

A Modern Approach to Volatility Surface Calibration:

From Implied to Local Volatility with Optimal Transport

Research Report

August 2025

For more information

info@genOTC.com

genOTC.com

©2025 genOTC. All rights reserved.

This whitepaper and its contents are the intellectual property of **genOTC**. No part of this document may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

For permission requests, contact:
Marketing at info@genOTC.com



Authored by Grégoire Loeper

Grégoire Loeper (PhD in Mathematics) is a French mathematician renowned as a global expert in Optimal Transport and an experienced derivatives trader. His groundbreaking research has been published in top-tier, peer-reviewed academic journals, where he bridges cutting-edge mathematical theory with practical applications in quantitative finance.

His work spans derivatives pricing, volatility modeling, stochastic calculus, and the use of optimal transport methods in financial modeling.

Grégoire has held senior positions including Senior Scientific Advisor at BNP Paribas, Director of the Monash Centre for Quantitative Finance and Investment Strategies, and Professor of Mathematics at Monash University in Australia.

[>Explore more of Grégoire's work](#)

[>Follow for updates on LinkedIn](#)

A Modern Approach to Volatility Surface Calibration: From Implied to Local Volatility with Optimal Transport

genotc Research

Abstract

Volatility modeling plays a central role in pricing and managing derivatives. The aim of this white paper is to introduce some basic concepts about volatility, explain the challenges of volatility calibration, and introduce the optimal transport based technique offered by genOTC.

1 Diffusion models

A diffusion model is a **stochastic evolution model** for an asset:

$$dS_t/S_t = \mu dt + \sigma dW_t \quad (1)$$

In practice, sigma is a highly stochastic process, reflecting the agitation of the market. In particular, in equity markets, sigma has a tendency to increase when the market goes down, thus explaining the volatility smile. On the contrary, for assets such as gold, or oil, sigma can increase when the asset value goes up. In any case, σ is a noisy random process and a proper diffusion model has to take into account the dynamic nature of the volatility. A market (i.e. a collection of options prices in our case) is arbitrage free if there exists a diffusion model such that the price of every instrument is given by the discounted expectation of its payoff under this model. This result is striking, fundamental, and at the heart of the derivatives pricing: taking an expectation is very easily done numerically by the so-called Monte-Carlo algorithm (simulating a large number of paths of the asset and taking the average of the instrument payoff over these realizations).

2 Implied volatility

The Implied Volatility (IV) is maybe the most familiar object for practitioners. It originates from the historical work of and Black, Scholes and Merton, which led to the celebrated **Black-Scholes formula**. Assume that a risky asset follows a Black-Scholes diffusion, with constant volatility σ , then calls and puts prices on this asset are given by the Black-Scholes formula. This formula is not useful for pricing options, as the constant volatility assumption is a strong over-simplification, however it is very useful when used in the opposite way, i.e. **going back from an observed option price to the equivalent constant volatility model that would have given this price. This is the implied volatility (IV)**. There is one value of the IV for each strike and maturity, and the way IV varies across strikes is known as the **volatility smile**. While useful for quoting and benchmarking, implied volatility tells us little about how an asset might behave dynamically. Implied volatility is what you see in the market — a price-derived measure of expected volatility for a specific option, at a specific strike and maturity. It reflects how the market values risk today. But **implied vol is just a collection of points, not a model, it is merely a parametrization of option prices**. In particular, it does not provide a diffusion model.

3 Calibration

If one wants to actually price derivatives with a model, the challenge now becomes to adjust the diffusion parameters from observed market prices, to guarantee the arbitrage free property. The model prices (i.e. the expectation of options payoffs) have to match with market prices. This exercise is an inverse problem because one goes from prices to the model. Going the other way is easy, as this is a simple stochastic simulation exercise, and a graduate student would be able to do it without difficulty. The calibration problem is notoriously hard, and this what genOTC tackles.

4 Local volatility

Going back to our model (1) the first observation is that in order to be arbitrage free, the drift μ has to be equal to the risk free rate, that we call r . Indeed the asset itself has to be equal to its discounted future value, therefore the drift term (i.e the expected value of the return) has to equal the discount term (we neglect dividends here for now). This is called the risk-neutral property of the model, it has to be satisfied by any arbitrage free model. The specification of the process μ is therefore very rigid. The volatility process in contrast has a lot of flexibility: many different volatility models can lead to the same set of options prices (say listed options, so with a set of strikes at given listed maturities). Now here is an important question: what are the common features of all models that coincide on a given set of options. We can even go further and ask: assume that call prices (and therefore put prices also from call put parity) are known for every strike K and maturity T up to a time T_{\max} .

