0^t \alpha_u(T, x) du + \sum_{n=1}^m \int_0^t \beta_u^n(T, x) dB_u^n, \end{cases} \quad (2.5)$$

where $B = (B^1, \dots, B^m)$ is a standard Brownian motion, $(\alpha_t)_{t \in [0, \bar{T}]}$ is a progressively measurable integrable stochastic process with values in \mathcal{B} , and, for each $n \in \{1, \dots, m\}$, $(\beta_t^n)_{t \in [0, \bar{T}]}$ is a progressively measurable square integrable stochastic process taking values in \mathcal{H} (cf. (5.4)).

Notice that (2.5) defines the dynamics of the code-book $(S_t, \kappa_t)_{t \in [0, \bar{T}]}$, but it does not ensure that it does, indeed, produce tangent models at each time t : in other words, there is no guarantee that (2.4) holds. Thus, additional ‘‘consistency’’ conditions have to be enforced to obtain models which are, indeed, tangent to the true underlying process. As shown in [4], this consistency is, in fact, equivalent to the fact that call prices generated by these tangent models are free of dynamic arbitrage. In order to present the main consistency result, we state the following regularity assumptions on β .²

Assumption 1. For each $n \leq m$, almost surely, for almost every $t \in [0, \bar{T}]$, we have:

RA1 $\sup_{T \in [t, \bar{T}]} \int_{-1}^1 |\beta_t^n(T, x)| dx < \infty$

RA2 For every $T \in [t, \bar{T}]$, the function $\beta_t^n(T, \cdot)$ is absolutely continuous on $\mathbb{R} \setminus \{0\}$.

RA3 For any $T \in [t, \bar{T}]$, $\int_{\mathbb{R}} (e^x - 1) \beta_t^n(T, x) dx = 0$.

Finally, we introduce some extra notation,

$$\bar{\beta}_t^n(T, x) := \int_{t \wedge T}^T \beta_t^n(u, x) du, \quad (2.6)$$

and formulate the consistency result, which is a simple corollary of Theorem 12 in [4].

²The conditions RA1–RA2 are of technical nature – they can be viewed as the regularity assumptions on β . The last condition, RA3, is a symmetry assumption on β : if this condition does not hold, one would have to introduce a continuous martingale component in the dynamics of S , in order to obtain a consistent model.

Theorem 1. (Carmona-Nadtochiy 2012) Assume that $(S_t)_{t \in [0, \bar{T}]}$ is a true martingale,³ β satisfies the above regularity assumptions RA1-RA3, and $\kappa_t(T, x) \geq 0$, almost surely for all $t \in [0, \bar{T}]$ and almost all $(T, x) \in [t, \bar{T}] \times \mathbb{R}$. Then the processes $(S_t, \kappa_t)_{t \in [0, \bar{T}]}$ satisfying (2.5) are consistent, in the sense that (2.4) holds, if and only if the following conditions hold almost surely for almost every $x \in \mathbb{R}$ and $t \in [0, \bar{T})$, and all $T \in (t, \bar{T}]$:

1. **Drift restriction:**

$$\alpha_t(T, x) = - \sum_{n=1}^m \left\{ \int_{\mathbb{R}} \bar{\beta}_t^n(T, y) \beta_t^n(T, x - y) dy - \bar{\beta}_t^n(T, x) \cdot \int_{\mathbb{R}} \beta_t^n(T, z) dz - \beta_t^n(T, x) \cdot \int_{\mathbb{R}} \bar{\beta}_t^n(T, z) dz \right\}. \quad (2.7)$$

2. **Compensator specification:** $K_t(x) = \kappa_t(t, x)$.

Theorem 1, along with equations (2.5) provide a general method for constructing a market-based model for call prices (i.e. an arbitrage-free dynamic model for implied volatility surface). Indeed, choosing $(\beta_t^1, \dots, \beta_t^m)_{t \in [0, \bar{T}]}$, we use the drift restriction in Theorem 1 and the second equation in (2.5) to generate the paths of $(\kappa_t)_{t \in [0, \bar{T}]}$. Finally, to generate the paths of $(S_t)_{t \in [0, \bar{T}]}$, one can use the compensator specification in Theorem 1 and the first equation in (2.5), after representing the random measure M through its compensator K and a Poisson random measure N (as shown in [4]). However, in the present paper we avoid simulating $(S_t)_{t \in [0, \bar{T}]}$ at all, by simply noticing that

$$\begin{aligned} \frac{1}{S_t} C_t^{S_t, \kappa_t}(T, x + \log S_t) &= \mathbb{E} \left[(\tilde{S}_T / S_t - e^x)^+ | \tilde{S}_t = S_t \right] = \mathbb{E} \left[(\tilde{S}_T - e^x)^+ | \tilde{S}_t = 1 \right] = C_t^{1, \kappa_t}(T, x), \\ \frac{1}{S_t} C_t^{S_t, bs}(T, x + \log S_t; \sigma) &= C_t^{1, bs}(T, x; \sigma), \end{aligned}$$

where $C_t^{S_t, bs}(T, x; \sigma)$ is the Black-Scholes price at time t of a call option with maturity T and strike e^x given that the level of underlying is at S_t and the volatility is σ . At any time t , regardless of the value of S_t , if we find the level of σ that makes the right hand sides of the two equations above coincide, then the option prices in the left hand sides have to coincide as well. This means that we can obtain the implied volatility of $C_t^{S_t, \kappa_t}$, in the maturity and log-moneyness variables, by computing the corresponding implied volatility of C_t^{1, κ_t} , for which we do not need to generate S_t .

2.2 Implied volatility simulation with tangent Lévy models

We first introduce the general framework of the simulation procedure. Our procedure has two stages, *estimation* and *simulation*. The estimation stage, where the additive density of the tangent process as well as its dynamics are fitted to market data, is performed in two steps:

- *Static fitting.* In static fitting, the additive density κ_t for each day t is obtained by least squares optimization which minimizes the squared difference between model prices and actual market prices. Notice that for any given day t , κ_t is fixed and there is no dynamics involved, which explains the term ‘static’.
- *Dynamic fitting.* In dynamic fitting, we recover the dynamics of the time series (κ_t) . In view of the drift restriction in Theorem 1, this boils down to determining the volatility terms $\{\beta_t^n\}_{n=1}^m$. This is done by applying the Principal Components Analysis to the time series of $(\kappa_t)_t$.

Once the estimation is completed, we generate the future paths of (κ_t) using Euler scheme Monte Carlo applied to the second equation in (2.5). From the simulated additive densities, we compute call prices C_t^{1, κ_t} and, then, implied volatilities by inverting the Black-Scholes formula.

Within the general framework, the simulation stage is generic, but the static part of the estimation stage can be quite different depending on the specific subclass of tangent Lévy densities $\kappa(u, x)$ that we fit to option price at any given time. In this section, we implement the procedure with the Lévy densities arising from the double exponential Lévy models proposed by Kou in [17]. The small number of parameters in double exponential models and the availability of an analytical pricing formula for call options make the resulting family of tangent Lévy models fairly easy to calibrate.

³The conditions on the input parameters which ensure that S is a martingale are presented in [4] – in Remark 5.3 and in the preceding paragraph.

2.3 Market data

We use *SPX* (S&P 500) call option prices provided by OptionMetrics, an option database containing historical prices of options and their underlying instruments. Throughout the paper, we use the option data from two time periods: Jan. 2007 - Aug. 2008 and Jan. 2011 - Dec. 2012. Table 1 gives a quick summary of the two periods. We cut off the first period at August 2008 to reduce the impact of the financial crisis.

On each day of a period, we only keep the options with time to maturity less than one year, whose best closing bid price and best closing offer price are both available, and take the average of the two prices as the option price. To ensure the validity of all prices, the contracts with zero open interest are excluded. As a result, there are roughly 10 to 80 call contracts with valid prices available for each maturity. The log-moneyness (more precisely, the put log-moneyness, defined as $\log(K/S_t)$) of these call options ranges roughly from -0.3 to 0.1, varying for different t and T . Our calibration also requires dividend and interest rate data available on OptionMetrics and the homepage of U.S. Department of Treasury, respectively. This dividend yield is recovered from option prices via put-call parity with the method proposed in [1]. On day t , we denote the dividend yield by q_t , and the risk-free rate between t and T by $r_{t,T}$. To simplify our implementation, we perform a simple transformation on the market data so that we can assume that the interest and dividend rates are both zero from now on:

$$C_t^{mkt}(T, x) = e^{q_t(T-t)} \bar{C}_t^{mkt}(T, \bar{x}), \quad \text{with } x = \bar{x} - (r_{t,T} - q_t)(T - t), \quad (2.8)$$

where $\bar{C}_t^{mkt}(T, \bar{x})$ is the market price of a call option with maturity T and strike $e^{\bar{x}}$. The adjusted call prices $C_t^{mkt}(T, x)$, corresponding to maturity T and strike e^x , are then consistent with the assumption of zero interest and dividend rates (i.e. they do not contain arbitrage under these assumptions). In a similar way, we define the adjusted bid and ask prices, $C_t^{mkt,b}$ and $C_t^{mkt,a}$.

In this section, we will perform the calibration of a tangent Lévy model on the time span from Jan. 3, 2007 to Dec. 31, 2007, denoted by $[t_0, \bar{T}]$. In Section 3, data from both periods will be used to test the performance of the tangent Lévy model.

2.4 Static fitting

Before we proceed with the static fitting, let us first have a quick review of the double exponential model. In such a model, the logarithm of underlying follows a pure jump Lévy process whose jump sizes have a double exponential distribution. More specifically, assuming no diffusion term, the dynamics of the underlying in the double exponential model are given by

$$d\hat{S}_t = \mu \hat{S}_{t-} dt + \hat{S}_{t-} d \left(\sum_{i=1}^{N_t} (\exp(Y_i) - 1) \right), \quad (2.9)$$

where μ is the drift term, N_t is a Poisson process with rate λ , $\{Y_i\}$ is a sequence of i.i.d. random variables with asymmetric double exponential distribution, independent of N_t . Note that the double exponential model is not exactly a tangent model that we will be using – rather, we will use it as a building block to construct our tangent model. This is why we use the notation \hat{S} to denote the dynamics under the double exponential model, whereas the tangent model dynamics are denoted by \tilde{S} . The density of an asymmetric double exponential distribution is given by

$$f_Y(y) = p \cdot \lambda_1 e^{-\lambda_1 y} \mathbf{1}_{y \geq 0} + q \cdot \lambda_2 e^{\lambda_2 y} \mathbf{1}_{y < 0}, \quad (2.10)$$

where $p, q \geq 0$, $p + q = 1$ represent the probabilities of positive and negative jumps, and $\lambda_1 > 1$, $\lambda_2 > 0$ are the parameters of the two exponential distributions. In other words, a double exponential model is a martingale model for the underlying whose logarithm is a pure jump Lévy process, with the Lévy density

$$\eta(x) = \lambda(p \cdot \lambda_1 e^{-\lambda_1 x} \mathbf{1}_{x \geq 0} + q \cdot \lambda_2 e^{\lambda_2 x} \mathbf{1}_{x < 0}). \quad (2.11)$$

One of the advantages of double exponential models is the availability of analytical pricing formulas for European call options, which could greatly simplify the calibration. [17] gives the pricing formula for double exponential models with a diffusion term. A minor modification of the derivation in [17] gives us the pricing formula in absence of the diffusion term, summarized in the lemma below (its proof is omitted, as it is, essentially, a repetition of the derivations in [17]).

Lemma 1. Under the assumptions of zero interest and dividend rates, assume, in addition, that the underlying process S follows a double exponential process with Lévy density given by (2.11), under the risk-neutral probability measure. Then, the price of a European call option with strike K and maturity T is given by

$$C_t^{\lambda, \lambda_1, \lambda_2, p}(T, \log K) = S_t \Psi(-\lambda \zeta, \lambda^*, p^*, \lambda_1^*, \lambda_2^*; \log(K/S_t), T-t) - K \Psi(-\lambda \zeta, \lambda, p, \lambda_1, \lambda_2; \log(K/S_t), T-t), \quad (2.12)$$

where

$$p^* = \frac{p}{1+\zeta} \cdot \frac{\lambda_1}{\lambda_1-1}, \quad \lambda_1^* = \lambda_1 - 1, \quad \lambda_2^* = \lambda_2 + 1, \\ \lambda^* = \lambda(\zeta + 1), \quad \zeta = \frac{p\lambda_1}{\lambda_1-1} + \frac{q\lambda_2}{\lambda_2+1} - 1,$$

and the function Ψ is given by:

$$\begin{aligned} & \Psi(\mu, \lambda, p, \lambda_1, \lambda_2; a, T) \\ &= \pi_0 \mathbf{1}_{a-\mu T \leq 0} + \sum_{n=1}^{\infty} \pi_n \sum_{k=1}^n P_{n,k} \left[\sum_{i=0}^{k-1} \frac{(\lambda_1(a-\mu T))^i}{i!} e^{-\lambda_1(a-\mu T)} \mathbf{1}_{a-\mu T \geq 0} + \mathbf{1}_{a-\mu T < 0} \right] \\ &+ \sum_{n=1}^{\infty} \pi_n \sum_{k=1}^n Q_{n,k} \left(1 - \sum_{i=0}^{k-1} \frac{(-\lambda_2(a-\mu T))^i}{i!} e^{\lambda_2(a-\mu T)} \right) \mathbf{1}_{a-\mu T < 0}, \end{aligned} \quad (2.13)$$

with

$$\pi_n = \frac{e^{-\lambda T} (\lambda T)^n}{n!}$$

and

$$\begin{aligned} P_{n,k} &= \sum_{i=k}^{n-1} \binom{n-k-1}{i-k} \binom{n}{i} \left(\frac{\lambda_1}{\lambda_1 + \lambda_2} \right)^{i-k} \left(\frac{\lambda_2}{\lambda_1 + \lambda_2} \right)^{n-i} p^i q^{n-i}, \quad 1 \leq k \leq n-1, \\ Q_{n,k} &= \sum_{i=k}^{n-1} \binom{n-k-1}{i-k} \binom{n}{i} \left(\frac{\lambda_1}{\lambda_1 + \lambda_2} \right)^{n-i} \left(\frac{\lambda_2}{\lambda_1 + \lambda_2} \right)^{i-k} p^{n-i} q^i, \quad 1 \leq k \leq n-1, \\ P_{n,n} &= p^n, \quad Q_{n,n} = q^n. \end{aligned}$$

For each T_l , with $l = 1, \dots, L$, we would like to find the set of parameters $\{\lambda, \lambda_1, \lambda_2, p\}$ that minimizes the difference between the market and the model prices. For practical reasons, we will work with time values instead of options prices. The market time value and the model time value are calculated as follows

$$\begin{aligned} V_t^{mkt,j}(T_l) &= C_t^{mkt}(T_l, e^{x_j}) - (S_t - e^{x_j})^+, \\ V_t^{\lambda, \lambda_1, \lambda_2, p, j}(T_l) &= C_t^{\lambda, \lambda_1, \lambda_2, p}(T_l, e^{x_j}) - (S_t - e^{x_j})^+. \end{aligned}$$

There are two reasons for working with time values. Firstly, the time values go to zero for very large and very small log-moneyness, which allows us to truncate the x -space with negligible numerical errors. Secondly, time values and option prices are often of different magnitudes, especially for in the money options, with option prices much greater than time values, hence, working with time values is likely to result in smaller numerical errors. For fixed time t and fixed maturity T_l , the optimization problem can be written as

$$\min_{\lambda > 0, \lambda_1 > 0, \lambda_2 > 0, p \in (0,1)} \sum_{j=1}^N \omega_j |V_t^{\lambda, \lambda_1, \lambda_2, p, j}(T_l) - V_t^{mkt,j}(T_l)|^2, \quad (2.14)$$

where $\omega_j = |C_t^{bid}(T_l, e^{x_j}) - C_t^{ask}(T_l, e^{x_j})|^{-2}$ are the weights we put on different options to take into account the difference in liquidity (measured by bid-ask spread). For every fixed maturity T_l , the solution of the above optimization problem, $(\lambda^l, \lambda_1^l, \lambda_2^l, p^l)$, yields the Lévy density $\eta_t(T_l, x)$ via (2.11). Then, we search for a function $\kappa_t(\cdot, \cdot)$, such that

$$\eta_t(T_l, x) = \frac{1}{T_l - t} \int_t^{T_l} \kappa_t(u, x) du, \quad (2.15)$$

for every maturity T_l and all $x \in \mathbb{R}$. The resulting tangent model on day t is defined as a martingale model for the underlying whose logarithm is a pure jump additive (non-homogeneous Lévy) process, with the Lévy density $\kappa_t(\cdot, \cdot)$. It is easy to see that the call prices produced by this model, for every maturity T_l and strike e^{x_j} , coincide with the prices produced by the double exponential model, $C_t^{\lambda_1, \lambda_1, \lambda_2, p^l}(T_l, e^{x_j})$. Thus, for a given t , the problem of static fitting is essentially a series of optimization problems (2.14), over all maturities T_l , along with the fitting problem (2.15).

