

FINANCIAL MATHEMATICS TEAM CHALLENGE

A collection of the four reports from the 2025
Financial Mathematics Team Challenge.



AIFMRM

AFRICAN INSTITUTE OF FINANCIAL MARKETS AND RISK MANAGEMENT



Preamble

One of the key aims of the FMTC is for South African postgraduate students in Financial and Insurance Mathematics to have the opportunity to focus on a topical, industry-relevant research project, while simultaneously developing links with international students and academics in the field. An allied objective is to bring a variety of international researchers to South Africa to give them a glimpse of the dynamic environment that is developing at UCT in the African Institute of Financial Markets and Risk Management. The primary goal, however, is for students to learn to work in diverse teams and to be exposed to a healthy dose of fair competition.

The Tenth Financial Mathematics Team Challenge was held from the 30th of June to the 11th of July 2025. The challenge brought together four teams of Master's and PhD students from Canada, Italy, and the United Kingdom to pursue intensive research in Financial Mathematics. Each team worked on a distinct research problem over the twelve days. Professional and academic experts from Canada, South Africa, France, and the United Kingdom mentored the teams; fostering teamwork and providing guidance. As they have in the past, the students applied themselves with remarkable commitment and energy.

This year's research included topical projects on (a) A Tale of Two Regions: A North and South Macroeconomic-Ecological Model, (b) Competition Between Market Makers, (c) Development and Analysis of a Uniswap v3 Agent-Based Model, and (d) xVA Greeks using the Likelihood Ratio Method within an American Monte Carlo Framework. These were either proposed directly by our academic/industry partners or chosen from areas of current relevance to the finance and insurance industry. To prepare the teams, guidance and preliminary reading was given to them a month before the meeting in Cape Town. During the final two days of the challenge, the teams presented their conclusions and solutions in extended seminar talks. The team whose research findings were adjudged to be the best was awarded a floating trophy. Each team authored a report containing a critical analysis of their research problem and the results that they obtained. This volume contains these four reports and will be available to future FMTC participants. It may also be of use and inspiration to Masters and PhD students in Financial and Insurance Mathematics.

To our astonishment, this year was the 10th anniversary edition of the FMTC. Who would have thought it! We celebrated in modest fashion by designing a new, special edition hoodie with a designer logo. The participants were super impressed.

We're taking a short hiatus next year, and then FMTC XI returns in 2027.

David Taylor, University of Cape Town

Andrea Macrina, University College London & University of Cape Town

Contents

1. **B. Badenhorst, K. Baldeo, K. Bopape, E. Kroell**
A Tale of Two Regions: A North and South Macroeconomic-Ecological Model
2. **N. Chilenje, M. Daba, D. Feleppa, C. Fellner**
Competition Between Market Makers¹
3. **D. Presta, S. Chen, T. Koen, D. Oosthuizen**
Development and Analysis of a Uniswap v3 Agent-Based Model
4. **A. Mathieu, G. Burgin, G. Matheri, L. Matlonya**
xVA Greeks using the Likelihood Ratio Method within an American Monte Carlo Framework

¹Winning team of the tenth Financial Mathematics Team Challenge

A Tale of Two Regions: A North and South Macroeconomic- Ecological Model

TEAM 1

BEN BADENHORST, University of Cape Town
KIASH BALDEO, University of Cape Town
KGAUGELO BOPAPE, University of Cape Town
EMMA KROELL, University of Toronto

Mentor:
MATHEUS GRASSELLI, McMaster University

African Collaboration for Quantitative Finance and Risk Research

Contents

1 Introduction	4
1.1 Literature Review	5
1.1.1 The DICE model	5
1.1.2 The SFC Approach	6
1.1.3 GEMMES and Related Models	7
1.2 Summary and Outlook	7
2 Methodology and Model	9
2.1 Climate-Economic Interactions	10
2.1.1 Production Before and After Damages	10
2.1.2 Abatement and Government Policy	12
2.2 Macro-Economic Module	13
2.2.1 Exchange Rates and Interest Rates	13
2.2.2 Nominal Profits, Investment, and Dividends	14
2.2.3 Labour, Inflation, and Wages	16
2.2.4 Consumption and Trade	17
2.3 Climate Module	19
2.4 Summary	21
3 Analysis of Results	22
3.1 Initial Calibration	22
3.2 Scenario Analysis	23
3.3 Case Study: Intervention by the North to Save Both Regions	28
3.3.1 Monte Carlo	31
4 Conclusion	34
4.1 Summary	34
4.2 Model Extensions and Future Work	35

A Additional Details	37
1.1 Stock-Flow Consistency Table	41
1.2 Data Collection and Estimation	42

Chapter 1

Introduction

The recent COP29 summit, held in Baku, Azerbaijan, brought renewed global attention to the deepening divide between developed and developing nations in addressing climate change. Branded the “Finance COP,” the conference resulted in a headline commitment by developed countries to mobilize 300 billion USD per year by 2035 to support climate adaptation and mitigation in developing nations (Wei et al., 2025). However, many emerging economies criticized the pledge as insufficient and delayed, given the immediate and escalating costs they face. These outcomes reaffirmed the structural imbalance in climate negotiations and underscored the need for models that capture the differentiated vulnerabilities and capacities across regions.

Traditional economic models often treat the global economy as a homogeneous entity, overlooking the stark disparities between developed and developing regions in terms of vulnerability to climate impacts, capacity for mitigation, and access to green technologies (Gallagher, 2014). This report seeks to bridge this gap by developing a two-region macroeconomic-ecological model that explicitly accounts for the divergent dynamics between the Global North (developed nations) and the Global South (developing nations).

Building on the foundations of the GEMMES (General Monetary and Multisectoral Macrodynamics for the Ecological Shift) model introduced by Bovari et al. (2018a), our study extends the framework to incorporate two interconnected economies linked through trade, investment, and exchange rate mechanisms. Unlike the GEMMES model, which aggregates the global economy into a single unit, our approach captures the asymmetric effects of climate policies and financial interactions between the two regions. This disaggregation allows us to explore how climate-

related damages, carbon pricing, and abatement subsidies in one region affect each other, offering a more nuanced understanding of the global climate-economy debate.

This report begins with an overview of the relevant literature, providing the context for our work. We then present a comprehensive description of the extended GEMMES model and the methodology used to calibrate and simulate it in a two-region framework. To explore the dynamic interactions between the North and South economies, particularly how their policies influence each other and the global climate, we implement a series of scenarios. These scenarios allow us to evaluate both the economic and environmental outcomes, offering insights into the trade-offs, spillover effects, and potential benefits of cooperation in regional climate policy.

1.1 Literature Review

This report connects several related macroeconomic modelling approaches. In this section, we give an overview of the different interconnected approaches and explain where our model fits in.

1.1.1 The DICE model

Numerous studies have been conducted to examine the dynamics between the global economy and climate, specifically, the effect of policies on greenhouse gas emissions and global temperature and the effect of global warming and climate change on an economy. Integrated Assessment Models (IAMs) of global climate change, model the effect of various interacting economic factors and policies on global climate over time.

The earliest IAM was the DICE (Dynamic Integrated Climate-Economy) model, first introduced by Nordhaus (1993), which models the dynamics of emissions and the resulting damages to the economy, as well as optimal emission reduction policies. As one of the first economic models including climate change impact, the DICE model has been adapted and extended extensively since the original 1993 paper and has proven so consequential that Nordhaus was awarded the 2018 Nobel prize for its development (Nobel Prize Committee / Nobel Media, 2018).

Nordhaus (1993) found that a modest carbon tax (starting at \$5 per ton of carbon,

rising to \$20 by 2100) is economically efficient, reducing emissions by 10-15% while minimizing costs. This work was a key advance in seeing the need for carbon pricing, but also criticized by others for low climate damage assumptions which led to more modest policy prescriptions. The results of this study were later updated in Nordhaus (2018), who predicts rapid climate change under minimal climate policies, with global temperatures rising significantly by 2100. These updated results show that the 2°C target set by international agreements such as the Paris Agreement is unlikely to be achieved.

1.1.2 The SFC Approach

Stock-Flow Consistent (SFC) modelling, introduced by Godley (1996) and Godley and Lavoie (2006), is a macroeconomic modelling approach that tracks all financial and real flows between different sectors of an economy using a balance sheet approach. These models capture monetary movement between real and financial economic variables across various sectors of an economy. An SFC model is characterised by a strict accounting framework executed by a balance sheet matrix indicating the stocks and a transaction flow matrix indicating the flows in the economy. Stock consistency refers to the property that every financial asset of some sector is a liability of another sector of the economy, flow consistency points to the property that every monetary flow is between two sectors in an economy, and stock-flow consistency refers to how every flow changes two stocks (Nikiforos and Zezza, 2017). Moreover, a set of behavioural equations beyond the accounting framework determine a particular stock flow consistent model.

SFC models offer several advantages in addition to models such as DICE, particularly in capturing macro-financial realism and sectoral interactions. Unlike DICE, which simplifies the economy into a single representative agent maximizing utility over time (Barrage and Nordhaus, 2024), SFC models explicitly represent multiple sectors— households, firms, government, and financial institutions— ensuring that all monetary flows and stock changes are accounted for using rigorous double-entry accounting. This allows SFC models to track debt accumulation, income distribution, and financial fragility, which are key channels through which climate change and mitigation policies affect the real economy (Godin and Yilmaz, 2020). Thus SFC models are particularly useful for policymakers assessing climate transition risks, distributional effects, and the macroeconomic feasibility of green investment strategies.

1.1.3 GEMMES and Related Models

Many IAMs use traditional macroeconomic models. In contrast, [Bovari et al. \(2018a\)](#) present an IAM that combines the climate module of the DICE model with the macroeconomic core of a Keen-based model which allows for rapid economic collapse ([Keen, 1995](#)). These models use ordinary differential equation (ODEs) to describe the interplay of the workers and the productive sector in an economy. For further work with these models, see [Grasselli and Costa Lima \(2012\)](#) and [Grasselli and Huu \(2015\)](#).

The GEMMES model of [Bovari et al. \(2018a\)](#) integrates the economic impact of climate change with the critical role of private debt, where the capital, debt, and employment dynamics are adapted from the Keen model described above. Their results highlight that both long-term climate change impacts and too rapid a transition to a low-carbon economy can lead to severe economic consequences, forcing the private sector to increase leverage and thus endangering financial stability. This model has been used for many subsequent expansions, including [Bovari et al. \(2018b\)](#), who introduce a government sector, and [Martin et al. \(2024\)](#), who update the climate module.

Sensitivity analyses of the GEMMES model are conducted by [Bolker et al. \(2021\)](#) and [Longaretti and Martin \(2025\)](#). In the former case, the authors identify inflation dynamics as being key to the outcome of the model. In the latter case, the authors find that the 2°C global warming target is already out of reach without negative emissions.

1.2 Summary and Outlook

While [Bovari et al. \(2018a\)](#) provides a comprehensive analysis of climate change impacts using the GEMMES model, its global aggregate approach could be refined to better capture regional disparities in climate vulnerability and policy effectiveness. A key limitation is the lack of disaggregation between developed and emerging economies, which differ significantly in growth trajectories, decarbonization capacities, and climate risk exposure. For instance, emerging markets often face higher adaptation costs and slower technological adoption ([Gallagher, 2014](#)), while developed nations contribute disproportionately to historical emissions ([Allen and Stocker, 2014](#)). Our goal is to address this gap by explicitly modelling financial flows, trade imbalances, and asymmetric policy impacts between the two region.

We do this by creating a new model which has two economic regions, who interact with each other and with a unified climate module. We introduce new exchange rate and trade dynamics, which link the economies together. We further estimate initial conditions for our two-region model from a variety of data sources and run several policy scenarios.

The remainder of the report is structured as follows. In [Chapter 2](#), we detail our model, explaining how we adapt the GEMMES model and add new exchange interactions. In [Chapter 3](#), we provide a detailed scenario analysis. Finally, in [Chapter 4](#), we give a brief conclusion and give some next steps for future research.

Chapter 2

Methodology and Model

In this section, we present our new two-region extension to the GEMMES model, where one region represents the economic South and the other the economic North. Our approach is to duplicate the economic module of the GEMMES model [Bovari et al. \(2018a\)](#), creating two regional economies. We then add economic interactions between the two regions. Both economic modules feed into a unified global climate module, based on the DICE model [\(Nordhaus, 2018\)](#). This is because the effects of each economy on the climate have global implications.

We link the economies mainly via an exchange rate mechanism, which connects the regions' domestic interest rates and affects trade. Trade between the two economies is modelled through an import-export mechanism, where both regions can use foreign goods for both investment and consumption. Demand for foreign goods depends on the exchange rate.

In this section, we provide an in-depth overview of how we calculate and determine basic economic and climate metrics. We first describe how the economy and climate are linked via climate damages and abatement, before describing the macroeconomic and climate modules in detail. In what follows, parameters that are set constant are denoted with a bar on top. [Table 2.1](#) provides a list of the 28 state variables used in the model and their respective meanings. Those listed with a subscript j exists in both regions, North and South, while those without a subscript are unique and defined for the entire world.

Variable	Description
K_j	Real capital stock in each region
a_j	Labour productivity in each region
D_j	Debt level in each region
p_j	Price level in each region
w_j	Wage level in each region
N_j	Population in each region
p_{Car_j}	Carbon price in each region
σ_j	Emission intensity in each region
$g\sigma_j$	Growth rate of emission intensity in each region
p_{BS_j}	Price of green (backstop) technology in each region
E_{land}	Exogenous land-use based emissions
CO_2^{AT}	CO ₂ -e concentration in the atmosphere
CO_2^{UP}	CO ₂ -e concentration in the upper ocean and biosphere
CO_2^{LO}	CO ₂ -e concentration in the lower ocean
T	Mean temp. of the atmosphere, land surface, and upper ocean
T_{LO}	Mean temp. of the deeper ocean
e^n	Nominal exchange rate
e_n^e	Expected exchange rate

Table 2.1: List of state variables used in the two-region ecological macroeconomic model, where $j \in \{N, S\}$ indicates two variables, one for each region.

2.1 Climate-Economic Interactions

2.1.1 Production Before and After Damages

The real output before damages in economy $j \in \{N, S\}$ is defined as:

$$Y_j^0 := \frac{K_j}{\bar{\nu}_j} = a_j L_j,$$

where K_j is the capital stock, L_j is the level of employed labour, a_j is labour productivity, and $\bar{\nu}_j$ is the capital–output ratio which we assume is constant. The labour productivity in each region grows exponentially according to the following ODE:

$$\frac{\dot{a}_j}{a_j} := \bar{\alpha}_j.$$

A key assumption of this model and other IAMs is that the real output Y^0 is not all available for investment, consumption, and trade, but is first reduced by the effects of climate change. There are two avenues for this reduction: climate damages, which directly reduce output due to climate change, and abatement, which reduces available economic output as firms seek to reduce the effects of climate change.

The available output after climate damages and abatement is given by

$$Y_j := (1 - \mathbf{D}_j)(1 - A_j)Y_j^0,$$

where $\mathbf{D}_j \in [0, 1]$ is the fraction of output lost due to climate damages and A_j is the abatement cost as a fraction of output.

Damages are calculated using a polynomial damage function. Opinions on the correct specification of the convexity of the damage curve differ, with three common opinions given by Nordhaus (2018), Weitzman (2012), and Dietz and Stern (2015). The general form for the damage curve is as follows:

$$\mathbf{D}_j := 1 - \frac{1}{1 + \bar{\pi}_{1,j}T + \bar{\pi}_{2,j}T^2 + \bar{\pi}_{3,j}T^{\bar{\xi}_j}},$$

where $\bar{\pi}_1$, $\bar{\pi}_2$, $\bar{\pi}_3$ and $\bar{\xi}_j$ are parameters which differ depending on the damage curve specification and T is the state variable giving the temperature deviation of the atmosphere, land, and upper ocean from pre-industrial levels. A comparison of the three damage curves is given in Figure 2.1. The Stern assumptions lead to a more convex damage curve, where global warming, measured in degrees Celsius




Figure 2.1: Comparison of three different damage curves, corresponding to different assumptions made in the literature by Nordhaus, Weitzman, and Stern.

from the preindustrial average temperature, leads to a more significant reduction in output, compared to the Weitzman or Nordhaus. In our model, we allow for any one of these three damage specifications, and will explore the impact of different damages specifications in [Section 3.2](#).

2.1.2 Abatement and Government Policy

In order to reflect that economies are introducing mechanisms to reduce emissions and limit temperature increases, we model the cost of mitigation through two channels: the evolution of backstop technology prices and carbon pricing, both of which influence abatement decisions in each region.

We assume there exists a zero-emission backstop production technology which firms can switch to. The price of the backstop technology in each region may be different and is assumed to be decreasing according to the ODE

$$\frac{\dot{p}_{BS_j}}{p_{BS_j}} := \bar{\delta}_{p_{BS_j}} \leq 0.$$

The abatement cost, A_j , gives the cost firms face in reducing their emissions. It is defined as

$$A_j := \frac{\sigma_j p_{BS_j}}{\bar{\theta}_j} n_j^{\bar{\theta}_j},$$

where σ_j is the region's emission intensity, n_j is the region's emissions reduction rate, p_{BS} is the price of backstop (green) technology, and $\bar{\theta}_j$ is a convexity parameter ([Bovari et al., 2018a](#)). We assume that the government may offer a subsidy as a fraction of abatement costs, $\bar{s}_{A_j} \in [0, 1)$, to cover a portion of the costs of the productive sector transitioning towards green production technologies. The total subsidy payment is

$$S_{f,j} := \bar{s}_{A_j} A_j Y^0.$$

The government may also impose a carbon tax of p_{Car_j} on industrial emissions E_{ind_j} to incentivize firms to reduce their emissions, with a total cost of

$$T_{f,j} := p_{Car_j} E_{ind_j}.$$

We assume the carbon price is set by the government according to the ODE

$$\frac{\dot{p}_{Car_j}}{p_{Car_j}} := \bar{\delta}_{p_{Car_j}} \geq 0.$$

As the carbon price level is an exogenous variable, we explore different carbon pricing pathways in the scenarios in [Section 3.2](#).

Given the subsidy and carbon tax amounts, a net transfer of $S_{f,j} - T_{f,j}$ takes place from the households to the productive sector via the government.

Following [Bovari et al. \(2018a\)](#), we assume that firms choose the emissions reduction rate n_j to minimize their abatement cost minus the net transfer they receive from the government. Thus firms in each region choose n_j such that

$$\min_{n_j \in [0,1]} \{A_j Y_j^0 + T_{f,j} - S_{f,j}\}.$$

One verifies by taking the first order conditions that this results in an emissions reduction rate of

$$n_j = \min \left\{ \left(\frac{p_{C_j}}{p_{BS_j}(1 - \bar{s}_{A_j})} \right)^{\frac{1}{\theta_j - 1}}, 1 \right\}.$$

2.2 Macro-Economic Module

In this subsection, we describe how the relevant macro-economic variables are calculated, based on the approach of [Bovari et al. \(2018a\)](#), extended to a two-economy set-up where each variable is indexed by $j \in \{N, S\}$ to reflect the respective region (North or South). The economies are interconnected through an exchange rate mechanism, which we describe first.

2.2.1 Exchange Rates and Interest Rates

Our approach to modelling exchange rates is based on the work of [Godin and Yilmaz \(2020\)](#). We first define the nominal exchange rate e^n , given in units of Southern currency per unit of Northern currency. Let i_S and i_N denote the inflation rate in the South and North, respectively, whose evolution will be described later (see [Equation \(2.2\)](#)). Furthermore, let Ex_S and Im_S denote the exports and imports of the South, respectively. Then we assume that the nominal exchange rate evolves as follows:

$$\frac{\dot{e}^n}{e^n} = (i_S - i_N) - \bar{\beta}_e \left(\frac{Ex_S - Im_S}{K_S} \right),$$

where $\bar{\beta}_e$ is a sensitivity parameter and we recall that K_S is the capital stock of the South. Thus the nominal exchange rate is influenced by the difference in inflation

between the South and the North as well as by the trade balance. When the Southern inflation rate exceeds the Northern one, or when imports exceed exports, the nominal exchange rate depreciates.

We also define the expected exchange rate, which gives the public's view on the future value of the nominal exchange rate e^n . We model the dynamics of the expected nominal exchange rate using the “backward-looking expectation structure” of Godin and Yilmaz (2020) which is given as follows:

$$\dot{e}_{exp}^n = \bar{\beta}_e^{ex} (e^n - e_n^e).$$

The change in the expected nominal exchange rate is based on how much the actual exchange rate e^n differs from what they expected e_n^e . The rate at which expectation of the nominal exchange rate is adjusted is determined by the adjustment parameter $\bar{\beta}_e^{ex}$.

The real exchange rate gives the difference in price levels between the economies. It is defined as follows:

$$e_R = \frac{p_N}{p_S} e_n,$$

where p_S and p_N denote the price levels in the South and North, respectively.

Having defined the exchange rates, we now discuss interest rates. We assume that the North sets a constant interest rate r_N . Then the interest rate in the South, r_S , is determined via an interest rate parity condition adapted from Godin and Yilmaz (2020):

$$r_S = r_N + (e_n^e + \dot{e}_n^e - e_n) + \bar{\iota}_S.$$

This condition links the interest rates of the two economies through the nominal exchange rate e_n , its expected future value e_n^e , and its rate of change $\dot{e}_n^e$. The term $\bar{\iota}_S$ represents the South's risk premium, which is assumed constant.

2.2.2 Nominal Profits, Investment, and Dividends

We assume nominal profit of the private sector in economy j is given by:

$$\Pi_j := p_j Y_j - w_j L_j - r_j D_j - p_j \bar{\delta}_j K_j + p_j (S_{f,j} - T_{f,j}),$$

where Π_j is the nominal profit in region j , p_j is the price level in region j , Y_j is real output, $w_j L_j$ is the total wage expense, $r_j D_j$ is the interest payment on outstanding private debt, $S_{f,j} - T_{f,j}$ is the net transfer from the government as defined in Section 2.1.2, K_j is the capital stock, and $\bar{\delta}_j$ is the capital depreciation rate.

We also define

$$\pi_j := \frac{\Pi_j}{p_j Y_j},$$

the ratio that measures total profit relative to the total nominal output of the economy.

The level of investment in each economy depends on private sector profits and output after damages and abatement according to

$$I_j := \kappa_j(\pi_j) Y_j,$$

where the function $\kappa(\cdot)$ is a smooth increasing function (Bovari et al., 2018a).

