POLITECNICO MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE E DELL'INFORMAZIONE

Executive Summary of the Thesis

Hawkes Process Models of Limit Order Books for Optimal Market Making: A Stochastic Control Approach

Laurea Magistrale in Mathematical Engineering – Quantitative Finance

Author: Andrea Tarditi

Advisor: Prof.ssa *Confortola Fulvia*

Academic year: 2024–2025

1. Context and motivation

Modern electronic markets are organized around limit order books (LOBs), where liquidity providers post limit orders and liquidity takers submit marketable orders that execute against the best quotes. At sub-second horizons, aggressive executions do not arrive as independent Poisson events: they come in bursts, display clustering, and show strong interaction between buy and sell activity. These empirical patterns matter directly for market making. When flow becomes intense, fill opportunities increase, but short-horizon adverse selection risk typically rises at the same time. A market maker therefore operates in an endogenous, fast-changing regime environment where both opportunity and risk are time-varying in event time.

This thesis addresses a concrete question: how should a liquidity provider dynamically choose bid and ask quote offsets to trade off spread capture and inventory risk when aggressive execution flow is self-exciting and state-dependent? The goal is not to build a full-fidelity market simulator, nor to forecast long-horizon prices. The goal is to construct a data-anchored, control-oriented pipeline that (i) captures clustered order arrivals, (ii) yields a mathematically coherent stochastic control problem, and (iii) produces an executable feedback quoting policy on tick-consistent grids.

2. Market model and control problem

The modelling backbone is an event-time description of aggressive order flow. Buy and sell market orders are modelled as a bivariate Hawkes process, capturing both self-excitation (bursts on the same side) and cross-excitation (spillovers across sides). Exponential kernels are used as a tractability device: they admit a finite-dimensional Markov embedding for the intensity state. This step is essential for control, because it turns a history-dependent excitation mechanism into a controlled Markov system where dynamic programming can be applied transparently.

The market maker is treated as a price taker. The control is the pair of bid/ask quote offsets $\delta_t = (\delta_t^a, \delta_t^b)$ (in ticks) around a reference midprice. The maker does not affect whether market orders arrive; instead, quoting affects whether incoming market orders execute against the maker. This is modelled through predictable thinning: a decreasing fill function maps quote distance into the intensity of executions against the maker. Short-horizon adverse selection is incorporated via predictable Bernoulli marking: conditional on market order arrivals, one-tick midprice moves occur with state-dependent probabilities linked to the order-flow state (for instance through intensity imbalance). Consistently with the thesis market model, the midprice is represented as a diffusion with additional one-tick jumps driven by the Bernoulli-marked market-order flow. These mechanisms (thinning for fills and marking for adverse moves) keep executions and toxicity inside a single compensator-based point-process framework, which is crucial for clean generator calculations and a consistent HJB derivation.

The maker's portfolio is described by inventory and cash processes updated by executions; terminal wealth is marked to mid and penalized for residual inventory. The objective is posed using exponential (CARA) utility:

$$V(t,z) = \sup_{\delta} \mathbb{E}\Big[-e^{-\gamma W_T^f} \mid Z_t = z\Big],$$
$$W_T^f = X_T + Q_T S_T - \frac{\eta}{2} Q_T^2.$$

The controlled state includes inventory, cash, the Hawkes intensities, and the price component (in simulation) or an exogenous midprice path (in replay). Controls are restricted to a compact, tick-consistent admissible set, matching microstructure constraints and numerical solvability.

3. Dynamic programming, HJB equation, and verification

A central contribution of the thesis is to make the bridge from Hawkes microstructure modelling to dynamic programming fully explicit and reproducible. The difficulty is not writing down an HJB formally, but making each step legitimate when order arrivals have stochastic predictable intensities and the control acts through predictable mechanisms (thinning and marking). To avoid "black-box" modelling shortcuts, the analysis constructs a clean semimartingale representation of the controlled state and derives the controlled generator from compensator calculus rather than heuristic infinitesimals.

