Is the minimum value of an option on variance generated by local volatility?

Mathias Beiglböck[*], Peter Friz[†] and Stephan Sturm[‡]

January 22, 2010

Abstract

We discuss the possibility of obtaining model-free bounds on volatility derivatives, given present market data in the form of a calibrated local volatility model. A counter-example to a wide-spread conjecture is given.

keywords: local vol, Dupire's formula; MSC: 91G99; JEL: G10.

1 Introduction

"... *it has been conjectured that the minimum possible value of an option on variance is the one generated from a local volatility model fitted to the volatility surface.*"; Gatheral [Gat06, page 155].

Leaving precise definitions to below, let us clarify that an *option on variance* refers to a derivative whose payoff is a convex function f of *total realized variance*. Turning from convex to concave, this conjecture, if true, would also imply that that the maximum possible value of a *volatility swap* ($f(x) = x^{1/2}$) is the one generated from a local volatility model fitted to the volatility surface. Given the well-documented model-risk in pricing volatility swaps, such bounds are of immediate practical interest.

The mathematics of local volatility theory (à la Dupire, Derman, Kani, ...) is intimately related to the following

Theorem 1 ([Gyö86])**.** *Assume $dY_t = \mu(t,\omega)\,dt + \sigma(t,\omega)\,dB_t$ is a multi-dimensional Itô-diffusion where μ, σ are progressively measurable, bounded and $\sigma\sigma^T \geq \varepsilon^2 I$ for some $\varepsilon > 0$. Then*

$$d\tilde{Y}_t = \mu_{loc}\left(t, \tilde{Y}_t\right) dt + \sigma_{loc}\left(t, \tilde{Y}_t\right) dB_t$$

has a unique weak solution, where

$$\begin{aligned} \mu_{loc}(t,y) &= E\left[\mu(t,\omega) \,|Y_t = y\right], \\ \sigma^2_{loc}(t,y) &= E\left[\sigma^2(t,\omega) \,|Y_t = y\right], \end{aligned}$$

and $\tilde{Y}_t \overset{law}{=} Y_t$ for all fixed t.

[*]University of Vienna, partially supported by the FWF under grant P21209.
[†]TU and WIAS Berlin, partially supported by MATHEON.
[‡]TU Berlin, supported by MATHEON.

A generic stochastic vol model (already written under the appropriate equivalent martingale measure and with suitable choice of numéraire) is of the form $dS = S\sigma dB$ where $\sigma = \sigma(t, \omega)$ is the (progressively measurable) instantenous volatility process. (It will suffice for our application to assume σ to be bounded from above and below by positive constants.) Arguing on log-price $X = \log S$ rather than S,

$$dX_t = \sigma(t, \omega)\, dB_t - \left(\sigma^2(t, \omega)/2\right) dt, \tag{1.1}$$

a classical application of theorem 1 yields the following Markovian projection result[1]: the (weak) solution to

$$\begin{aligned} d\tilde{X}_t &= \sigma_{loc}\left(t, \tilde{X}\right) dB_t - \left(\sigma_{loc}^2\left(t, \tilde{X}_t\right)/2\right) dt, & (1.2)\\ \text{(with } \sigma_{loc}^2(t, x) &= E\left[\sigma^2(t, \omega)\,|X_t = x\right]\text{)} & (1.3) \end{aligned}$$

has the one-dimensional marginals of the original process X_t. Equivalently[2], the process $\tilde{S} = \exp \tilde{X}$,

$$d\tilde{S}_t = \sigma_{loc}(t, \tilde{S}_t)\tilde{S}_t\, dB_t,$$

known as (Dupire's) *local volatility model*, gives rise to identical prices of all European options $C(T, K)$. It easily follows that $\sigma_{loc}^2(t, \tilde{S})$ is given by *Dupire's formula*

$$\sigma_{loc}^2(T, \tilde{S})|_{\tilde{S}=K} = 2\frac{\partial_T C}{K^2 \partial_{KK} C}. \tag{1.4}$$

Volatility derivatives are options on realized variance; that is, the payoff is given by some function f of realized variance. The latter is given by

$$W_T := \langle \log S \rangle_T = \langle X \rangle_T = \int_0^T \sigma^2(t, \omega)\, dt;$$

in the model $dS = \sigma(t, \omega)\, S\, dB$ and by

$$\tilde{W}_T := \langle \log \tilde{S} \rangle_T = \langle \tilde{X} \rangle_T = \int_0^T \sigma_{loc}^2(t, \tilde{X}_t)\, dt$$

in the corresponding local volatility model.