As we extend the model from a single-economy to a two-economy framework, each region is now able to purchase a portion of its investment goods from the foreign economy, with the remainder purchased domestically. The foreign component of investment, I_j^f , is governed by the real exchange rate through an allocation function $\gamma_{I_j}(e_R) \in [0, 1]$, such that

$$I_j^f := \gamma_{I_j}(e_R) \cdot I_j.$$

The domestic component is then given as:

$$I_j^D := (1 - \gamma_{I_j}(e_R)) \cdot I_j.$$

The function $\gamma_{I_j}(e_R)$ is decreasing in the real exchange rate for each region as follows

$$\begin{aligned} \gamma_{I_S}(e_R) &:= \min \{1, \max \{0, \bar{\gamma}_{I_S} - \bar{\beta}_{I_S} \cdot (e_R - \bar{e}_R)\} \}, \\ \gamma_{I_N}(e_R) &:= \min \{1, \max \{0, \bar{\gamma}_{I_N} - \bar{\beta}_{I_S} \cdot (1/e_R - 1/\bar{e}_R)\} \}, \end{aligned}$$

where e_R is the real exchange rate, $\bar{e}_R$ is the equilibrium exchange rate, $\bar{\gamma}_{I_S}$, $\bar{\gamma}_{I_N}$ are the baseline investment share in the foreign goods, and $\bar{\beta}_{I_S}$, $\bar{\beta}_{I_N}$ are sensitivity parameters.

Investment determines the accumulation of capital over time. Given the amount of investment made, the change in capital stock K_j in each region is governed by the following ODE:

$$\dot{K}_j = I_j - \bar{\delta}_j K_j,$$

where we recall that $\bar{\delta}_j$ is the capital depreciation rate.

Once investment has been made, the productive sector allocates its remaining profits either to dividend payments or to retained earnings. Dividends are distributed

to households, while retained earnings remain within the productive sector. This distribution is modelled as follows:

$$\begin{aligned}\Pi_{d,j} &:= \Delta(\pi_j) p_j Y_j, \\ \Pi_{r,j} &:= \Pi_j - \Pi_{d,j},\end{aligned}$$

where $\Pi_{d,j}$ is the amount dividends paid out by the productive sector and $\Pi_{r,j}$ is the remaining profits retained by the productive sector after the payment of dividends. Here, $\Delta(\cdot)$ is an increasing, smooth function bounded between 0 and 1. We note that the size of dividends paid out in region j depends on both the nominal profit share π_j and the nominal output $p_j Y_j$, which already accounts for damages due to climate change and abatement costs.

Having now specified how profits are distributed between dividends and retained earnings, and how investment affects capital accumulation, we have all the necessary components to model the change in nominal debt. The dynamics of change in debt accumulation in region j are governed by the following ODE:

$$\dot{D}_j = p_j I_j - \Pi_{r,j} - p_j \bar{\delta}_j K_j. \quad (2.1)$$

Debt increases when firms pay dividends ($\Pi_{d,j}$) and invest ($p_j I_j$), as these activities require borrowing. Conversely, debt decreases when firms generate nominal profits (Π_j) and account for capital depreciation ($p_j \bar{\delta}_j K_j$), which reduces the need for borrowing.

2.2.3 Labour, Inflation, and Wages

In this section, we discuss how the workforce evolves over time, what drives inflation, and how this impacts the changes in wages in economy j . We also define key macroeconomic variables.

First, we define the working-age population in each region N_j , and give its evolution over time. The change in economy j 's workforce N_j is assumed to follow a sigmoidal function, where $\bar{q}_j$ is the speed of convergence of the workforce N_j to its upper limit $\bar{P}_j^N$:

$$\frac{\dot{N}_j}{N_j} := \bar{q}_j \left(1 - \frac{N_j}{\bar{P}_j^N} \right).$$

Next, in order to analyse the performance of the each economy, we define the same economic ratios as in [Bovari et al. \(2018a\)](#):

$$\lambda_j := \frac{L_j}{N_j}, \quad \omega_j := \frac{w_j L_j}{p_j Y_j}, \quad \text{and} \quad d_j := \frac{D_j}{p_j Y_j}.$$

The employment rate λ_j is the ratio of total employment relative to the size of the workforce in economy j , ω_j represents the wage share in economy j , which is the proportion of total output paid to workers, and d_j measures the proportion of debt D_j compared to the total output of the economy.

These ratios are important indicators of the macroeconomic health of each economy. For example, if the ratio λ_j is approaching zero, this signals rising unemployment, which may indicate that the economy is experiencing significant distress or is on the verge of collapse.

Inflation in each region is driven by the wage share as follows:

$$i_j := \frac{\dot{p}_j}{p_j} := \bar{\eta}_{p_j} (\bar{\xi}_j \omega_j - 1). \quad (2.2)$$

Here $\bar{\eta}_{p_j}$ is a speed adjustment term that controls how quickly inflation responds to changes in wage shares. The term $\bar{\xi}_j$ represents a markup factor that the private sector applies to prices. If $\bar{\xi}_j \omega_j > 1$, then there are rising prices in the economy and thus inflation. Conversely, if $\bar{\xi}_j \omega_j < 1$, then prices are decreasing, leading to deflation.

Finally, we define the short-run wage dynamics in each region j , based on [Bolker et al. \(2021\)](#):

$$\frac{\dot{w}_j}{w_j} := \phi(\lambda_j) + \bar{\gamma}_j i_j.$$

The first term is an increasing function $\phi(\cdot)$ which represents the Phillips curve. It implies that as λ_j (the employment rate in economy j) increases, the bargaining power of workers also increases, resulting in upward pressure on wages. The second term is an inflation indexation component, where i_j is the inflation rate and $\bar{\gamma}_j \in [0, 1]$ reflects the degree of money illusion. A higher $\bar{\gamma}_j$ indicates that wages are more strongly adjusted in response to inflation to maintain purchasing power.

2.2.4 Consumption and Trade

In our two-economy model, part of the domestic economy's consumption is driven by the demand for goods from abroad. This demand is, once again, influenced by exchange rate dynamics. We model the proportion of consumption imported from the foreign economy using a similar γ function to the one introduced in the investment case. Let

$$C_S = C_S^f + C_S^d \quad \text{where} \quad C_j^f = \gamma_{C_j}(e_R) \cdot C_j,$$

where the allocation functions are defined as follows

$$\begin{aligned}\gamma_{C_S}(e_R) &:= \min \left\{ 1, \max \left\{ 0, \bar{\gamma}_{C_S} - \bar{\beta}_{C_S} \cdot (e_R - \bar{e}_R) \right\} \right\}, \\ \gamma_{C_N}(e_R) &:= \min \left\{ 1, \max \left\{ 0, \bar{\gamma}_{C_N} - \bar{\beta}_{C_N} \cdot (1/e_R - 1/\bar{e}_R) \right\} \right\}.\end{aligned}$$

Here, $\bar{\gamma}_{C_j}$ is the baseline propensity to import for consumption in each region, $\bar{\beta}_{C_j}$ are speed parameters, and we recall that e_R and $\bar{e}_R$ are the real and equilibrium real exchange rates, respectively.

We can now determine the imports and exports for the respective economies. In the South, they import their foreign-sourced consumption and investment and export the North's foreign-sourced consumption and investment:

$$Im_S = C_S^f + I_S^f, \quad \text{and} \quad Ex_S = C_N^f + I_N^f.$$

Given that there are only two economies in the model, this construction implies that:

$$Im_N = Ex_S \quad \text{and} \quad Ex_N = Im_S.$$

To determine exports and imports in terms of consumption and investment, we first need to solve for the consumption levels of the respective economies. This is done using the following two national income identity equations:

$$\begin{aligned}Y_N &= I_N + C_N + Ex_N - Im_N, \\ Y_S &= I_S + C_S + Ex_S - Im_S.\end{aligned}$$

We define the system of equations for consumption using a matrix representation:

$$\Gamma = \begin{bmatrix} 1 - \gamma_{C_N} & \gamma_{C_S} \\ \gamma_{C_N} & 1 - \gamma_{C_S} \end{bmatrix}, \quad Y = \begin{bmatrix} Y_N \\ Y_S \end{bmatrix}, \quad I = \begin{bmatrix} I_N - I_N^f + I_S^f \\ I_S - I_S^f + I_N^f \end{bmatrix},$$

where Y_N, Y_S are the outputs in the North and South, I_N, I_S are total investments in each region, I_N^f, I_S^f are foreign investments received by the North and South, respectively. Let C be the vector of total consumption in the North and South. Then we compute consumption in North and South as

$$C = \Gamma^{-1} \cdot (Y - I).$$

With this, we have fully characterised all macroeconomic quantities in our two-region model.

2.3 Climate Module

The climate module is based on Bovari et al. (2018a)'s continuous time version of the DICE model (Nordhaus, 2018). We do not differentiate between regional climates for the North and South economies, as climate impacts, such as temperature change and CO₂ concentrations are modelled globally. This reflects the assumption that if one region pollutes, it will affect the climate for both regions.

Industrial emissions in each region depend on raw production Y_j^0 , the emissions intensity σ_j , and the emissions reduction rate n_j :

$$E_{ind_j} := Y_j^0 \sigma_j (1 - n_j).$$

Total global emissions are then determined as the sum of industrial emissions in each region and land use emissions, E_{land} , the latter of which is an exogenous variable in our model:

$$E := E_{ind_N} + E_{ind_S} + E_{land}.$$

To capture how emissions evolve over time, we introduce dynamics for both land-use emissions and emission intensity. Land-use emissions decrease exogenously at a constant rate, while emission intensity and its rate of change are modelled endogenously. These dynamics are specified as follows:

$$\begin{aligned} \frac{\dot{E}_{land}}{E_{land}} &:= \bar{\delta}_{E_{land}} < 0, \\ \frac{\dot{\sigma}_j}{\sigma_j} &:= g_{\sigma,j} \quad \text{and} \quad \frac{\dot{g}_{\sigma,j}}{g_{\sigma,j}} := \bar{\delta}_{g_{\sigma,j}}. \end{aligned}$$

The term $\dot{\sigma}_j$ governs the change in emission intensity, which may differ between regions. The term $\dot{g}_{\sigma,j}$ governs the change in the rate of emissions reduction.

The rest of the climate module is identical to the continuous-time DICE model. The carbon cycle is represented by Equation (2.3), which describes how carbon emissions in the atmosphere, upper ocean, and deep ocean evolve over time. This process is governed by a system of ordinary differential equations, whereby CO₂ is redistributed across the three levels according to the diffusion matrix Φ (Bovari et al., 2018a).

$$\begin{pmatrix} \dot{CO}_2^{AT} \\ \dot{CO}_2^{UP} \\ \dot{CO}_2^{LO} \end{pmatrix} := \begin{pmatrix} E \\ 0 \\ 0 \end{pmatrix} + \Phi \begin{pmatrix} CO_2^{AT} \\ CO_2^{UP} \\ CO_2^{LO} \end{pmatrix}, \quad (2.3)$$

where Φ is given by

$$\Phi := \begin{pmatrix} -\bar{\phi}_{12} & \bar{\phi}_{12}\bar{C}_{UP}^{AT} & 0 \\ \bar{\phi}_{12} & -\bar{\phi}_{12}\bar{C}_{UP}^{AT} - \bar{\phi}_{23} & \bar{\phi}_{23}\bar{C}_{LO}^{UP} \\ 0 & \bar{\phi}_{23} & -\bar{\phi}_{23}\bar{C}_{LO}^{UP} \end{pmatrix},$$

where

$$\bar{C}_{UP}^{AT} := \frac{\bar{C}_{pind}^{AT}}{\bar{C}_{pind}^{UP}}, \quad \bar{C}_{LO}^{UP} := \frac{\bar{C}_{pind}^{UP}}{\bar{C}_{pind}^{LO}},$$

and $\bar{\phi}_{12}, \bar{\phi}_{23}, \bar{C}_{pind}^{AT}, \bar{C}_{pind}^{UP}$ and $\bar{C}_{pind}^{LO}$ are given parameters.

Equation (2.4) and Equation (2.5) calculate radiative forcing, which is a measure of the difference between the amount of energy that enters the Earth's atmosphere and the amount of energy that leaves it (MIT Climate Portal, n.d.). This is due to the build-up of greenhouse gases, which affects energy levels in the atmosphere (Bovari et al., 2018a).

$$F := F_{ind} + F_{exo}, \quad (2.4)$$

$$F_{ind} := \frac{\bar{F}_{2 \times CO_2}}{\log(2)} \log \left(\frac{CO_2^{AT}}{\bar{C}_{pind}^{AT}} \right), \quad (2.5)$$

If $F > 0$, there will be a warming effect, and if $F < 0$, there will be a cooling effect. We distinguish between forcing caused by industrial CO_2 emissions, such as fossil fuels, and exogenous forcing, which originates outside the model. In Equation (2.5), $\bar{F}_{2 \times CO_2}$ is a parameter that represents the increase in radiative forcing resulting from a doubling of the pre-industrial CO_2 concentration. The atmospheric CO_2 concentration, CO_2^{AT} , is obtained from the carbon cycle ODE, whose dynamics are presented. Note, $\bar{C}_{pind}^{AT}$ is the pre-industrial levels of CO_2 emissions in the atmosphere (Bovari et al., 2018a).

The impacts of forcing on global temperatures are showcased with the following temperature dynamics:

$$\dot{T} := \frac{1}{\bar{C}} (F - \rho T - \bar{\gamma}^*(T - T_{LO})), \quad (2.6)$$

$$\dot{T}_{LO} := \frac{\bar{\gamma}^*}{\bar{C}_{LO}} (T - T_{LO}). \quad (2.7)$$

Equation (2.6) is the ordinary differential equation (ODE) that governs the change in global temperature across the land, atmosphere, and upper ocean. Equation (2.7) represents the change in the temperature of the deep ocean. These two ODEs interact with each other: as radiative forcing increases, the surface temperature T rises. However, part of this heat is transferred to the deep ocean, which reduces the rate




Figure 2.2: A diagram of the flows between key variables in our two-region climate-economic model

of surface warming. This exchange is governed by the exchange rate coefficient γ^* , while ρ represents the rate of heat loss to space. Therefore, as the atmosphere, land, and upper ocean warm, heat is gradually transferred to the deeper ocean, which in turn causes the deep ocean temperature T_{LO} to rise over time (Bovari et al., 2018a).

2.4 Summary

All together, we have a two region climate-economic model with 28 state variables: 10 for each region and 8 for the world. Figure 2.2 gives an overview of the interactions in our model. It shows how the two economic models for North and South both feed into the climate model via emissions E , and interact with each other via the exchange rate e^n , interest rate parity r_S , and exports Ex and imports Im .

Finally, a stock-flow consistency table for our model is given in the Appendix in Table A.4. The table indicates the main transactions between firms, households, and governments in the two regions, ensuring all financial flows are accounted for and balanced using a double entry system. It tracks key transactions like consumption, investment, wages, taxes, and importantly trade (imports/exports), with each entry showing how resources move between sectors while maintaining a zero-sum equilibrium.

Chapter 3

Analysis of Results

3.1 Initial Calibration

We compare the simulation results of our model over a 9-year period (from 2016 to 2024) to real-world data over the same period. The model is calibrated using parameter estimates from 2016, to match [Bovari et al. \(2018a\)](#), who used parameter estimates from the 2016 DICE model. Since the DICE model is forward-looking, we can compare real-world data from 2016 onward to the performance of the model. If the simulation results closely match the observed data it will indicate the model is appropriately calibrated and analysis of the models can be performed.

[Figure 3.1](#) plots the simulation path given by the model (represented by the lines) against the observed data (represented by points) for industrial emissions E_{ind} , population N , temperature T and real GDP, Y . As observed from the plots, our model is able to closely match the observed data, and almost perfectly matches population. For real GDP, we match the overall trend, although in the latter years we slightly overestimate the level. This is primarily due to the global reduction in real output in 2020 caused by the COVID-19 pandemic and associated lockdowns. Regarding industrial emissions, E_{ind} , our model captures the downward trajectory in the North, driven primarily by European investment in green technologies and the introduction of carbon pricing. Conversely, the model also captures the rise in industrial emissions in the South, slightly overestimating the level, which can again be attributed to the impact of the COVID-19 pandemic. For temperature, the model closely matches observed levels in the early years but underestimates the level in more recent years, as global temperatures have risen rapidly since the beginning of the decade. From this analysis, we conclude that our model has been appropriately




Figure 3.1: Simulated model trajectory (line) vs. observed data (points), 2016–2024.

calibrated to the data and is capable of replicating observed real-world dynamics.

Given this initial calibration and the estimates from real data as described in [Appendix 1.2](#), we compute a set of initial conditions and model parameters which we use for all of the following scenarios. [Table A.1](#) gives the initial conditions used in our simulations while [Table A.2](#) and [Table A.3](#) give the parameter values.

3.2 Scenario Analysis

We implement a different set of scenarios for the two regions with the aim to examine how they impact each other and their respective economies. In our framework, a different scenario represents a set of factors we can vary, namely how large the damages are in each economy; this can be adjusted between Nordhaus, Weitzman, and Stern. Another scenario we vary is the level of carbon pricing in each economy, which can vary from none, low, or high. Lastly, as an additional policy, each region

can introduce a subsidy for the abatement cost ranging from 0 to 1. We initially look at three policy scenarios.

The first scenario is the one which is most applicable to what we are currently observing, and hence it is named *business as usual*. The second scenario is called *best case scenario* which reflects the policy framework where both economies take climate change risk seriously. Finally, the third policy is when neither economy takes climate risk seriously and is called *worst case scenario*. The three policies are summarised in [Table 3.1](#). What our model allows us to now see, which was not previously possible, is how policy decisions in one economy affect the other and vice versa.

Table 3.1: Summary of Scenario Settings Across Scenarios

Policy	Scenario 1 (Business-as-Usual)	Scenario 2 (Best Case)	Scenario 3 (Worst Case)
Damage Function (North)	Nordhaus	Nordhaus	Weitzman
Damage Function (South)	Stern	Stern	Stern
Carbon Pricing (North)	Low	High	None
Carbon Pricing (South)	None	Low	None
Subsidy (North)	0.25	0.50	0
Subsidy (South)	0	0.25	0

[Figure 3.2](#) shows the performance of key economic and climate variables through to the year 2175 under **Scenario 1**, which, as discussed above, represents the *business-as-usual* scenario. In this configuration, climate damages in the North are modelled using the Nordhaus damage function, while damages in the South follow the Stern specification. Nordhaus corresponds to low damages from climate change, whereas Stern means there are high damages from climate. The motivation behind this low-high configuration is that climate impacts on the Southern economy are more severe than in the North, largely because the North possesses more resources to respond to climate disasters and reduce its impact.

The carbon pricing regime is set to *low* in the North and *none* in the South. This reflects the observation that developed economies have implemented some carbon pricing schemes (mainly European nations). In contrast, most developing countries, particularly in the Global South, have yet to introduce carbon pricing in any significant way ([Maplecroft, 2024](#)). Lastly, the North is assumed to subsidize 25% of abatement costs, while the South provides no such subsidy.

Under this *business as usual* scenario, the trajectory of key variables is deeply concerning. As shown in the real GDP (Y) plot, the Southern economy essentially




Figure 3.2: Scenario 1: business as usual

collapses by 2175. This is accompanied by mass deflation, evidenced by the inflation rate (red line) going negative. Similarly, the employment ratio (λ) also falls to zero. Furthermore, the nominal exchange rate (e^n) collapses to zero. This outcome is driven by the severe deflation and collapse of the Southern economy, rendering its currency effectively worthless relative to the North. On the climate front, the global temperature exceeds 3.5°C by 2100 and reaches over 5°C by 2150, entirely missing the internationally agreed target of limiting warming to 2°C. In contrast, macro-economic indicators in the North appear relatively resilient. Real GDP continues to grow, the employment ratio λ oscillates at a stable level, and inflation stays stable. This suggests that the economic activity of the Northern economy is stable despite the massive increases in temperature and the Southern economy collapsing.

These findings raise critical concerns. Our model suggests that developed economies, under their current climate policy frameworks, can potentially reduce emissions to near zero before 2100 and still sustain economic growth even while the economies of developing regions collapse. This raises a serious question of whether or not the North will be willing to support the South's transition to cleaner green technologies

and take climate change seriously if it won't impact their economies as much.

In [Figure 3.3](#) are the results of the key climate and macro-economic variables under **Scenario 2** which represents the *best case scenario*. Under this scenario, we have followed the same damages function as in policy one. However, in this policy regime, we assume that both regions implement carbon pricing at a high level and introduce subsidies for abatement costs which is set at 0.5 in the North and 0.25 in the South, respectively.

When looking at plots we can see that if aggressive carbon pricing and relatively high subsidies are implemented both regions reduce their industrial emissions to zero, with the Northern economy reaching zero emissions by 2050 and South after 2100. Furthermore, global temperature rise plateaus by 2175 around a 3.25 °C increase. Unfortunately, the goal of maintaining 2 °C increase before the year 2100 is not achieved.

The macro-economic variables also perform well as both regions' real GDP continue to grow over the period of simulation. The employment ratio λ is stable;




Figure 3.3: Scenario 2: best case

meanwhile, the size of the variation in the employment ratio of the South also stabilizes (note: the oscillations represent the business cycle). The exchange rate keeps growing between the two regions, which is due to inflation in the South increasing.

The results from **Scenario 2** suggest that if both economies take climate change seriously and implement carbon pricing and subsidies, it is possible to reduce and eventually halt the rise in temperatures. However, the target of limiting warming to 2.5°C before 2100 is still missed. Interestingly, the economic outcomes for the South are significantly better than in the *business as usual* case. Not only does the Southern economy avoid collapse, but it also experiences growth in real output. The Northern economy also performs well under this scenario. Therefore, based on these results, it is possible for the Southern economy not only to avoid collapse but also to grow. However, this would require the implementation of subsidies and carbon pricing, which may be fiscally and politically challenging for the region due to constrained resources.




Figure 3.4: Scenario 3: worst case

Figure 3.4 showcases the results of the climate and macro-economic variables under the *worst-case scenario* where neither economy does anything about climate change and in the North the damages of climate change are worse than expected as

represented by the damages function set to Weitzman which represents a medium level of damages. It is reasonable to expect in the worst-case scenario that the ability of the North to use its resources to mitigate the impact of climate disasters is not as effective as previously thought, which is modeled using the Weitzman function.