Technically, this is achieved by embedding the Hawkes-driven counting channels into a Poisson random measure formulation and using a triangular (step-by-step) construction to guarantee well-posedness under the model's admissibility and integrability conditions. This yields a transparent bookkeeping of jump events (fills, adverse price moves, and intensity updates) and provides the exact Dynkin/Itô identities needed for dynamic programming. Within this framework, a stopping-time Dynamic Programming Principle (DPP) is proved for the finite-horizon exponential-utility objective, leading to a nonlocal Hamilton–Jacobi–Bellman equation whose jump terms correspond one-to-one to economically meaningful events in the LOB model.

The CARA structure is then used in a constructive way: exponential utility implies a factorization that removes cash and the absolute price level from the PDE, reducing the HJB to a tractable equation in $(t, q, \lambda^+, \lambda^-)$. A verification theorem is stated (and proved under the required regularity/integrability conditions) to certify that a sufficiently regular solution of the reduced HJB coincides with the value function and yields an optimal feedback quoting policy. In this sense, the resulting strategy is a white-box algorithm: its decisions are the explicit output of a calibrated model and a solved control problem, not a black-box predictor.

A stopping-time Dynamic Programming Principle (DPP) is established for the finite-horizon terminal-utility criterion. The DPP leads to a Hamilton–Jacobi–Bellman equation in which the supremum is taken over admissible quote offsets and the nonlocal terms account for both (i) inventory/cash jumps induced by executions and (ii) intensity jumps induced by excitation. A fully general PDE treatment can be formulated in a viscosity-solution framework; however, the thesis follows a classical verification route: for a sufficiently regular candidate solution, a verification theorem shows that it coincides with the value function and yields an optimal feedback control.

Exponential (CARA) utility plays a structural role. The value function factorizes in a way that removes cash and the absolute price level from the PDE, leading to a reduced HJB in $(t, q, \lambda^+, \lambda^-)$. This reduction is not merely algebraic convenience: it makes the numerical problem feasible on bounded grids and makes explicit that the optimal policy depends on inventory and on the order-flow regime encoded by the Hawkes intensities (including their imbalance).

4. Numerical solution and empirical implementation

The reduced HJB is solved numerically on bounded grids for inventory and intensities, with a discrete (tick-consistent) control set so that the maximization becomes a finite search at each grid point. Time stepping is designed to remain stable under fast Hawkes mean reversion, and post-jump states are handled by interpolation and clamping to grid bounds. The solver exports an executable feedback policy mapping

$$(t, q, \lambda^+, \lambda^-) \;\mapsto\; (\delta_b^\star, \delta_a^\star),$$

which can be used directly inside a backtesting engine.
The empirical pipeline connects data to control in a reproducible way. LOBSTER-style messages are mapped into a multivariate event representation, then reduced to a two-dimensional aggressive execution flow for Hawkes calibration. The Hawkes model is fitted on intraday windows to reflect local stationarity while allowing regime changes through the day; subcriticality checks based on the branching matrix provide safeguards for simulation. Reduced-form microstructure components are calibrated to map quote distance into fill propensity and to estimate short-horizon adverse-move probabilities at event level.
Policy evaluation is carried out in two complementary regimes. In replay mode, the historical midprice path is fixed, so performance is not confounded by price-model mismatch and the test focuses on execution and inventory control under realized price moves. In simulation mode, events are generated from the calibrated Hawkes model together with the calibrated fill/marking mechanisms, enabling controlled stress tests and sensitivity analyses. Numerically, the optimal quotes widen and skew more aggressively when Hawkes intensities and imbalance signal elevated short-horizon toxicity, while reverting to tighter and more symmetric quotes in calmer regimes.

5. Scope and outlook

The modelling choices are deliberately control-oriented. The framework aims to capture the dominant high-frequency mechanism that matters for market making, namely clustered aggressive flow and its link to fill opportunities and adverse selection, while preserving a tractable Markov state and a numerically solvable HJB. Natural extensions include richer marking (multi-tick moves and variable horizons), explicit queue dynamics and latency/priority effects, multi-scale kernels (sums of exponentials) to capture longer memory while preserving a Markov state, and a full viscosity/weak-DPP formulation that removes smoothness assumptions

Overall, the thesis delivers an end-to-end, reproducible pipeline in which each component has an explicit role: LOB data are mapped into event streams, Hawkes and microstructure primitives are calibrated, the reduced HJB is solved on tick-consistent grids, and the resulting policy is exported and evaluated in replay and simulation regimes.