Common choices of f are $f(x) = x$, the *variance swap*, $f(x) = x^{1/2}$, the volatility swap, or simply $f(x) = (x - K)^+$, a *call-option on realized variance.* See [FG05] for instance. As is well-known, see e.g. [Gat06], the pricing of a variance swap, assuming continuous dynamics of S such as those specified above, is model free in the sense that it can be priced in terms of a log-contract; that is, a European option with payoff $\log S_T$. In particular, it follows that

$$E\left[\tilde{W}_T\right] = E\left[W_T\right].$$

Of course this can also be seen from (1.3), after exchanging E and integration over $[0, T]$. Passing from W_T to $f(W_T)$ for general f this is not true, and the resulting differences are known in the industry as *convexity adjustment.* We can now formalize the conjecture given in the first lines of the introduction[3].

[1] Let us quickly remark that Markovian projection techniques have led recently to a number of new applications (see [Pit06], for instance).

[2] The abuse of notation, by writting both $\sigma_{loc}\left(t, \tilde{X}\right)$ and $\sigma_{loc}\left(t, \tilde{S}\right)$, will not cause confusion.

[3] It is tacitly assumed that $f(W_T), f\left(\tilde{W}_T\right)$ are integrable.

Conjecture 1. *For any convex f one has* $E\big[f\big(\tilde{W}_T\big)\big] \le E\big[f\big(W_T\big)\big]$.

Our contribution is twofold: first we discuss a simple (toy) example which provides a counterexample to the above conjecture; secondly we refine our example using a 2-dimensional Markovian projection (which may be interesting in its own right) and thus construct a perfectly sensible Markovian stochastic volatility model in which the conjectured result fails. All this narrows the class of possible dynamics for S for which the conjecture can hold true and so should be a useful step towards positive answers.

2 Idea and numerical evidence

Example 2. *Consider a Black–Scholes "mixing" model $dS = S\sigma dB, S_0 = 1$ with time horizon $T = 3$ in which $\sigma^2(t,\omega)$ is given by $\sigma_+^2(t)$ or $\sigma_-^2(t)$,*

$$\sigma_+^2(t) := \begin{cases} 2 & \textit{if } t \in [0,1], \\ 3 & \textit{if } t \in]1,2], \\ 1 & \textit{if } t \in]2,3], \end{cases} \qquad \sigma_-^2(t) := \begin{cases} 2 & \textit{if } t \in [0,1], \\ 1 & \textit{if } t \in]1,2], \\ 3 & \textit{if } t \in]2,3], \end{cases}$$

depending on a fair coin flip $\epsilon = \pm 1$ (independent of B). Obviously $W = W_3 = \int_0^3 \sigma^2\,dt \equiv 6$ in this example, hence $E\big[(W-6)^+\big] = (W-6)^+ = 0$. On the other hand, the local volatility is explicitly computable (cf. the following section) and one can see from simple Monte Carlo simulations that for $\tilde{W} = \tilde{W}_3$

$$E\left[(\tilde{W} - 6)^+\right] \approx 0.026 > 0$$

thereby (numerically) contradicting conjecture 1, with $f(x) = (x-6)^+$.

Our analysis of this toy model is simple enough: in section 3 below we prove that $P[\tilde{W} = 6] \neq 1$. Since $E[\tilde{W}] = E[W] = 6$ and $(x-6)^+$ is strictly convex at $x = 6$, Jensen's inequality then tells us that $E[(\tilde{W}-6)^+] > 0 = E\left[(W-6)^+\right]$.

3 Analysis of the toy example

We recall that it suffices to show that $\tilde{W} = \int_0^3 \sigma_{loc}^2(t, \tilde{X}_t)\,dt$ is not a.s. equal to $W \equiv 6$. The distribution of X_t is simply the mixture of two normal distributions. More explicitly, $X_t = I_{\{\epsilon=+1\}} X_{t,+} + I_{\{\epsilon=-1\}} X_{t,-}$,

$$X_{t,\pm} = \int_0^t \sigma_\pm(s)\,dB_s + \tfrac{1}{2}\int_0^t \sigma_\pm^2(s)\,ds \quad \sim \quad N\Big(\tfrac{1}{2}\Sigma_\pm(t), \Sigma_\pm(t)\Big),$$