This scenario is very concerning as temperature rises to 6 °C and looks to keep on increasing beyond the year 2175 with emissions still increasing in the North but emissions decreasing in the South due to its economy collapsing.

Looking at the macro-economic variables we see that the real GDP as in scenario 1 goes to zero for the South as well the employment share, wage share going to zero; there is also serious deflation. However, in this scenario we can see that the Northern economy starts to experience deflation after 2150 and its real GDP starts to plateau, suggesting that its economy is slowing and can potentially collapse.

Scenario 3 shows that if the North and South do nothing about rising temperatures, they will both be impacted, with the South decreasing and the North experiencing very worrying key macro-economic trends such as deflation.

3.3 Case Study: Intervention by the North to Save Both Regions

We now allow for a new scenario: the North may incentivise the South to enact climate policies by means of a subsidy, which the North pays to the South. In particular, the South's abatement subsidy is now paid as a result of a reduction in the consumption of the North, i.e.,

$$\tilde{C}_N = C_N - s_{A_S} A_S Y_S^0,$$

where $\tilde{C}_N$ denotes the North's consumption after paying the subsidy.

We now present a scenario where if the North does not help the South via subsidies, the collapse in the South could impact them. In this setup, both regions face high climate damages. The North adopts a high carbon tax and subsidizes abatement at a rate of 0.5, while the South implements neither a carbon tax nor an abatement subsidy. The simulation results are presented in [Figure 3.5](#).


From the climate perspective, this scenario mirrors the business-as-usual case: global temperature begins to plateau only around 5.25°C. This outcome is primarily driven by continued high emissions in the South, even as the North successfully reduces its emissions to zero.

The macroeconomic results are particularly interesting. When both regions are exposed to high damages, real GDP (Y) for both economies starts to decline sharply around 2100. This suggests that the collapse of the Southern economy imposes spillover effects that drag down the Northern economy as well. Further evidence of systemic failure can be seen in the collapse of price levels in both regions (inflation goes negative), indicating deflationary pressure consistent with broad economic breakdown.


Looking at [Figure 3.5b](#), we examine the scenario titled *North helps South*. In this scenario, the Northern economy actively supports the South by subsidizing its abatement efforts. The policy settings are as follows: the North adopts a high carbon tax and subsidizes its own abatement at a rate of 0.5, while the South implements a low carbon tax and receives a subsidy of 0.75 from the North. Both regions are subject to high climate damages.

The simulation results provide strong support for this policy configuration. As with the best-case scenario, global emissions fall to zero by 2075, and global temperatures plateau at approximately 2.75°C by 2175. On the economic front, both regions experience robust growth. Real GDP in the North reaches 1250 by the end of the simulation period, the highest among all policy scenarios, while the South's GDP climbs to 750, also a peak value across all five scenarios.

Other macroeconomic indicators further reinforce the strength of this policy. The North maintains a stable employment ratio, while the South shows a clear reduction in the volatility of employment over time. Profits in the productive sectors remain healthy: in the North, they stabilize around 0.3, and in the South, they show signs of convergence. Inflation in both regions follows a cyclical pattern that remains bounded and stable throughout the simulation.



(a) North in isolation



(b) North helps South

Figure 3.5: Case study: intervention by the North to save both regions




Figure 3.6: Consumption of the Northern economy

Taken together, these results position the *North helps South* scenario as the most favourable among those tested, delivering strong outcomes both from an economic and climate viewpoint. Importantly, the North’s act of subsidising Southern abatement does not impose a noticeable economic burden. In fact, analysing [Figure 3.6](#), we observe that consumption, which we use as a proxy for living standards, remains relatively stable and performs well in the non-cooperative scenario where the North does not subsidize the South. This suggests that the North’s living standards are not significantly compromised even without providing support. Furthermore, in this scenario, although consumption is slightly lower in the short run, it eventually surpasses that of the non-cooperative case as consumption collapses in that case. This indicates that mutual cooperation not only supports Southern abatement but also leads to higher long-term living standards in the North. In fact, the North performs better under this cooperative policy than in scenarios where it acts alone. This makes a compelling case for Northern support of Southern climate action, both from the standpoint of economic efficiency and from a moral perspective.

3.3.1 Monte Carlo

To further our analysis, we perform Monte Carlo simulations for our case study. Following [Bovari et al. \(2018b\)](#) and [Bolker et al. \(2021\)](#), we assume certain key parameters are drawn from distributions. We focus on two groups of parameters: the wage-inflation parameters $\bar{\eta}_p$ and $\bar{\gamma}_j$ and the climate parameters $\bar{S}$ and $\bar{C}_{pind}^{UP}$.

In particular, we assume the speed adjustment term from the inflation equation, $\bar{\eta}_p$, is drawn from a normal distribution with mean 0.4 and standard deviation 0.12, while money illusion parameter $\bar{\gamma}_j$ is drawn from a generalized Gamma distribution with shape parameter $s = 6.2327$, scale parameter $m = 0.0033$, and family

parameter $f = 0.3158$. In terms of the climate parameters, we assume the equilibrium climate sensitivity $\bar{S}$ is drawn from a log-normal distribution with log mean 1.107 and log standard deviation 0.264 and the intermediate climate reservoir $\bar{C}_{pind}^{UP}$ is also drawn from a log-normal distribution, with log mean 5.89 and log standard deviation 0.25.




Figure 3.7: Monte Carlo simulation output for two scenarios: the North has only internal climate policy and no inter-region subsidy (green), and the North pays a subsidy to the South for them to enact carbon pricing (purple)

Figure 3.7 shows the result of 1000 Monte Carlo runs for our key intervention scenario, with the simulation run until 2100. The plot shows the median of the runs for the employment rate in North and South, inflation rate in North and South, the total global emission, and the temperature deviation. The upper 95th and lower 5th quantiles are shown by the filled areas. We observe that crashes in both the North and South are much more likely when the North focuses its climate action only on itself, even on the shortened timescale of 84 years. This gives further evidence

of the key effect of the North's intervention in stabilizing global temperatures and avoiding an economic collapse in both regions.

Chapter 4

Conclusion

4.1 Summary

In this report, we extended the GEMMES framework to a two-region setting by introducing mechanisms that allow for interaction between regions through investment, trade, and exchange rate dynamics. This was achieved by incorporating two additional differential equations governing the nominal and expected exchange rates. We then aggregated data from 34 countries into two representative regions and calibrated the model using these initial conditions.

We then implemented a range of policy scenarios to examine how the economic performance of each region influences the other, as well as their joint impact on the climate. The scenarios included business as usual, best case, and worst case. In the business as usual scenario, which reflects current global policy trends, the 2°C temperature target is comfortably missed by 2100. The Southern economy collapses, while the Northern economy remains relatively resilient. This outcome raises a critical moral question: if the collapse of the South does not materially affect the North, will the North have sufficient incentive to support the South's transition to a low-carbon economy?

We also examined a best case scenario, where both regions take aggressive action through carbon pricing and abatement subsidies and a worst case scenario, where neither region implements meaningful climate policies. In the best case scenario, both economies continue to grow and temperature increases are significantly slowed, although the 2°C target is still missed. In contrast, the worst case scenario leads to economic collapse in both regions, driven by runaway temperature in-

creases and widespread systemic breakdown.

As a final set of scenarios, we examined two important cases. In the first, we explored what happens when the effects of climate change on the North are more severe than previously anticipated, and the North acts in isolation without supporting the South. In this case, although the North initially reduces its own emissions, the continued high emissions from the South lead to global temperature rise. As the Southern economy collapses under the pressure of climate damages, it generates systemic spillovers that ultimately trigger economic decline and collapse in the North as well.

In contrast, the second scenario considers a cooperative approach where the North subsidizes the South's abatement costs. This leads to significantly better outcomes. Emissions fall to zero, temperature stabilizes at the same level as in the best case scenario, and both regional economies experience strong, sustained growth. This highlights that supporting global mitigation efforts not only delivers climate benefits but also protects long-term economic stability in the North.

4.2 Model Extensions and Future Work

We now discuss a possible extension for future work on this project. One area for improvement is to strengthen the linkages between the two economies. Currently, they are connected through an exchange rate mechanism as well as exports and imports. This also feeds into an interest rate mechanism. This is a significant improvement on the GEMMES model, and allows us to capture interdependencies in both the economies and the climate of the two regions. However, the interconnections could be strengthened to better capture the complexity of real-world economic interdependencies.

In our model, each economy allocates imports a portion of its consumption, both based on the exchange rate. However, this structure, modeled using a gamma function, allows the allocation of investment and consumption to vary continuously between 0 and 1. This implies that if the foreign economy collapses, the domestic economy will fully reallocate its consumption to domestic sources.

The issue with this assumption is that it overlooks product differentiation and import dependence. In reality, many goods are not available domestically and must be imported. Therefore, if the foreign economy collapses, it is unrealistic to assume that all consumption can simply shift to domestic markets without any loss. Some goods and services would become unavailable, leading to reduced consumption.

Future research should incorporate these constraints. Doing so would enhance the realism of the model and deepen our understanding of economic interdependencies under climate stress.

Appendix A

Additional Details

Variable	Value	Description
North		
Y_N	43.41	Real output in the North
d_N	0.82	Debt ratio in the North
λ_N	0.70	Employment level in the North
ω_N	0.56	Wage share in the North
N_N	0.60	Working-age population level in the North
g_{σ_N}	-0.0152	Growth rate of emission intensity in the North
n_N	0.03	Emissions reduction rate in the North
p_N	1	Price level in the North
E_{ind_N}	10.96	Industrial emissions in the North
pBS_N	446.67	Price of green (backstop) technology in the North
South		
Y_S	21.00	Real output in the South
d_S	0.80	Debt ratio in each region
λ_S	0.65	Employment level in the South
ω_S	0.54	Wage share in the South
N_S	2.48	Working-age population level in the South
g_{σ_S}	-0.0152	Growth rate of emission intensity in the South
n_S	0.03	Emissions reduction rate in the South
p_S	1	Price level in the South
E_{ind_S}	21.61	Industrial emissions in the South
pBS_S	600.00	Price of green (backstop) technology in the South
World		
E_{land}	2.6 Gt CO ₂ -e	Exogenous land-use based emissions
CO_2^{AT}	851 Gt C	CO ₂ -e concentration in the atmosphere
CO_2^{UP}	460 Gt C	CO ₂ -e concentration in the upper ocean and biosphere
CO_2^{LO}	1740 Gt C	CO ₂ -e concentration in the lower ocean
T	0.85°C	Mean temp. of the atmosphere, land surface, and upper ocean
T_{LO}	0.0068°C	Mean temp. of the deeper ocean
e^n	6.6036	Nominal exchange rate
e^{exp}	6.6036	Expected nominal exchange rate

Table A.1: List of initial conditions used for simulations in our two-region ecological macroeconomic model. These are converted to the state variables used in the model.

Parameter	Value	Description
North		
$\bar{P}_N^N$	1	Ceiling on Northern workforce, in billions
$\bar{q}_N$	0.01	Speed of growth of Northern workforce dynamics
$\bar{\alpha}_N$	0.02	Growth rate of Northern labour productivity (const spec)
$\bar{\delta}_N$	0.04	Depreciation rate of capital
$\bar{\nu}_N$	2.7	Capital to output ratio
r_N	0.02	Real interest rate
$\bar{\eta}_{pN}$	0.2	Adjustment speed for prices
ξ_N	1.6	Mark-up rate for inflation
$\bar{\gamma}_N$	0.9	Money illusion
$\bar{\phi}_{0N}$	-0.292	Constant of linear Phillips curve
$\bar{\phi}_{1N}$	0.469	Slope of linear Phillips curve
$\bar{\kappa}_{0N}$	0.0318	Constant of investment function
$\bar{\kappa}_{1N}$	0.575	Slope of investment function
$\bar{\kappa}_{minN}$	0	Min of investment function
$\bar{\kappa}_{maxN}$	0.3	Max of investment function
$\bar{\Delta}_{0N}$	-0.078	Constant of dividend function
$\bar{\Delta}_{1N}$	0.553	Slope of dividend function
$\bar{\Delta}_{minN}$	0	Min of dividend function
$\bar{\Delta}_{maxN}$	0.3	Max of dividend function
$\bar{\delta}_{pBSN}$	-0.0051	Growth rate of the backstop tech price
$\delta_{g\sigma N}$	-0.001	Variation rate of the growth of emission intensity
$\bar{\gamma}_{IN}$	0.2	Baseline investment share in foreign goods
$\bar{\beta}_{IN}$	0.01	Foreign investment sensitivity parameter
$\bar{e}_R$	0.6667	Equilibrium exchange rate in foreign propensity functions
$\bar{\gamma}_{CN}$	0.1	Baseline consumption share of foreign goods
$\bar{\beta}_{CN}$	0.01	Foreign consumption sensitivity parameter
South		
$\bar{P}_S^N$	4	Ceiling on Southern workforce, in billions
$\bar{q}_S$	0.02	Speed of growth of Southern workforce dynamics
$\bar{\alpha}_S$	0.02	Growth rate of Southern labour productivity (const spec)
$\bar{\delta}_S$	0.04	Depreciation rate of capital
$\bar{\nu}_S$	2.7	Capital to output ratio
$\bar{\eta}_{pS}$	0.2	Adjustment speed for prices
ξ_S	1.6	Mark-up rate for inflation
$\bar{\gamma}_S$	0.9	Money illusion
$\bar{\phi}_{0S}$	-0.292	Constant of linear Phillips curve
$\bar{\phi}_{1S}$	0.469	Slope of linear Phillips curve
$\bar{\kappa}_{0S}$	0.1	Constant of investment function
$\bar{\kappa}_{1S}$	1	Slope of investment function
$\bar{\kappa}_{minS}$	0	Min of investment function
$\bar{\kappa}_{maxS}$	0.5	Max of investment function

Table A.2: List of parameters used for simulations in our two-region ecological macroeconomic model.

Variable	Value	Description
$\bar{\Delta}_{0S}$	0.3	Constant of dividend function
$\bar{\Delta}_{1S}$	0.5	Slope of dividend function
$\bar{\Delta}_{minS}$	0	Min of dividend function
$\bar{\Delta}_{maxS}$	0.5	Max of dividend function
$\bar{\delta}_{PBS_S}$	-0.0051	Growth rate of the backstop tech price
$\bar{\delta}_{g\sigma_N}$	-0.001	Variation rate of the growth of emission intensity
$\bar{\gamma}_{IS}$	0.5	Baseline investment share in foreign goods
$\bar{\beta}_{IS}$	0.05	Foreign investment sensitivity parameter
$\bar{e}_R$	1.5	Equilibrium exchange rate in foreign propensity functions
$\bar{\gamma}_{C_N}$	0.15	Baseline consumption share of foreign goods
$\bar{\beta}_{C_N}$	0.01	Foreign consumption sensitivity parameter
World		
$\bar{\beta}_e$	0.05	Exchange rate sensitivity parameter
$\bar{\beta}_e^{ex}$	0.95	Expected exchange rate sensitivity parameter
$\bar{t}_S$	0.01	Southern risk premium
$\bar{C}_{preind}^{AT}$	588	Preind. concentration of CO ₂ in the atmosphere, in Gt C
$\bar{C}_{preind}^{UP}$	360	Preind. concentration of CO ₂ in the upper ocean, in Gt C
$\bar{C}_{preind}^{LO}$	1720	Preind. concentration of CO ₂ in the lower ocean, in Gt C
$\bar{\phi}_{23}$	0.0013409	Transfer coefficient for carbon from UP to LO
$\bar{\delta}_{E_{land}}$	-0.022	Growth rate of land use change CO ₂ -e emissions
$\bar{F}_{dbl}$	3.6813	Change in radiative forcing from doubling of preind CO ₂
$\bar{F}_{exo}^{start}$	0.5	Initial value of exogenous radiative forcing
$\bar{F}_{exo}^{end}$	1	Initial value of exogenous radiative forcing
$\bar{T}_{preind}$	13.74	Preindustrial temperature, in degrees Celsius
$\bar{C}$	49.76	Heat capacity of AT and UP
$\bar{C}_{LO}$	3.52	Heat capacity of the lower ocean layer
$\bar{\gamma}^*$	0.0176	Heat exchange coefficient between temperature layers, in SI
$\bar{S}$	3.1	Equilibrium climate sensitivity, in degrees Celsius
$\bar{\xi}_1$	0	Damage function parameter
$\bar{\xi}_2$	0.00236	Damage function parameter
$\bar{\xi}_3$	0.00000507	Damage function parameter - Weitzman
$\bar{\xi}_3$	0.0000819	Damage function parameter - Dietz and Stern
$\bar{\theta}$	2.6	Abatement cost function parameter

Table A.3: List of parameters used for simulations in our two-region ecological macroeconomic model.

1.1 Stock-Flow Consistency Table

	NORTH			
	FIRMS		HOUSEHOLDS	GOVERNMENT
	CURRENT	CAPITAL		
Transactions				
Consumption	$C_N^d + C_S^f$		$-C_N$	0
Investment	$I_N^d + I_S^f$	$-I_N$		0
Wage Cost	$-W_N L_N$		$W_N L_N$	0
Capital Depreciation	$-\delta_N K_N$	$\delta_N K_N$		0
Carbon Tax	$-T_{fN}$		T_{fN}	0
Abatement Subsidies	S_{A_N}			$-S_{A_N}$
Dividends	$-\Pi_{dN}$		Π_{dN}	0
Imports	$-I_N^f$		$-C_N^f$	$-I_N^f - C_N^f$
Exports	I_S^f		C_S^f	$I_S^f + C_S^f$
SOUTH				
Transactions				
Consumption	$C_S^d + C_N^f$		$-C_S$	0
Investment	$I_S^d + I_N^f$	$-I_S$		0
Wage Cost	$-W_S L_S$		$W_S L_S$	0
Capital Depreciation	$-\delta_S K_S$	$\delta_S K_S$		0
Carbon Tax	$-T_{fS}$		T_{fS}	0
Abatement Subsidies	S_{A_S}			$-S_{A_S}$
Dividends	$-\Pi_{dS}$		Π_{dS}	0
Imports	$-I_S^f$		$-C_S^f$	$-I_S^f - C_S^f$
Exports	I_N^f		C_N^f	$I_N^f + C_N^f$
Sum				0

Table A.4: Stock-flow consistency table for the North-South model

1.2 Data Collection and Estimation

This study uses four main databases to gather time series data for 34 countries, grouped into two regions (North and South). The data sources are given in [Table A.5](#). These countries follow the same sample used in [Bovari et al. \(2018a\)](#), with the addition of Brazil. The division of the countries into the Global North and the Global South is as follows:

- **North** (developed economies): Canada, Denmark, Finland, France, Germany, Greece, Hungary, Ireland, Israel, Italy, Japan, Netherlands, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, United Kingdom, United States.
- **South** (emerging markets): Argentina, Brazil, Chile, China, India, Indonesia, Malaysia, Mexico, South Africa, Thailand, Turkey.

We note that exchange rates are converted from the local currency unit per US\$ to US\$ per local currency unit. The U.S. CPI (base year 2010) is adjusted to obtain a base 2015 index. Since the Penn World Table data ends in 2019, labour share data for 2020–2024 is sourced from the United Nations database.

Next, from the raw data, the following variables are constructed:

$$\lambda = \frac{\text{Employment-to-population \%} \times (\text{Population 65+ \%} + \text{Population 15–64 \%})}{\text{Population 15–64 \%}}$$

$$N = \text{Population 15–64 \%} \times \text{Total Population}$$

$$L = \lambda \times N$$

$$wL = \text{Wage Share} \times \text{GDP}$$

$$w = \frac{wL}{L}$$

Regional nominal exchange rates (US\$ per LCU) are calculated as GDP-weighted averages of country-level exchange rates. These are used to derive a cross-region exchange rate (South per North currency unit). Inflation rates, nominal interest rates, and real interest rates are also computed at the regional level.

The backstop price levels are set such that the price in the South is assumed level of 600 in 2016, while the price in the North is set so that the 2016 global GDP-weighted average is 550. Both levels decline annually at a rate of 0.5%, following [Bovari et al. \(2018a\)](#).

Table A.5: Data Sources

Source	Description
World Bank	Provides macroeconomic variables including GDP (constant 2015 US\$), total population, population aged 15–64 (% of total), population aged 65+ (% of total), employment-to-population ratio (15+, total, %, modeled ILO estimate), gross fixed capital formation (% of GDP), nominal exchange rate (LCU per US\$, period average), consumer price index (base 2010), final consumption expenditure (% of GDP), imports (% of GDP), exports (% of GDP), and real exchange rate index. data.worldbank.org
Penn World Table	Provides the share of labour compensation in GDP at national prices until 2019. rug.nl/ggdc/productivity/pwt
UN DESA (SDG Database)	Provides SDG Indicator 10.4.1: Labour share of GDP (used for 2020–2024). unstats.un.org/sdgs/dataportal/database
Bank for International Settlements (BIS)	Contains data on labour share in GDP and central bank interest rates. data.bis.org
EDGAR	Supplies greenhouse gas emissions for all countries (up to 2023). edgar.jrc.ec.europa.eu/report_2024
NASA (Temperature Data)	Provides global temperature data. climate.nasa.gov/vital-signs/global-temperature

Bibliography

- Allen, M.R., Stocker, T.F., 2014. Impact of delay in reducing carbon dioxide emissions. *Nature Climate Change* 4, 23–26. doi:[10.1038/nclimate2077](https://doi.org/10.1038/nclimate2077).
- Barrage, L., Nordhaus, W., 2024. Policies, projections, and the social cost of carbon: Results from the DICE-2023 model. *Proceedings of the National Academy of Sciences* 121, e2312030121. doi:[10.1073/pnas.2312030121](https://doi.org/10.1073/pnas.2312030121).
- Bolker, B.M., Grasselli, M.R., Holmes, E., 2021. Short communication: Sensitivity analysis of an integrated climate-economic model. *SIAM Journal on Financial Mathematics* 12, SC44–SC57. doi:[10.1137/21M1404120](https://doi.org/10.1137/21M1404120).
- Bovari, E., Giraud, G., Mc Isaac, F., 2018a. Coping with collapse: A stock-flow consistent monetary macrodynamics of global warming. *Ecological Economics* 147, 383–398. doi:[10.1016/j.ecolecon.2018.01.034](https://doi.org/10.1016/j.ecolecon.2018.01.034).
- Bovari, E., Lecuyer, O., Mc Isaac, F., 2018b. Debt and damages: What are the chances of staying under the 2°C warming threshold? *International Economics* 155, 92–108. doi:[10.1016/j.inteco.2018.02.002](https://doi.org/10.1016/j.inteco.2018.02.002). special issue on ‘Social values of carbon and climate policy signals in the post-COP21 context’.
- Dietz, S., Stern, N., 2015. Endogenous growth, convexity of damage and climate risk: How nordhaus’ framework supports deep cuts in carbon emissions. *The Economic Journal* 125, 574–620. doi:[10.1111/econj.12188](https://doi.org/10.1111/econj.12188).
- Gallagher, K.S., 2014. *The globalization of clean energy technology: Lessons from China*. MIT press.
- Godin, A., Yilmaz, S.D., 2020. *Modelling small open developing economies in a financialized world: A stock-flow consistent prototype growth model*. Technical Report. Agence française de développement.
- Godley, W., 1996. *Money, finance and national income determination: an integrated approach*. Technical Report. URL: <https://www.econstor.eu/bitstream/10419/186846/1/wp167.pdf>.