At the first glance, the optimization in (2.14) seems to have four parameters. However, the following constraints will reduce the number of parameters to two in our calibration:

- To improve the stability of small-jumps intensity over time, we would like the Lévy density $\eta(T_l, x)$ to be continuous in x . The continuity at $x = 0$ requires

$$p \cdot \lambda_1 = (1 - p) \cdot \lambda_2 \Leftrightarrow \lambda_2 = \frac{p}{1 - p} \lambda_1. \quad (2.16)$$

- In view of the results in Section 2.1, we have to impose the symmetry condition RA3 on β^n 's. A simple application of Itô's lemma shows that, for the symmetry condition RA3 to hold, it suffices to choose every κ_t , so that

$$\int_{\mathbb{R}} (e^x - 1) \kappa_t(T, x) dx$$

is a deterministic function of $T - t$, for all times $0 \leq t < T \leq \bar{T}$. To achieve this, in view of (2.15), we need to choose every $\eta_t(T_l, \cdot)$ so that the symmetry index

$$\Xi(T - t) := \int_{\mathbb{R}} (e^x - 1) \eta_t(T, x) dx = \lambda \left(\frac{p}{\lambda_1 - 1} - \frac{1 - p}{\lambda_2 + 1} \right) \quad (2.17)$$

is a deterministic function of $T - t$. This yields:

$$p = \frac{-(1 + \Xi(T - t)/\lambda)(\lambda_1 - 1)}{\Xi(T - t)/\lambda(\lambda_1 - 1)^2 - 2(\lambda_1 - 1) - 1}, \quad (2.18)$$

where Ξ is a fixed (estimated a priori) function.

With the two constraints, our calibration takes only two variables: λ and λ_1 . The condition $p \in (0, 1)$ transforms to the following condition on λ_1 :

$$\lambda_1 \in \begin{cases} (1, \infty), & \text{if } \Xi(T - t) \leq 0, \\ \left(1, 1 + \frac{1}{\Xi(T - t)}\right), & \text{if } \Xi(T - t) > 0. \end{cases} \quad (2.19)$$

As a result, the optimization problem (2.14) can be re-written as

$$\min_{\lambda > 0, \lambda_1 \in I_{\lambda_1}} \sum_{j=1}^N \omega_j |V_t^{\lambda, \lambda_1, j}(T_l) - V_t^{mkt, j}(T_l)|^2, \quad (2.20)$$

where I_{λ_1} is the interval defined in (2.19). The symmetry index function $\Xi(\tau)$, for all $\tau \in \mathbb{R}_+$, can be obtained on the first calibration day $t = 0$, solving a three-variable optimization problem,

$$\min_{\lambda > 0, \lambda_1 > 1, p \in (0, 1)} \sum_{j=1}^N \omega_j |V_0^{\lambda, \lambda_1, p, j}(T_l) - V_0^{mkt, j}(T_l)|^2, \quad (2.21)$$

and setting

$$\Xi(T_l) = \lambda \left(\frac{p}{\lambda_1 - 1} - \frac{1 - p}{\lambda_2 + 1} \right), \quad (2.22)$$

for every maturity T_l , and, finally, interpolating linearly between every T_{l-1} and T_l . We summarize the calibration procedure for $\{\eta_t(T_l, \cdot)\}$ in the following algorithm:

Algorithm 1: Algorithm for calibrating $\{\eta_t(T_l, \cdot)\}$

- 1 Preprocess the market data according to (2.8);
 - 2 For $t = 0$, run the three-variable optimization (2.21), without the symmetry condition, for all maturities, and compute $\Xi(\cdot)$ by (2.22) and linear interpolation;
 - 3 For the subsequent days $t \in (0, \bar{T}]$, run the two-variable optimization (2.20), with already estimated Ξ , to obtain the time series of Lévy densities $(\eta_t)_{t \in [0, \bar{T}]}$.
-

Below are the calibration results. The Lévy densities η on Jan. 3, 2007 – the first day of calibration – are obtained by the three-variable optimization (2.21). From the calibrated parameters, we compute the symmetry index Ξ via (2.22), which is shown in Figure 1. With the symmetry index Ξ , we run the two-variable optimization (2.20) on the following day, Jan. 4, 2007, and obtain the Lévy densities η shown in Figure 2. The corresponding time values are shown in Figure 3, and the implied volatilities are shown in Figure 4.⁴ We can see that the calibration results are quite precise for the strike values close to the spot. The quality of the calibration deteriorates for very large and very small strike values. However, for the experiment presented in the subsequent sections, we are only concerned with the strike values close to the spot. The full calibration results are presented here in order to illustrate the use of tangent Lévy models, and not in order to argue that the double exponential model provides a perfect fit of the market prices.⁵ As for the calibrated Lévy densities η , their values tend to decrease as the time to maturity increases (cf. Figure 2). The magnitude of Ξ (which measures the “asymmetry” of the Lévy measure) is decreasing with maturity as well. Both results are in line with empirical findings on jump intensities and volatility skews. In order to see how the calibration parameters change over time, in Figure 5, we plot the time series of calibrated λ and λ_1 , corresponding to the first two maturities (recall that the optimization in (2.20) is performed over (λ, λ_1)). We can see that these parameters change significantly over time (even the average between the first and the second maturity values changes significantly), hence, it is not realistic to use the double exponential model (with fixed parameters) for analyzing the future dynamics of options’ prices.

Next, for every day t , we need to find κ_t that satisfies (2.15). Notice that, if $\eta_t(T, x)$ is differentiable in T , we obtain:

$$\eta_t(T, x) + (T - t) \frac{\partial \eta_t(T, x)}{\partial T} = \kappa_t(T, x), \quad (2.23)$$

for each $x \in \mathbb{R}$. The relationship (2.23) can be used to back out the additive densities $(\kappa_t)_{t \in [0, \bar{T}]}$ from the calibrated Lévy densities $(\eta_t)_{t \in [0, \bar{T}]}$. However, the calibrated densities $\eta_t(T, \cdot)$ are only defined for $T = T_l$, hence, we need to interpolate them across maturities. An analysis of the calibrated Lévy densities shows that $\eta_t(T, x)$ generally exhibits one of the following two patterns as a function of T .

- For small jump sizes x , $\eta_t(T, x)$ decreases rapidly as T increases. To ensure that the recovered κ is non-negative, we used a combination of exponential function and power function

$$\eta_t(T, x) = c_1(T - t)^{c_2} + c_3(T - t) \exp(-c_4(T - t)) + c_5 \quad (2.24)$$

to fit η , for any fixed x . The corresponding Lévy density κ can then be computed as

$$\kappa_t(T, x) = c_1(c_2 + 1)(T - t)^{c_2} + \exp(-c_4(T - t))(2c_3(T - t) - c_3c_4(T - t)^2) + c_5. \quad (2.25)$$

For every day t , we run an optimization algorithm, searching for constants c_i ’s that would provide the best approximation of the calibrated $\{\eta_t(T_l, x)\}_l$ via (2.24). This optimization is subject to the following constraints: $c_1 > 0$, $c_2 \in (-1, 0)$, $c_3 > 0$, $c_4 \in (0, 2/\bar{T})$ and $c_5 > 0$. It is easy to see that these constraints ensure the positivity of κ_t , defined by (2.25).

- For large jump sizes x , $\eta_t(T, x)$ increases as T increases. The function we used to fit this scenario is a simple polynomial function

$$\eta_t(T, x) = c_1(T - t)^4 + c_2(T - t)^3 + c_3(T - t)^2 + c_4(T - t) + c_5. \quad (2.26)$$

⁴In Figure 4, we assume that the market implied volatility is zero whenever the bid price of the call option falls below its payoff value. Recall that, as described in Subsection 2.3, the market prices of call options have been adjusted so that they are consistent with the assumption of zero interest and dividend rates. Hence, the implied volatility of such an option is not well defined if its value is below the payoff, and we have to set it to zero. Notice also that it does not lead to any arbitrage opportunities, due to the presence of a bid-ask spread.

⁵If additional precision of calibration is required, one can simply use a family of Lévy models with a larger number of parameters. We, however, do not do it in the present paper, due to somewhat limited computing power available to us.

Then, κ is computed as

$$\kappa_t(T, x) = 5c_1(T - t)^4 + 4c_2(T - t)^3 + 3c_3(T - t)^2 + 2c_4(T - t) + c_5. \quad (2.27)$$

As before, for every day t , we run an optimization algorithm, searching for constants c_i 's that would provide the best approximation of the calibrated $\{\eta_t(T_l, x)\}_l$ via (2.26). However, in this case, it is not easy to obtain explicit bounds on the constants c_i that would guarantee the positivity of κ_t , given by (2.27), while preserving enough flexibility and the tractability of the optimization problem. Therefore, we simply construct a very fine partition of the interval $[0, \bar{T}]$ (with the diameter of one day) and introduce a system of obvious (linear) constraints on $\{c_i\}$, which ensure that $\kappa_t(T, x)$, given by (2.27), is positive for the values of $T - t$ in the chosen partition.⁶

An illustration of the two scenarios together with an example of the reconstructed κ is shown in Figure 6.

Remark 1. *It is worth commenting on the potential drawbacks of the proposed calibration method. First of all, as the Kou's model has very few parameters, the resulting curve $\eta_t(T_l, \cdot)$ does not provide a very good fit of the option prices with very large and very small strikes. In addition, the fact that each $\eta_t(T_l, \cdot)$ is fitted to the options' prices with maturity T_l , separately for each T_l , introduces possible inconsistencies in this family of surfaces for large T_l . Namely, the larger is the range of maturities, the worse is the performance of the above interpolation of $\eta_t(\cdot, x)$: the algorithm struggles to maintain a high quality of fit of the actual, calibrated, $\eta_t(T_l, x)$ while preserving the nonnegativity of κ_t . Finally, the proposed calibration method is computationally expensive: it takes over an hour to calibrate κ_t to a surface with 10 maturities and 70 strikes.⁷ These are some of the reasons why we restrict the empirical analysis in the subsequent sections to a set of options with a fairly small range of strikes and maturities. We believe that there may exist a more flexible family of tangent Lévy densities and/or a more efficient calibration algorithm, however, even the present specification allows us to illustrate the advantages of the market-based approach.*

2.5 Dynamic fitting

Recall that, in view of (2.5), the Lévy density κ has the following dynamics:

$$d\kappa_t(T, x) = \alpha_t(T, x)dt + \sum_{n=1}^m \beta_t^n(T, x)dB_t^n. \quad (2.28)$$

In the dynamic fitting, we need to assume that the time increments of κ are stationary, which is only natural if we work with the time to maturity $\tau = T - t$ instead of the maturity T . Namely, we define $\hat{\kappa}_t(\tau, x) = \kappa_t(t + \tau, x)$ and its dynamics

$$d\hat{\kappa}_t(\tau, x) = \hat{\alpha}_t(\tau, x)dt + \sum_{n=1}^m \hat{\beta}_t^n(\tau, x)dB_t^n. \quad (2.29)$$

A simple application of Itô's formula shows that

$$\hat{\alpha}_t(\tau, x) = \alpha_t(t + \tau, x) + \frac{\partial \kappa_t(t + \tau, x)}{\partial T} \quad \text{and} \quad \hat{\beta}_t^n(\tau, x) = \beta_t^n(t + \tau, x). \quad (2.30)$$

To simulate future implied volatility surfaces, all we need are the diffusion terms $\hat{\beta}^n$'s, because the drift term $\hat{\alpha}$ can be computed from $\hat{\beta}^n$'s. We assume that $\hat{\beta}_t^n(\tau, x)$'s are deterministic and constant as functions of t , for any (τ, x) (from a finite family of points).⁸ The values of $\hat{\kappa}$ are observed at the discrete times $\{t_i\}$ (which correspond to days in our experiment). Then, every increment $\Delta \hat{\kappa}_{t_i} = \hat{\kappa}_{t_{i+1}} - \hat{\kappa}_{t_i}$ is a sum of a Gaussian random vector, corresponding to the diffusion part, and a vector that corresponds to the drift term (we view every surface as a vector whose entries correspond to different values of (τ, x)). Notice that the distribution of the Gaussian component is completely determined by its covariance matrix, hence, we will aim to choose $\hat{\beta}^n$'s to match the estimated covariance matrix.

⁶Of course, in theory, the values of $\kappa_t(\tau, x)$ may become negative for the values of τ between the partition points (albeit unlikely, if the diameter is small). However, this causes no problem for our experiment, because, in what follows, we restrict our analysis to a finite number of points (τ, x) on the surface $\kappa_t(\cdot, \cdot)$, and we may (and will) choose these points so that τ belongs to the given partition.

⁷Of course, calibration to a smaller set of option's prices (e.g. single maturity and multiple strikes) can be done significantly faster.

⁸Notice that this assumption implies that $\hat{\beta}^n$'s are the same under both the physical and the risk-neutral measures.

Assuming that the drift term is bounded, it is easy to notice that the standard estimate of the covariance of $\Delta\hat{\kappa}_{t_i}$ also provides a consistent estimate of the covariance of the aforementioned Gaussian vector, asymptotically, as the length of the time increments converges to zero. In the actual computations, we use daily increments – these are small compared to the time span of the entire sample, which is one year. To fit $\hat{\beta}^n$'s to the estimated covariance matrix, it is natural to use the Principal Components Analysis (PCA), which finds the directions that explain most of the variance in the increments $\{\Delta\hat{\kappa}_{t_i}\}$. However, the PCA cannot be applied directly because the number of points on the surface is close to the sample size, which is 251 (i.e. the number of days in the sample): for each t , we have call prices for 10 maturities and 21 jump sizes, which gives us 210 points on the $\hat{\kappa}$ surface after static fitting. To reduce the number of points, we pick every other maturity and the 7 jump sizes whose intensities are larger than others across time t . This gives us $5 * 7 = 35$ points on the reduced surface of $\Delta\hat{\kappa}_{t_i}$.

Applying PCA to the reduced surface, we see that the first three eigenmodes $\{f^n(\tau, x)\}_{n=1}^3$ explain over 93% of the daily variance of $\{\Delta\hat{\kappa}_{t_i}\}$, as shown in Figure 7. To extend the values of the eigenmodes to other points (i.e. other jump sizes and maturities), we simply perform a linear interpolation. The first three eigenmodes have very unique characteristics. The first eigenmode takes the most prominent feature of $\Delta\hat{\kappa}$ – the large values are concentrated around small jumps at very short time to maturity. This eigenmode can be understood as a combination of the “level” factor and the “slope” factor (appearing in a typical PCA result for yield curve dynamics) along both the maturity and the jump size directions. The second eigenmode shows the curvature along the jump size direction, and the third eigenmode shows the curvature along the time to maturity direction. As the eigenmodes $\{f^n(\tau, x)\}_{n=1}^3$ are normalized, to obtain the diffusion terms $\hat{\beta}^n$'s, we need to multiple the eigenmodes by the loading factors:

$$\hat{\beta}_t^n(\tau, x) = \sqrt{\lambda_n} \cdot f^n(\tau, x), \quad n = 1, 2, 3, \quad (2.31)$$

where λ_n is the eigenvalue of the covariance matrix corresponding to the eigenmode f^n . Once we have $\hat{\beta}^n$'s, we change the variables to pass to β^n 's and calculate the drift term α according to (2.7). Figure 8 shows the drift term α computed according to (2.7). Notice that, due to our assumptions on $\hat{\beta}^n$'s and according to (2.7), the drift $\hat{\alpha}_t$ is deterministic and constant in t . Notice that $\hat{\alpha}$ can then be computed as

$$\hat{\alpha}_t(\tau, x) = \alpha_t(t + \tau, x) + \frac{\partial \kappa_t(t + \tau, x)}{\partial T}, \quad (2.32)$$

where we have no problem with evaluating the partial derivative, as, in the static fitting stage, κ_t was interpolated across maturities.