References

[1] Frédéric Abergel, Marouane Anane, Anirban Chakraborti, Aymen Jedidi, and Ioane Muni Toke. *Limit Order Books*. Cambridge University Press, 2016.

[2] Marco Avellaneda and Sasha Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008. doi: 10.1080/14697680701381228.

[3] Emmanuel Bacry, Iacopo Mastromatteo, and Jean-François Muzy. Hawkes processes in finance. *Market Microstructure and Liquidity*, 1(1):1550005, 2015. doi: 10.1142/S2382626615500057.

[4] Pierre Brémaud. *Point Processes and Queues: Martingale Dynamics*. Springer, 1981. doi: 10.1007/978-1-4684-9477-8.

[5] Pierre Brémaud and Laurent Massoulié. Stability of nonlinear hawkes processes. *The Annals of Probability*, 24(3):1563–1588, 1996. doi: 10.1214/aop/1065725193.

[6] Álvaro Cartea, Sebastian Jaimungal, and José Penalva. *Algorithmic and High-Frequency Trading*. Cambridge University Press, 2015.

[7] Rama Cont, Sasha Stoikov, and Rishi Talreja. A stochastic model for order book dynamics. *Operations Research*, 58(3):549–563, 2010. doi: 10.1287/opre.1090.0806.

[8] L. Cui, A. Hawkes, and H. Yi. An elementary derivation of moments of hawkes processes. *Advances in Applied Probability*, 52 (1):102–137, 2020. doi: 10.1017/apr.2019.53.

[9] D. J. Daley and D. Vere-Jones. *An Introduction to the Theory of Point Processes. Volume I: Elementary Theory and Methods*. Springer, 2 edition, 2003. doi: 10.1007/978-0-387-49835-5.

[10] Wendell H Fleming and H Mete Soner. *Controlled Markov Processes and Viscosity Solutions*, volume 25 of *Stochastic Modelling and Applied Probability*. Springer Science & Business Media, 2 edition, 2006. doi: 10.1007/0-387-31071-1.

[11] Olivier Guéant, Charles-Albert Lehalle, and Joaquin Fernandez-Tapia. Dealing with the inventory risk: a solution to the market making problem. *Mathematics and Financial Economics*, 7:477–507, 2013. doi: 10.1007/s11579-012-0087-0.

[12] Alan G. Hawkes. Spectra of some self-exciting and mutually exciting point processes. *Biometrika*, 58(1):83–90, 1971. doi: 10.1093/biomet/58.1.83.

[13] Alan G. Hawkes and David Oakes. A cluster process representation of a self-exciting process. *Journal of Applied Probability*, 11(3):493–503, 1974. doi: 10.1017/S0021900200096273.

[14] W. Huang, Charles-Albert Lehalle, and Mathieu Rosenbaum. Simulating and analyzing order book data: The queue-reactive model. *Journal of the American Statistical Association*, 110(509):107–122, 2015. doi: 10.1080/01621459.2014.982278.

[15] Konark Jain, Nick Firoozye, Jonathan Kochems, and Philip Treleaven. Limit order book simulations: A review, 2024. revised Mar 2024.

[16] Théo Leblanc. Exponential moments for hawkes processes under minimal assumptions. *Electronic Communications in Probability*, 2024. doi: 10.1214/24-ECP625. Also available as arXiv:2505.15253.

[17] Xiaofei Lu and Frédéric Abergel. High-dimensional hawkes processes for limit order books: modelling, empirical analysis and numerical calibration. *Quantitative Finance*, 18(2):249–264, 2018. doi: 10.1080/14697688.2017.1403142.

[18] Yosihiko Ogata. On lewis' simulation method for point processes. *IEEE Transactions on Information Theory*, 27(1):23–31, 1981. doi: 10.1109/TIT.1981.1056305.

[19] Huyèn Pham. *Continuous-time Stochastic Control and Optimization with Financial Applications*. Springer, 2009. doi: 10.1007/978-3-540-89500-8.

[20] Eric Smith, J. Doyne Farmer, László Gillemot, and Supriya Krishnamurthy. Statistical theory of the continuous double auction. *Quantitative Finance*, 3(6):481–514, 2003. doi: 10.1088/1469-7688/3/6/307.