Executive Summary

In this thesis we examine efficient Monte Carlo methods for option pricing and risk management. We focus in particular on extending the multilevel approach of Giles (2008) to more complex exotic derivatives featuring periodic cash flows. For this, we provide a rigorous proof that Theorem 3.1 of Giles (2008) can be generalized to target accrual redemption forwards.

We compare empirically different multilevel pricing methods on calibrated local and stochastic volatility models. We conclude that although for Lipschitz payoffs the method proposed by Giles (2008) has a complexity of $O\left(\varepsilon^{-2}\left(\log\varepsilon^{-1}\right)^2\right)$ and Lemaire et al. (2017) has a lower complexity of just $O\left(\varepsilon^{-2}\left(\log\varepsilon^{-1}\right)\right)$ (where ε is the imposed upper RMSE bound), the ad-hoc stopping criterion in the former results in superior performance for realistic situations compared to the latter. We also prove empirically that the antithetic multilevel Monte Carlo method proposed in Giles et al. (2014) dominates both approaches, even when the cost model differs (e.g. when the generation of random variates does not dominate the running time of the simulation); such situations are encountered frequently in industry, where random variates are pre-generated outside trading hours.

The above empirical study is necessary as lower theoretical *asymptotic* complexity in Monte Carlo simulation often comes at the cost of a more complex simulation method (e.g. Milstein vs. Euler-Maruyama). Especially for $\varepsilon \in [10^{-4}; 10^{-5}]$ (which are values of practical interest), it is not *a priori* clear that the asymptotic complexity is what determines the ultimate performance of the methods. There are numerous other technical factors that can (and do) result in lower performance of an otherwise theoretically superior method.

Finally, we provide the technical details of a risk system based on algorithmic differentiation that is used in conjunction with Monte Carlo simulations. We show that the model sensitivities thus computed are more accurate. Furthermore, the method can be readily applied to multilevel Monte Carlo simulations, resulting in large computation gains for complex models. Just as before, we show that the performance of the system does not only depend on the theoretical efficiency of automatic differentiation; rather, the internal data structures play a disproportionate role in achieving this theoretical performance.