2.6 Monte Carlo simulation of implied volatility surfaces

Once all the terms in the right hand side of (2.29) are estimated, we can, for example, apply an explicit Euler scheme to simulate the future Lévy densities $\hat{\kappa}_t$. However, we need to ensure that the simulated $\hat{\kappa}_t$'s stay nonnegative at all times. Inspired by [4], we incorporate a scaling factor in (2.29) as follows:

$$d\hat{\kappa}_t(\tau, x) = \gamma_t^2 \hat{\alpha}_t(\tau, x) dt + \gamma_t \sum_{n=1}^m \hat{\beta}_t^n(\tau, x) dB_t^n, \quad (2.33)$$

where

$$\gamma_t = \frac{1}{\epsilon} \left(\inf_{\tau \in [0, \bar{\tau}], x \in \mathbb{R}} \hat{\kappa}_t(\tau, x) \wedge \epsilon \right), \quad (2.34)$$

with $\epsilon = 1e^{-6}$ and $\bar{\tau} = 1$. Of course, this modification changes the diffusion term of $\hat{\kappa}_t$, which was estimated from historical data. However, the value of ϵ is chosen to be so small that, in the historical sample, γ_t is always equal to one.⁹ Hence, if we use the $\hat{\beta}^n$'s chosen in the previous subsection, the resulting dynamics are still consistent with the past observations. It is also easy to see that, since γ_t is a scalar, the drift restriction (2.7) is satisfied by the new drift and volatility of κ . Finally, this modification ensures that $\hat{\kappa}_t$ is almost surely nonnegative for any t .

To simulate the future values of κ , we apply the explicit Euler scheme to (2.33), to obtain

$$\hat{\kappa}_{t_{i+1}}(\tau, x) = \hat{\kappa}_{t_i}(\tau, x) + \gamma_{t_i}^2 \hat{\alpha}_{t_i}(\tau, x) \Delta t_i + \gamma_{t_i} \sum_{n=1}^m \hat{\beta}_{t_i}^n(\tau, x) \Delta B_{t_i}^n, \quad (2.35)$$

⁹Due to the natural restrictions on the values of the calibration parameters, it is guaranteed that all the calibrated $\hat{\kappa}_t$ are strictly positive, hence, there always exists a small enough ϵ , such that, in the historical sample, γ_t is equal to one.

with $t_{i+1} = t_i + \Delta t_i$ and Δt_i being one day. Having simulated future $\hat{\kappa}_{t_i}$, we compute η_{t_i} via (2.15). Then, for every fixed maturity T , the option prices in the model given by the Lévy density $\eta_{t_i}(T, \cdot)$ can then be computed, for example, using the methods proposed in [5] or [21]. These methods are based on Fourier transform and can be implemented efficiently via numerical integration.¹⁰ In particular, in our computation, we set the initial value of the associated exponential Lévy process to one, and use the following formula to calculate future option prices:

$$C_t^{1, \kappa_t}(T, x) = 1 - \frac{e^{x/2}}{\pi} \int_0^\infty \frac{du}{u^2 + \frac{1}{4}} \operatorname{Re} \left[\exp(-iux) \phi_t \left(T, u - \frac{i}{2} \right) \right], \quad (2.36)$$

where ϕ_t is the characteristic function of an exponential Lévy process with the Lévy density $\eta_t(T, \cdot)$ and with initial value one:

$$\phi_t(T, u) = \exp \left[-iu(T-t) \int_{\mathbb{R}} \eta_t(T, x)(e^x - 1)dx + (T-t) \int_{\mathbb{R}} (e^{iux} - 1)\eta_t(T, x)dx \right].$$

From the above option price, $C_t^{1, \kappa_t}(T, x)$, we can easily calculate the implied volatility by inverting the Black Scholes formula, assuming that the spot level is at one and that the interest and dividend rates are zero. As discussed at the very end of Subsection 2.1, this value is the same as the value of implied volatility of a call option for any spot level S_t , strike $S_t e^x$, and maturity T , regardless of what the level of S_t is (hence, we don't need to simulate it). Using this method, we simulate the implied volatility surfaces five days into the future starting from Dec. 13, 2007, as shown in Figures 9 and 10.

3 Empirical analysis of the performance of tangent Lévy models

In this section, we discuss the importance of modeling derivatives prices in a way that is consistent with their historical time series. As we know, an investment manager's portfolio or a trader's trading book often contains multiple financial derivatives written on the same underlying. As a simple example, an equity trader might hold a calendar spread and a butterfly spread at the same time. To properly manage the risk, one needs to understand the joint dynamics of these derivatives. In turn, to construct a reliable model, it is crucial to have a modeling framework which produces feasible combinations of option prices (i.e. is arbitrage-free) and which can be fitted both to the present and to the historical values of these prices. Tangent Lévy models (as any market-based model) are built to achieve this goal precisely. This is due to the fact that not only present but also historical information contained in the time series of options' prices is used in the estimation of model dynamics. Classical stochastic volatility models cannot capture the historical evolution of options' prices, hence, there is a reason to believe that market-based models would lead to better performance in portfolio management. To show that tangent Lévy models do indeed work better, here, we test the model implemented in Section 2, using the portfolio choice problem described below. The results are, then, compared against one of the most popular volatility models in the industry – the Stochastic Alpha Beta Rho (SABR) model. In addition, we compare the results produced by the tangent Lévy model to the ones produced by the method of direct simulation of implied volatility, described in [7], which also makes use of historical options' prices. However, the latter method produces arbitrage in the simulated options' prices, which makes it impossible to use this method in many applications (e.g. pricing illiquid securities) and which brings it beyond the scope of arbitrage-free models, within which we aim to show the superiority of market-based models. We conduct an additional experiment, aimed at estimating Value at Risk of an options' portfolio, in order to illustrate the negative consequences of the presence of arbitrage in the direct simulation method.

3.1 The variance-minimizing portfolio choice problem

This example is a simplified Markowitz-type portfolio optimization problem. Consider a portfolio manager who needs to decide how he/she should balance a portfolio of SPX options so that its risk is minimized. Among the many definitions of portfolio risk, we adopt the one used in the classic Markowitz problem (for example, see Section 6.6 of [22]) – namely, the standard deviation of the portfolio return over a given (future) time period. Notice that this is not a typical Markowitz portfolio problem, given we are not considering the trade-off between return and risk as a typical Markowitz problem would. As a matter of fact, we assume that the portfolio manager lives

¹⁰Please note that we cannot use (2.12) to calculate option prices, because, even though the calibrated Lévy densities $\{\eta_t(T, \cdot)\}$ are double exponential, there is no reason to believe that the simulated η 's remain double exponential.

in a risk-neutral world, so that the expected return is normalized. We admit that lacking excess return might make the example less exciting, but it helps us compare the model performance in an apples-to-apples fashion. With the normalized return, there is no need to worry about the impact of different market views portfolio managers might build into the investment decisions. Of course, without such a trade-off, there is a trivial solution to the portfolio choice problem – do not invest at all, reducing the risk to zero. To make the problem non-trivial, we require that the value of the portfolio at the time when it is constructed must be equal to a fixed number. Such a restriction is relevant if the manager makes profits off the commission, proportional to the size of the investment portfolio he/she manages. For example, an option market maker might want to know the optimal inventory so that he/she can adjust the quoting strategy accordingly to reach the portfolio composition with minimal inventory risk. Or, a broker dealer might need to know her optimal position in options over the next several days to meet the risk and capital requirements.

We now formulate this problem mathematically. Let us assume that there are n options with the same maturity T but with different strikes K_1, \dots, K_n in the portfolio.¹¹ Let $C_t(K_i)$ be the time- t price of the K_i -struck option, and let ω_i be the quantity of this option in the portfolio, with a negative ω_i representing the short-selling. The weights ω_i have to be determined at the initial time t . The portfolio value at any future time $t + u$ is simply $V_{t+u} = \sum_{i=1}^n \omega_i C_{t+u}(K_i)$, and the return over the time period $[t, t + u]$ is $R_u^t = V_{t+u}/V_t$. For simplicity, we assume that the risk-free rate and the dividend yield are both zero, so the expectation of R_u^t , computed at time t , is simply 1.¹² For a given $t \in [0, T)$ and $u \in (0, T - t]$, to determine the portfolio weights, we need to solve the following convex optimization problem:

$$\begin{aligned} \min_{\omega \in \mathbb{R}^n} \quad & \mathbb{E}_t(R_u^t - 1)^2 = \min_{\omega \in \mathbb{R}^n} \quad \mathbb{E}_t(V_{t+u} - 1)^2 \\ \text{s.t.} \quad & V_t = \omega^T C_t = 1, \end{aligned}$$

where the initial value of the portfolio is one. This is equivalent to

$$\begin{aligned} \min_{\omega \in \mathbb{R}^n} \quad & \omega^T \Lambda_u^t \omega \\ \text{s.t.} \quad & V_t = \omega^T C_t = 1, \end{aligned} \tag{3.1}$$

where $\Lambda_u^t = \mathbb{E}_t[(C_{t+u} - C_t)^2]$ is the time- t covariance matrix of the (absolute) returns of the options in the time period $[t, t + u]$. It is easy to see that the closed-form solution to the quadratic optimization (3.1) is

$$\omega^t = \frac{(\Lambda_u^t)^{-1} C_t}{C_t^T (\Lambda_u^t)^{-1} C_t}. \tag{3.2}$$

Thus, as it is well-known, the key to solving this optimization problem is to estimate the covariance matrix Λ_u^t . To do this, we compute the sample covariance matrix using the time- $(t + u)$ option prices simulated under each model at time t . Note that, to obtain a fair comparison, the parameters of each model are only estimated using the options data prior to day t . Then, given N samples of the time- $(t + u)$ options' prices,

$$C_t^{(j)} = [C_{t+u}^{(j)}(K_1), \dots, C_{t+u}^{(j)}(K_n)]^T, \quad j = 1, \dots, N,$$

the sample covariance matrix is estimated as

$$\Lambda_u^t = \frac{1}{N-1} \sum_{j=1}^N (C_t^{(j)} - C_t)(C_t^{(j)} - C_t)^T. \tag{3.3}$$

Different models generate different simulated paths of options' prices, which then lead to different optimal weights. Naturally, how these optimal weights perform in the real world serves as an indicator of the *model performance*. To be more specific, a better model should be able to generate portfolios with smaller standard deviation in the returns. To estimate the standard deviation of portfolio returns, we define the figure of merit Q as the *average realized deviation* of the portfolio return in the testing period, i.e.

$$Q = \sqrt{\frac{1}{N_{test}} \sum_{k=1}^{N_{test}} (R_u^{t_k} - 1)^2}, \tag{3.4}$$

¹¹The strikes will be chosen so that they always stay between 95% and 105% of the spot value (evaluated at the time when the portfolio is constructed).

¹²This assumption is justified by the very short time horizon for each portfolio return – 8 days – and it is confirmed empirically.

where N_{test} is the number of trials (i.e. the number of days in the testing period), and $R_u^{t_k}$ is the actual portfolio return (given by market data) over the time period $[t_k, t_k + u]$, with the optimal weights ω^{t_k} , obtained by (3.2) on day t_k . Every t_k correspond to the k th day in the testing period, i.e.

$$R_u^{t_k} = \sum_{i=1}^n \omega_i^{t_k} C_{t_k+u}(K_i). \quad (3.5)$$

Recall that, by assumption, the mean of $R_u^{t_k}$, computed at t_k , should always be 1. To make this assumption be consistent with the data, we choose a relatively small time horizon u .

3.2 Simulation algorithms

As mentioned in the previous subsection, to find the optimal portfolio, we need to estimate the covariance matrix using simulated option prices. In this section, we describe the simulation algorithms for each model.

- **Double exponential tangent Lévy model.** For this experiment, we need to simulate both the underlying process S and the non-homogeneous Lévy density κ . For the double exponential tangent Lévy model, in particular, we need to complete the following two steps to move one step ahead from t_i to t_{i+1} :

- Step 1: Simulate the underlying process by

$$S_{t_{i+1}} = S_{t_i} \exp\left\{-\int_{\mathbb{R}} (e^x - 1) \kappa_{t_i}(t_i, x) dx \Delta t_i + \sum_{k=1}^{N_{t_i}} J_k\right\},$$

with $t_{i+1} = t_i + \Delta t_i$ and Δt_i being one day. Here $\kappa_{t_i}(t_i, x)$ is the additive density for immediate maturity $T = t_i$, N_{t_i} is the number of jumps during the $(t_i, t_{i+1}]$ period, which has a Poisson distribution with parameter $\lambda \Delta t_i$, where $\lambda = \int_{\mathbb{R}} \kappa_{t_i}(t_i, x) dx$, and J_k 's are the jump sizes having the distribution $\frac{1}{\lambda} \kappa_{t_i}(t_i, x) dx$. Notice that we approximate the jump component of $\log S$ with a compound Poisson process, which is reasonable given that the jump activity is finite in our setting.

- Step 2: Simulate the Lévy density $\kappa_{t_{i+1}}$ via (2.35).

Simulating u days ahead requires repeating the two steps u times. We can then use the Fourier transform methods, as described in Subsection 2.6, to calculate time- u option prices, and estimate the covariance matrix to obtain optimal weights.

- **SABR model.** The simulation based on SABR model is slightly easier. SABR model, as proposed by Hagan et al. in [13], describes the dynamics of the forward price F and the volatility α as follows:

$$\begin{aligned} dF_t &= \alpha_t F_t^\beta dB_t^1, \\ d\alpha_t &= \nu \alpha_t dB_t^2, \end{aligned} \quad (3.6)$$

where F and α are correlated through $dB_t^1 dB_t^2 = \rho dt$. [13] provides the following asymptotic formula for the time- t implied volatility under the SABR model:¹³

$$\begin{aligned} \sigma_t(K, T, F_t, \alpha_t) &\approx \frac{\alpha_t}{(F_t K)^{(1-\beta)/2} \left\{1 + \frac{(1-\beta)^2}{24} \log^2 F_t/K + \frac{(1-\beta)^4}{1920} \log^4 F_t/K\right\}} \cdot \left(\frac{z}{x(z)}\right) \\ &\left\{1 + \left[\frac{(1-\beta)^2}{24} \frac{\alpha_t^2}{(F_t K)^{(1-\beta)}} + \frac{1}{4} \frac{\rho \beta \nu \alpha_t}{(F_t K)^{(1-\beta)/2}} + \frac{2-3\rho^2}{24} \nu^2\right] (T-t)\right\}, \end{aligned} \quad (3.7)$$

where K is the strike value, T is the maturity, F_t is the current spot level, and z and $x(z)$ are defined as

$$\begin{aligned} z &= \frac{\nu}{\alpha_t} (F_t K)^{(1-\beta)/2} \log \frac{F_t}{K}, \\ x(z) &= \log \left\{ \frac{\sqrt{1-2\rho z + z^2} + z - \rho}{1-\rho} \right\}. \end{aligned} \quad (3.8)$$

¹³Note that 3.7 is only an approximation for the true implied volatility in SABR model. This approximation speeds up the computations significantly, and, as mentioned in [13], approximates the true implied volatility well for reasonable parameter values.

As for the parameters' values, [13] suggests that β can be fixed in advance and [29] verifies empirically that this is a reasonable assumption. In our example, we will use two values of β : $\beta = 1$ and $\beta = 0.7$. $\beta = 1$ is probably the most natural choice for equity market as it mimics a log-normal model most closely, and $\beta = 0.7$ is widely used on trading desks as it provides better results for risk management. The other parameters – the current volatility α_t , the volatility of volatility ν and the correlation ρ – will be calibrated to market prices by minimizing the sum of squared differences between the market call prices and those produced by the model, calculated with (3.7). With the parameters calibrated on the initial day d_k , the forward price and the volatility can be simulated as follows:

$$\begin{aligned} F_{t_{i+1}} &= F_{t_i} e^{-0.5\alpha_{t_i}^2 \Delta t_i + \alpha_{t_i} \Delta B_{t_i}^1}, \\ \alpha_{t_{i+1}} &= \alpha_{t_i} e^{-0.5\nu^2 \Delta t_i + \nu(\rho \Delta B_{t_i}^1 + \sqrt{1-\rho^2} \Delta B_{t_i}^2)}, \end{aligned} \quad (3.9)$$

where B^1 and B^2 are independent Brownian motions. The time- t_i implied volatilities and option prices can then be computed via (3.7), with the simulated spot F_{t_i} and volatility α_{t_i} . Figure 11 shows the results of calibrating SABR model with $\beta = 1$ to SPX options with a fixed (short) maturity, on the first trading day of 2008. We can see that the model fits market prices very well, at least, if the range of strikes and the maturity are small enough.