What can we say about the models that are consistent with this price surface ? They share the same local volatility which we define as:

$$\text{Local Volatility (LV)} = \sigma(t, s)$$

is the expected value of the instantaneous volatility sigma at time t conditional to the asset price being s. There is even an explicit formula to obtain the LV from options prices, the Dupire Formula:

$$\frac{1}{2}\sigma^2(T, K) = \frac{\partial_T C(K, T) + rK\partial_K C(K, T)}{K^2\partial_{KK} C(K, T)} \quad (2)$$

These observations have a deep implication: given an arbitrage free European price surface, i.e. given $C_{\text{market}}(K, T)$ for every K and T such that no arbitrage exists, one can build a local volatility model:

$$dS/S = rdt + \sigma(t, S)dW$$

such that the price of all European options under this model are calibrated to the market surface:

$$C_{\text{model}}(K, T) = \mathbb{E}[\text{Discount}(0, T)(S_T - K)^+] = C_{\text{market}}(K, T)$$

The local volatility model is clearly not the most realistic model, it is not based on empirical observation of statistical properties of prices. It is built uniquely by requesting a model to be:

- arbitrage free
- consistent with a given European price surface
- Markovian

The last assumption is clearly the most constraining one, it says that there is no «memory» in the process, the future evolution of the price is only affected by the current price, not by recent or less recent events. Taking into account this memory effect requires a statistical approach that is not necessary for the sole purpose of calibrating to European options.

5 The challenge of building local volatility surfaces

Although Dupire's formula suggests that it can be done directly via a simple analytic formula, constructing a local volatility surface is not trivial. Indeed, what is given by the market is only a discrete version of the surface $C(K, T)$, with bid-ask spreads, and possibly asynchronous quotes. Therefore one can only have an approximation of the true surface $C(K, T)$, and taking derivatives of an approximation can quite often lead to difficulties. A common answer to the problem of approximating $C(K, T)$ is to take a parametric form the implied volatility surface $IV(K, T, \text{Theta})$, where the parameter Theta lives in a relatively low dimensional space. This has indeed the effect off stabilizing the computation of the local volatility, at the expense of relying on a parametric model to fit the IV surface, therefore restricting the range of attainable market configurations.

6 Optimal Transport (OT) in a nutshell

Over the past two decades, OT has become a foundational tool in mathematics, statistics, machine learning, and now quantitative finance. At genOTC, we apply Optimal Transport to the task of building arbitrage-free local volatility surfaces. Originally proposed by Gaspard Monge in 1781 and reformulated by Leonid Kantorovich in the 1940s, Optimal Transport provides a mathematically rigorous way to displace a distribution in a constrained way, while minimizing a transport cost. Originally stated in the context of civil engineering, where actual matter is displaced, and the constraints are an initial and final distribution of matter, we will look here at moving a probability distribution, under a diffusion process compatible with risk neutral dynamics, and the constraints will be the options prices. In contrast with Dupire's approach, where a first important step of extrapolation/approximation of the surface $C(K,T)$ has to be done, the OT approach deals directly with the natural set of constraints: a given list of options prices within a given bid-ask spread. Under these constraints, the algorithm seeks to find the least costly compatible model. But what does "least costly" mean in this context? In Optimal Transport, the "cost" refers to the mathematical effort required to shift probability mass from one distribution to another. In our context of volatility modeling, the cost will more be seen as a regularization mechanism, penalizing volatilities that are too extreme and irregular. This process replaces Dupire's numerical derivatives with a global optimization that:

- Is arbitrage-free by construction
- Handles sparse and noisy data
- Avoids making any parametric assumptions
- Works across any asset class

7 genOTC: bringing it all together

genOTC operationalizes this mathematical innovation in a cross-asset SaaS platform designed for banks, hedge funds, asset managers, and exchanges. Key features include:

- Exact calibration to market prices
- Fast Real-time computation
- Asset-agnostic framework — works seamlessly across equities, FX, crypto, rates, commodities
- No integration or heavy setup required
- Production-ready volatility surfaces for pricing, hedging, and risk analysis

By avoiding the use of rigid parametric models and directly generating marked to market arbitrage free models, genOTC offers a next-generation tool for mastering market volatility.

If there is a fit, we will find it.

8 Conclusion

Volatility calibration is at the core of derivatives trading and hedging. It is not just a numerical problem — it is a probabilistic and geometric challenge. Optimal Transport provides the natural mathematical framework for solving it. genOTC transforms this insight into a powerful, scalable solution that delivers robust and consistent volatility surfaces — ready for real market conditions and coherent risk management.



Better Pricing | Better insights

Join the revolution.

genOTC is a fintech specialized in quantitative finance. We democratize access to, and empower clients with robust and arbitrage-free calibration methods. Ready to transform your OTC experience?

Contact us today to learn more about our solutions: info@genotc.com

LinkedIn: <https://www.linkedin.com/company/genotc>

Learn more: <https://genotc.com/optimal-transport-technology/>

Try it now: <https://genotc.com/solution>