- Godley, W., Lavoie, M., 2006. Monetary economics: an integrated approach to credit, money, income, production and wealth. Springer. doi:[10.1007/978-1-137-08599-3](https://doi.org/10.1007/978-1-137-08599-3).
- Grasselli, M.R., Costa Lima, B., 2012. An analysis of the keen model for credit expansion, asset price bubbles and financial fragility. *Mathematics and Financial Economics* 6, 191–210. doi:[10.1007/s11579-012-0071-8](https://doi.org/10.1007/s11579-012-0071-8).
- Grasselli, M.R., Huu, A.N., 2015. Inflation and speculation in a dynamic macroeconomic model. *Journal of Risk and Financial Management* 8, 285–310. URL: <https://www.mdpi.com/1911-8074/8/3/285>, doi:[10.3390/jrfm8030285](https://doi.org/10.3390/jrfm8030285).
- Keen, S., 1995. Finance and economic breakdown: Modeling minsky's "financial instability hypothesis". *Journal of Post Keynesian Economics* 17, 607–635. URL: <http://www.jstor.org/stable/4538470>.
- Longaretti, P.Y., Martin, H.A., 2025. A variance-based sensitivity analysis of a goodwin-keen type economic model. URL: <https://hal.science/hal-04930500/document>.
- Maplecroft, V., 2024. North–south divergence on carbon policy driving uncertainty for business. URL: <https://www.maplecroft.com/products-and-solutions/climate-and-environment/insights/north-south-divergence-on-carbon-policy-driving-uncertainty-for-business/>, accessed: 2025-07-07.
- Martin, H.A., Quiquet, A., Nicolas, T., Giraud, G., Charbit, S., Roche, D.M., 2024. Climate-induced economic damages can lead to private-debt tipping points. URL: https://hal.science/hal-04224077v2/file/clim_eco_v2.pdf.
- MIT Climate Portal, n.d. What is radiative forcing? URL: <https://climate.mit.edu/explainers/radiative-forcing>, accessed: 2025-07-07.
- Nikiforos, M., Zezza, G., 2017. Stock-flow consistent macroeconomic models: A survey. *Journal of Economic Surveys* 31, 1204–1239. doi:<https://doi.org/10.1111/joes.12221>.
- Nobel Prize Committee / Nobel Media, 2018. William D. Nordhaus – facts. URL: <https://www.nobelprize.org/prizes/economic-sciences/2018/nordhaus/facts/>, accessed: 2025-07-09.
- Nordhaus, W., 2018. Projections and uncertainties about climate change in an era of minimal climate policies. *American Economic Journal: Economic Policy* 10, 333–60. URL: <https://www.aeaweb.org/articles?id=10.1257/pol.20170046>, doi:[10.1257/pol.20170046](https://doi.org/10.1257/pol.20170046).

- Nordhaus, W.D., 1993. Rolling the 'dice': an optimal transition path for controlling greenhouse gases. *Resource and Energy Economics* 15, 27–50. doi:[https://doi.org/10.1016/0928-7655\(93\)90017-0](https://doi.org/10.1016/0928-7655(93)90017-0).
- Wei, J., Jiang, T., Ménager, P., Kim, D.G., Dong, W., 2025. COP29: Progresses and challenges to global efforts on the climate crisis. *The Innovation* 6.
- Weitzman, M.L., 2012. GHG targets as insurance against catastrophic climate damages. *Journal of Public Economic Theory* 14, 221–244. doi:[10.1111/j.1467-9779.2011.01539.x](https://doi.org/10.1111/j.1467-9779.2011.01539.x).

Competition Between Market Makers

TEAM 2

NEHELO CHILENJE, University of Cape Town
MANDISA DABA, University of Cape Town
DAVIDE FELEPPA, Sapienza University of Rome
CAMERON FELLNER, University of Cape Town

Supervisor:

LEANDRO SÁNCHEZ-BETANCOURT, University of Oxford

Contents

1 Introduction	3
2 The model	8
1 Theoretical framework	8
2 Maximization problem	12
3 Approximate closed-form solution	14
4 Leader objective function	16
3 Numerical analysis	21
1 Comparative results between the Euler scheme and the closed form approximation	21
2 Numerical solution of the Stackelberg equilibrium	23
3 Including time dependencies	29
4 Conclusion	32
A Proofs	36
1 Proof of Proposition 2.1	36
2 Proof of Theorem 3.1	38
3 Proof of Proposition 4.3	40
B Algorithms	42
C Approximate closed form	43

Chapter 1

Introduction

Modern financial markets operate primarily through sophisticated electronic trading systems known as limit order books (LOBs). These systems serve as the technological backbone enabling efficient price discovery and trade execution in today's markets. A LOB represents an electronic record keeping system used by exchanges to organize and match buyers and sellers of financial instruments, acting essentially as a digital marketplace where trading orders are systematically arranged and executed according to well-established priority rules [Gould et al. \(2013\)](#).

The LOB operates as a continuous two-sided auction mechanism where buy orders, known as bids, are ranked by price with the highest offers positioned at the top of the queue, while sell orders, termed asks, are similarly organized with the lowest prices receiving execution priority. When a buyer's maximum acceptable price meets or exceeds a seller's minimum price requirement, the electronic matching system automatically facilitates the transaction, see [Cartea et al. \(2015\)](#) for further details. The numerical difference between the highest bid price and the lowest ask price constitutes what market participants refer to as the bid-ask spread can be seen as the cost of immediate execution, see e.g. [Amihud and Mendelson \(1986\)](#), which serves as a proxy of market liquidity.

Within this electronic framework, two primary categories of orders interact to create market dynamics. Limit orders specify both a desired quantity and a maximum (minimum) purchase price for buy (sell) orders, effectively providing liquidity to the market. Conversely, market orders execute immediately at the best available prices in the book, consuming the liquidity that limit orders have previously supplied. The depth and overall liquidity of any given market can be assessed by examining the thickness of the limit order book, which reflects the aggregate number of orders waiting at each discrete price level [Ho and Stoll \(1983\)](#). An example of how a LOB works in practice is shown in [Figure 1.1](#), where market orders of both small and large sizes are simulated, highlighting how it works in practice.




Figure 1.1: Simulations of market order executions for small and large amount of shares.

Liquidity providers (LPs) emerge as crucial market participants whose primary function involves enhancing overall market liquidity by ensuring continuous availability of both buyers and sellers across various price points. These entities, most prominently represented by market makers, maintain a persistent presence in financial markets by continuously quoting bid and ask prices while standing ready to transact specific quantities of securities at their posted prices.

Financial regulators maintain ongoing oversight of liquidity provision activities to ensure fair and orderly market operations, adequate liquidity availability during periods of market stress, prevention of manipulative trading practices, and protection of investor interests across all market segments. As an example, the EU introduced MiFiD II Regulatory Requirements, in the US market makers must be compliant with Rule 7E, in South Africa the Financial Markets Act regulates the market-making business, and so on.

The fundamental business model of liquidity providers centers on their ability to

earn the spread between their bid and ask quotes while carefully managing their exposure to adverse price movements, commonly referred to as inventory risk. This creates an ongoing optimization problem where market makers must balance their desire to post prices that generate substantial profits against the necessity of positioning their quotes to effectively manage their accumulated inventory positions. Additionally, liquidity providers face adverse selection risk when trading against informed market participants who possess superior information about an asset's true underlying value.

The foundations of inventory risk modeling in market making trace back to [Garmann \(1976\)](#) seminal work on market microstructure, which first distinguished between dealership and auction market models and characterized the inventory problems faced by market makers managing security and cash holdings. [Ho and Stoll \(1981\)](#) provided the first systematic treatment of dealer inventory risk using stochastic dynamic programming. The framework modeled dealer behavior under uncertainty with stochastic demand processes and return risk, deriving optimal bid-ask prices that maximize expected utility of terminal wealth. The theoretical framework was extended in [Ho and Stoll \(1983\)](#), where the authors analyze competing dealers.

[Glosten and Milgrom \(1985\)](#) introduced the adverse selection component to bid-ask spreads, demonstrating that positive spreads exist even with risk-neutral specialists making zero expected profits, while [Kyle \(1985\)](#) the continuous auction model further developed the understanding of information incorporation in market making, analyzing the dynamic interaction between informed traders and market makers. The modern era of inventory risk modeling began with [Avellaneda and Stoikov \(2008\)](#), which formalized the [Ho and Stoll \(1983\)](#) framework mathematically for limit order book markets.

This work introduced a rigorous stochastic control approach to market making optimization, modeling the reference price as Brownian motion with order arrival rates dependent on distance to the reference price. [Guéant et al. \(2013\)](#) provided significant mathematical innovations by transforming the Hamilton-Jacobi-Bellman equations into systems of linear ordinary differential equations, and in a similar setting [Cartea et al. \(2015\)](#) models the objective as a linear-quadratic functional form. Most research in market microstructure considers inventory risk and asymmetric information risk separately, yielding highly significant findings; however, few studies have investigated the interaction among multiple market makers and their impact on market liquidity.

As a matter of fact, competition among multiple liquidity providers can create complex strategic interactions that significantly influence market dynamics. [Boyce et al.](#)

(2025) were among the first to consider the interaction between a reference and competing market makers, with competition modeled through exponential dependence on spread differences between market makers. Guo et al. (2024) established connections between the market making problem, price impact, and the optimal execution problem in a Nash equilibrium framework. In Bergault and Sánchez-Betancourt (2025), the authors study a mean field game between brokers and informed traders, finding out that the Nash equilibrium helps informed traders in decision making and brokers in externalizing (or internalizing) order flow.

This work represents a first attempt to extend the theoretical framework in Boyce et al. (2025) by considering a Stackelberg equilibrium between two LP entities: the leader, representing the aggregate of market makers, and the follower, which is the reference market maker.

A Stackelberg equilibrium represents a fundamental solution concept in hierarchical game theory, originally introduced by Von Stackelberg (2010) (English translation of the original German volume) and subsequently formalized in dynamic game theory in Simaan and Cruz Jr (1973), while extensively developed in Başar and Olsder (1998). The equilibrium captures the strategic interaction between players in a hierarchical decision-making framework where one player, designated as the leader, possesses a structural advantage through the ability to commit to a strategy before other players, termed followers, make their decisions. The critical feature distinguishing the Stackelberg equilibrium from the simultaneous Nash equilibrium lies in the information structure and timing of decisions. The leader possesses the strategic advantage of commitment power, enabling anticipation of followers' rational responses, while followers observe the leader's committed strategy before making their own decisions.

This hierarchical information pattern creates an asymmetric game structure where the leader can effectively manipulate the subsequent equilibrium to optimize his own objective. In stochastic differential game settings, as analyzed in Moon and Başar (2018), the Stackelberg equilibrium concept extends naturally to dynamic environments. The existence and uniqueness of Stackelberg equilibria depend critically on the structural properties of the game, including the convexity of strategy spaces, continuity of payoff functions, and the uniqueness of followers' best response correspondences.

In linear-quadratic (LQ) frameworks, as demonstrated by Yong (2002) and Xu and Zhang (2015), the Stackelberg solution often admits explicit characterization through coupled systems of differential equations, where the leader's optimal strategy incorporates the anticipated trajectory of the followers' optimal responses. In the

same LQ framework, many financial applications have been proposed, such as [Zheng and Shi \(2020\)](#) and the references therein. The fundamental requirement for a well-defined Stackelberg equilibrium is that the follower must know ex ante that leader will observe the committed strategy, and followers must possess no means of committing to strategies that deviate from their best responses to the leader's announcement, as stated in [Başar and Olsder \(1998\)](#).

In our theoretical framework, two different entities act in the financial market as LPs. Firstly, the leader is constituted by the set of market makers, and therefore is the agent that can influence the market by quoting bid and ask spreads. This can be considered as the market maker that quotes the smallest spread at every instant in the considered time horizon. The second entity is the follower, who observes the spreads in the market and defines his optimal strategy as a function of these, quoting his spreads accordingly. In contrast, the leader anticipates the follower's response and optimizes the minimum spread to maximize its utility function.

These are a number of novelties in this work. First of all, the theoretical framework in [Boyce et al. \(2025\)](#) is generalized by allowing the leader to precommit to a strategy that takes into account the follower's inventory. Furthermore, the coefficients of the spreads quoted by the leader are, in full generality, deterministic time dependent functions. Moreover, in [Boyce et al. \(2025\)](#), the authors analyze the market-making competition exclusively from the follower's perspective. Considering a Stackelberg equilibrium, instead, this work also aims to analyze the leader's point of view, and therefore to verify the existence of an optimal strategy for the leader.

The work is organized as follows: in Chapter [2](#) the theoretical framework is presented and approximated closed-form solutions are obtained, Chapter [3](#) contains an extensive numerical analysis and comparisons with other existing frameworks, then Chapter [4](#) concludes.

Chapter 2

The model

1 Theoretical framework

We consider a financial market operating under a limit order book mechanism where liquidity provision is facilitated by competing market makers. We model a competitive environment with heterogeneous market participants consisting of a leader market maker (denoted by “he”), characterized as the participant offering the best aggregate available price at any given time, and a follower market maker (denoted by “she”), who responds strategically to the leader’s pricing decisions. This leader-follower structure captures the realistic asymmetry observed in modern electronic markets, where certain market makers (often due to technological advantages, information, or scale) consistently provide the best quotes, while others adapt their strategies in response [Cartea et al. \(2015\)](#).

Let $(W_t)_{0 \leq t \leq T}$ be a standard Brownian motion, where $T > 0$ represents the finite trading horizon. Additionally, we define independent Poisson processes for modeling liquidity taking orders: $(M_t^a)_{0 \leq t \leq T}$ for ask-side market orders with intensity $\lambda^a > 0$ and $(M_t^b)_{0 \leq t \leq T}$ for bid-side market orders with intensity $\lambda^b > 0$.

The fundamental asset price follows the following equation,

$$S_t = S_0 + \sigma_S W_t, \quad S_0 > 0. \quad (2.1)$$

We denote by $\delta^L := \{\delta_t^{a,L}, \delta_t^{b,L}\}_{0 \leq t \leq T}$ the ask and bid depths for the leader and by $\delta^F := \{\delta_t^{a,F}, \delta_t^{b,F}\}_{0 \leq t \leq T}$ the ask and bid depths for the follower, where $\delta_t^{a,i}, \delta_t^{b,i}$ represent the depths at which players place their ask and bids quotes relative to the fundamental asset price S_t . The depths represent the price “half spreads” that each market maker offers relative to the fundamental price, determining their competitiveness and expected profit margins.

The observable market mid-price of the leader is given by

$$S_t^{\text{mid},L} = \frac{1}{2} \left(2S_t + \delta_t^{a,L} - \delta_t^{b,L} \right),$$

while the mid-price of the whole market is given by

$$S_t^{\text{mid}} = \frac{1}{2} \left(S_t + \min\{\delta_t^{a,L}, \delta_t^{a,F}\} \right) + \frac{1}{2} \left(S_t - \min\{\delta_t^{b,L}, \delta_t^{b,F}\} \right).$$

The filled orders for each market maker are represented by the point processes

$$(N_t^{a,L})_{0 \leq t \leq T} \text{ and } (N_t^{b,L})_{0 \leq t \leq T},$$

$$(N_t^{a,F})_{0 \leq t \leq T} \text{ and } (N_t^{b,F})_{0 \leq t \leq T}.$$

Along the lines of [Boyce et al. \(2025\)](#), the intensities for order fills incorporate competitive effects through exponential functions:

$$\Lambda_t^{a,L} := \lambda^a \left(1 - \frac{\Lambda_t^{a,F}}{\lambda^a} \right),$$

$$\Lambda_t^{b,L} := \lambda^b \left(1 - \frac{\Lambda_t^{b,F}}{\lambda^b} \right),$$

$$\Lambda_t^{a,F} := \lambda^a \min \left\{ \exp \left\{ -\kappa \left(\delta_t^{a,F} - \delta_t^{a,L} + \iota \right) \right\}, 1 \right\},$$

$$\Lambda_t^{b,F} := \lambda^b \min \left\{ \exp \left\{ -\kappa \left(\delta_t^{b,F} - \delta_t^{b,L} + \iota \right) \right\}, 1 \right\},$$

where $\kappa > 0$ is the decay rate parameter, $\iota > 0$ is the minimum tick size, and the $\min\{1, \cdot\}$ operation ensures intensities remain bounded. From now on we will assume that ι is absorbed in the leader's spread, such that δ^L stands for a tick more generous than the leader bid-ask quotes.

The intensity functions satisfy monotonicity properties, such as $\Lambda_t^{a,L}$ is decreasing in $\delta_t^{a,L}$ and increasing in $\delta_t^{a,F}$, while $\Lambda_t^{a,F}$ is decreasing in $\delta_t^{a,F}$ and increasing in $\delta_t^{a,L}$, with similar properties holding for bid-side intensities. In accordance with the economic intuition, offering better prices (smaller spreads) increases the probability of order execution. Furthermore, probability conservation holds, such as for any side $s \in \{a, b\}$, we have $\Lambda_t^{s,L} + \Lambda_t^{s,F} \leq \lambda^s$, ensuring that all market orders are filled by either the leader or follower.

In a general framework where λ^a and λ^b are fixed, the execution probability of both market agents depends exclusively on the difference in quoted spreads between them. This feature is not realistic, since higher spreads correspond to a lower frequency of incoming market orders, even if the spread difference between leader

and follower is negligible. Therefore, a possible solution could be consider the arrival frequency of orders as dependent on the base spread level. In this case, when quoted spreads are very distant from the base price of the stock, market orders decrease, introducing a trade-off for operators between trade frequency and revenue per trade.

We return to the point below under the stated assumptions, the market making model is well-posed in the sense that the intensity processes $(\Lambda_t^{s,i})$ are well-defined and bounded, the order fill processes $(N_t^{s,i})$ exist as point processes with the specified intensities, for $i = \{L, F\}$. The inventories dynamics are given by

$$dQ_t^i = dN_t^{b,i} - dN_t^{a,i}, \quad Q_0^i = q_0^i, \text{ for } i = \{L, F\}. \quad (2.2)$$

Let the inventory space of both agents be $\mathcal{Q}^i = \{\underline{q}^i, \underline{q}^i + 1, \dots, \bar{q}^i\}$ with cardinality $n^i = \bar{q}^i - \underline{q}^i + 1$, for $i = \{L, F\}$, such that $\underline{q}^i \leq Q_t^i \leq \bar{q}^i$.

The cash account is

$$dX_t^i = (S_t + \delta_t^{a,i}) dN_t^{a,i} - (S_t - \delta_t^{b,i}) dN_t^{b,i}, \quad X_0^i = x_0^i, \text{ for } i = \{L, F\}. \quad (2.3)$$

Extending the existing literature, see e.g. [Boyce et al. \(2025\)](#), we assume the leader's strategy depending on his own inventory and on the follower's one. In our framework, the leader precommits his quoting strategy in $t = 0$ and the follower can observe it in the market without noise, such as

$$\begin{aligned} \delta_t^{a,L} &= a(t) - \beta(t)Q_t^L - \theta(t)Q_t^F, \\ \delta_t^{b,L} &= b(t) + \beta(t)Q_t^L + \theta(t)Q_t^F, \end{aligned} \quad (2.4)$$

where $a(t)$, $b(t)$, $\beta(t)$ and $\theta(t)$ are deterministic functions of time. In order to pre-commit, the leader should know the follower inventory in $t = 0$, so he can observe the transactions on the LOB and extrapolate the inventory of the counterparty. To this end, from now on we will always assume $q_0^L = q_0^F = 0$.

A rigorous definition of the probability space in this framework is necessary, reiterating the argument in [Barucci et al. \(2025\)](#). Consider first Ω^d as the set of increasing piecewise constant càdlàg functions from $[0, T]$ into $\mathbb{N}$ with jumps equal to one, and let Ω^c denote the set of continuous functions from $[0, T]$ into $\mathbb{R}$. The sample space is then defined as $\Omega = \Omega^c \times \Omega_d^4$, where the components correspond to the Brownian motion driving the fundamental asset price and the point processes for ask and bid market orders arriving. Let

$$\left(W_t, \overline{N}_t^{a,L}, \overline{N}_t^{b,L}, \overline{N}_t^{a,F}, \overline{N}_t^{b,F} \right)_{t \in [0, T]}$$

be the canonical process on Ω , where $(W_t)_{t \in [0, T]}$ represents the fundamental price innovation and $(\bar{N}_t^{s, i})_{t \in [0, T]}$ for $s \in \{a, b\}$ and $i \in \{L, F\}$ are the canonical counting processes. The associated filtration is given by $\bar{\mathcal{F}} = (\mathcal{F}_t^c \otimes \mathcal{F}_t^{a, L} \otimes \mathcal{F}_t^{b, L} \otimes \mathcal{F}_t^{a, F} \otimes \mathcal{F}_t^{b, F})_{t \in [0, T]}$, where $(\mathcal{F}_t^c)_{t \in [0, T]}$ is the right-continuous completed filtration generated by $(W_t)_{t \in [0, T]}$, and $(\mathcal{F}_t^{s, i})_{t \in [0, T]}$ are the right-continuous filtrations generated by the respective point processes for market maker i .