3.3 Direct simulation of implied volatility

In this subsection, we describe another method of simulating the implied volatility surface, mentioned in the introduction. This method is described in [7], and it is based on the PCA analysis of the increments of the logarithm of implied volatility. Namely, on any past observation day t_i , we consider the market implied volatility surface $IV_{t_i}(\tau, x)$, for different values of the times to maturity τ and the (negative of) log-moneyness $x = K/S_{t_i}$,¹⁴ where K is the strike of an option and S_{t_i} is the underlying value at time t_i . As mentioned earlier, we use market prices of SPX options and deduce the implied volatility from these prices using an inverse of the Black-Scholes formula (note that we apply discounting to the options' prices and strikes, so that the inverse function is used with zero interest and dividend rates). In order to obtain the value of $IV_{t_i}(\tau, x)$ for an arbitrary (τ, x) , we interpolate linearly across the observed values: first, in τ , then, in x (i.e. first, we construct continuous curves $IV_{t_i}(\tau, x_j)$, for every observed x_j , then, we interpolate in x for every value of τ). We fix 15 pairs of $\{\tau_j, x_j\}$, located in the region that is relevant for our experiment (in particular, all $\{x_j\}$ are between 0.9 and 1.1). In fact, for the numerical experiment described in the next section, we only need the prices (or implied volatilities) of the options with short maturities (around 1 month), hence, we choose all τ_j to be equal to one month (the implied volatilities for shorter maturities are simply assumed to be equal to the ones with maturity one month). Thus, we obtain a time-series of vectors,

$$\{IV_{t_i} = (IV_{t_i}(x_1), \dots, IV_{t_i}(x_{15}))\},$$

where t_i is the i th day in a given historical sample.

As in [7], our standing assumption is that the log-increments of this time series $\{\Delta \log IV_{t_i} = \log(IV_{t_{i+1}}/IV_{t_i})\}$ have the same Gaussian distribution. In other words,

$$\Delta \log IV_{t_i}(x_j) = \mu(x_j) \Delta t_i + \sum_{n=1}^m \Sigma^n(x_j) \Delta B_{t_i}^n, \quad (3.10)$$

where Δt_i is one day, B^n 's are independent Brownian motions, and $\mu = (\mu(x_1), \dots, \mu(x_{15}))$ and each

$$\Sigma^n = (\Sigma^n(x_1), \dots, \Sigma^n(x_{15}))$$

is a constant vector. The vector μ is estimated as a sample mean, and we estimate the vectors $\{\Sigma^n\}$ via PCA, as it is done in [7]. The results of this analysis, carried out using the market data in the time period Jan. – Dec. 2011, are presented in Figure 12. In particular, we see that the first two eigenmodes $\{f^n\}_{n=1}^2$ explain over 98% of the daily variance of $\{\Delta \log IV_{t_i}\}$. As the eigenmodes $\{f^n\}_{n=1}^2$ are normalized, to obtain the vectors Σ^n 's, we need to multiple the eigenmodes by the loading factors:

$$\Sigma^n = \sqrt{\lambda_n} \cdot f^n, \quad n = 1, 2,$$

¹⁴With a slight abuse of notation, in this paper, we refer to x as the “log-moneyness”, while it is, in fact, a negative of log-moneyness.

where λ_n is the eigenvalue of the covariance matrix corresponding to the eigenmode f^n . The estimate of μ shows that its values are negligible (in view of the small size of the time increment Δt_i). Having estimated the parameters, we can simulate the future values of implied volatility directly using (3.10). In order to simulate the underlying values, we use the following model

$$\Delta S_{t_i} = IV_{t_i}(0) \left(\sum_{n=1}^m \rho_n \Delta B_{t_i}^n + \rho_0 \Delta B_{t_i}^0 \right),$$

where B^n 's are the same Brownian motions that appear in (3.10), B^0 is an independent Brownian motion, and ρ_n 's are the correlation parameters, to be estimated. Recall that we have found that $m = 2$. Due to the orthogonality of the eigenmodes of $\Delta \log IV$, it is easy to extract the time series $\{B_{t_i}^n\}$, for $n = 1, 2$, from the time series $\{\Delta \log IV_{t_i}\}$ (cf. [7]). Then, ρ_n 's, for $n = 1, 2$, can be estimated using the standard methods – as the correlation between the increments of B^n and $\log S$. We obtain: $\rho_1 = -0.865$ and $\rho_2 = 0.22$. The remaining coefficient is computed in an obvious way: $\rho_0 = \sqrt{1 - \rho_1^2 - \rho_2^2} \approx 0.45$.

Despite the fact the model (3.10) is very simple to implement, it suffers from one serious drawback – it is not arbitrage-free. As we discussed at the beginning of this paper, there are two types of arbitrage in options market: the static arbitrage, in which certain relations between options' prices for different strikes and maturities are violated, and the dynamic one, in which the time changes in options' prices do not have the correct pattern (i.e. do not have the correct drifts). From a practical point of view, the presence of the latter type of arbitrage may not lead to any obvious trading opportunities (especially, for short time horizons), while the static arbitrage opportunities are very easy to exploit and, hence, are not typically observed in the market. It turns out that the model implied by (3.10) does produce static arbitrage in the simulated options' prices. Indeed, there is no particular reason to believe that the shapes of implied volatility produced by (3.10) are such that the shapes of the associated options' prices satisfy the desired static no-arbitrage conditions. This observation has been documented in several existing publications (cf. [2]), and it serves as one of the main arguments for using the market-based models. From a practical point of view, the presence of static arbitrage in the model may or may not lead to a problem, depending on how the model is used. In fact, it is shown in Subsection 3.4.2 that the direct simulation method performs reasonably well for the problem at hand, which, effectively, reduces to the computation of the covariance matrix of options' returns. However, the presence of static arbitrage may create unnatural (and, hence, unreliable) results when computing other measures of risk of a portfolio of options. To illustrate the latter, in Subsection 3.4.2, we conduct an additional experiment, which compares the performance of the direct simulation method and of the DETL model in estimating the Value at Risk (VaR) of a portfolio consisting of two call options. In particular, we show that the estimates of VaR produced by the direct simulation are unrealistic, while the DETL model does a fairly good job at this task.

3.4 Results of empirical analysis

In this section, we go through the test procedure in detail and present the test results for the following models:

- Double exponential tangent Lévy model (DETL).
- SABR model with $\beta = 1$.
- SABR model with $\beta = 0.7$.
- Direct simulation of implied volatility.

The DETL and SABR models will be run in two periods: (I) Jan. 2007- Aug. 2008 and (II) Jan. 2011 - Dec. 2012. The direct simulation of implied volatility will only be tested in period (II). For each period, we use the first year's data as a training sample, to estimate the parameters of the tangent Lévy model, and we use the rest of the data as the testing sample, to compute the figure of merit Q defined in (3.4). The division between training and testing samples is shown in Table 2. Please note that we cut off the first period at August 2008 to reduce the impact of the financial crisis. The tests will be run on a portfolio of call options and underlying – referred to as a “(C + S) portfolio” – with three, four and five strikes. In each case, we pick every other strike starting from the strike closest to the underlying spot value (in other words, closest to at-the-money) at the moment when the portfolio is constructed. We pick these options because their market prices are most accurate. Assuming the set of available strikes is $K_1 < K_2 < \dots < K_n$ and the spot S satisfies $K_{i-1} < S < K_i$, Table 3 illustrates the strikes used in each

case. The values of the strikes chosen this way always stay between 95% and 105% of the spot value (evaluated at the time when the portfolio is constructed).

As announced earlier, we also investigate the performance of the method based on the direct simulation of implied volatility. As the latter model is not arbitrage-free (and, hence, is not the main emphasis of the present investigation), we only provide the analysis of its performance for a smaller number of experiments: i.e. for the case of five strikes and only in period (II). The parameters of this model (i.e. the vectors Σ^n 's and correlations ρ^n 's) are estimated using the 2011 data, and the model is tested (i.e. the portfolios are constructed and the figure of merit is computed) on the 2012 data.

For all portfolios, we use a simulation horizon u equal to 8 days, and, at the time t_k , when the portfolio is constructed, the options have maturity equal to $t_k + u$ plus additional 30 days, so that their time-to-maturity becomes 30 days when the given simulation period ends. We also assume the budget constraint $M = 1$. In addition to the figure of merit Q , we also check the *average predicted deviation* defined as

$$P = \sqrt{\frac{1}{N_{test}} \sum_{k=1}^{N_{test}} (\omega^{t_k})^T \Lambda_u^{t_k} \omega^{t_k}}, \quad (3.11)$$

where ω^{t_k} is the set of optimal weights obtained via (3.2) on the day t_k . The difference between Q and P is another measure of the accuracy of a model's prediction. Besides the predicted and realized deviation, one may be interested in how much the optimal portfolio weights fluctuate across the initial days t_k . To measure this fluctuation, we define the *average quantity oscillation* index K :

$$K = \frac{1}{n} \sum_{i=1}^n \left(\frac{1}{N_{test}} \sum_{k=1}^{N_{test}-1} \left| \omega_i^{t_{k+1}} - \omega_i^{t_k} \right| \right), \quad (3.12)$$

where $\omega_i^{t_k}$ is the quantity of the K_i -struck option in the optimal portfolio constructed on day t_k .

3.4.1 Period I

For Period I, the estimation of the parameters of the tangent Lévy model is described in Section 2. Following the simulation algorithm outlined in Subsection 3.2, for every initial day t_k in the testing sample, we simulate 500 sample paths of the underlying and the options' prices, using each model, and starting with the actual prices observed on day t_k . In the simulation for the tangent Lévy model, we use the drift α and the volatility β estimated from the training sample, so that the testing is performed out-of-sample (this issue is irrelevant for SABR model, as it does not allow for any use of past option prices). Using the simulated prices, for each model, we calculate the average predicted deviation P , according to (3.11), and estimate the optimal portfolio weights ω^{t_k} via (3.2). Using these weights, we construct the corresponding portfolio and record its value at time $t_k + u$ using the *actual market prices*. Collecting the results for all initial days t_k , we compute the average realized deviation Q via (3.4).

The results are shown in Table 4. It is easy to see that, for a portfolio with 5 strikes, DETL model produces much smaller values of Q than those produced by SABR model, indicating that the tangent Lévy models do a much better job at finding the minimal-variance portfolio. This can also be seen in Figure 13, which shows that the distribution of realized returns is much more concentrated around 1 under the DETL model than under the SABR model. Furthermore, if we look at the difference between Q and P , we can see that it is much smaller for DETL than for SABR model. This suggests that the tangent Lévy models produce a more reliable prediction of the risk of an options' portfolio (as measured by the standard deviation of its return) than the SABR model. Besides a small return deviation, another nice feature of tangent Lévy models is the stability of optimal option quantities across the initial days t_k . Figure 14 shows the optimal quantities of options and underlying index, in the portfolio with 5 strikes, across all initial days in the testing period, for every model. Similarly, Table 5 shows the average quantity oscillation K (defined in (3.12)) for all portfolios and all models. It is easy to see that the portfolio weights constructed via DETL model are much more stable than those constructed using SABR model. This can be explained by the fact that the parameters of a tangent Lévy model are estimated from both the present and the historical options' prices, while a classical stochastic volatility model, such as SABR, can only be calibrated to the options' prices available on day t_k . It is well known (and obvious intuitively) that an estimate based on a larger sample is more robust. Thus, the ability of tangent Lévy models to be fitted to the historical options prices makes their output (in this case, the optimal portfolio weights) more stable. This advantage, of course, comes at a cost of higher computational complexity. The Monte Carlo simulations and the computation of optimal portfolio weights

for 5 strikes, using DETL model, take around 10 minutes for each initial day. Analogous computations using SABR model take around 10 seconds.¹⁵ Nevertheless, in many cases, 10 minutes is a reasonable time to wait in order to obtain higher quality results (e.g. if the computations are performed daily).

Tables 4 and 5 also show that the difference between the performances of DETL and SABR models shrinks as the number of strikes in the portfolio decreases. This is not a surprise: as the number of strikes decreases, the number of degrees of freedom in the dynamics of options' prices, which have to be captured by the model, decreases as well. Eventually, for a very small number of strikes, the SABR model does relatively well. However, even in the case of 3 strikes, the tangent Lévy model does at least as good as SABR (although at a higher computational cost). Of course, the real benefit of using tangent Lévy models is only visible when the number of options in the portfolio is relatively large. Figure 15 provides a visual explanation for DETL's outperformance. It shows the 500 simulated call option prices, as functions of strike, at the end of the simulation period, under DETL model and under SABR model with $\beta = 1$. It is easy to see that the SABR model only allows for very limited shapes of the simulated call price curves, while the tangent Lévy model is able to generate a much wider variety of shapes. It is the lack of variety of different scenarios for the joint evolution of call prices (not merely the lack of parameters in the model) that prohibits the classical stochastic volatility models, such as SABR, from capturing the true dynamics of options' prices (or, of implied volatility surface) contained in the historical data.

3.4.2 Period II

Herein, we repeat the same analysis for Period II. The main purpose of this analysis is to show that the outperformance of tangent Lévy models is not due to our choice of a testing period, but that it is a persistent property. In addition, we investigate the performance of the method based on direct simulation of implied volatility.

First, we need to estimate the parameters of DETL model using the data of year 2011. The estimation procedure is exactly the same as the one described in Section 2, so we only present the main results here. In particular, the PCA shows that the first three eigenmodes explain over 93% of the variance. The eigenvalues and the eigenmodes are shown in Figure 16, and the corresponding drift term α is shown in Figure 17.a. Comparing to Figures 7 and 8, we see that these results are almost the same as for the year 2007, suggesting that this model is very robust.

Once the estimation is completed, we can repeat the same simulation and testing procedures as in Subsection 3.4.1, to obtain the results shown in Tables 6 and 7, as well as in Figures 18 and 19. These results confirm the finding of Subsection 3.4.1: for sufficiently many strikes in the portfolio, the tangent Lévy model does a much better job at finding a portfolio with smallest variance, its prediction for the variance is more reliable, and the portfolio weights are more stable.

In the case of direct simulation of implied volatility, we proceed similarly. Namely, we estimate the parameters of the model (i.e. the vectors Σ^n 's and the correlations ρ^n 's) using the PCA analysis of the time series of implied volatilities from 2011, and, then, conduct the forward 8-day simulation starting from every day in 2012 (using the market implied volatility as the initial condition, and reusing the same vectors Σ^n 's and the same correlations ρ^n 's). As shown in Tables 6–7 and in Figures 18–19, the method of direct simulation of implied volatility does slightly worse than the DETL model, but it is quite close. From the results of this particular experiment, one can argue that the performance of the direct simulation is comparable to the DETL performance. This can be explained by the fact that the direct simulation method also takes into account both the present and the historical values of the implied volatility. In addition, this method is significantly faster than the other two models: the Monte Carlo simulations and the computation of optimal portfolio weights for 5 strikes take around 0.15 seconds for each initial day.