where $\Sigma_\pm(t) := \int_0^t \sigma_\pm^2(s)\,ds$. Thus $\sigma_{loc}^2(t,x) = E[\sigma^2(t,\omega)|X_t = x]$ is given by[4]

$$\sigma_{loc}^2(t,x) = \frac{\frac{\sigma_+^2(t)}{\sqrt{\Sigma_+(t)}}\exp\left[-\frac{(x+\Sigma_+(t)/2)^2}{2\Sigma_+(t)}\right] + \frac{\sigma_-^2(t)}{\sqrt{\Sigma_-(t)}}\exp\left[-\frac{(x+\Sigma_-(t)/2)^2}{2\Sigma_-(t)}\right]}{\frac{1}{\sqrt{\Sigma_+(t)}}\exp\left[-\frac{(x+\Sigma_+(t)/2)^2}{2\Sigma_+(t)}\right] + \frac{1}{\sqrt{\Sigma_-(t)}}\exp\left[-\frac{(x+\Sigma_-(t)/2)^2}{2\Sigma_-(t)}\right]}. \quad (3.1)$$

[4] More general expression for local volatility are found in [BM06, Chapter 4] and [Lee01, Lab09]. Note the necessity to keep $\sigma^2(.,\omega)$ constant on some interval $[0,\varepsilon]$, for otherwise the local vol surface is not Lipschitz in x, uniformly as $t \to 0$.

Since $\sigma_{loc} = \sigma_{loc}(s,x)$ is bounded, measurable in t and Lipschitz in x (uniformly w.r.t. t) and bounded away from zero it follows from [SV06, Theorem 5.1.1] that the SDE

$$d\tilde{X}_t = \sigma_{loc}\left(t, \tilde{X}\right) dB_t - \tfrac{1}{2}\sigma_{loc}^2\left(t, \tilde{X}_t\right) dt$$

has a unique strong solution (started from $\tilde{X}_0 = 0$, say). Since σ_{loc} is uniformly bounded away from 0 it follows that the process $(\tilde{X}_t)$ has *full support*, i.e. for every continuous $\varphi : [0,3] \to \mathbb{R}$, $\varphi(0) = 0$ and every $\varepsilon > 0$

$$P[\|\tilde{X}_t - \varphi(t)\|_{\infty;[0,T]} \leq \varepsilon] > 0.$$

Indeed, there a various ways to see this: one can apply Stroock–Varadhan's support theorem, in the form of [Pin95, Theorem 6.3] (several simplifications arise in the proof thanks to the one-dimensionality of the present problem); alternatively, one can employ localized lower heat kernel bounds (à la Fabes–Stroock [FS86]) or exploit that the Itô-map is continuous here (thanks to Doss–Sussman, see for instance [RW00, page 180]) and deduce the support statement from the full support of B.

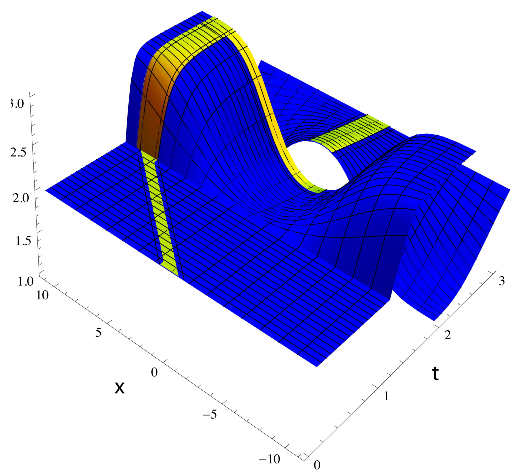

Figure 1: Time evolution of local variance ${\sigma_{loc}}^2(t,x)$ in dependence of log-moneyness. The bright strip indicates a set of paths with realized variance strictly larger than 6.

Figure 1 illustrates the dependence of ${\sigma_{loc}}^2(t,x)$ on time t and log-moneyness x. To gain our end of proving that $\tilde{W}(\omega) = \int_0^3 \sigma_{loc}^2(t, \tilde{X}_t)\, dt$ is not constantly equal 6, we can determine a set of paths $(\tilde{X}_t(\omega))$ for which $\tilde{W}$ is strictly larger than 6. In view of Figure 1 it is natural to consider paths which are large, i.e. $\tilde{X}_t(\omega) \in [8,10]$, for $t \in]1, 2 - \frac{1}{10}]$ and small, i.e. $|\tilde{X}_t(\omega)| \leq 1$, on the interval $]2,3]$. A short *mathematica*-calculation reveals that $\tilde{W}(\omega) \gtrsim 6.65 > 6$ for each such path and according to the full-support statement the set of all such paths has positive probability, hence $\tilde{W}$ is indeed not deterministic.