We denote by $\mathbb{P}_0$ the probability measure on $(\Omega, \mathcal{F})$ such that: (1) $(\bar{M}_s^{a, L} = \bar{N}_s^{a, L} - s, \bar{M}_s^{b, L} = \bar{N}_s^{b, L} - s, \bar{M}_s^{a, F} = \bar{N}_s^{a, F} - s, \bar{M}_s^{b, F} = \bar{N}_s^{b, F} - s)_{s \in [0, T]}$ are $\mathcal{F}$ -martingales; (2) $(W_t)_{t \in [0, T]}$ is a standard Brownian motion; and (3) all processes are mutually independent.

The actual point processes $(N_t^{a, L}, N_t^{b, L}, N_t^{a, F}, N_t^{b, F})_{t \in [0, T]}$ representing filled orders for both market makers can be defined as solutions to the coupled system

$$\begin{aligned} dN_t^{a, L} &= \mathbb{1}_{\{Q_{t-}^L > \underline{q}^L\}} d\bar{N}_t^{a, L}, & dN_t^{b, L} &= \mathbb{1}_{\{Q_{t-}^L < \bar{q}^L\}} d\bar{N}_t^{b, L}, \\ dN_t^{a, F} &= \mathbb{1}_{\{Q_{t-}^F > \underline{q}^F\}} d\bar{N}_t^{a, F}, & dN_t^{b, F} &= \mathbb{1}_{\{Q_{t-}^F < \bar{q}^F\}} d\bar{N}_t^{b, F}, \end{aligned}$$

where $[\underline{q}^i, \bar{q}^i]$ for $i \in \{L, F\}$ define the inventory constraints for each market maker.

Let $\mathcal{U}$ denote the set of admissible control processes $\delta = (\delta_t^{a, L}, \delta_t^{b, L}, \delta_t^{a, F}, \delta_t^{b, F})_{t \in [0, T]} \in L^2(\Omega \times [0, T])$ such that $\delta \geq -\delta_\infty$ and the processes are $\mathcal{F}$ -predictable, where $\delta_\infty = (\delta_\infty^{a, L}, \delta_\infty^{b, L}, \delta_\infty^{a, F}, \delta_\infty^{b, F})$ characterizes the lower bounds of the controls.

For $\delta \in \mathcal{U}$, we define the state processes $(S_t, Q_t^L, Q_t^F, X_t^L, X_t^F)$, where the dynamics are given by equations (2.1), (2.2), and (2.3). For any control $\delta \in \mathcal{U}$, we define the probability measure $\mathbb{P}^\delta$ via the Radon-Nikodym derivative $\frac{d\mathbb{P}^\delta}{d\mathbb{P}_0} = \mathcal{L}_T^\delta$, where $\mathcal{L}_t^\delta$ is the Doléans-Dade exponential.

Since the control processes $(\delta_t^{a, L}, \delta_t^{b, L}, \delta_t^{a, F}, \delta_t^{b, F})_{t \in [0, T]}$ are bounded from below and the intensities are uniformly bounded by construction through the $\min\{1, \cdot\}$ operation, the process $(\mathcal{L}_t^\delta)_{t \in [0, T]}$ is a $\mathbb{P}_0$ -martingale by Novikov's criterion. Under $\mathbb{P}^\delta$, the compensated processes $M_t^{s, i, \delta} := N_t^{s, i} - \int_0^t \Lambda_u^{s, i} du$, for $s \in \{a, b\}$ and $i \in \{L, F\}$, are true $\mathcal{F}$ -martingales, and by Watanabe's characterization theorem, all point processes $(N_t^{s, i})_{t \in [0, T]}$ for $s \in \{a, b\}$ and $i \in \{L, F\}$ are $\mathcal{F}$ -Poisson processes with the specified stochastic intensities.

Throughout the analysis, we work on the probability space $(\Omega, \mathcal{F}, \mathbb{P}^\delta)$, which satisfies the usual conditions of right-continuity and completeness, ensuring that all stochastic processes are well-defined and the market making optimization problems for both the leader and follower are mathematically well-posed.

2 Maximization problem

The follower aim is to maximize their objective function

$$\begin{aligned}
J(\delta^F) &= \mathbb{E} \left[X_T^F + Q_T^F S_T^{\text{mid},L} - \gamma(Q_T^F)^2 - \phi \int_0^T (Q_s^F)^2 ds \middle| \mathcal{F}_0 \right] \\
&= \mathbb{E} \left[X_T^F + Q_T^F \left(S_T + \frac{a(T) - b(T)}{2} - \beta(T)Q_T^L - \theta(T)Q_T^F \right) \right. \\
&\quad \left. - \gamma(Q_T^F)^2 - \phi \int_0^T (Q_s^F)^2 ds \middle| \mathcal{F}_0 \right], \tag{2.5}
\end{aligned}$$

under the cash [2.3](#) and inventory [2.2](#) constraints where δ^F is the set of $\mathcal{F}$ -adapted and integrable strategies for the follower. By setting $a(t), b(t), \beta(t) \in \mathbb{R}$ as fixed $\forall t \in [0, T]$ and $\theta(t) = 0, \forall t \in [0, T]$, we obtain the model in [Boyce et al. \(2025\)](#) without noise, while setting also $\delta^L = 0$ and $\iota = 0$ our framework reduces in the Avellaneda and Stoikov model.

Starting at $t = 0$, during the time horizon the follower will quote her spreads, and hopefully these will be filled by random market orders. At the end of the time horizon, at $t = T$, the follower's wealth will be equal to cash X_T^F plus the marked-to-market inventory $Q_T^F S_T^{\text{mid},L}$. Since we consider the follower as a risk-averse agent, we include in the objective function a penalization term on the square of the terminal inventory and a running penalization on the square of the inventory. These terms ensure that the follower is not excessively exposed to inventory risk, and therefore reduces her exposure to the stock market.

The intuition behind the terms in the HJB equation [\(A.4\)](#) follows the same argument and interpretation made by [Boyce et al. \(2025\)](#).

The value function differences capture

$$\begin{aligned}
v(t, s, x, q^F, q^L - 1) - v(t, s, x, q^F, q^L) &\quad \text{ask side leader execution} \\
v(t, s, x, q^F, q^L + 1) - v(t, s, x, q^F, q^L) &\quad \text{bid side leader execution} \\
v(t, s, x + s + \delta^a, q^F - 1, q^L) - v(t, s, x, q^F, q^L - 1) &\quad \text{ask side follower versus leader} \\
v(t, s, x - s + \delta^b, q^F + 1, q^L) - v(t, s, x, q^F, q^L + 1) &\quad \text{bid side follower versus leader}
\end{aligned}$$

These differences are weighted by: (1) the intensities λ^a and λ^b for market order arrivals, and (2) the follower's fill probabilities, for both bid and ask sides:

$$\begin{aligned}
&\min\{e^{-\kappa(\delta^a - a(t) + \beta(t)q^L + \theta(t)q^F)}, 1\}, \\
&\min\{e^{-\kappa(\delta^b - b(t) - \beta(t)q^L - \theta(t)q^F)}, 1\}.
\end{aligned}$$

Proposition 2.1. Consider the problem (2.5). Then, for any $t \in [0, T]$, the value function is

$$v(t, s, x, q^F, q^L) := x + q^F s - \frac{1}{2}(\beta(t) + \theta(t)) (q^F)^2 - \beta(t)q^F q^L + g(t, q^F, q^L), \quad (2.6)$$

where $g(t, q^F, q^L)$ satisfies the following equation

$$\begin{aligned} 0 = & + \frac{\partial g}{\partial t}(t, q^F, q^L) - \frac{1}{2} \left(\frac{\partial \beta}{\partial t}(t) + \frac{\partial \theta}{\partial t}(t) + 2\phi \right) (q^F)^2 - \frac{\partial \beta}{\partial t}(t) q^F q^L \\ & + \lambda^a (g(t, q^F, q^L - 1) - g(t, q^F, q^L) + \beta(t)q^F) \\ & + \lambda^b (g(t, q^F, q^L + 1) - g(t, q^F, q^L) - \beta(t)q^F) \\ & + \sup_{c^{a,F}} \left\{ \lambda^a \min \left\{ e^{-\kappa(c^{a,F} + \frac{1}{2}(\beta(t) + \theta(t)) - a(t))}, 1 \right\} \right. \\ & \left. [c^{a,F} + g(t, q^F - 1, q^L) - g(t, q^F, q^L - 1)] \right\} \mathbb{1}_{q^F > \underline{q}^F} \\ & + \sup_{c^{b,F}} \left\{ \lambda^b \min \left\{ e^{-\kappa(c^{b,F} + \frac{1}{2}(\beta(t) + \theta(t)) - b(t))}, 1 \right\} \right. \\ & \left. [c^{b,F} + g(t, q^F + 1, q^L) - g(t, q^F, q^L + 1)] \right\} \mathbb{1}_{q^F < \bar{q}^F}, \end{aligned} \quad (2.7)$$

with terminal condition

$$g(T, q^F, q^L) = \frac{a(T) - b(T)}{2} q^F - \left(\gamma + \frac{\theta(T) - \beta(T)}{2} \right) (q^F)^2.$$

The optimal follower quoted spreads are

$$\begin{aligned} \delta^{a,F*}(t, q^F, q^L) &= \max \left\{ \hat{\delta}^{a,F*}(t, q^F, q^L), a(t) - \beta(t)q^L - \theta(t)q^F \right\}, \\ \delta^{b,F*}(t, q^F, q^L) &= \max \left\{ \hat{\delta}^{b,F*}(t, q^F, q^L), b(t) + \beta(t)q^L + \theta(t)q^F \right\}. \end{aligned} \quad (2.8)$$

where

$$\begin{aligned} \hat{\delta}^{a,F*}(t, q^F, q^L) &= \frac{1}{\kappa} - g(t, q^F - 1, q^L) + g(t, q^F, q^L - 1) \\ &\quad - \beta(t)q^L - \theta(t)q^F + \frac{1}{2}(\beta(t) + \theta(t)) \\ \hat{\delta}^{b,F*}(t, q^F, q^L) &= \frac{1}{\kappa} - g(t, q^F + 1, q^L) + g(t, q^F, q^L + 1) \\ &\quad + \beta(t)q^L + \theta(t)q^F + \frac{1}{2}(\beta(t) + \theta(t)). \end{aligned} \quad (2.9)$$

Proof. See Appendix 1 □

The follower's optimal quoted depths align with [Boyce et al. \(2025\)](#), but in extend formulation, given the introduction of the time-dependent parameters $a(t), b(t), \beta(t)$, and an additional term $\theta(t)$. This new parameter $\theta(t)$ directly captures how the follower's inventory q^F affects the leader's quote adjustments.

This assumes that the leader has information relating to the inventory of the follower at each point in time and posts his quotes relative to such information. The economic interpretation for θ fixed is intuitive. Consider the follower has a large amount of inventory q^F , so the leader might increase the inventory risk for the follower by decreasing his depths on the ask side so that the follower has higher probability to be filled in that side of the LOB. We will delve on this aspect in the next chapter, when analyzing the equilibrium.

3 Approximate closed-form solution

Here we propose an approximate solution for the problem [\(2.5\)](#). We assume that the unrestrained maximizers [\(2.9\)](#) are always greater (or equal) than the leader depths and $\beta(t)$ fixed for $t \in [0, T]$.

In this case, [\(2.7\)](#) simplifies as

$$\begin{aligned} 0 &= \frac{\partial g}{\partial t}(t, q^F) - \frac{1}{2} \left(\frac{\partial \theta}{\partial t}(t) + 2\phi \right) (q^F)^2 \\ &\quad + \left(\lambda^a - \lambda^b \right) \beta q^F \\ &\quad + \frac{\lambda^a}{\kappa} e^{-1 - \kappa(\frac{1}{2}(\beta + \theta(t)) - a(t))} \exp \left\{ -\kappa [g(t, q^F) - g(t, q^F - 1)] \right\} \mathbb{1}_{q^F > \underline{q}^F} \\ &\quad + \frac{\lambda^b}{\kappa} e^{-1 - \kappa(\frac{1}{2}(\beta + \theta(t)) - b(t))} \exp \left\{ -\kappa [g(t, q^F) - g(t, q^F + 1)] \right\} \mathbb{1}_{q^F < \bar{q}^F}, \\ g(T, q^F) &= \frac{a(T) - b(T)}{2} q^F - \left(\gamma + \frac{\theta(T) - \beta}{2} \right) (q^F)^2. \end{aligned}$$

Theorem 3.1. Assume $\underline{q}^F = -\bar{q}^F$. Assume also the unrestrained maximizers [\(2.9\)](#) to be equal or greater than the one tick more generous leader's spread.

Given the time interval $[0, T]$, we define its discretization as $\tau = \{t_n, n = 0, 1, \dots, N\}$, where $\Delta t = t_n - t_{n-1}$ and $N = \frac{T}{\Delta t}$.

Define the function $\omega : \tau \times \mathcal{Q}^F \rightarrow (0, \infty)$, the tridiagonal set of matrices

$(\mathbf{A}_{i,q,n})_{\underline{q}^F \leq i, q \leq \bar{q}^F, 0 \leq n \leq N}$ and the vector $(\mathbf{B}_q)_{\underline{q}^F \leq q \leq \bar{q}^F}$. Then the function $v : \tau \times (0, \infty) \times \mathbb{R} \times \mathcal{Q}^F \times \mathbb{Z}$ is given by

$$v(t_n, s, x, q^F, q^L) := x + q^F s - \frac{1}{2}(\beta + \theta(t_n)) (q^F)^2 - \beta q^F q^L + \frac{1}{\kappa} \log(\omega(t_n, q^F)), \quad (2.10)$$

where

$$\omega(t_n, q) = \left(\exp \left\{ - \sum_{m=n}^{N-1} \mathbf{A}_m \cdot dt \right\} \mathbf{B} \right) q, \quad (2.11)$$

and

$$\mathbf{A}_{i,q,n} = \begin{cases} \frac{\kappa}{2} \left(\frac{\partial \theta}{\partial t}(t_n) + 2\phi \right) q^2 - \beta \kappa (\lambda^a - \lambda^b) q & \text{if } i = q, \\ -\lambda^a \exp \left(-1 - \kappa \left(\frac{\beta + \theta(t_n)}{2} - a(t_n) \right) \right) & \text{if } i = q - 1, \\ -\lambda^b \exp \left(-1 - \kappa \left(\frac{\beta + \theta(t_n)}{2} - b(t_n) \right) \right) & \text{if } i = q + 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$\mathbf{B}_q = \exp \left\{ \kappa \left[\frac{a(T) - b(T)}{2} q - \left(\gamma + \frac{\theta(T) - \beta}{2} \right) q^2 \right] \right\}.$$

If the following conditions are satisfied, (2.10) solves problem 2.5.

$$\begin{aligned} \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F - 1)} \right) \right] + \frac{1}{2}(\beta + \theta(t_n)) &\geq a(t_n), \quad t_n \in \tau, \quad q^F \in \{\underline{q}^F + 1, \dots, \bar{q}^F\} \\ \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F + 1)} \right) \right] + \frac{1}{2}(\beta + \theta(t_n)) &\geq b(t_n), \quad t_n \in \tau, \quad q^F \in \{\underline{q}^F, \dots, \bar{q}^F - 1\}. \end{aligned} \quad (2.12)$$

Then the candidate optimizers satisfy

$$\begin{aligned} \hat{\delta}^{a, F^*}(t_n, q^F, q^L) &= \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F - 1)} \right) \right] - \beta q^L - \theta(t_n) q^F + \frac{1}{2}(\beta + \theta(t_n)) \\ \hat{\delta}^{b, F^*}(t_n, q^F, q^L) &= \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F + 1)} \right) \right] + \beta q^L + \theta(t_n) q^F + \frac{1}{2}(\beta + \theta(t_n)). \end{aligned} \quad (2.13)$$

Proof. See Appendix 2 □

Proposition 3.2. Under the assumptions of Theorem 3.1. Consider $a(t), b(t), \theta(t) \in \mathbb{R}$, to be constant values denoted by $a, b, \theta \in \mathbb{R}$.

Define the function $\omega : [0, T] \times \mathcal{Q}^F \rightarrow (0, \infty)$, the tridiagonal matrix $(\mathbf{A}_{i,q})_{\underline{q}^F \leq i, q \leq \bar{q}^F}$ and the vector $(\mathbf{B}_q)_{\underline{q}^F \leq q \leq \bar{q}^F}$.

Then, for any $t \in [0, T]$, the value function $v : [0, T] \times (0, \infty) \times \mathbb{R} \times \mathcal{Q}^F \times \mathbb{Z}$ is

$$v(t, s, x, q^F, q^L) := x + q^F s - \frac{1}{2}(\beta + \theta) (q^F)^2 - \beta q^F q^L + \frac{1}{\kappa} \log(\omega(t, q^F)), \quad (2.14)$$

where

$$\omega(t, q) = (\exp \{-\mathbf{A}(T - t)\} \mathbf{B})_q, \quad (2.15)$$

with

$$\mathbf{A}_{i,q} = \begin{cases} \kappa \phi q^2 - \beta \kappa (\lambda^a - \lambda^b) q & \text{if } i = q, \\ -\lambda^a \exp \left(-1 - \kappa \left(\frac{\beta + \theta}{2} - a \right) \right) & \text{if } i = q - 1, \\ -\lambda^b \exp \left(-1 - \kappa \left(\frac{\beta + \theta}{2} - b \right) \right) & \text{if } i = q + 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$\mathbf{B}_q = \exp \left\{ \kappa \left[\frac{a - b}{2} q - \left(\gamma + \frac{\theta - \beta}{2} \right) q^2 \right] \right\}.$$

Proof. The proof comes directly from the solution of the ODEs system when the matrix $\mathbf{A}$ is fixed. $\square$

Obtaining the matrix form of the approximated solution of $\omega(t, q^F)$ makes the computational process immediate. This also gives insight into the follower's response strategy and how the parameters affect its value function. Since $-\mathbf{A}$ is essentially positive, in the sense that all off-diagonal elements are nonnegative, $\omega(t, q^F)$ takes values in $(0, \infty)$.

We show in Appendix [C](#) the goodness of the approximated closed form for the optimal quoted spreads via numerical experiments. With a reasonable set of parameters, violations of [\(2.12\)](#) appear in less than 0.01% of the total simulated runs.

4 Leader objective function

In order to characterize a Stackelberg game, we must define an objective function for the leader. The latter can influence the average market prices, so he might be incentivized to quote high spreads so that his final inventory is valued at a much higher price compared to the base price. For this reason, we consider an objective function slightly different from the follower's one, where the final inventory is valued at the base price S_T . This ensures that the leader has no arbitrage opportunities by increasing the quoted spreads.

As mentioned earlier, it is also necessary to introduce a trade-off for the leader so that he is not incentivized to raise the quoted spreads excessively. For this purpose, we consider the intensity of market order arrival processes as functions of the base spread level, such as $\lambda^a(t) = \lambda_0^a e^{-\kappa a(t)}$, $\lambda^b(t) = \lambda_0^b e^{-\kappa b(t)}$. In this way, when the market is less liquid, that is when spreads are higher, market orders arrive at a lower frequency.

We set the leader's objective function given by

$$\mathbb{E} \left[X_T^L + Q_T^L S_T - \gamma (Q_T^L)^2 - \phi \int_0^T (Q_s^L)^2 ds \middle| \mathcal{F}_0 \right]. \quad (2.16)$$

Denote $u \in C^{1,2}([0, T] \times (0, \infty) \times \mathbb{R} \times \mathcal{Q}^L \times \mathbb{Z})$ as

$$u(t, s, x, q^L, q^F) = \mathbb{E} \left[X_T^L + Q_T^L S_T - \gamma (Q_T^L)^2 - \phi \int_0^T (Q_s^L)^2 ds \middle| \mathcal{F}_t \right], \forall t \in [0, T]. \quad (2.17)$$

The following proposition aims to characterize the leader's objective function using the Feynman-Kac formula, permitting a partial-integro differential equation representation.

Proposition 4.1. Consider the leader's objective function (2.16). Recalling the conditional expectation in (2.17), the u function solves the following equation

$$\begin{aligned} 0 = & \frac{\partial u}{\partial t}(t, s, x, q^L, q^F) + \frac{\sigma_S^2}{2} \frac{\partial^2 u}{\partial s^2}(t, s, x, q^L, q^F) - \phi (q^L)^2 \\ & + \lambda^a (u(t, s, x, q^L, q^F - 1) - u(t, s, x, q^L, q^F)) \\ & + \lambda^b (u(t, s, x, q^L, q^F + 1) - u(t, s, x, q^L, q^F)) \\ & + \Lambda^{a,L} (u(t, s, x + s + \delta^{a,L}, q^L - 1, q^F) - u(t, s, x, q^L, q^F - 1)) \mathbb{1}_{q^L > \underline{q}^L} \\ & + \Lambda^{b,L} (u(t, s, x - s + \delta^{b,L}, q^L + 1, q^F) - u(t, s, x, q^L, q^F + 1)) \mathbb{1}_{q^L < \bar{q}^L} \end{aligned} \quad (2.18)$$

with terminal condition

$$u(T, s, x, q^L, q^F) = x + q^L s - \gamma (q^L)^2. \quad (2.19)$$

Proof. The argument follows the Feynman-Kac representation theorem, see [Zhu et al. \(2008\)](#). $\square$

Through Proposition 4.1, it is possible to solve the equation (2.18) using numerical techniques, which allows computing the expectation of the leader's objective function without the need to use Monte Carlo methods. Direct computation of the expectation is impossible because of the stochastic nature of the intensities $\Lambda^{a,s}$ and $\Lambda^{b,s}$ where $s \in \{F, L\}$. Indeed, by solving equation (2.18) backward in time and plugging the initial conditions into the function u , one obtains (2.16).

Before proceeding with the characterization of the u function, it is useful to observe how the filling probability for the leader does not depend on his own inventory, but only on the parameters of his precommitted strategy and on the follower's inventory.

Remark 4.2. The filling intensities for the leader $\Lambda^{a,L}$ and $\Lambda^{b,L}$ depend on his own precommitted strategy and on the inventory of the follower

$$\begin{aligned}\Lambda^{a,L}(t, q^F) &= \lambda^a \left[1 - \min \left\{ \exp \left[\frac{\beta + \theta(t)}{2} + \frac{1}{\kappa} \left(1 + \log \frac{\omega(t, q^F)}{\omega(t, q^F - 1)} \right) - a(t) \right], 1 \right\} \right], \\ \Lambda^{b,L}(t, q^F) &= \lambda^b \left[1 - \min \left\{ \exp \left[\frac{\beta + \theta(t)}{2} + \frac{1}{\kappa} \left(1 + \log \frac{\omega(t, q^F)}{\omega(t, q^F + 1)} \right) - b(t) \right], 1 \right\} \right].\end{aligned}$$

According to this remark, the follower, through her optimal response, determines her proximity to the leader in the LOB, affecting both filling probabilities. The leader, by pre-committing, only sets a baseline for his own fill probability.