However, the direct simulation method does not constitute an arbitrage-free model: in particular, the simulated options' prices, typically, contain static arbitrage, as discussed in Subsection 3.3. The presence of static arbitrage implies certain obvious shortcomings, which make it impossible to use the same model for other tasks, in addition to the one considered in this experiment. To illustrate this, we conduct an additional experiment, in which we use the direct simulation method and the DETL model to estimate the VaR of a so-called “digital spread” portfolio. This portfolio consists of a long position in a European call option with a smaller strike and short position in a co-maturing call with a larger strike, scaled for convenience: $(C(K_4) - C(K_5))/(K_4 - K_5)$, with $K_4 < K_5$. On each day t_k in 2012, we construct the digital spread portfolio by choosing K_4 and K_5 to be the two largest strikes used in the main experiment (i.e. the variance minimization problem) and by choosing the same maturity as in the main experiment (defined at the beginning of Subsection 3.4). Recall that VaR is a quantile of the distribution of the future value of the portfolio. Hence, for each day t_k , we use the direct simulation method and the DETL model

¹⁵These computations are performed on 2.2 GHz Intel Core i7, 8 GB 1600 MHz DDR3.

to simulate 500 future sample values of this portfolio, for 16 days ahead. In each of the two samples, we compute the 0.01-quantile, thus, obtaining two time series of quantiles. These time series are plotted in Figure 20. Notice that the actual portfolio values have to be nonnegative at all times, due to the obvious static no-arbitrage condition. Therefore, a reasonable estimate of a quantile has to be nonnegative as well. In agreement with this observation, the quantiles produced by the DETL model are nonnegative, because this model is designed to exclude the possibility of static arbitrage in the simulated prices. However, the direct simulation method does not take the issue of static arbitrage into account. As a result, the quantiles produced by the direct simulation often take negative values of significant magnitude (relative to the quantiles produced by the DETL model). In addition, the quantile estimates of the direct simulation method exhibit strong oscillation over time, while the estimates produced by the DETL model do not. This demonstrates the additional advantage of the DETL model, which is the stability of the quantile estimate. On every day t_k , we also compute the actual (realized) future value of the digital spread portfolio, for 16 days ahead. These values, as expected, remain positive at all times. Part (b) of Figure 20 shows the time period with the smallest realized future values of the digital spread portfolio. We can see that these values remain positive and above the quantiles produced by the DETL model, while they do fall below the quantiles produced by the direct simulation method, in two instances. The latter observation, itself, does not imply that the quality of the quantile estimates produced by direct simulation is low.¹⁶ However, it does indicate that, for two days in the sample, the direct simulation method fails to provide the appropriate capital requirements for the given portfolio, even at a 99% confidence level. We conclude that the DETL model performs significantly better than the direct simulation method at the task of computing VaR of the digital spread portfolio. In practice, one would like to use the same model for various purposes, hence, it is highly desirable that the model performs reasonably well for a wide range of problems. The simplistic example of the VaR computation, described above, shows that using a model that is not free of arbitrage may have unpredictable consequences.¹⁷

4 Conclusion

In this paper, we implement and test a market-based model for European-type options. This model is a numerically tractable specification of the family of tangent Lévy models proposed in [4] and [3]. Such models, in particular, provide a method for generating Monte Carlo samples of future implied volatility surfaces, in a way that is consistent with their past and present values. We estimate the parameters of this model using real market data, for two periods: 2007-2008 and 2011-2012. The estimation procedure is described in detail, so that it can be reproduced by any interested reader.

In addition, we use the estimated model and the real market data to conduct an empirical study, whose main goal is to compare the performance of a market-based model with the performance of a classical stochastic volatility model. We choose the problem of minimal-variance portfolio choice to compare the performance of the tangent Lévy model with the SABR model. Our study demonstrates that the tangent Lévy model does a much better job at finding a portfolio with smallest variance. In addition, its prediction of the future return variance is more reliable, and the portfolio weights are more stable. To the best of our knowledge, this is the first example of empirical analysis, based on real market data, which provides a convincing evidence of the outperformance of the market-based models for European options, as compared to the classical spot models.

We also compare the DETL model to the method of direct simulation of implied volatility proposed in [7]. We find that the performance of direct simulation is quite close to that of the DETL model (only slightly worse), for the problem of portfolio selection. However, unlike the DETL and SABR models, the method of direct simulation does not belong to the class of arbitrage-free models. As a result, it may produce a priori impossible scenarios in the Monte Carlo simulations and, in turn, unrealistic estimates. We illustrate this fact by estimating VaR of a given portfolio of call options using both the DETL model and the direct simulation method. This analysis illustrates some of the advantages of arbitrage-free models in general, and the market-based models in particular.

Our work is subject to certain limitations, which suggest directions for future research. One of the biggest

¹⁶For example, if the sample members are independent, it is likely that one of the members falls below the 0.01-quantile. In the present case, the members are not independent, as the testing time periods overlap, hence, the probability of such an event is lower. Nevertheless, it may be non-negligible.

¹⁷The consequences may be much more severe if, for example, the model is used for pricing exotic instruments, such as options written on portfolios of call options. From a theoretical point of view, this type of phenomena are well understood, and, hence, we do not provide a detailed discussion of this in the present paper. It may be interesting to construct a particular example in which the absence of arbitrage in the direct simulation causes serious deficiencies in pricing exotic options, thus, demonstrating the benefits of arbitrage-free models. However, this requires an entirely new numerical experiment, and, hence, we leave it for future research.

challenges in implementing a market-based model is the numerical complexity and potential instability of the static fitting (see Remark 1). To mitigate some of these issues, we chose to work with a parametric family of Lévy densities. Although this increases the stability of computations, we still have to rely on the convergence of a generic optimization algorithm, which is applied to a non-convex problem and takes a long time to converge. In addition, the restriction to a parametric form of the density also implies that we may not be able to fit option prices with a required precision. Finally, the transition between Lévy density and option prices, at least, as it is implemented herein, requires substantial (although reasonable) time. As a result, the DETL model performs slower than the competing models, and its applications are limited to the problems that do not require frequent re-runs (e.g. daily runs would be appropriate). Herein, we do not aim to show that the present choice and implementation of a market-based model are the most efficient ones: rather, our goal is to demonstrate the practical advantages of the general principle of market-based modeling. However, it is important to, ultimately, design more efficient classes of market-based models. To achieve this using tangent Lévy models, one has to find a family of Lévy densities that is rich enough, to approximate market prices of the options with sufficient precision, and, at the same time, not too large, so that the transition between Lévy density and options' prices (and vice versa) is fast and stable.¹⁸ This is a balance that seems hard to find. One can go even further along these lines and search for other families of tangent models – not necessarily based on Lévy processes. This, in turn, motivates the search for other families of models, which can always fit an arbitrary family of arbitrage-free options' prices.¹⁹ An example of such a family is provided in [6], but the existence and description of consistent dynamics within this family of tangent models remains an open question.

5 Appendix A

Here, we define the Banach spaces associated with tangent Lévy processes.

- \mathcal{B}_0 is a Banach space of Borel measurable functions satisfying

$$\|f\|_{\mathcal{B}_0} := \int_{\mathbb{R}} (|x| \wedge 1) |x| (1 + e^x) |f(x)| dx < \infty. \quad (5.1)$$

- \mathcal{B} is a Banach space of absolutely continuous functions $f : [0, \bar{T}] \rightarrow \mathcal{B}_0$ satisfying

$$\|f\|_{\mathcal{B}} := \|f(0)\|_{\mathcal{B}_0} + \int_0^{\bar{T}} \left\| \frac{d}{du} f(u) \right\|_{\mathcal{B}_0} du < \infty. \quad (5.2)$$

- \mathcal{H}_0 is a Hilbert space of Borel measurable functions $f : \mathbb{R} \rightarrow \mathbb{R}$ satisfying

$$\|f\|_{\mathcal{H}_0}^2 := \int_{\mathbb{R}} |x|^4 (1 + e^x)^2 |f(x)|^2 dx < \infty. \quad (5.3)$$

- \mathcal{H} is a Hilbert space of absolutely continuous functions $f : [0, \bar{T}] \rightarrow \mathcal{H}_0$ satisfying

$$\|f\|_{\mathcal{H}}^2 := \|f(0)\|_{\mathcal{H}_0}^2 + \int_0^{\bar{T}} \left\| \frac{d}{du} f(u) \right\|_{\mathcal{H}_0}^2 du < \infty. \quad (5.4)$$

- $C([0, \bar{T}])$ is a Banach space of continuous functions $f : [0, \bar{T}] \rightarrow \mathbb{R}$ satisfying

$$\sup_{x \in [0, \bar{T}]} |f(x)| < \infty. \quad (5.5)$$

- $W^{1,2}([0, \bar{T}])$ is a Hilbert space of absolutely continuous functions $f : [0, \bar{T}] \rightarrow \mathbb{R}$ satisfying

$$|f(0)|^2 + \int_0^{\bar{T}} \left| \frac{d}{du} f(u) \right|^2 du < \infty. \quad (5.6)$$

¹⁸Ideally, we would like to be able to propose an algorithm for the associated optimization problem that is guaranteed to converge at a known rate.

¹⁹In fact, one has to go beyond Lévy processes to do this. For example, it is not hard to find a combination of arbitrage-free prices of three call options, with the same maturity and different strikes, which cannot be approximated with an arbitrary precision (simultaneously) by any exponential Lévy model. This means that, in principle, if there are more than two strikes traded in the market, the associated call prices (even with the bid-ask spreads) may be such that there is no tangent Lévy model that can match them.

- \mathcal{B}_d is a Banach space of absolutely continuous functions $f : [0, \bar{T}] \rightarrow \mathbb{R}$ satisfying

$$\|f\|_{\mathcal{B}_d} := |f(0)| + \int_0^{\bar{T}} \left| \frac{d}{du} f(u) \right| du < \infty. \quad (5.7)$$

Here the subscript d is used to indicate the “discrete” models.

- \mathcal{H}_d is the Hilbert space of absolutely continuous functions $f : [0, \bar{T}] \rightarrow \mathbb{R}$ satisfying

$$\|f\|_{\mathcal{H}_d}^2 := |f(0)|^2 + \int_0^{\bar{T}} \left| \frac{d}{du} f(u) \right|^2 du < \infty. \quad (5.8)$$

We know that $\mathcal{H}_0 \subset \mathcal{B}_0$, $\mathcal{H} \subset \mathcal{B}$, $W^{1,2}([0, \bar{T}]) \subset C([0, \bar{T}])$ and $\mathcal{H}_d \subset \mathcal{B}_d$. In addition, it is not hard to see that the completion of \mathcal{H}_0 is \mathcal{B}_0 with respect to the norm $\|\cdot\|_{\mathcal{B}_0}$. Similarly, the completion of \mathcal{H} is \mathcal{B} with respect to $\|\cdot\|_{\mathcal{B}}$, the completion of $W^{1,2}([0, \bar{T}])$ is $C([0, \bar{T}])$ with respect to the “sup” norm, and the completion of \mathcal{H}_d is \mathcal{B}_d with respect to the $\|\cdot\|_{\mathcal{B}_d}$ norm. Hence, we conclude that the couples $(\mathcal{H}, \mathcal{B})$, $(W^{1,2}([0, \bar{T}]), C([0, \bar{T}]))$, and $(\mathcal{H}_d, \mathcal{B}_d)$ are all conditional Banach spaces (see III 5.3 in [18] for definition).

6 Appendix B

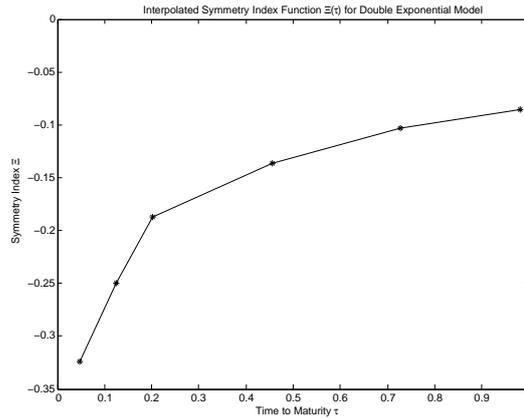


Figure 1: Symmetry index Ξ as a function of time to maturity, in DETL model

Table 1: Time periods

	Jan. 2007 - Aug. 2008	Jan. 2011 - Dec. 2012
# of days	419	502
Range of SPX spot price	\$1214.9 - \$1565.2	\$1099.2 - \$1465.8

Table 2: Testing periods

Period	Training period	Testing period
I	Jan. 2007 - Dec. 2007	Jan. 2008 - Aug. 2008
II	Jan. 2011 - Dec. 2011	Jan. 2012 - Dec. 2012

Table 3: Strikes used in each portfolio

# of strikes	Strikes used
5	$K_{i-3}(\text{call}), K_{i-1}(\text{call}), K_{i+1}(\text{call}), K_{i+3}(\text{call}), K_{i+5}(\text{call})$
4	$K_{i-3}(\text{call}), K_{i-1}(\text{call}), K_{i+1}(\text{call}), K_{i+3}(\text{call})$
3	$K_{i-1}(\text{call}), K_{i+1}(\text{call}), K_{i+3}(\text{call})$

Table 4: Average deviation of (C + S) portfolio in period I

	# of strikes	DETL	SABR ($\beta = 1$)	SABR ($\beta = 0.7$)
Average realized deviation Q	5	0.55%	84.97%	111.42%
	4	0.54%	4.69%	24.43%
	3	0.64%	2.18%	10.50%
Average predicted deviation P	5	0.87%	0.19%	9.33%
	4	0.88%	0.30%	9.66%
	3	1.05%	0.53%	10.29%

Table 5: Average quantity oscillation K (as defined in (3.12)) in (C + S) portfolio in Period I

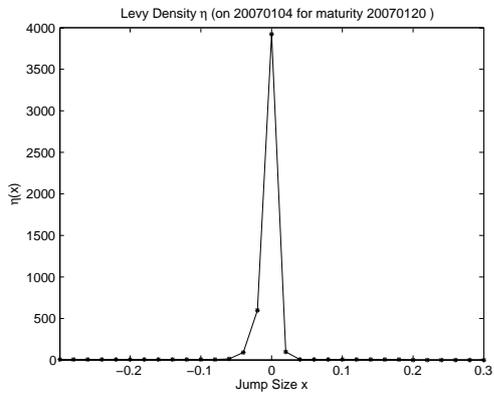
# of strikes	DETL	SABR ($\beta = 1$)	SABR ($\beta = 0.7$)
5	0.0039	1.1747	2.3846
4	0.0038	0.1339	0.4629
3	0.0027	0.0263	0.0807

Table 6: Average deviation of (C + S) portfolio in Period II

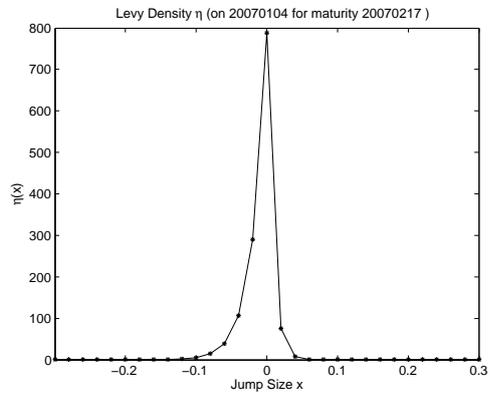
	# of strikes	DETL	SABR ($\beta = 1$)	SABR ($\beta = 0.7$)	Direct Simulation
Average realized deviation Q	5	0.41%	9.07%	33.22%	0.6%
	4	0.42%	3.51%	17.61%	-
	3	0.42%	0.90%	5.22%	-
Average predicted deviation P	5	0.79%	0.36%	7.98%	0.43%
	4	0.79%	0.43%	8.11%	-
	3	0.94%	0.62%	8.46%	-

Table 7: Average quantity oscillation K (as defined in (3.12)) of (C + S) portfolio with 5 strikes in Period II

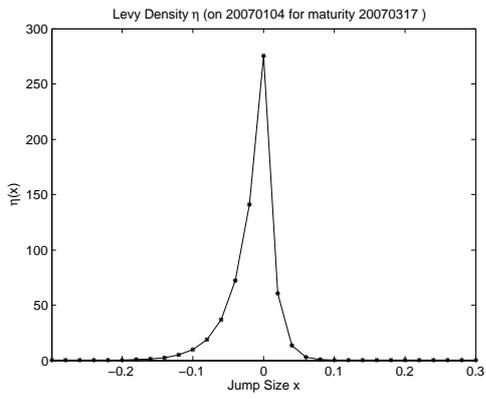
# of strikes	DETL	SABR ($\beta = 1$)	SABR ($\beta = 0.7$)	Direct Simulation
5	0.0011	0.1410	0.6642	0.0065
4	0.0012	0.0537	0.2736	-
3	0.0011	0.0145	0.0474	-