Using elementary analysis it is not difficult to turn numerical evidence into rigorous mathematics. Making (3.1) explicit yields that $\sigma_{loc}^2(t,x) \equiv 2$ for $t \in$

$[0,1]$ and that

$$\sigma^2_{loc}(t+1,x) = \frac{\frac{3}{\sqrt{2+3t}} e^{-\frac{(2x+2+3t)^2}{8(2+3t)}} + \frac{1}{\sqrt{2+t}} e^{-\frac{(2x+2+t)^2}{8(2+t)}}}{\frac{1}{\sqrt{2+3t}} e^{-\frac{(2x+2+3t)^2}{8(2+3t)}} + \frac{1}{\sqrt{2+t}} e^{-\frac{(2x+2+t)^2}{8(2+t)}}} \tag{3.2}$$

$$\sigma^2_{loc}(t+2,x) = \frac{\frac{1}{\sqrt{5+t}} e^{-\frac{(2x+5+t)^2}{8(5+t)}} + \frac{3}{\sqrt{3+3t}} e^{-\frac{(2x+3+3t)^2}{8(3+3t)}}}{\frac{1}{\sqrt{5+t}} e^{-\frac{(2x+5+t)^2}{8(5+t)}} + \frac{1}{\sqrt{3+3t}} e^{-\frac{(2x+3+3t)^2}{8(3+3t)}}} \tag{3.3}$$

for $t \in]0,1]$. We fix $\varepsilon \in]0,1]$ and observe that it is simple to see that $\lim_{x\to\infty} \sigma^2_{loc}(t+1,x) = 3$, uniformly w.r.t. $t \in [\varepsilon, 1]$, and that $\lim_{x\to 0} \sigma^2_{loc}(t+1,x) \geq 2$, uniformly w.r.t. $t \in]0,1]$. It follows that there exists some $\delta > 0$ such that

$$\begin{aligned} &\sigma^2_{loc}(t+1,x) \geq 3-\varepsilon \text{ for } x > \tfrac{1}{\delta}, t \in [\varepsilon,1] \text{ and} \\ &\sigma^2_{loc}(t+1,x) \geq 2-\varepsilon \text{ for } |x| < \delta, t \in]0,1]. \end{aligned}$$

Thus we obtain

$$\tilde{W}(\omega) = \int_0^3 \sigma^2_{loc}(t, \tilde{X}_t(\omega))\, dt \geq 1\cdot 2 + (1-2\varepsilon)\cdot(3-\varepsilon) + 1\cdot(2-\varepsilon) \tag{3.4}$$

for every path $\tilde{X}(\omega)$ satisfying $\tilde{X}_t(\omega) > \frac{1}{\delta}$ for $t \in [1+\varepsilon, 2-\varepsilon]$ and $|\tilde{X}_t(\omega)| < \delta$ for $t \in [2,3]$. This set of paths $\tilde{X}(\omega)$ has positive probability and the quantity on the right side of (3.4) is strictly larger than 6 provided that ε was chosen sufficiently small. Hence we find that $\tilde{W}$ is not constantly equal to 6 as required.

For what it's worth, the example can be modified such that volatility is adapted to the filtration of the driving Brownian motion.

The trick is to choose a random sign $\hat{\epsilon}, P(\hat{\epsilon} = +1) = P(\hat{\epsilon} = -1) = \frac{1}{2}$ depending solely on the behavior of $(B_t)_{0\leq t\leq 1}$ and in such a way that S_1 is *independent* of $\hat{\epsilon}$. For instance, if we let $m(s)$ be the unique number satisfying $P(S_{1/2} > m(s)|S_1 = s) = P(S_{1/2} \leq m(s)|S_1 = s) = \frac{1}{2}$, it is sensible to define $\hat{\epsilon} := +1$ if $S_{1/2} > m(S_1)$ and $\hat{\epsilon} := -1$ otherwise.