The equation (2.18) consists of high dimension PIDE, for which its resolution usually requires computationally expensive numerical techniques. The aim is therefore to transform (2.18) into a linear system of ODEs in matrix form, in a similar fashion to what done in Theorem 3.1. Due to the dependence of function u on both inventories q^L and q^F , a closed form is not obtainable, but the following proposition permits us to use computationally efficient numerical techniques, typically used for solving Sylvester equations.

Proposition 4.3. Consider the differential equation (2.18) characterizing the evolution of the leader objective function with the terminal condition (2.19). Consider also the following ansatz for the function u :

$$u(t, s, x, q^L, q^F) = x + q^L s + h(t, q^L, q^F), \quad (2.20)$$

with $h(T, q^L, q^F) = -\gamma (q^L)^2$.

Let $\underline{q}^L = -\bar{q}^L$, $\underline{q}^F = -\bar{q}^F$ and $\underline{q}^L = q^L$. We represent the function $h(t, q^L, q^F)$ as a matrix $(\mathbf{H}_{i,j}(t))_{\underline{q}^L \leq i \leq \bar{q}^L, \underline{q}^F \leq j \leq \bar{q}^F}$, where the entry $\mathbf{H}_{i,j}(t)$ corresponds to $h(t, q_i^L, q_j^F)$.

The PIDE for h can be written as

$$\frac{\partial \mathbf{H}}{\partial t}(t) = \mathbf{A}^1(t) \mathbf{H}(t) + \mathbf{H}(t) \mathbf{A}^2(t) + \mathbf{B}(t), \quad (2.21)$$

with the terminal condition $(\mathbf{H}_{i,j}(T)) = -\gamma (q_i^L)^2$, $\forall i \in \mathcal{Q}^L, j \in \mathcal{Q}^F$, where

$$(\mathbf{A}_{i,k}^1(t)) = \begin{cases} \lambda^a + \lambda^b & \text{if } i = k, \\ -\Lambda^{a,L}(t, q^F) & \text{if } k = i - 1, \\ -\Lambda^{b,L}(t, q^F) & \text{if } k = i + 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$(\mathbf{A}_{j,\ell}^2(t)) = \begin{cases} \Lambda^{a,L}(t, q_\ell^F) - \lambda^a & \text{if } \ell = j - 1, \\ \Lambda^{b,L}(t, q_\ell^F) - \lambda^b & \text{if } \ell = j + 1, \\ 0 & \text{otherwise,} \end{cases}$$

$$(\mathbf{B}_{i,j}(t)) = -\phi(q_i^L)^2 + \Lambda^{a,L}(t, q_j^F) \delta^{a,L}(t, q_i^L, q_j^F) + \Lambda^{b,L}(t, q_j^F) \delta^{b,L}(t, q_i^L, q_j^F).$$

Proof. See Appendix [3](#) □

The proposition [4.3](#) allows us to write the partial integro-differential equation for the leader's objective into a matrix differential equation format. Recognizing the discrete nature of the inventory state spaces for both the leader and follower allows us to represent the value function as a finite-dimensional matrix $\mathbf{H}(t)$, where each entry corresponds to a specific combination of leader and follower inventory levels.

The matrix decomposition reveals the underlying structure of the competitive market making dynamics through three distinct components. The matrix $\mathbf{A}^1(t)$ captures the transitions in the leader's inventory dimension and operates through left multiplication on $\mathbf{H}(t)$. Its tridiagonal structure reflects the fact that inventory can only change by discrete units when orders are filled, with the diagonal elements representing the combined arrival intensities of ask and bid orders, while the super- and sub-diagonal elements encode the inventory transition effects through the trading intensities. The matrix $\mathbf{A}^2(t)$ governs transitions in the follower's inventory dimension through right multiplication and embodies the competitive coupling between the two market makers. Its entries capture how changes in the follower's inventory position affect the leader's value function; while the source matrix $\mathbf{B}(t)$ incorporates all the deterministic components of the system.

The formulation [\(2.21\)](#) transforms a high-dimensional equation into a standard Sylvester equation. The computational advantage of this representation is substantial, as, when considering a discretization of the time interval τ , it allows the application of well-established numerical algorithms such as the Bartels-Stewart method, proposed in [Bartels and Stewart \(1972\)](#), or Hessenberg-Schur decomposition for efficient solution, provided that the matrices $\mathbf{A}^1(t)$ and $-\mathbf{A}^2(t)$ satisfy the non-resonance condition of having no common eigenvalues, which ensures the

uniqueness and stability of the solution.

In the Chapter 3 we will numerically solve the maximization problem for the leader, thus solving the whole Stackelberg game. The idea is to find the optimal parameters that maximize the expected value of the leader's objective. The leader, indeed, being able to observe the follower response on the market quotes, can select ad hoc parameters for his precommitted strategy in order to maximize his own future expected value.

Chapter 3

Numerical analysis

1 Comparative results between the Euler scheme and the closed form approximation

In this Chapter we proceed to carry out an extensive series of numerical analyses in order to validate the theoretical results obtained, confirm the economic intuitions of the previous chapters, and verify the existence of a solution for Stackelberg equilibrium by evaluating the leader's objective as a function of his optimization possibilities¹.

First of all, we present the results in Boyce et al. (2025), setting $a, b, \beta \in \mathbb{R}$, fixed and $\theta(t) = 0, \forall t \in [0, T]$. The parameter value for this instance are in the Table 3.1. We also fix $x_0^F, x_0^L, q_0^F, q_0^L = 0$.

a	b	β	θ	ϕ	γ	s_0	σ_S	λ^a	λ^b	κ
0.1	0.1	0.05	0	0.1	0.03	100	1	10	10	2

Table 3.1: Parameter values for the Boyce et al. (2025) model.

In Figure 3.1 a single simulation in such framework is shown. As we can expect, the bid-ask spread of the leader is fixed during the whole trading period, while the follower's is decreasing due to the penalization terms.

As time horizon approaches, in fact, the follower needs to liquidate her own inventory (buy if it is negative), so she is willing to give up part of the profit derived from the spread for this purpose. When the follower's inventory is negative near expiration, the optimal bid spread reduces drastically in order to increase the chances of being filled by a market sell order, while conversely the ask spread increases. As

¹We provide the algorithm sketches in Appendix B




Figure 3.1: Simulation of a single path using the approximate closed-form solution (2.13). Parameter values are in Table 3.1.

for the leader, instead, if his inventory near expiration is negative, he is willing to buy with a negative spread (at a discount), rather than being penalized at $t = T$.

In the same framework, a Monte Carlo procedure was carried out in order to estimate the expected value of the leader's objective function, comparing it with the approximate closed form $v(0, x_0^F, s_0, q_0^F, q_0^L)$. The histograms in Figure 3.2 show how the theoretical value (3.243) is well approximated both in the case of using numerical methods (Euler schemes) (3.188) for solving the Hamilton-Jacobi-Bellman equation, and in the case of calculating the optimal strategy through the matrix form (3.199).

In Table 3.2 we provide a summary of the Monte Carlo procedure for both the leader and the follower. Due to the risk aversion of the agents, the average final inventory Q_T is close to zero, with rather limited variability. The spread on average shows the same behavior previously observed in Figure 3.1, such as decreasing in time.




Figure 3.2: Distribution of the terminal objective function of the follower over 1 million simulations. Euler approximation (left), approximate closed form (2.13) (right). Parameter values are in Table 3.1.

	Follower		Leader	
	Mean	Std	Mean	Std
Q_T	-0.0013	1.8695	-0.0040	3.7013
X_T	3.5813	186.9584	1.4625	370.1303
Obj	3.1878	1.8487	1.0677	2.6807
$\bar{\delta}^a + \bar{\delta}^b$	1.1290	0.0310	0.2000	0.0000
$\delta_T^a + \delta_T^b$	1.0602	2.0161×10^{-8}	0.2000	0.0000

Table 3.2: Summary of the results from 1 million simulations using the Euler scheme for the value function approximation. Parameter values are in Table 3.1.

2 Numerical solution of the Stackelberg equilibrium

In order to verify the existence of a Stackelberg equilibrium, the values of the leader's objective function at maturity are shown in Figure 3.3 as the parameters a and b vary. As mentioned earlier, it is necessary to introduce a trade-off for the leader so that he is not incentivized to raise the quoted spreads excessively. For this purpose, we consider the intensity of market order arrival processes as functions of the base spread level, such as $\lambda^a(t) = \lambda_0^a e^{-\kappa a(t)}$, $\lambda^b(t) = \lambda_0^b e^{-\kappa b(t)}$. In this way, when the market is less liquid, that is when spreads are higher, market orders arrive at a lower frequency. In order to optimize his objective, the leader must find a point, or a region, of optimality, managing the trade-off between revenues and trading frequency.




Figure 3.3: Average leader's terminal objective function over 100 thousands simulations. Parameter values are in Table 3.1 except from a and b .

When a , b are fixed during the trading period, it is evident that there exists one (or more) combination of these parameters that maximizes the leader's expected return. This confirms from a numerical point of view the existence of a Stackelberg equilibrium. It is interesting to note that the region where the objective is maximum is symmetric in the parameters a , b , with value in the interval $[0.2, 0.4]$.

Interesting details can derive from evaluating the effect of θ on the follower's quoted spreads. In Figure 3.4, the spread levels are shown as the follower's inventory varies. For the bid, the spread is not quoted when $q^F = \bar{q}^F$, conversely for the ask it is not quoted when $q^F = \underline{q}^F$. As can be noticed, when the inventory is high, at the beginning of the trading period, the follower needs to liquidate the shares in portfolio, and therefore his bid spread reduces as θ decreases.




Figure 3.4: Optimized follower bid δ_0^{b, F^*} (bottom panel) and ask δ_0^{a, F^*} (top panel) spreads as a function of the θ , at the initial point in time $t = 0$ (left panel) and at the terminal $t = T$ (right panel). Each line corresponds to a different inventory level for the follower. Parameter values are in Table [3.1](#).

The parameter θ exhibits distinctly asymmetric effects on ask and bid spreads, reflecting the strategic nature of the game where the leader's consideration of follower inventory affects buying and selling sides differently. For ask spreads (δ_0^{a, F^*}), when θ is negative, high inventory positions are associated with higher spreads, while for bid spreads (δ_0^{b, F^*}), the relationship appears inverted. This asymmetry captures the economic intuition that a leader who accounts for follower inventory will strategically adjust spreads to either compete or second with the follower's position.

Around $\theta = 0$, where the leader largely ignores follower inventory, spreads converge regardless of inventory levels, creating a reference case. However, as $|\theta|$ increases, the spreads diverge significantly based on inventory.

One of the most striking feature is the temporal evolution from $t = 0$ to $t = 1$, where the spread patterns become significantly more pronounced as the trading period progresses. At $t = 0$, the spreads converge relatively smoothly around $\theta = 0$, but at $t = 1$, we observe more dramatic divergence and steeper gradients, indicating that inventory pressure intensifies as the end of the trading period approaches.




Figure 3.5: Optimized follower bid δ^{b, F^*} (right panel) and ask δ^{a, F^*} (left panel) spreads as a function of time, when θ takes different values. Each line corresponds to a different inventory level for the follower. Parameter values are in Table [3.1](#).

In Figure [3.5](#) the time evolution of optimal leader spreads is shown, when the parameter θ takes negative ($\theta = -0.1$), zero or positive ($\theta = 0.1$) values. When such parameter is not negative we observe a systematic convergence of spreads as time approaches $t = 1$, reflecting the intensifying urgency to liquidate inventory positions as the trading period concludes. This convergence pattern demonstrates that regardless of initial inventory levels, the finite time horizon creates a natural force that drives the follower to quote near the leader.

The parameter θ exhibits profound asymmetric effects on the equilibrium structure, with the benchmark case of $\theta = 0.00$ providing a symmetric equilibrium where both ask and bid spreads converge to identical values at $t = 1$. This symmetry breaks down dramatically when θ deviates from zero, creating distinct patterns for positive and negative values. When $\theta = -0.10$, spreads show greater sensitivity to inventory levels throughout the trading period, and in particular when approaching the terminal date.

The inventory-dependent trajectories, represented by the color gradient from blue (negative inventory) to yellow (positive inventory), demonstrate that extreme inventory positions require more dramatic spread adjustments in the late stages of the trading period.

This asymmetry captures the strategic nature of the leader's response to follower inventory, where positive θ values create competition on the market while negative values generate higher spreads.




Figure 3.6: Average leader's terminal objective function over 100 thousands simulations as a function of θ , for different values of β . Parameter values are in Table 3.1 except from β .

Indeed, if we take as an example the case where the follower purchases a share, that is, her inventory $q^F \rightarrow q^F + 1$, her optimal strategy at $t + dt$ would be to decrease the bid spread and increase the ask spread, in order to be more competitive on the buying side and be filled with lower probability by a market order in the opposite direction. If the value of θ were positive, the leader at $t + dt$ decreases the relative ask spread (increasing the bid spread), creating greater competition in the market, which will therefore be more liquid and with smaller spreads. Conversely, if the value of θ is negative, the leader will put pressure on the follower by raising the market ask spread (and decreasing the bid), thus increasing the chances that the follower will be filled again on the buy side. As can be seen in Figure 3.6, when the leader creates greater competition in the market by lowering spreads, the expected value of his profit at maturity is positive but of limited amount. Conversely, when he tries to work against the follower, raising spreads in a manner contrary to what the latter desires, his expected objective increases considerably.

Figure 3.7 shows the expected value of the leader objective at maturity as a function of a and b when θ takes different values. In this way it is possible to understand what is the effect of the follower's inventory on the leader's optimal strategy and on his objective function.




Figure 3.7: Average leader's terminal objective function over 100 thousands simulations when $\theta = -0.5$ (top-left panel), $\theta = -0.05$ (top-right panel), $\theta = 0.05$ (bottom-left panel) and $\theta = 0.5$ (bottom-right panel). Parameter values are in Table 3.1.

We can observe how as the magnitude of parameter θ increases, the expected value of the objective for the leader increases, and the optimal base spread combination $\{a, b\}$ tends to move along the main diagonal, maintaining symmetry between bid and ask sides. However, when θ is positive, such as when the leader creates liquidity in the market, he must compensate by increasing the base spread levels, so the optimal values of a and b increase, and the expected value of the terminal objective grows linearly. Conversely, for negative values of θ , the optimal combination $\{a, b\}$ tends to zero and the expectation of final profit grows more than linearly. Relevant insights derive from the analysis of the risk aversion of the two agents




Figure 3.8: Average leader's terminal objective function over 100 thousands simulations when he is risk tolerant (left panel) and risk averse (right panel). We fix the leader parameters $\gamma = 0.01$, $\phi = 0.02$ for the left panel and $\gamma = 0.03$, $\phi = 0.1$ for the right panel. Parameter values are in Table 3.1 except from γ and ϕ .

in the market. As observable in Figure 3.8, although there is no marked relationship between the follower's risk aversion and the expected value of the leader's objective, an interesting trend is noted. When the leader is risk tolerant, that is, he slightly penalizes his running and terminal square inventory, his expected value is greater when the follower is particularly risk averse. Conversely, when the leader is more risk averse, he would prefer the follower to be risk tolerant. This behavior stems from the nature of the Stackelberg equilibrium itself, as the leader's objective is greater when he is opposite to the follower, being able to take advantage of his dominant position.

3 Including time dependencies

In this section we investigate the behavior of the equilibrium when the parameters $a(t)$ and $b(t)$ are time dependent. Through Monte Carlo simulations (1 million runs per scenario), we examine four distinct cases of parameter evolution: positive exponential, negative exponential, positive linear, and negative linear. All other parameters remain constant across all scenarios, and their value is shown in the Table 3.1

Scenario	Functional Form
Positive Exponential	$a(t) = b(t) = 0.1 + 0.1e^t$
Negative Exponential	$a(t) = b(t) = 0.1 - 0.1e^t$
Positive Linear	$a(t) = b(t) = 0.1 + 0.1t$
Negative Linear	$a(t) = b(t) = 0.1 - 0.1t$

Table 3.3: Selected time dependent functional form for the base bid $b(t)$ and ask $a(t)$ spreads, for $t \in [0, T]$.

The simulation results are summarized in Table 3.4, showing the mean both Follower and Leader strategies across key metrics.

Scenario	Metric	F	L	$F - L$
Positive Exp.	Q_T	-0.0018	0.0017	-0.0035
	X_T	3.5409	1.4201	2.1208
	Obj	3.0000	1.5925	1.4075
Negative Exp.	Q_T	0.0043	-0.0041	0.0084
	X_T	2.9137	-0.7564	3.6701
	Obj	2.9511	-1.1665	4.1176
Positive Linear	Q_T	-0.0050	-0.0143	0.0093
	X_T	3.8376	2.8654	0.9722
	Obj	2.9484	1.4373	1.5111
Negative Linear	Q_T	-0.0015	-0.0035	0.0020
	X_T	3.4747	1.0667	2.4080
	Obj	2.9331	0.7427	2.1904

Table 3.4: Summary of the results from 1 million simulations using approximate closed form 2.13 for the follower and numerical techniques for the leader, with time-dependent parameters $a(t)$, $b(t)$. Parameter values are given in Table 3.1 and Table 3.3.

The positive exponential case yields the highest expected objective for the follower, while in contrast negative exponential parameters produce particularly adverse outcomes for the leader strategy, showing negative mean performance metrics. Linear parameter changes demonstrate more moderate results, where positive linear trends favor the leader's terminal wealth while negative linear trends lead to more balanced performance.

A possible cause of this can be attributed to the fact that, by decreasing spreads exponentially, the leader has no particular profit margins, especially in the late stages, and for this reason is forced to trade at a loss in the market. In contrast, the follower maintains profit margins by quoting higher spreads; she will have a smaller number of transactions but with high expected profit. In general, when the trend of base

spreads $a(t)$, $b(t)$ increases over time, the expected objective of the leader is greater: the lower frequency of trades in the market, in fact, is offset by higher profits per trade.

These results show that the time evolution of parameters $a(t)$ and $b(t)$ significantly impacts system behavior, with exponential changes creating more extreme outcomes than linear ones.




Figure 3.9: Spreads difference $\delta^a - \delta^b$ in time for the leader (orange line) and the follower (blue line) for a single simulation. Positive exponential scenario (left panel) and positive linear scenario (right panel). Parameter values are given in Table 3.1 and Table 3.3.

In order to better understand the evolution of spreads over time, Figure 3.9 shows the spreads for a single simulation for both agents, when $a(t)$ and $b(t)$ are time dependent functions.

Chapter 4

Conclusion

In this work we have extended the framework in [Boyce et al. \(2025\)](#) by considering complex interactions between two LPs in the capital markets. The set of LPs is defined as the leader, while a single entity is the follower, who observes the quotes on the market and optimizes his strategy accordingly. A modeling of this framework through a Stackelberg equilibrium comes naturally, as the market, in turn, observes the follower acting, monitoring the LOB, and optimizes the parameters of its own precommitted strategy in order to maximize his own expected utility.

The leader precommits a strategy with a minimum level of spread, which can vary over time, and depends linearly on his own inventory and on that of the follower. While the dependence of the strategy on the leader's own inventory had already been analyzed, the introduction of the term θ can provide relevant details on how the leader can exploit the available information on the market and his dominant position to his advantage. When the leader acts against the follower, such as $\theta < 0$, the expectation on the terminal objective is much higher than when he accommodates the follower by providing liquidity to the market (lowering spreads).

The follower's optimal strategy was obtained in approximate closed form, while the leader problem was characterized through a system of Sylvester equations, solvable through numerical techniques well known in the literature. The existence of a Stackelberg equilibrium was proven numerically, as there exists one (or more) combinations of parameters that maximize the leader's terminal objective. This combination varies as the sensitivity parameters to the respective inventories vary, and always remains symmetric with respect to bid and ask. In accordance with the principle of non-cooperative games, moreover, the leader benefits when the follower assumes opposite behavior, that is when the levels of risk aversion between the two agents are contrary.

Finally, the impact of the time dependence of the base bid and ask spreads on the objectives of both leader and follower was evaluated. When the base spread grows over time, the expected profit grows accordingly, even though liquidity in the market decreases and market orders arrive with lower frequency. Conversely, if spreads decrease over time, the leader might find himself having to execute orders with negligible profit margins, or even at a loss.

This work is a first step in the study of interaction between market makers, introducing novelties in the literature. Possible developments of the same could be considering more complex interactions between market agents, or the development of a Nash equilibrium between multiple homogeneous LPs. Furthermore, a Stackelberg equilibrium framework is well-suited to the context of automated market making, so much can be explored in this direction.

Bibliography

- Amihud, Y., Mendelson, H., 1986. Asset pricing and the bid-ask spread. *Journal of financial Economics* 17, 223–249.
- Avellaneda, M., Stoikov, S., 2008. High-frequency trading in a limit order book. *Quantitative Finance* 8, 217–224.
- Bartels, R.H., Stewart, G.W., 1972. Algorithm 432 [c2]: Solution of the matrix equation $ax + xb = c$ [f4]. *Communications of the ACM* 15, 820–826.
- Barucci, E., Mathieu, A., Sánchez-Betancourt, L., 2025. Market making with fads, informed, and uninformed traders. *arXiv preprint arXiv:2501.03658*.
- Başar, T., Olsder, G.J., 1998. *Dynamic noncooperative game theory*. SIAM.
- Bergault, P., Sánchez-Betancourt, L., 2025. A mean field game between informed traders and a broker. *SIAM Journal on Financial Mathematics* 16, 358–388.
- Boyce, R., Herdegen, M., Sánchez-Betancourt, L., 2025. Market making with exogenous competition. *SIAM Journal on Financial Mathematics* 16, 692–706.
- Cartea, Á., Jaimungal, S., Penalva, J., 2015. *Algorithmic and high-frequency trading*. Cambridge University Press.
- Garman, M.B., 1976. Market microstructure. *Journal of financial Economics* 3, 257–275.
- Glosten, L.R., Milgrom, P.R., 1985. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of financial economics* 14, 71–100.
- Gould, M.D., Porter, M.A., Williams, S., McDonald, M., Fenn, D.J., Howison, S.D., 2013. Limit order books. *Quantitative Finance* 13, 1709–1742.
- Guéant, O., Lehalle, C.A., Fernandez-Tapia, J., 2013. Dealing with the inventory risk: a solution to the market making problem. *Mathematics and financial economics* 7, 477–507.