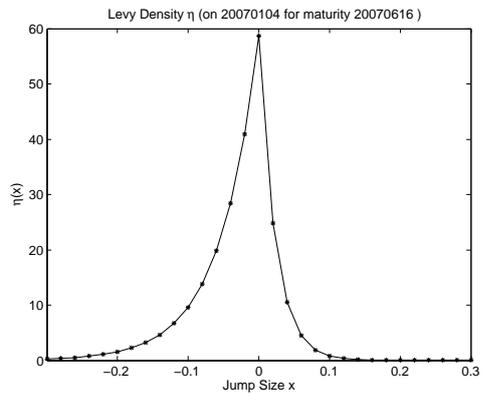
(a) Maturity Jan.20,2007



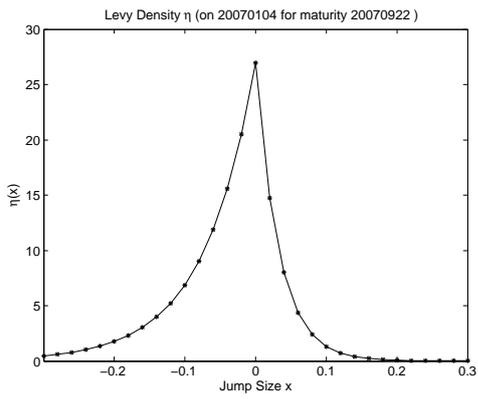
(b) Maturity Feb.17,2007



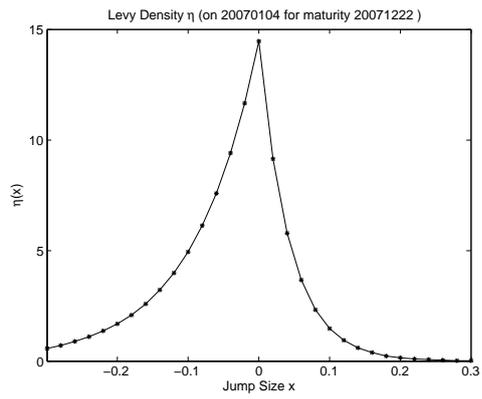
(c) Maturity Mar.17,2007



(d) Maturity Jun.16,2007

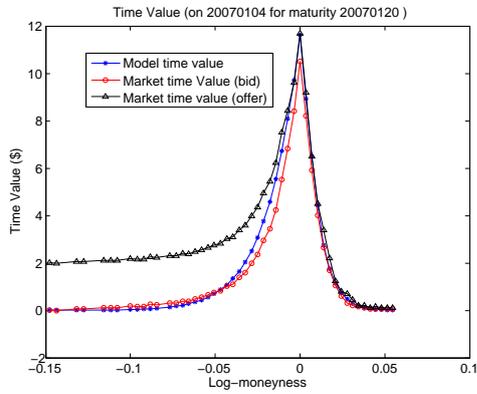


(e) Maturity Sep.22,2007

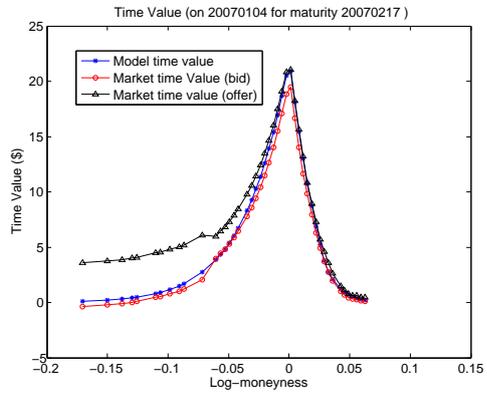


(f) Maturity Dec.22,2007

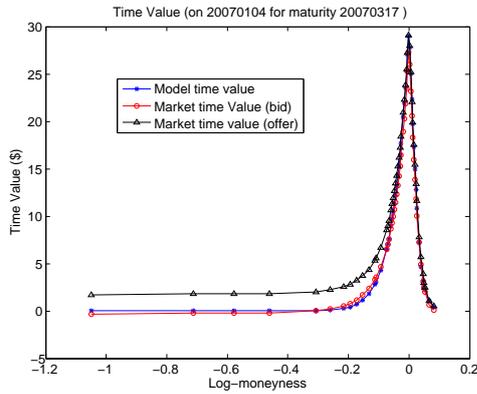
Figure 2: Calibrated densities η for DETL model on the second day, Jan. 4, 2007



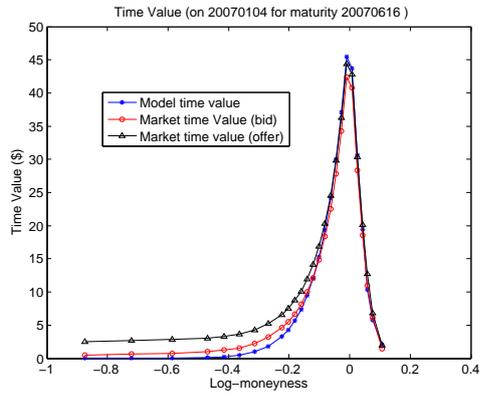
(a) Maturity Jan.20,2007



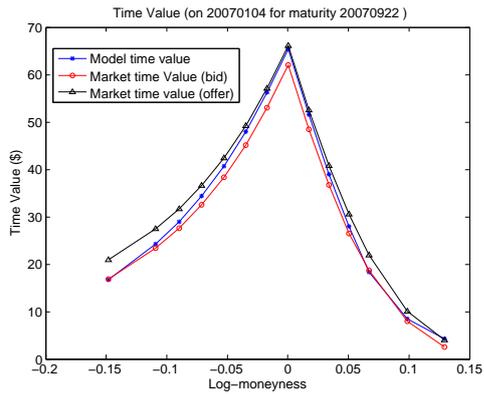
(b) Maturity Feb.17,2007



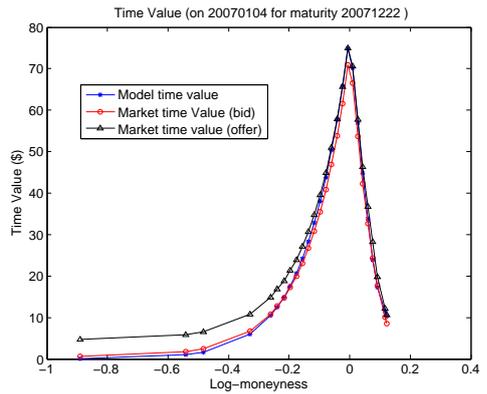
(c) Maturity Mar.17,2007



(d) Maturity Jun.16,2007

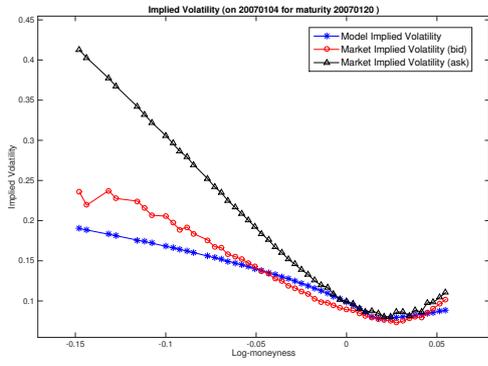


(e) Maturity Sep.22,2007

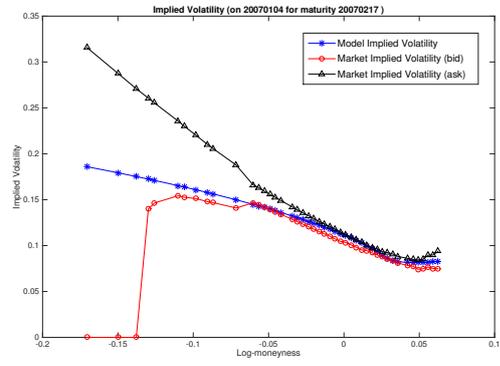


(f) Maturity Dec.22,2007

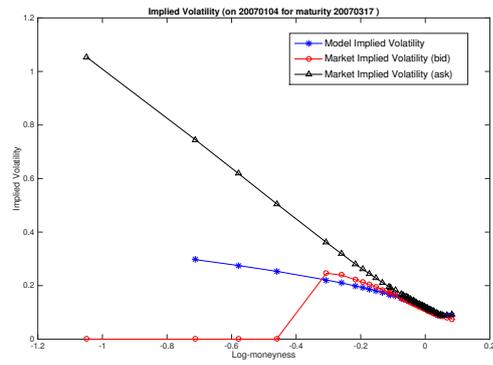
Figure 3: Calibrated time values for DETL model on the second day, Jan. 4, 2007



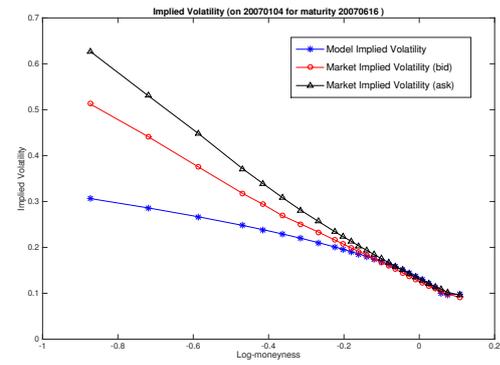
(a) Maturity Jan.20,2007



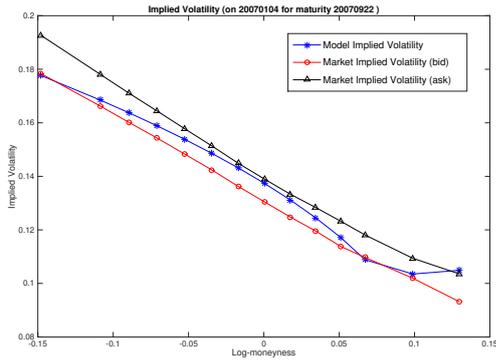
(b) Maturity Feb.17,2007



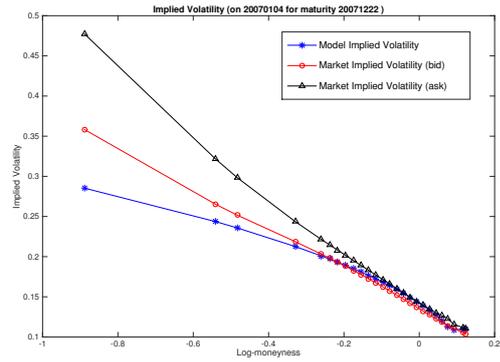
(c) Maturity Mar.17,2007



(d) Maturity Jun.16,2007

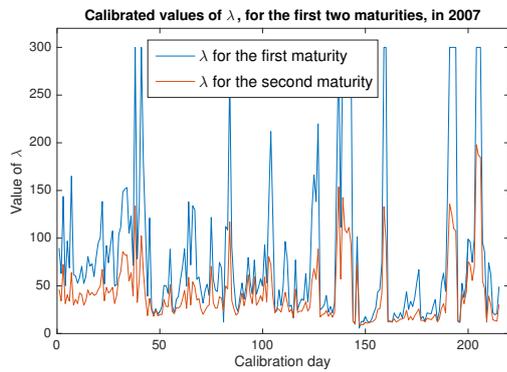


(e) Maturity Sep.22,2007

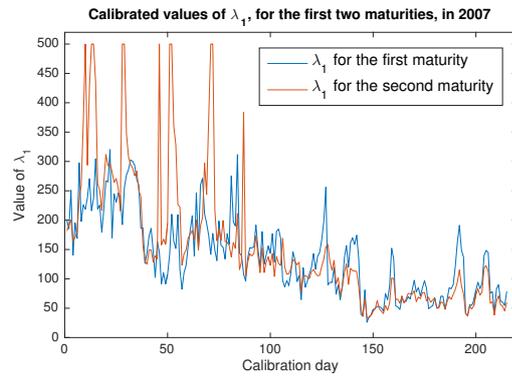


(f) Maturity Dec.22,2007

Figure 4: Calibrated Implied Volatilities for DETL model on the second day, Jan. 4, 2007

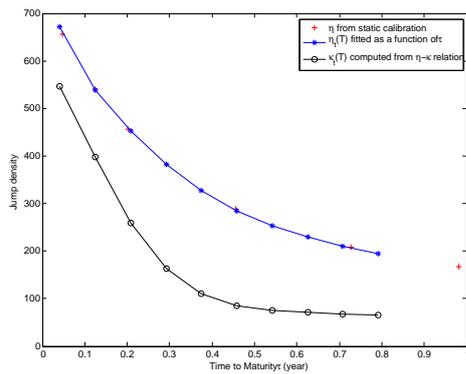


(a) Calibrated values of λ

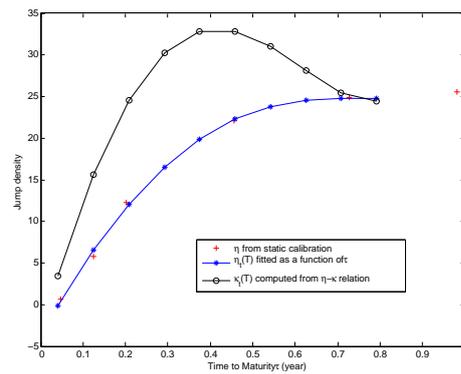


(b) Calibrated values of λ_1

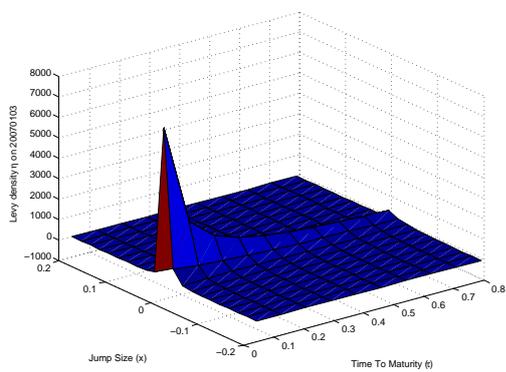
Figure 5: Calibrated values of λ and λ_1 , for the first two maturities, in Jan. – Dec. 2007



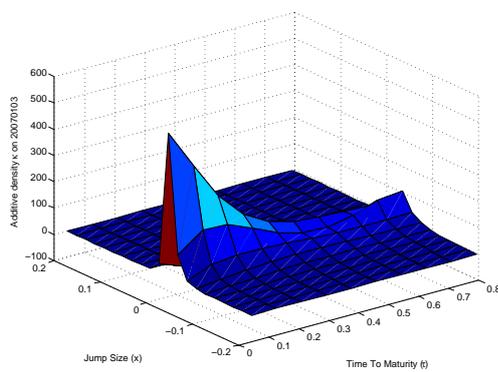
(a) η and κ for small jumps ($x = -0.02$)



(b) η and κ for large jumps ($x = -0.08$)

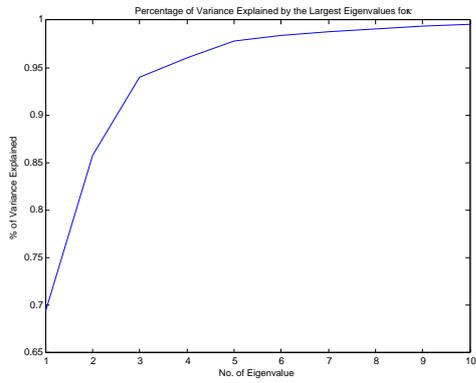


(c) Lévy density η

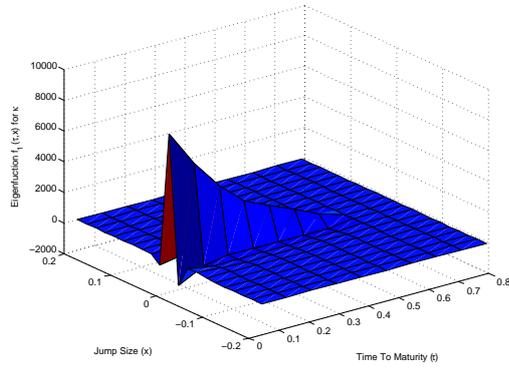


(d) Additive density κ reconstructed from η in (c)

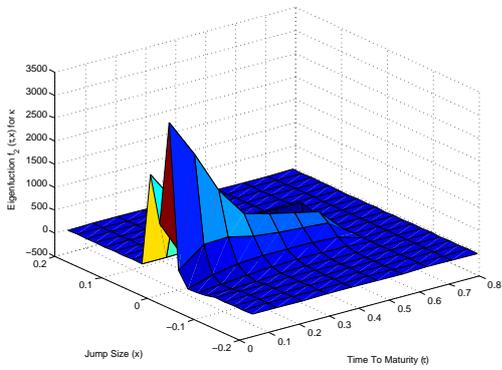
Figure 6: Calculating Lévy density κ from η