We then leave the stock price process unchanged on $[0,1]$, i.e. we define $\hat{\sigma}^2(t) = \sigma^2(t) = 2$ and $\hat{S}_t = S_t$ for $t \in [0,1]$. On $]1,2]$ resp. $]2,3]$ we set $\hat{\sigma}^2(t) := 2 + \hat{\epsilon}$ resp. $\hat{\sigma}^2(t) := 2 - \hat{\epsilon}$ and define $\hat{S}_t, t \in]1,3]$ as the solution of the SDE

$$d\hat{S}_t = \hat{\sigma}(t)\hat{S}_t\, dB_t, \;\; \hat{S}_1 = S_1. \tag{3.5}$$

Here (3.5) depends only on S_1 and the process $(B_t - B_1)_{1\leq t\leq 3}$; since both are independent of $\hat{\epsilon}$, we obtain that $(\hat{S}_t)_{1\leq t\leq 3}$ and $(S_t)_{1\leq t\leq 3}$ are equivalent in law. It follows that $\hat{W} = \int_0^3 \hat{\sigma}^2(t,\omega)\, dt \equiv 6$ and since $\hat{S}_t$ and S_t have the same law for each $t \in [0,3]$, they induce the same local volatility model and in particular the same (non deterministic) $\tilde{W}$.

4 Counterexample for a Markovian stochastic volatility model

Recall that X denotes the log-price process of a general stochastic volatility model;

$$dX_t = \sigma(t,\omega)\, dB_t - \left(\sigma^2(t,\omega)/2\right) dt,$$

where $\sigma = \sigma(t,\omega)$ is the (progressively measurable) instantenous volatility process. Recall also our standing assumption that σ is bounded from above and below by positive constants.) We would like to apply theorem 1 to the 2D-diffusion (X,a) where $da = \sigma^2 dt$ keeps track of the running realized variance[5]. We can only do so after elltiptic regularization. That is, we consider

$$\begin{aligned} dX_t &= \sigma(t,\omega)\, dB_t - \left(\sigma^2(t,\omega)/2\right) dt, \\ da_t^\varepsilon &= \sigma^2(t,\omega)\, dt + \varepsilon^{1/2} dZ_t \end{aligned}$$

where Z is a Brownian motion, independent of $\sigma\left(B,\sigma^2\right)$. It follows that the following "double-local" volatility model

$$\begin{aligned} d\tilde{X}_t^\varepsilon &= \sigma_{dloc}\left(t,\tilde{X}_t^\varepsilon,\tilde{a}_t^\varepsilon\right) dB_t - \left(\sigma_{dloc}^2\left(t,\tilde{X}_t^\varepsilon,\tilde{a}_t^\varepsilon\right)/2\right) dt, \\ d\tilde{a}_t^\varepsilon &= \sigma_{dloc}^2\left(t,\tilde{X}_t^\varepsilon,\tilde{a}_t^\varepsilon\right) dt + \varepsilon^{1/2} dZ_t, \\ (\text{with } \sigma_{dloc}^2(t,x,a) &= E\left[\sigma^2(t,\omega)\,|X_t = x, a_t^\varepsilon = a\right]) \end{aligned}$$

has the one-dimensional marginals of the original process (X_t, a_t^ε). That is, for all fixed t and ε,

$$X_t \overset{law}{=} \tilde{X}_t^\varepsilon \text{ and } \tilde{a}_t^\varepsilon \overset{law}{=} a_t^\varepsilon.$$

Let us also note that the law of a_t^ε is the law of $a_t = a_t^0$ convolved with a standard Gaussian of mean 0 and variance ε. Let us also note that the log-price processes X and $\tilde{X}^\varepsilon$ induces the same local volatility surface. To this end, just observe that $X_t \overset{law}{=} \tilde{X}_t^\varepsilon$ implies identical call option prices for all strikes and maturities and hence (by Dupire's formula) the same local volatility:

$$\sigma_{loc}^2(t,x) = E\left[\sigma^2(t,\omega)\,|X_t = x\right] = E\left[\sigma_{dloc}^2\left(t,\tilde{X}_t^\varepsilon,\tilde{a}_t^\varepsilon\right)|\tilde{X}_t^\varepsilon = x\right].$$

Since the law of a time inhomogoneous Markov process is fully specified by its generator, it follows that the law of the local volatility process associated to (X) has the same law as the local volatility process associated to $\left(\tilde{X}^\varepsilon\right)$.