- Guo, I., Jin, S., Nam, K., 2024. Macroscopic market making games. arXiv preprint arXiv:2406.05662 .
- Ho, T., Stoll, H.R., 1981. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial economics* 9, 47–73.
- Ho, T.S., Stoll, H.R., 1983. The dynamics of dealer markets under competition. *The Journal of finance* 38, 1053–1074.
- Kyle, A.S., 1985. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society* , 1315–1335.
- Moon, J., Başar, T., 2018. Linear quadratic mean field stackelberg differential games. *Automatica* 97, 200–213.
- Pham, H., 2009. Continuous-time stochastic control and optimization with financial applications. volume 61. Springer Science & Business Media.
- Simaan, M., Cruz Jr, J.B., 1973. On the stackelberg strategy in nonzero-sum games. *Journal of Optimization Theory and Applications* 11, 533–555.
- Von Stackelberg, H., 2010. Market structure and equilibrium. Springer Science & Business Media.
- Walter, W., 2013. Ordinary differential equations. volume 182. Springer Science & Business Media.
- Xu, J., Zhang, H., 2015. Sufficient and necessary open-loop stackelberg strategy for two-player game with time delay. *IEEE transactions on cybernetics* 46, 438–449.
- Yong, J., 2002. A leader-follower stochastic linear quadratic differential game. *SIAM Journal on Control and Optimization* 41, 1015–1041.
- Zheng, Y., Shi, J., 2020. A stackelberg game of backward stochastic differential equations with applications. *Dynamic Games and Applications* 10, 968–992.
- Zhu, C., Yin, G., Baran, N.A., 2008. Feynman-kac formulas for regime-switching jump diffusions and their applications. *Mathematical Biosciences and Engineering* 5, 645–660. doi:[10.3934/mbe.2008.5.645](https://doi.org/10.3934/mbe.2008.5.645).

Appendix A

Proofs

1 Proof of Proposition 2.1

The problem 2.5 satisfies the classical dynamic programming condition, see e.g. Pham (2009), so can be expressed in terms of the value function as follows

$$v(t, s, x, q^F, q^L) = \sup_{\delta^F \in \mathcal{A}^F} J(\delta^F). \quad (\text{A.1})$$

We guess a solution $v(t, s, x, q^F, q^L) \in C^{1,2}([0, T] \times (0, \infty) \times \mathbb{R} \times \mathcal{Q}^F \times \mathbb{Z})$ of the form

$$v(t, s, x, q^F, q^L) := x + q^F s - \frac{1}{2}(\beta(t) + \theta(t)) (q^F)^2 - \beta(t) q^F q^L + g(t, q^F, q^L), \quad (\text{A.2})$$

where $g : [0, T] \times \mathcal{Q}^F \times \mathbb{Z} \rightarrow \mathbb{R}$ and $g \in C^{1,2}([0, T] \times \mathcal{Q}^F \times \mathbb{Z})$. The Hamilton-Jacobi-Bellman equation (HJB) associated with (2.5) with under the cash 2.3 and inventory 2.2 constraints is

$$\begin{aligned} -\frac{\partial v}{\partial t}(t, s, x, q^F, q^L) - \sup_{\delta^F \in \mathcal{A}^F} \mathcal{L}v(t, s, x, q^F, q^L) &= 0, \\ v(T, s, x, q^F, q^L) &= x + q^F s + \frac{a(T) - b(T)}{2} - (\gamma + \theta(T)) (q^F)^2 - \beta(T) q^F q^L, \end{aligned} \quad (\text{A.3})$$

where $\mathcal{L}v(t, s, x, q^F, q^L)$ stands for the infinitesimal generator of the value function, such that

$$\begin{aligned}
\mathcal{L}v(t, s, x, q^F, q^L) &= \frac{\sigma_S^2}{2} \frac{\partial^2 v}{\partial s^2}(t, s, x, q^F, q^L) - \phi(q^F)^2 \\
&+ \lambda^a [v(t, s, x, q^F, q^L - 1) - v(t, s, x, q^F, q^L)] \\
&+ \lambda^b [v(t, s, x, q^F, q^L + 1) - v(t, s, x, q^F, q^L)] \\
&+ \sup_{\delta^{a, F} \in \mathcal{A}^F} \left\{ \lambda^a \min \left\{ e^{-\kappa(\delta^{a, F} - a(t) + \beta(t)q^L + \theta(t)q^F)}, 1 \right\} \right. \\
&\quad \left. [v(t, s, x + s + \delta^{a, F}, q^F - 1, q^L) - v(t, s, x, q^F, q^L - 1)] \right\} \mathbb{1}_{q^F > \underline{q}^F} \\
&+ \sup_{\delta^{b, F} \in \mathcal{A}^F} \left\{ \lambda^b \min \left\{ e^{-\kappa(\delta^{b, F} - b(t) - \beta(t)q^L - \theta(t)q^F)}, 1 \right\} \right. \\
&\quad \left. [v(t, s, x - s + \delta^{b, F}, q^F + 1, q^L) - v(t, s, x, q^F, q^L + 1)] \right\} \mathbb{1}_{q^F < \bar{q}^F}
\end{aligned} \tag{A.4}$$

By plugging the ansatz [A.2](#) in [A.4](#) we obtain

$$\begin{aligned}
0 &= + \frac{\partial g}{\partial t}(t, q^F, q^L) - \frac{1}{2} \left(\frac{\partial \beta}{\partial t}(t) + \frac{\partial \theta}{\partial t}(t) + 2\phi \right) (q^F)^2 - \frac{\partial \beta}{\partial t}(t) q^F q^L \\
&+ \lambda^a (g(t, q^F, q^L - 1) - g(t, q^F, q^L) + \beta(t)q^F) \\
&+ \lambda^b (g(t, q^F, q^L + 1) - g(t, q^F, q^L) - \beta(t)q^F) \\
&+ \sup_{\delta^{a, F} \in \mathcal{A}^F} \left\{ \lambda^a \min \left\{ e^{-\kappa(\delta^{a, F} - a(t) + \beta(t)q^L + \theta(t)q^F)}, 1 \right\} \right. \\
&\quad \left. \left[\delta^{a, F} + \beta(t)q^L + \theta(t)q^F - \frac{1}{2} (\beta(t) + \theta(t)) + g(t, q^F - 1, q^L) - g(t, q^F, q^L - 1) \right] \right\} \mathbb{1}_{q^F > \underline{q}^F} \\
&+ \sup_{\delta^{b, F} \in \mathcal{A}^F} \left\{ \lambda^b \min \left\{ e^{-\kappa(\delta^{b, F} - b(t) - \beta(t)q^L - \theta(t)q^F)}, 1 \right\} \right. \\
&\quad \left. \left[\delta^{b, F} - \beta(t)q^L - \theta(t)q^F - \frac{1}{2} (\beta(t) + \theta(t)) + g(t, q^F + 1, q^L) - g(t, q^F, q^L + 1) \right] \right\} \mathbb{1}_{q^F < \bar{q}^F},
\end{aligned}$$

with

$$g(T, q^F, q^L) = \frac{a(T) - b(T)}{2} q^F - \left(\gamma + \frac{\theta(T) - \beta(T)}{2} \right) (q^F)^2. \tag{A.5}$$

Considering an adjustment term for the depths,

$$\begin{aligned}
c^{a, F}(t) &= \delta^{a, F}(t) + \beta(t)q^L + \theta(t)q^F - \frac{1}{2} (\beta(t) + \theta(t)), \\
c^{b, F}(t) &= \delta^{b, F}(t) - \beta(t)q^L - \theta(t)q^F - \frac{1}{2} (\beta(t) + \theta(t)),
\end{aligned}$$

we have the following PIDE

$$\begin{aligned}
0 = & + \frac{\partial g}{\partial t}(t, q^F, q^L) - \frac{1}{2} \left(\frac{\partial \beta}{\partial t}(t) + \frac{\partial \theta}{\partial t}(t) + 2\phi \right) (q^F)^2 - \frac{\partial \beta}{\partial t}(t) q^F q^L \\
& + \lambda^a (g(t, q^F, q^L - 1) - g(t, q^F, q^L) + \beta(t) q^F) \\
& + \lambda^b (g(t, q^F, q^L + 1) - g(t, q^F, q^L) - \beta(t) q^F) \\
& + \sup_{c^{a,F}} \left\{ \lambda^a \min \left\{ e^{-\kappa(c^{a,F} + \frac{1}{2}(\beta(t) + \theta(t)) - a(t))}, 1 \right\} [c^{a,F} + g(t, q^F - 1, q^L) - g(t, q^F, q^L - 1)] \right\} \mathbb{1}_{q^F > \underline{q}^F} \\
& + \sup_{c^{b,F}} \left\{ \lambda^b \min \left\{ e^{-\kappa(c^{b,F} + \frac{1}{2}(\beta(t) + \theta(t)) - b(t))}, 1 \right\} [c^{b,F} + g(t, q^F + 1, q^L) - g(t, q^F, q^L + 1)] \right\} \mathbb{1}_{q^F < \bar{q}^F},
\end{aligned} \tag{A.6}$$

where the terminal condition is [\(A.5\)](#).

Optimizing over $c^{a,F}$ and $c^{b,F}$ we obtain

$$\begin{aligned}
c^{a,F*}(t, q^F, q^L) &= \max \left\{ \hat{c}^{a,F}(t, q^F, q^L), a(t) - \frac{\beta(t) + \theta(t)}{2} \right\}, \\
c^{b,F*}(t, q^F, q^L) &= \max \left\{ \hat{c}^{b,F}(t, q^F, q^L), b(t) - \frac{\beta(t) + \theta(t)}{2} \right\},
\end{aligned}$$

where

$$\begin{aligned}
\hat{c}^{a,F}(t, q^F, q^L) &= \frac{1}{\kappa} - g(t, q^F - 1, q^L) + g(t, q^F, q^L - 1), \\
\hat{c}^{b,F}(t, q^F, q^L) &= \frac{1}{\kappa} - g(t, q^F + 1, q^L) + g(t, q^F, q^L + 1).
\end{aligned}$$

In terms of optimal control strategy, such as follower's quoted depths, we have

$$\begin{aligned}
\delta^{a,F*}(t, q^F, q^L) &= \max \left\{ \hat{\delta}^{a,F*}(t, q^F, q^L), a(t) - \beta(t) q^L - \theta(t) q^F \right\}, \\
\delta^{b,F*}(t, q^F, q^L) &= \max \left\{ \hat{\delta}^{b,F*}(t, q^F, q^L), b(t) + \beta(t) q^L + \theta(t) q^F \right\}.
\end{aligned}$$

where

$$\begin{aligned}
\hat{\delta}^{a,F*}(t, q^F, q^L) &= \frac{1}{\kappa} - g(t, q^F - 1, q^L) + g(t, q^F, q^L - 1) - \beta(t) q^L - \theta(t) q^F + \frac{1}{2}(\beta(t) + \theta(t)) \\
\hat{\delta}^{b,F*}(t, q^F, q^L) &= \frac{1}{\kappa} - g(t, q^F + 1, q^L) + g(t, q^F, q^L + 1) + \beta(t) q^L + \theta(t) q^F + \frac{1}{2}(\beta(t) + \theta(t)).
\end{aligned}$$

2 Proof of Theorem [3.1](#)

The argument follows the proofs of Proposition 1 and Theorem 2 in [Guéant et al. \(2013\)](#). Define the function $\omega : [0, T] \times \mathcal{Q}^F \times \mathbb{Z} \rightarrow (0, \infty)$ and $\omega \in C^{1,2}([0, T] \times \mathcal{Q}^F \times \mathbb{Z})$.

The ansatz for the value function $v(t, s, x, q^F, q^L) \in C^{1,2}([0, T] \times (0, \infty) \times \mathbb{R} \times \mathcal{Q}^F \times \mathbb{Z})$ associated to the problem (2.5) with unrestrained maximizers equal or greater than the leader quoted spread is

$$v(t, s, x, q^F, q^L) := x + q^F s - \frac{1}{2}(\beta(t) + \theta(t)) (q^F)^2 - \beta(t) q^F q^L + \frac{1}{\kappa} \log(\omega(t, q^F)).$$

The value function solves the following approximate Hamilton-Jacobi-Bellman equation

$$\begin{aligned} -\frac{\partial v}{\partial t}(t, s, x, q^F, q^L) - \sup_{\delta^F \in \mathcal{A}^F} \mathcal{L}v(t, s, x, q^F, q^L) &= 0, \\ v(T, s, x, q^F, q^L) &= x + q^F s + \frac{a(T) - b(T)}{2} - (\gamma + \theta(T)) (q^F)^2 - \beta(T) q^F q^L, \end{aligned} \quad (\text{A.7})$$

where

$$\begin{aligned} \mathcal{L}v(t, s, x, q^F, q^L) &= \frac{\sigma_S^2}{2} \frac{\partial^2 v}{\partial s^2}(t, s, x, q^F, q^L) - \phi(q^F)^2 \\ &\quad + \lambda^a [v(t, s, x, q^F, q^L - 1) - v(t, s, x, q^F, q^L)] \\ &\quad + \lambda^b [v(t, s, x, q^F, q^L + 1) - v(t, s, x, q^F, q^L)] \\ &\quad + \sup_{\delta^{a, F} \in \mathcal{A}^F} \left\{ \lambda^a \min \left\{ e^{-\kappa(\delta^{a, F} - a(t) + \beta(t)q^L + \theta(t)q^F)}, 1 \right\} \right. \\ &\quad \left. [v(t, s, x + s + \delta^{a, F}, q^F - 1, q^L) - v(t, s, x, q^F, q^L - 1)] \right\} \mathbb{1}_{q^F > \underline{q}^F} \\ &\quad + \sup_{\delta^{b, F} \in \mathcal{A}^F} \left\{ \lambda^b \min \left\{ e^{-\kappa(\delta^{b, F} - b(t) - \beta(t)q^L - \theta(t)q^F)}, 1 \right\} \right. \\ &\quad \left. [v(t, s, x - s + \delta^{b, F}, q^F + 1, q^L) - v(t, s, x, q^F, q^L + 1)] \right\} \mathbb{1}_{q^F < \bar{q}^F}. \end{aligned} \quad (\text{A.8})$$

Considering the first ansatz (A.2) and by setting $\omega(t, q^F) = \exp\{\kappa g(t, q^F)\}$, we obtain

$$\begin{aligned} 0 &= + \frac{\partial \omega}{\partial t}(t, q^F) \\ &\quad + \kappa \left[-\frac{1}{2} \left(+ \frac{\partial \theta}{\partial t}(t) + 2\phi \right) (q^F)^2 + (\lambda^a - \lambda^b) \beta(t) q^F \right] \omega(t, q^F) \\ &\quad + \lambda^a e^{-1 - \kappa(\frac{1}{2}(\beta + \theta(t)) - a(t))} \omega(t, q^F - 1) \mathbb{1}_{q^F > \underline{q}^F} \\ &\quad + \lambda^b e^{-1 - \kappa(\frac{1}{2}(\beta + \theta(t)) - b(t))} \omega(t, q^F + 1) \mathbb{1}_{q^F < \bar{q}^F}, \\ \omega(T, q^F) &= \exp \left\{ \kappa \left[\frac{a(T) - b(T)}{2} q^F - \left(\gamma + \frac{\theta(T) - \beta}{2} \right) (q^F)^2 \right] \right\}. \end{aligned}$$

Given the time interval $[0, T]$, we define its discretization as $\tau = \{t_n, n = 0, 1, \dots, N\}$, where $dt = t_n - t_{n-1}$ and $N = \frac{T}{dt}$.

Let us also assume $\underline{q}^F = -\bar{q}^F$. We define the tridiagonal set of matrices $(\mathbf{A}_{i,q,n})_{\underline{q}^F \leq i, q \leq \bar{q}^F, 0 \leq n \leq N}$

$$\mathbf{A}_{i,q,n} = \begin{cases} \frac{1}{2} \left(\frac{\partial \theta}{\partial t}(t_n) + 2\phi \right) q^2 - \beta \kappa (\lambda^a - \lambda^b) q & \text{if } i = q, \\ -\lambda^a \exp \left(-1 - \kappa \left(\frac{\beta + \theta(t_n)}{2} - a(t_n) \right) \right) & \text{if } i = q - 1, \\ -\lambda^b \exp \left(-1 - \kappa \left(\frac{\beta + \theta(t_n)}{2} - b(t_n) \right) \right) & \text{if } i = q + 1, \\ 0 & \text{otherwise,} \end{cases}$$

and the vector $(\mathbf{B}_q)_{\underline{q}^F \leq q \leq \bar{q}^F}$

$$\mathbf{B}_q = \exp \left\{ \kappa \left[\frac{a(T) - b(T)}{2} q - \left(\gamma + \frac{\theta(T) - \beta}{2} \right) q^2 \right] \right\}.$$

It is straightforward to prove that

$$\omega(t_n, q) = \left(\exp \left\{ - \sum_{m=n}^{N-1} \mathbf{A}_m \cdot dt \right\} \mathbf{B} \right) q, \quad (\text{A.9})$$

where $-\mathbf{A}_n$ is essentially positive, for further details see [Boyce et al. \(2025\)](#) and [Walter \(2013\)](#). The optimal controls can be computed as

$$\begin{aligned} \hat{\delta}^{a, F^*}(t_n, q^F, q^L) &= \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F - 1)} \right) \right] - \beta q^L - \theta(t_n) q^F + \frac{1}{2}(\beta + \theta(t_n)) \\ \hat{\delta}^{b, F^*}(t_n, q^F, q^L) &= \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F + 1)} \right) \right] + \beta q^L + \theta(t_n) q^F + \frac{1}{2}(\beta + \theta(t_n)). \end{aligned} \quad (\text{A.10})$$

where the unrestrained controls must be equal or greater the one-tick more generous spreads, such as parameters must satisfy the following conditions

$$\begin{aligned} \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F - 1)} \right) \right] + \frac{1}{2}(\beta + \theta(t_n)) &\geq a(t_n), \quad t_n \in \tau, \quad q^F \in \{\underline{q}^F + 1, \dots, \bar{q}^F\}, \\ \frac{1}{\kappa} \left[1 + \log \left(\frac{\omega(t_n, q^F)}{\omega(t_n, q^F + 1)} \right) \right] + \frac{1}{2}(\beta + \theta(t_n)) &\geq b(t_n), \quad t_n \in \tau, \quad q^F \in \{\underline{q}^F, \dots, \bar{q}^F - 1\}. \end{aligned}$$

3 Proof of Proposition [4.3](#)

Consider the PIDE [\(2.18\)](#) characterizing the evolution of the leader objective function with the terminal condition [\(2.19\)](#). We make an ansatz for $u(t, s, x, q^L, q^F)$ of the form

$$u(t, s, x, q^L, q^F) = x + q^L s + h(t, q^L, q^F), \quad (\text{A.11})$$

with $h(T, q^L, q^F) = -\gamma (q^L)^2$.

The (2.18) can be written as

$$\begin{aligned}
-\frac{\partial h}{\partial t}(t, q^L, q^F) &= -\phi(q^L)^2 + \lambda^a[h(t, q^L, q^F - 1) - h(t, q^L, q^F)] \\
&\quad + \lambda^b[h(t, q^L, q^F + 1) - h(t, q^L, q^F)] \\
&\quad + \Lambda^{a,L}(t, q^F)[h(t, q^L - 1, q^F) - h(t, q^L, q^F - 1)]\mathbb{1}_{q^L > \underline{q}^L} \\
&\quad + \Lambda^{b,L}(t, q^F)[h(t, q^L + 1, q^F) - h(t, q^L, q^F + 1)]\mathbb{1}_{q^L < \bar{q}^L} \\
&\quad + \Lambda^{a,L}(t, q^F)\delta^{a,L}(t, q^L, q^F)\mathbb{1}_{q^L > \underline{q}^L} + \Lambda^{b,L}(t, q^F)\delta^{b,L}(t, q^L, q^F)\mathbb{1}_{q^L < \bar{q}^L}, \\
h(T, q^L, q^F) &= -\gamma (q^L)^2.
\end{aligned} \tag{A.12}$$

Let $\underline{q}^L = -\bar{q}^L$, $\underline{q}^F = -\bar{q}^F$ and $\underline{q}^L = \underline{q}^F$. We can represent the function $h(t, q^L, q^F)$ as a matrix $(\mathbf{H}_{i,j}(t))_{\underline{q}^L \leq i \leq \bar{q}^L, \underline{q}^F \leq j \leq \bar{q}^F}$, where the entry $\mathbf{H}_{i,j}(t)$ corresponds to $h(t, q_i^L, q_j^F)$.

We also introduce three different matrices, $(\mathbf{A}_{i,k}^1(t))_{\underline{q}^L \leq i, k \leq \bar{q}^L}$, $(\mathbf{A}_{j,\ell}^2(t))_{\underline{q}^F \leq j, \ell \leq \bar{q}^F}$ and $(\mathbf{B}_{i,j}(t))_{\underline{q}^L \leq i \leq \bar{q}^L, \underline{q}^F \leq j \leq \bar{q}^F}$, each of them involving a different component in the differential operator:

$$\begin{aligned}
(\mathbf{A}_{i,k}^1(t)) &= \begin{cases} \lambda^a + \lambda^b & \text{if } i = k, \\ -\Lambda^{a,L}(t, q^F) & \text{if } k = i - 1, \\ -\Lambda^{b,L}(t, q^F) & \text{if } k = i + 1, \\ 0 & \text{otherwise,} \end{cases} \\
(\mathbf{A}_{j,\ell}^2(t)) &= \begin{cases} \Lambda^{a,L}(t, q_\ell^F) - \lambda^a & \text{if } \ell = j - 1, \\ \Lambda^{b,L}(t, q_\ell^F) - \lambda^b & \text{if } \ell = j + 1, \\ 0 & \text{otherwise,} \end{cases}
\end{aligned}$$

$$(\mathbf{B}_{i,j}(t)) = -\phi(q_i^L)^2 + \Lambda^{a,L}(t, q_j^F)\delta^{a,L}(t, q_i^L, q_j^F) + \Lambda^{b,L}(t, q_j^F)\delta^{b,L}(t, q_i^L, q_j^F).$$

We can rewrite (A.12) in matrix formulation as

$$\frac{\partial \mathbf{H}}{\partial t}(t) = \mathbf{A}^1(t)\mathbf{H}(t) + \mathbf{H}(t)\mathbf{A}^2(t) + \mathbf{B}(t), \tag{A.13}$$

with the terminal condition $(\mathbf{H}_{i,j}(T)) = -\gamma(q_i^L)^2, \forall i \in \mathcal{Q}^L, j \in \mathcal{Q}^F$.