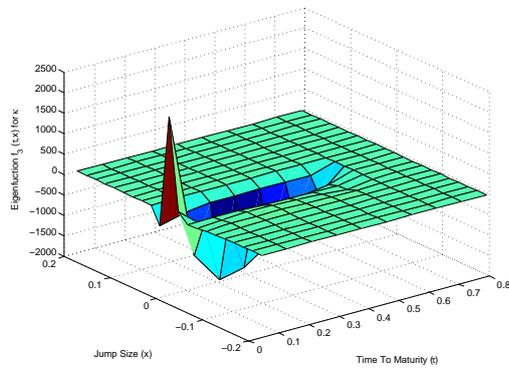
(a) Percentage of variance explained by the eigenmodes



(b) The first eigenmode scaled by $\sqrt{\lambda_1}$



(c) The second eigenmode scaled by $\sqrt{\lambda_2}$



(d) The third eigenmode scaled by $\sqrt{\lambda_3}$

Figure 7: Eigenvalues and eigenmodes of $\Delta \hat{\kappa}$ for DETL model

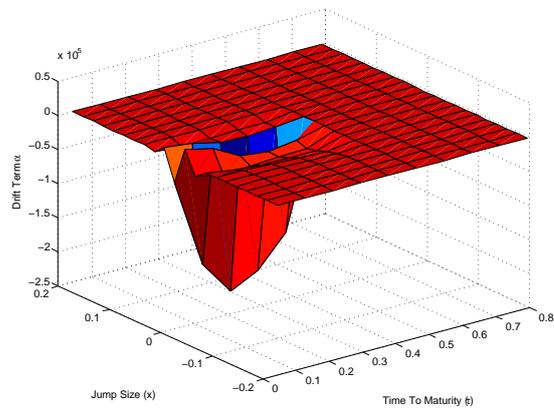
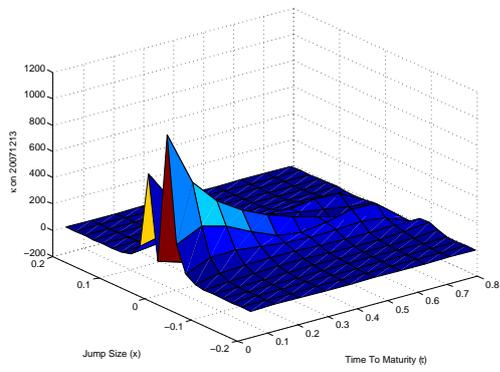
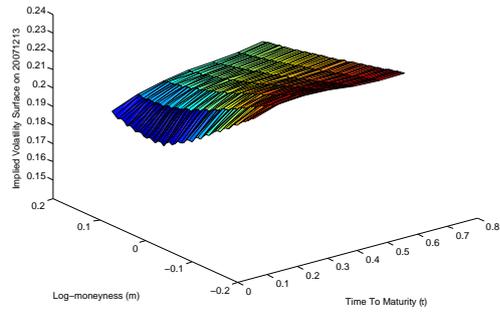


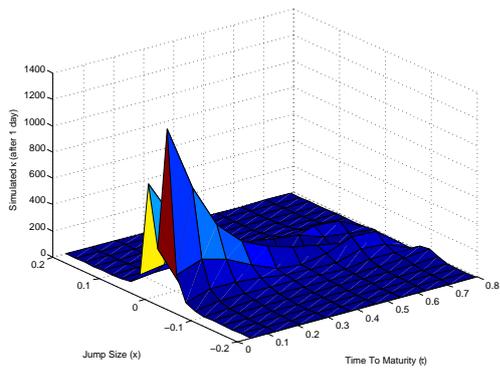
Figure 8: The drift term α for DETL model



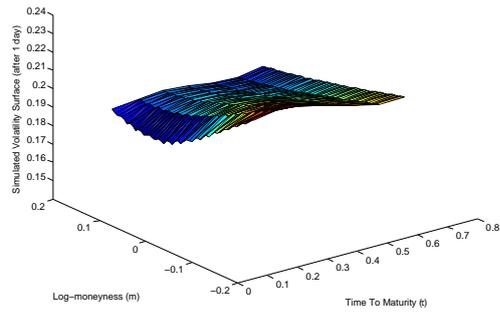
(a) Calibrated κ



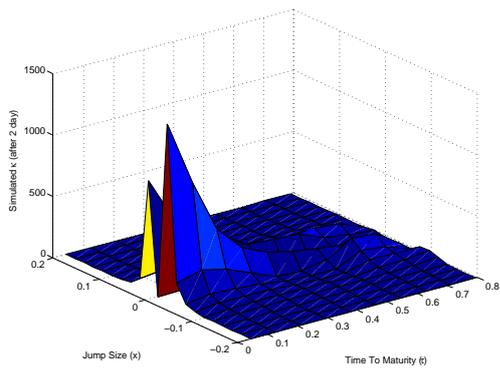
(b) Calibrated Implied Volatility Surface



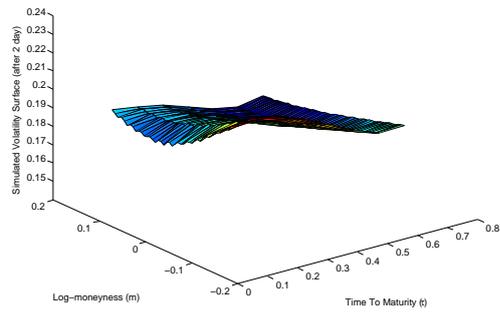
(c) Simulated κ (1st day)



(d) Simulated Implied Volatility Surface (1st day)

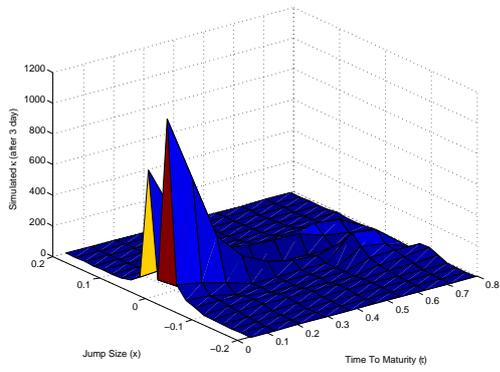


(e) Simulated κ (2nd day)

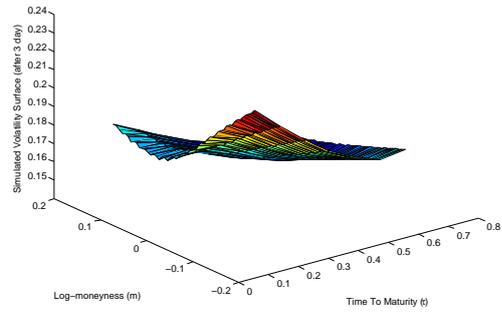


(f) Simulated Implied Volatility Surface (2nd day)

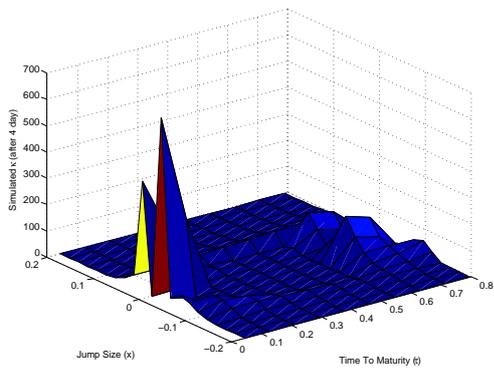
Figure 9: Simulated κ 's and implied volatility surfaces using DETL model (1)



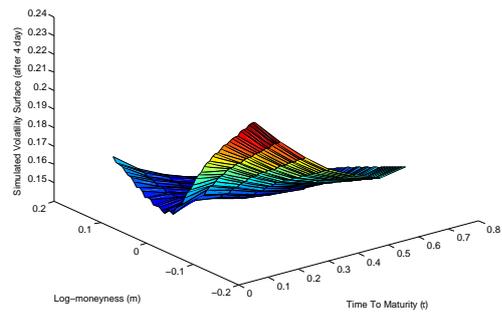
(a) Simulated κ (3rd day)



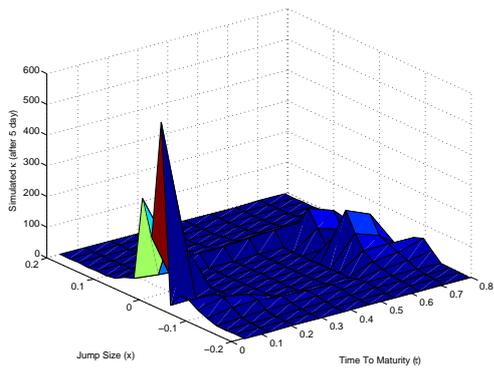
(b) Simulated Implied Volatility Surface (3rd day)



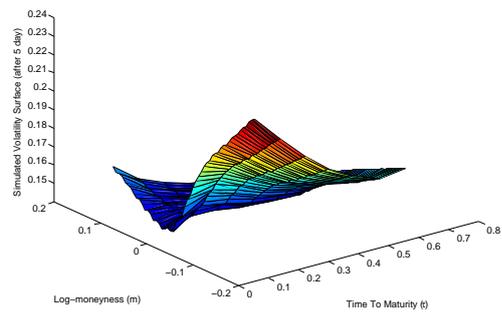
(c) Simulated κ (4th day)



(d) Simulated Implied Volatility Surface (4th day)

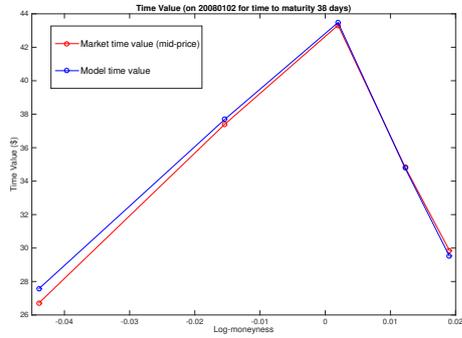


(e) Simulated κ (5th day)

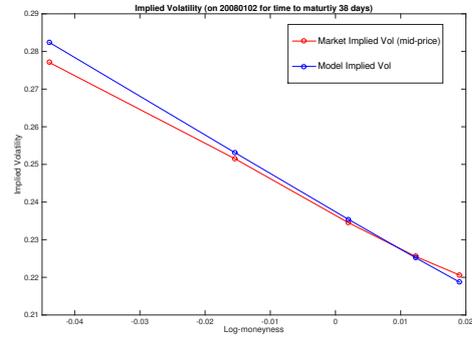


(f) Simulated Implied Volatility Surface (5th day)

Figure 10: Simulated κ 's and implied volatility surfaces using DETL model (2)

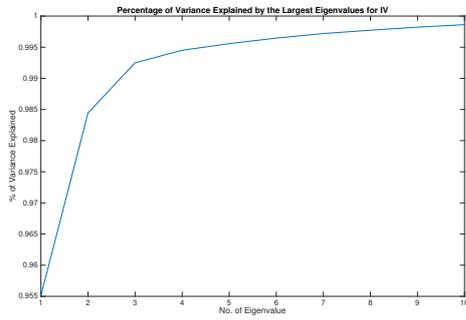


(a) Time values of the market and model prices

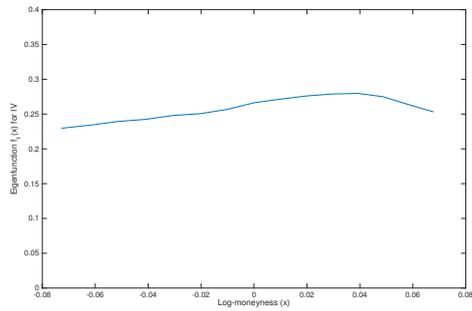


(b) Implied Volatilities of the market and model prices

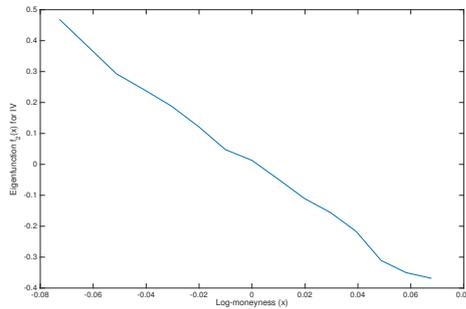
Figure 11: Results of calibrating SABR model with $\beta = 1$ to the market prices of SPX options on Jan. 2, 2008



(a) Percentage of variance explained by the eigenmodes

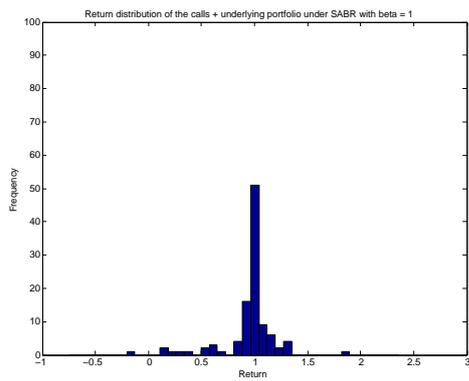


(b) The first eigenmode scaled by $\sqrt{\lambda_1}$

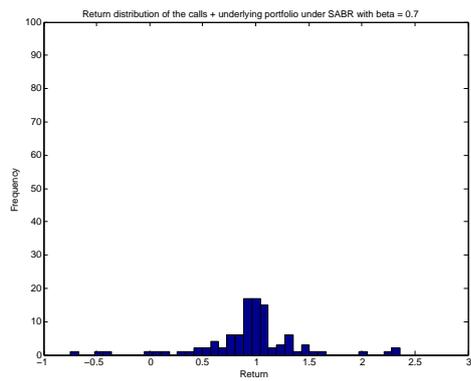


(c) The second eigenmode scaled by $\sqrt{\lambda_2}$

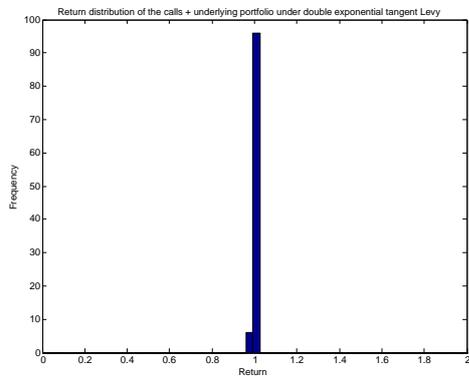
Figure 12: Eigenvalues and eigenmodes of $\Delta \log IV$



(a) Under SABR model with $\beta = 1$

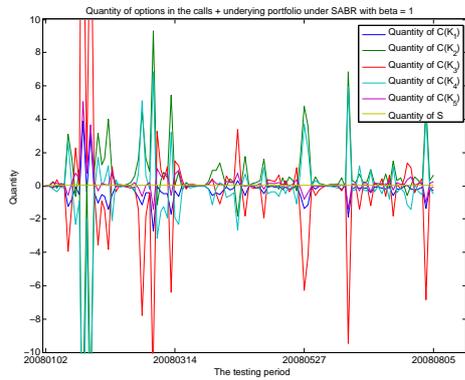


(b) Under SABR model with $\beta = 0.7$

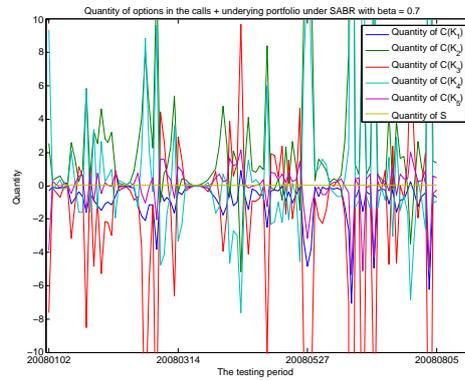


(c) Under double exponential tangent Lévy model

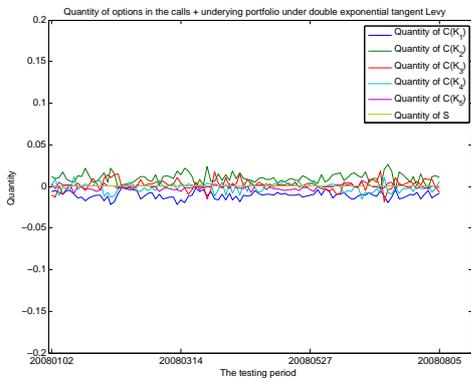
Figure 13: Distribution of the 8-day returns of (C + S) portfolio with 5 strikes in Period I. Different scales are used to show more details.