We apply this to the toy model discussed ealier. Recall that in this example, with $T=3$

$$a_T = W_T = \int_0^T \sigma^2(t,\omega)\, dt = 6$$

whereas realized variance under the corresponding local vol model,

$$\tilde{W}_T = \int_0^T \sigma_{loc}^2\left(t,\tilde{X}_t\right) dt$$

[5] In other words,

$$W_T = \int_0^T \sigma^2(t,\omega)\, dt = a_T,$$

provided $a_0 = 0$ which we shall assume from here on.

was seen to be random (but with mean $a_T = W_T$, thanks to the matching variance swap prices). As a particular consequence, using Jensen

$$
\begin{aligned}
E\Big(\int_0^T \sigma_{loc}^2\big(t,\tilde{X}_t\big)dt-6\Big)^+ \quad &> \quad \Big(E\int_0^T \sigma_{loc}^2\big(t,\tilde{X}_t\big)dt-6\Big)^+ = \\
&= \quad \Big(E\int_0^T \sigma^2\left(t,\omega\right)dt-6\Big)^+ = (a_T-6)^+.
\end{aligned}
$$

We claim that this persists when replacing the abstract stochastic volatility model (X) by $\left(\tilde{X}^\varepsilon\right)$, the first component of a 2D Markov diffusion, for any $\varepsilon > 0$. Indeed, thanks to the identical laws of the respective local volatility processes the left-hand side above does not change when replacing $(\tilde{X}_t)$ by the local volatility process associated to $\left(\tilde{X}^\varepsilon\right)$. On the other hand

$$
\begin{aligned}
E\int_0^T \sigma_{dloc}^2\big(t,\tilde{X}_t^\varepsilon,\tilde{a}_t^\varepsilon\big)dt \quad &= \quad E(\tilde{a}_T^\varepsilon-\varepsilon^{1/2}Z_T) \\
&= \quad E\left(\tilde{a}_T^\varepsilon\right) = E\left(a_T^\varepsilon\right) \\
&= \quad E\big(a_T+\varepsilon^{1/2}Z_T\big) = a_T.
\end{aligned}
$$

Thus, insisting again that the process $\tilde{X}$ is (in law) the local volatility model associated to $\left(\tilde{X}^\varepsilon\right)$ we see that

$$
E\Big(\int_0^T \sigma_{loc}^2\big(t,\tilde{X}_t\big)\,dt-6\Big)^+ > \Big(E\int_0^T \sigma_{dloc}^2\big(t,\tilde{X}_t^\varepsilon,\tilde{a}_t^\varepsilon\big)dt-6\Big)^+ = 0.
$$

In other words, the double-local vol model constitues an example of a Markovian stochastic volatility model, where stochastic volatility is a function of both state variables, in which conjecture 1 fails.

References

[BM06] D. Brigo and F. Mercurio. *Interest rate models—theory and practice*. Springer Finance. Springer-Verlag, Berlin, second edition, 2006. With smile, inflation and credit.

[FS86] E. B. Fabes and D. W. Stroock. A new proof of Moser's parabolic Harnack inequality using the old ideas of Nash. *Arch. Rational Mech. Anal.*, 96(4):327–338, 1986.

[FG05] P. Friz and J. Gatheral. Valuation of volatility derivatives as an inverse problem. *Quant. Finance*, 5(6):531–542, 2005.

[Gat06] J. Gatheral. The volatility surface: a practitioner's guide. *WILEY*, 2006.

[Gyö86] I. Gyöngy. Mimicking the one-dimensional marginal distributions of processes having an Itô differential. *Probab. Theory Relat. Fields*, 71(4):501–516, 1986.

[Lab09] P. H. Labordére. Calibration of local stochastic volatility models to market smiles. *Risk magazine*, 04 Sept. 2009.

[Lee01] R. Lee. Implied and local volatilities under stochastic volatility. *International Journal of Theoretical and Applied Finance*, vol 4:45–89, 2001.

[Pin95] R. G. Pinsky. *Positive harmonic functions and diffusion*, volume 45 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, Cambridge, 1995.

[Pit06] V. Piterbarg. Markovian projection method for volatility calibration. *Available at SSRN: http://ssrn.com/abstract=906473*, 2006.

[RW00] L. C. G. Rogers and D. Williams. *Diffusions, Markov processes, and martingales. Vol. 1.* Cambridge Mathematical Library. Cambridge University Press, Cambridge, 2000. Foundations, Reprint of the second (1994) edition.

[SV06] D. W. Stroock and S. R. Srinivasa Varadhan. *Multidimensional diffusion processes.* Classics in Mathematics. Springer-Verlag, Berlin, 2006. Reprint of the 1997 edition.