(a) Under SABR model with $\beta = 1$

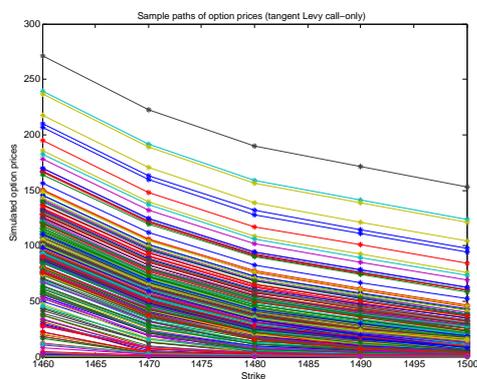


(b) Under SABR model with $\beta = 0.7$

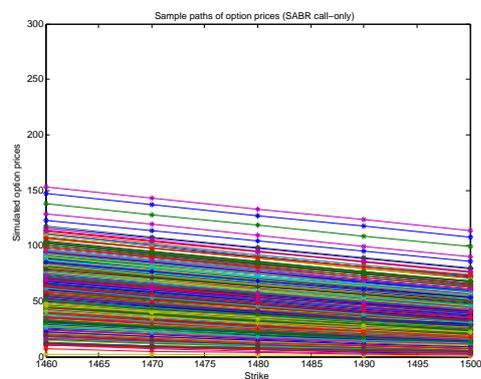


(c) Under double exponential tangent Lévy model

Figure 14: Option quantities in (C + S) portfolio with 5 strikes in Period I. Different scales are used to show more details

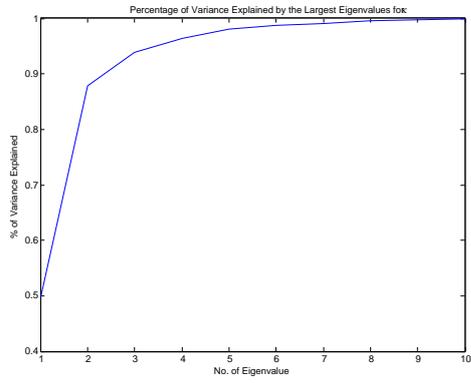


(a) Under double exponential tangent Lévy model

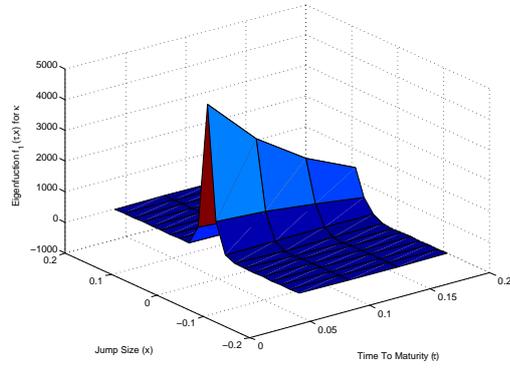


(b) Under SABR model with $\beta = 1$

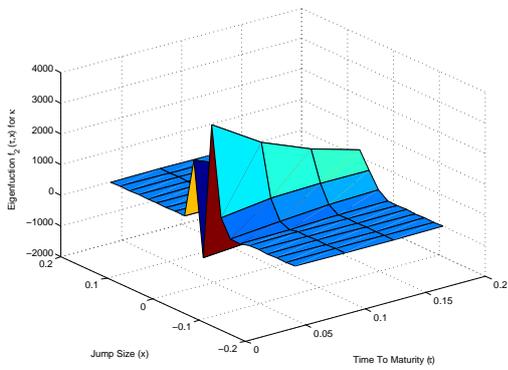
Figure 15: Terminal option prices in (C + S) portfolio, as functions of strike, simulated using 500 sample paths



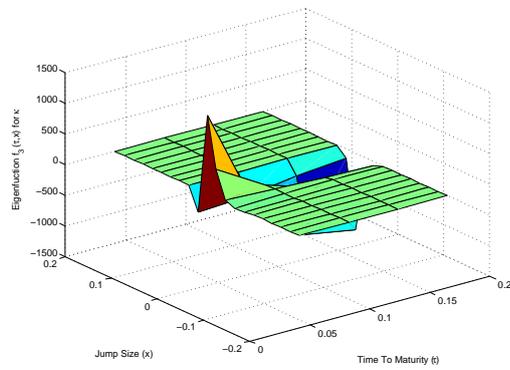
(a) Percentage of variance explained by the first eigenmodes



(b) The first eigenmode scaled by $\sqrt{\lambda_1}$



(c) The second eigenmode scaled by $\sqrt{\lambda_2}$



(d) The third eigenmode scaled by $\sqrt{\lambda_3}$

Figure 16: Eigenvalues and eigenmodes of $\Delta \hat{k}$ under DETL, estimated using 2011 data

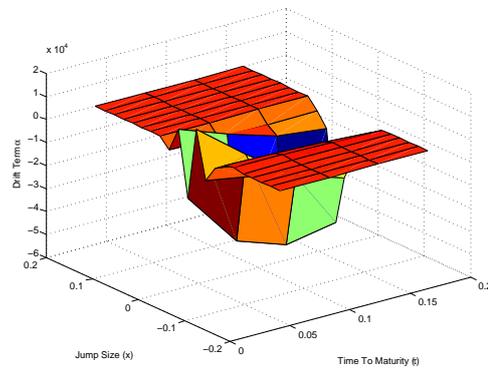
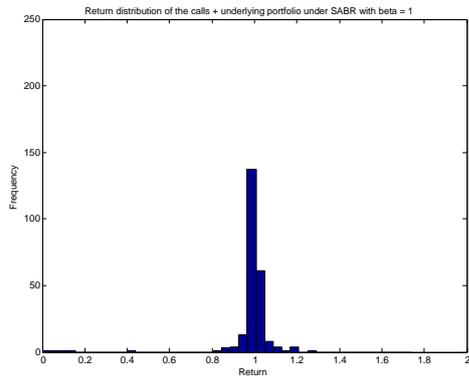
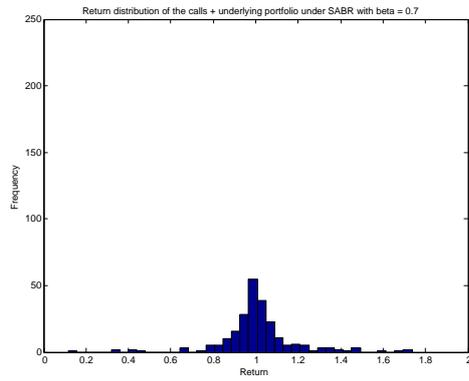


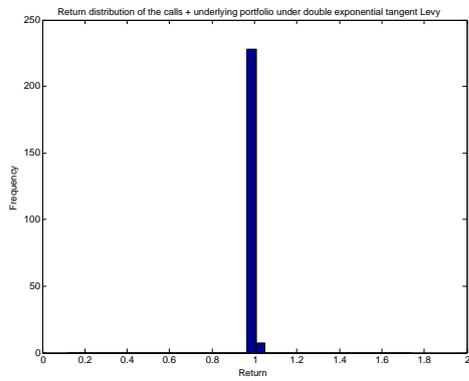
Figure 17: The drift term α in DETL model, estimated using 2011 data



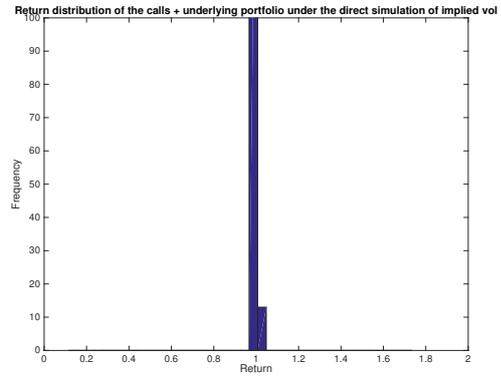
(a) Under SABR model with $\beta = 1$



(b) Under SABR model with $\beta = 0.7$

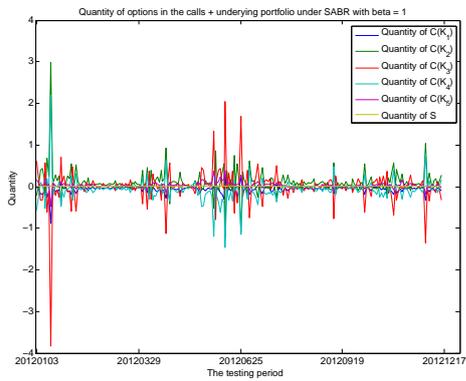


(c) Under double exponential tangent Lévy model

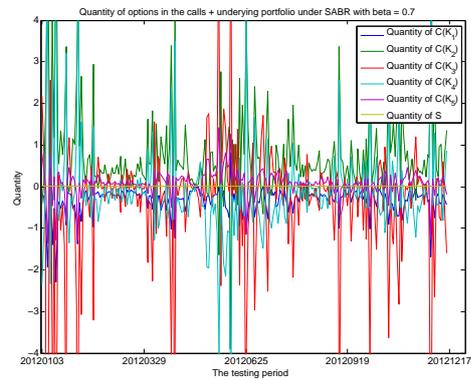


(d) Under the direct simulation of implied volatility

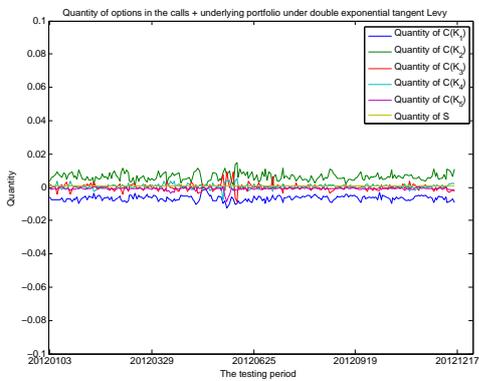
Figure 18: Distribution of the 8-day returns of (C + S) portfolio with 5 strikes in Period II



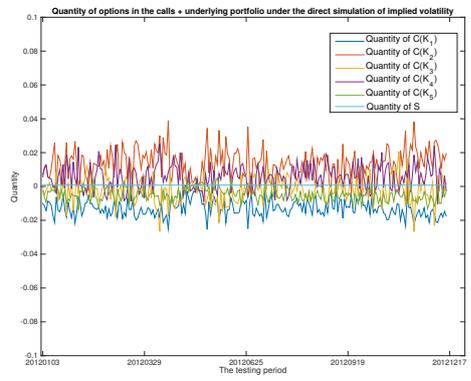
(a) Under SABR model with $\beta = 1$



(b) Under SABR model with $\beta = 0.7$

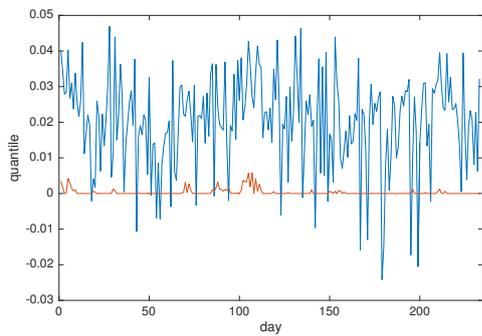


(c) Under double exponential tangent Lévy model

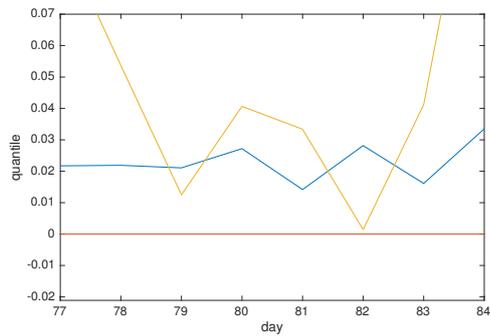


(d) Under the direct simulation of implied volatility

Figure 19: Option quantities in (C + S) portfolio with 5 strikes in Period II. Different scales are used to show more details



(a) Quantiles produced by the simulations



(b) Quantiles and realized portfolio values

Figure 20: Estimated 0.01-quantiles of the future (16 days ahead) values of the “digital spread” portfolio, produced by the direct simulation (blue) and by DETL model (red), on every day in 2012. Part (b) contains the same quantiles on a smaller time interval, with the realized future (16 days ahead) values of the portfolio (yellow).

References

- [1] Y. Ait-Sahalia and A. Lo. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance*, 53:499–547, 1997.
- [2] R. Carmona and S. Nadtochiy. Local volatility dynamic models. *Finance and Stochastics*, 13(1):1–48, 2009.
- [3] R. Carmona and S. Nadtochiy. Tangent models as a mathematical framework for dynamic calibration. *International Journal of Theoretical and Applied Finance (IJTAF)*, 14(01):107–135, 2011.
- [4] R. Carmona and S. Nadtochiy. Tangent Lévy market models. *Finance and Stochastics*, 16(1):63–104, 2012.
- [5] P. Carr and D. Madan. Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2:61–73, 1999.
- [6] P. Carr and S. Nadtochiy. Local variance gamma and explicit calibration to option prices. *Mathematical Finance*, published online, DOI: 10.1111/mafi.12086, 2014.
- [7] R. Cont and J. da Fonseca. Dynamics of implied volatility surfaces. *Quantitative Finance*, 2:45–60, 2002.
- [8] R. Cont, J. da Fonseca, and V. Durrleman. Stochastic models of implied volatility surfaces. *Economic Notes*, 31(2):361–377, 2002.
- [9] R. Cont and P. Tankov. *Financial modelling with jump processes*. Chapman & Hall/CRC, 2004.
- [10] L. Cousot. Conditions on option prices for absence of arbitrage and exact calibration. *Journal of Banking & Finance*, 31(11):3377 – 3397, 2007.
- [11] E. Derman and I. Kani. Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility. *International Journal of Theoretical and Applied Finance*, 01(01):61–110, 1998.
- [12] D. Filipovic, L.P. Hughston, and A. Macrina. Conditional density models for asset pricing. *International Journal of Theoretical and Applied Finance (IJTAF)*, 15:1–24, 2012.
- [13] P. Hagan, D. Kumar, A. Lesniewski, and D. Woodward. Managing smile risk. *Wilmott Magazine*, pages 84–108, 2002.
- [14] D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica: Journal of the Econometric Society*, pages 77–105, 1992.
- [15] Jan Kallsen and Paul Krühner. On a Heath-Jarrow-Morton approach for stock options. Technical report, Kiel University, 2010.
- [16] S. Karlsson. *Consistent dynamic equity market code-books from a practical point of view*. PhD thesis, University of Vienna, 2011.
- [17] S. Kou. A jump-diffusion model for option pricing. *Management Science*, 48:1086–1101, August 2002.
- [18] H. Kuo. *Gaussian Measures in Banach Spaces*. Springer-Verlag Berlin, 1975.
- [19] E. Leclercq. *Three essays on asset pricing*. PhD thesis, École Polytechnique Fédérale de Lausanne, 2014.
- [20] R. Lee. Implied volatility: Statics, dynamics, and probabilistic interpretation. In *Recent Advances in Applied Probability*, pages 241–268. Springer, 2005.
- [21] A. Lewis. A simple option formula for general jump-diffusion and other exponential Lévy processes. <http://optioncity.net/pubs/ExpLevy.pdf/>.
- [22] D. Luenberger. *Investment Science*. Oxford University Press, 1997.
- [23] R. Rebonato. *Volatility and Correlation: The Perfect Hedger and the Fox*. John Wiley & Sons, Ltd., 2004.
- [24] A. Richter and J. Teichmann. Discrete time term structure theory and consistent recalibration models. *arXiv:1409.1830*, preprint, 2014.

- [25] P. Schönbucher. A market model for stochastic implied volatility. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 357(1758):2071–2092, 1999.
- [26] P. Schönbucher. Portfolio losses and the term structure of loss transition rates: a new methodology for the pricing of portfolio credit derivatives. Technical report, ETH Zürich, December 2005.
- [27] M. Schweizer and J. Wissel. Arbitrage-free market models for option prices: The multi-strike case. *Finance and Stochastics*, 12(4):469–505, 2008.
- [28] M. Schweizer and J. Wissel. Term structures of implied volatilities: Absence of arbitrage and existence results. *Mathematical Finance*, 18(1):77–114, 2008.
- [29] G. West. Calibration of the SABR model in illiquid markets. *Applied Mathematical Finance*, 12(4):371–385, 2005.
- [30] J. Zhao. Parametric arbitrage-free models for implied smile dynamics. Master’s thesis, University of Oxford, June 2010.