Appendix B

Algorithms

In this appendix we provide sketches for the numerical algorithm developed in the Chapter [3](#). The full code is available in the [GitHub directory](#).

Algorithm 1 Market-Making Game Simulation

```
1: for  $i = 1$  to  $steps - 1$  do
2:    $t \leftarrow time(i)$ 
3:   Calculate intensities:
4:    $\lambda_a \leftarrow \lambda_0 \exp(-\kappa a(t)), \lambda_b \leftarrow \lambda_0 \exp(-\kappa b(t))$ 
5:   Follower's optimal spreads:
6:    $(\delta_a^F, \delta_b^F) \leftarrow \text{calculate\_deltas}(t, T, q_i^F, q_i^L, \dots)$ 
7:   Update stock price:
8:    $S_{i+1} \leftarrow S_i + \sigma \sqrt{\Delta t} Z \quad Z \sim \mathcal{N}(0, 1)$ 
9:   Leader's quotes:
10:   $\delta_a^L \leftarrow a(t) - \beta q^L - \theta q^F, \delta_b^L \leftarrow b(t) + \beta q^L + \theta q^F$ 
11:  Simulate order arrivals:
12:   $M_a \sim \text{Poisson}(\lambda_a \Delta t), M_b \sim \text{Poisson}(\lambda_b \Delta t)$ 
13:  Compute fill probabilities:
14:   $p_a^F = e^{-\kappa(\delta_a - \delta_a^L)} \wedge 1$ , similar for b
15:  Determine executed orders:
16:   $fill_a^F = (U_a \leq p_a^F) \wedge (q^F \neq \underline{q} + 1), U \sim \mathcal{U}[0, 1]$ , similar for b
17:   $fill_a^L = (U_a \leq 1 - p_a^F) \wedge (q^L \neq \underline{q} + 1), U \sim \mathcal{U}[0, 1]$ , similar for b
18:  Update positions:
19:   $q^F \leftarrow q^F - \mathbb{I}_{fill_a^F} \cdot N_a + \mathbb{I}_{fill_b^F} \cdot N_b, q^L \leftarrow q^L - \mathbb{I}_{fill_a^L} \cdot N_a + \mathbb{I}_{fill_b^L} \cdot N_b$ 
20:   $X_{t+\Delta t}^F = X_t^F + \mathbb{I}_{fill_a^F} \cdot M_t^a \cdot (S_t + \delta_a^F) - \mathbb{I}_{fill_b^F} \cdot M_t^b \cdot (S_t - \delta_b^F)$ 
21:   $X_{t+\Delta t}^L = X_t^L + \mathbb{I}_{fill_a^L} \cdot M_t^a \cdot (S_t + \delta_a^L) - \mathbb{I}_{fill_b^L} \cdot M_t^b \cdot (S_t - \delta_b^L)$ 
22: end for
```

Appendix C

Approximate closed form

In this appendix testing results are shown in order to evaluate the goodness of the approximate closed form solution (2.13) and the time independent form derived in Proposition 3.2. Out of ten rounds of ten thousand runs each, we obtained at most one violation of the condition (2.12) each. This means that, given a reasonable set of parameters, the violation rate is $\leq 0.01\%$, which means a high goodness of approximation for optimal controls.

In the Figure C.1 one of the violations obtained is shown.




Figure C.1: Random violation of condition (2.12) for the closed form approximation. Parameter value are in Table 3.1.

Development and Analysis of a Uniswap v3 Agent-Based Model

TEAM 3

DANIEL PRESTA, McMaster University
SHAOHONG CHEN, University of Cape Town
THOMAS KOEN, University of Cape Town
DE WET OOSTHUIZEN, University of Cape Town

Supervisor:
ADRIEN NGUYEN-HUU, Montpellier University

Contents

1	Introduction	4
2	Literature review	5
2.1	Introduction	5
2.2	Uniswap v3	5
2.3	Agent-based modeling	6
2.4	Ecology	7
2.5	Conclusion	9
3	The model	10
3.1	Uniswap v3	10
3.1.1	Price discretisation via ticks	10
3.1.2	Liquidity and virtual reserves	10
3.1.3	Fee accounting	11
3.1.4	Swap mechanism within ticks	11
3.1.5	Swap mechanism across ticks	12
3.1.6	Liquidity provision within a specified price range	12
3.2	Agent-based model	13
3.3	Ecology of agents	13
3.3.1	Noise trader	13
3.3.2	Trend follower	14
3.3.3	Fundamentalist	14
4	Results & Discussion	16
4.1	Visual comparison of market dynamics	16
4.1.1	Market dynamics with two liquidity providers	16
4.1.2	Market dynamics with three liquidity providers	21
4.1.3	Market dynamics with an active liquidity provider	23
4.2	Comparison of performance metrics	25
5	Tree-based sensitivity analysis	27
5.1	Theory of regression trees	27
5.2	Random forests and gradient boosting	29
5.2.1	Theory of random forests	30
5.2.2	Theory of gradient boosting	31
5.2.3	Variable importance	32

5.3	Tree-based sensitivity results	33
6	Conclusion	39

1 Introduction

In a world where comfort and efficiency have become two of the most prized social assets, the emergence of Decentralised Finance (DeFi) comes as no surprise. However, while DeFi enhances comfort and efficiency, it also introduces greater complexity and, much like traditional financial markets, carries inherent risks. As DeFi develops further, it is of the utmost importance to fully comprehend its workings and to stay up to date with its new platforms and implementations.

Auer R. and F. (2023) define DeFi as “a competitive, contestable, composable and non-custodial financial ecosystem built on technology that does not require a central organisation to operate.” Perhaps a simpler definition could be even more instructive: DeFi is a system in which people can use financial services (such as trading or lending) without the need for banks or other financial intermediaries, through the use of blockchain technology and smart contracts.

The largest active DeFi platform is Uniswap v3, a Decentralised Exchange (DEX) built on the same Automated Market Maker (AMM) foundations as its earlier versions, but with a handful of new features that give liquidity providers (LPs) and traders a lot more to consider when acting on the market (Adams et al. 2021). AMMs are systems that pool liquidity and make it available to traders using a set algorithm (Othman 2018). Some of these new features include concentrated liquidity, a flexible fee structure, and protocol fee governance. Despite Uniswap v3’s innovations, little is known about how the behaviours of LPs and traders collectively shape market outcomes. The introduction of concentrated liquidity and flexible fees have sparked a new debate regarding liquidity provision strategies and their abilities to optimise returns, as well as how different LP strategies affect the market. Furthermore, the ecology of traders in the DeFi market is just as prevalent as in the more common centralised market. There exists a gap in the literature, however, regarding the behavioural effects of traders within a DeFi market.

The aim of this paper is to fill this gap by analysing how LPs and different types of traders interact with each other within a Uniswap v3 framework, by investigating the concentrated liquidity feature introduced in Adams et al. (2021), and by modelling the individual behaviours and beliefs of agents in the Uniswap market. This is accomplished by building an agent-based-model (ABM) that simulates the features of a Uniswap v3 platform, thereby allowing us to seamlessly study the impact of the various agents on the model. This ABM enables us to capture the complex dynamics that are present on these DeFi platforms – a feature that is often missed when using more traditional models. Understanding this ecology is not only academically valuable, but also practically relevant, as it can help design better incentives for LPs, improve price stability, and enhance overall market effi-

ciency.

The remainder of this paper is structured as follows: Section 2 reviews related literature on Uniswap v3, agent-based modeling, and the ecology of LPs and traders. Section 3 explores the theoretical considerations of our model by providing a thorough explanation of the Uniswap v3 framework, the structure of our ABM, and the ecology of traders and LPs in the model. Section 4 presents and discusses our results, while Section 5 presents the theory of tree-based sensitivity analysis, and a handful of key findings in the context of the Uniswap v3 model. Lastly, Section 6 concludes by summarizing our most important results and their implications, as well as providing suggestions for future work.

2 Literature review

2.1 Introduction

Decentralised Finance (DeFi) is reshaping financial markets by removing financial intermediaries, and a prime example of this is the Uniswap v3 framework. Built on the same automated market maker (AMM) logic as the earlier Uniswap v1 and v2 systems, this newer decentralised exchange (DEX) lets liquidity providers (LPs) concentrate their liquidity within pre-specified price ranges and set flexible fee tiers. The additional flexibility gives LPs far more control than before, but also leads to an increase in system complexity. These features raise new strategic questions for both LPs and traders, sparking interest in how their choices influence market behaviour. This review dives into three key areas of literature: the foundational work of Uniswap v3, the workings of agent-based models (ABMs) in the world of finance, and the study of the ecology of LPs and traders in general DeFi markets, as well as the Uniswap v3 framework. The purpose of this literature review is to place our paper, in which we use an agent-based model to examine the ecology of LPs and traders in a Uniswap v3 model, within the broader body of research across these three areas.

2.2 Uniswap v3

The foundational paper on which this research is built is the paper by Adams et al. (2021). In their paper they present the Uniswap v3, a new type of AMM that lets liquidity providers choose the specific price ranges where their funds are active, while keeping liquidity fragmentation and gas costs relatively low. AMMs are systems that pool liquidity and provide it to traders using a set algorithm (Othman, 2018). Their paper follows the work by Adams et al. (2020), wherein they provide a

clear explanation of how the Uniswap v2 system works, which is essential to study the newer Uniswap v3. [Adams et al. \(2021\)](#) explain that Uniswap v3 builds on the same constant product formula used in previous versions, but introduces several important new features such as concentrated liquidity, flexible fees, protocol fee governance, and an improved price oracle and liquidity oracle. As mentioned earlier, an important feature in the Uniswap market is that of constant product markets. [Angeris et al. \(2021\)](#) provide a simple formal analysis of constant product markets and proceed to show that Uniswap satisfies several desirable properties for a market.

[Aigner and Gurvinder Dhaliwal \(2021\)](#) define the Uniswap as a DEX and mention how DEXs are replacing the traditional order books used for trading. Based on this idea, [Lehar and Parlour \(2023\)](#) go on to provide a comparison between AMMs and limit order books (LOBs). They argue that AMMs are ideal when assets have low volatility, trades are small and frequent, and noise trading is quite prevalent, whereas LOBs are ideal when assets are more volatile and their prices change rapidly or very often.

2.3 Agent-based modeling

As mentioned in the introduction, this research paper implements an ABM in the context of Uniswap v3. [Axtell and Farmer \(2025\)](#) define an ABM as a new computational methodology for representing the behaviour of individuals in order to study social phenomena. Similarly, [Datseris et al. \(2024\)](#) define agent-based modelling as a method for simulating a system wherein autonomous agents (LPs and traders, in our case) interact with their environment and with each other, given a set of rules defined at the start of the simulation. In their work, [Axtell and Farmer \(2025\)](#) review ABMs in the areas of finance and economics and proceed to emphasise the fact that it is useful in easing the usual assumptions in standard economic models. They go on to present a vision for how ABMs could be utilised to build more realistic economic models in later work, while emphasising potential issues with this approach.

In order to simulate an ABM in which LPs and traders interact with each other in a Uniswap v3 environment, we use the `Agents.jl` package in `Julia`, presented by [Datseris et al. \(2024\)](#). Due to the fact that ABMs are not described by simple and clear mathematical equations, the code that generates ABMs is usually very complicated and slow. The software presented by [Datseris et al. \(2024\)](#) provides an ABM analysis platform with minimal code complexity. By comparing their software to other ABM packages in other languages, they find that it outperforms all of them in terms of efficiency and simplicity, while still offering the same features as the rest.

Moving on from the implementation of ABMs, another important consideration is model validation. [Fagiolo et al. \(2019\)](#) review recent progress in this area and categorise validation techniques into three main types: comparing simulated and real-world data, calibrating model parameters, and exploring the parameter space. The paper also points out key challenges, such as the need for improved hypothesis testing, as well as the importance of ensuring that properties such as stationarity and ergodicity are satisfied. These issues are particularly relevant in financial applications of agent-based modelling.

Furthermore, [Hommes \(2005\)](#) reviews the research on dynamic heterogeneous agent models (HAMs) in economics and finance, focusing mainly on simpler models that can be explored using a combination of analytical and computational techniques. These models often feature agents who are not fully rational but instead rely on rules of thumb or heuristics that, while imperfect, are effective in practice. In particular, HAMs are able to reproduce key patterns seen in real financial markets, such as excessive price swings, large trading volumes, trend following, and volatility clustering.

Finally, we see how ABMs can be applied to decentralised finance (DeFi) through [Cong et al. \(2024\)](#), who introduce agent-based modelling for research on decentralised autonomous organisations (DAOs) and DeFi. They use agent-based modelling to study how different design choices affect the concentration of governance tokens in DAOs and DeFi systems. Their study shows that ABMs are powerful and flexible tools for analysing and improving governance structures in DeFi and DAOs.

Further relevant work is produced by [Wang and Kampakis \(2024\)](#), who introduce Tokenlab, an ABM framework designed to study price dynamics and speculative behaviour in token-based markets. The framework breaks down complex systems into simple agent interactions, allowing for the simulation of various speculative strategies and their impacts on price formation. A key feature is its controller mechanism, which models different types of speculators and their interactions. The study provides useful insights into how speculative activity shapes market sentiment and token price movements, making it highly relevant for analysing dynamics in DeFi markets.

2.4 Ecology

An interesting and expanding field of literature in the world of DeFi, specifically for Uniswap v3, is the ecology of LPs and traders in the market. There have been a few studies analysing the different types of behaviours and strategies for both

LPs and traders. [Fan et al. \(2024\)](#) discuss how LPs can either be active or passive when it comes to providing liquidity. Active LPs strategically reallocate liquidity constantly in order to achieve the maximum return, whereas passive LPs allocate liquidity once and then wait their position out, possibly reallocating once or twice if necessary. [Fan et al. \(2024\)](#) provide a new dynamic liquidity strategy that allows LPs to adjust their positions over time, which thus enables them to reduce potential losses by reallocating liquidity using capital earned through previous trades. Their results show that dynamic liquidity provision strategies outperform more general passive strategies.

Furthermore, as explained by [Adams et al. \(2021\)](#), one of the new features provided by Uniswap v3 is concentrated liquidity. This feature introduces a new avenue for LPs to exhibit differences in behaviour, as the framework now includes LPs who are either risk-averse or have an appetite for risk. Risk-averse LPs opt for a wider price range when providing liquidity in order to ensure that they still receive fees from trading, even if the price volatility is high. LPs with a larger risk appetite opt for a smaller price range, hoping to obtain a greater fee amount due to an increased proportion of liquidity. [Bar-On and Mansour \(2023\)](#) explain that this feature introduces the issue of locating the optimal strategy for choosing price intervals. They formalise this issue as an online learning problem with non-random rewards. They apply regret-minimisation methods to present a liquidity provision strategy that insures a lower bound on the reward. Furthermore, [Álvaro Cartea et al. \(2024\)](#) propose a strategy for selecting optimal liquidity range widths, which outperforms existing LP approaches based on real Uniswap v3 data.

[Heimbach et al. \(2022\)](#) raise the concern that the complexity of concentrated liquidity might cause issues due to the increased amount of considerations for new LPs, suggesting that a high level of skill and understanding of the Uniswap v3 system is needed in order to generate meaningful returns. However, [Fan et al. \(2022\)](#) suggest that increased flexibility of price range options is beneficial to both LPs and traders, considering the improvement of efficiency and returns.

Finally, we turn to another key participant in the DeFi ecosystem by examining the literature on traders. This paper studies the relationship between three types of traders in the Uniswap v3 environment: noise traders, trend followers, and fundamentalists. [Bottazzi et al. \(2005\)](#) investigate how market dynamics are influenced by changing the ecology of individual traders, while maintaining the governing institutional structures. They go on to analyse two different agent typologies in the form of trend followers and noise traders. We expand on their work by adding fundamentalist traders and by implementing an ABM on a simulated Uniswap v3 market, rather than the general financial market.

Additional literature on the ecology of traders is produced by Hommes (2005), who explores dynamic models where traders differ in expectations and trading strategies. They examine trend followers, noise traders, and fundamentalists, and find that markets with such heterogeneous agents can exhibit complex dynamics such as excess volatility, bubbles, and crashes.

Furthermore, Boswijk et al. (2007) estimate an asset pricing model that changes between fundamentalist and trend-following regimes. They find that trend-following regimes are more dominant during price run-ups, whereas fundamentalist regimes are more dominant after bubbles. Finally, Wang and Kampakis (2024) use the TokenLab ABM to show how different speculative strategies influence token prices. They simplify the market to include different types of interacting traders (representing different behaviours), namely short-term speculators, trend followers, and liquidity hoarders.

2.5 Conclusion

A clear progression in research emerges when exploring the current literature surrounding DeFi. The research begins by unpacking the functionality of AMMs like the Uniswap v3 framework, before constructing ABMs to simulate the interactions between different agents. Lastly, there is plenty of literature that analyses the wide range of behaviours and strategies used by participants in these markets. Uniswap v3 builds on previous editions through the introduction of concentrated liquidity and a flexible fee structure, which ultimately reshapes how LPs and traders interact with each other. After engaging with the literature, we conclude that agent-based modelling is a powerful method to capture these interactions by simulating a system in which LPs and traders make decisions within a DeFi market. This paper builds on the various fields of literature that have been discussed in this section, with the ultimate goal of using an ABM to deepen our understanding of agent interactions in a Uniswap v3 market.

3 The model

3.1 Uniswap v3

3.1.1 Price discretisation via ticks

Uniswap v3 discretises the continuous price space into evenly spaced units called *ticks*. Each tick index i corresponds to a price level defined by:

$$p(i) = 1.0001^i, \quad \sqrt{p(i)} = 1.0001^{i/2} \quad (1)$$

The inverse formula for obtaining a tick index from a given price is:

$$i = \lfloor \log_{\sqrt{1.0001}} \sqrt{p} \rfloor \quad (2)$$

Each tick represents a price movement of approximately 0.01%. A price range is specified by a pair of signed integer tick indices: a lower tick i_l and an upper tick i_u . In this model, we set the tick spacing to 1.

3.1.2 Liquidity and virtual reserves

Unlike earlier AMMs, Uniswap v3 does not explicitly maintain reserve balances for `token0` (x) and `token1` (y). Instead, it tracks two core state variables:

- The active liquidity, L
- The current square root price, $\sqrt{P}$

The liquidity L remains constant within a tick interval $[i_l, i_u]$ and only changes when crossing a tick boundary, or through the addition or removal of liquidity by the liquidity provider. In contrast, the price $\sqrt{P}$ evolves continuously during swaps. These variables are sufficient to calculate the level of virtual reserves and to facilitate efficient updates to the pool state. The core relationships between liquidity and price are defined as:

$$L = \sqrt{xy} \quad (3)$$

$$\sqrt{P} = \sqrt{\frac{y}{x}}, \quad (4)$$

where x and y represent the virtual reserves of `token0` and `token1`, respectively.

To express changes in token quantities as a function of liquidity, we rearrange the above relationships as follows:

$$\Delta y = \Delta \sqrt{P} \cdot L \quad (5)$$

$$\Delta x = \Delta \left(\frac{1}{\sqrt{P}} \right) \cdot L, \quad (6)$$

where Δy and Δx represent the changes in the `token1` and `token0` quantities, respectively. By further rearranging Equation (5) and Equation (6), we can express the changes in price:

$$\Delta\sqrt{P} = \sqrt{P_{\text{new}}} - \sqrt{P_{\text{old}}} = \frac{\Delta y}{L} \quad (7)$$

$$\Delta\left(\frac{1}{\sqrt{P}}\right) = \frac{1}{\sqrt{P_{\text{new}}}} - \frac{1}{\sqrt{P_{\text{old}}}} = \frac{\Delta x}{L}. \quad (8)$$

Here, $\sqrt{P_{\text{old}}}$ and $\sqrt{P_{\text{new}}}$ denote the square root of the pool price before and after the swap, respectively.

3.1.3 Fee accounting

This model assumes a fixed swap fee of 0.3%, denoted by $\gamma = 0.003$. Gas costs are excluded for simplicity, while protocol fees ϕ are also excluded, thus distributing all protocol fees to LPs. In addition, we note:

- x_{in} : The amount of `token0` swapped into the pool, prior to fee deduction.
- y_{in} : The amount of `token1` swapped into the pool, prior to fee deduction.
- $f_{g,0}, f_{g,1}$: Global fee growth variables for `token0` and `token1`.

Fees accumulate according to:

$$\Delta f_{g,0} = x_{\text{in}} \cdot \gamma \quad (9)$$

$$\Delta f_{g,1} = y_{\text{in}} \cdot \gamma \quad (10)$$

3.1.4 Swap mechanism within ticks

As discussed in the previous section, the active liquidity level L remains constant within a tick range. When a trader initiates a swap, the protocol adjusts the square root price $\sqrt{P}$ based on the direction of the trade. After applying the fixed swap fee γ , the change in tokens is given by:

$$\Delta y = y_{\text{in}} \cdot (1 - \gamma) \quad (\text{swap token1 for token0}) \quad (11)$$

$$\Delta x = x_{\text{in}} \cdot (1 - \gamma) \quad (\text{swap token0 for token1}). \quad (12)$$

For swaps exchanging `token1` for `token0`, the trader first provides an initial amount of `token1` y_{in} , which enables the algorithm to compute the corresponding price change $\Delta\sqrt{P}$ via Equation (7), after first computing Δy in Equation (11). The outflow Δx is then computed via Equation (6). Conversely, for swaps exchanging `token0` for `token1`, the model uses the known input x_{in} to compute the quantity change Δx and the price change $\Delta(1/\sqrt{P})$ through Equations (12) and (8). The outflow Δy is then computed in Equation (5).

3.1.5 Swap mechanism across ticks

A key feature of Uniswap v3 is determining when a swap crosses a tick boundary, as this directly affects which liquidity ranges become active during execution. The following flowchart is originally from [Adams et al. \(2021\)](#) and illustrates the swap control logic.




Figure 1: Swap control flow within and across ticks ([Adams et al., 2021](#)).

If the tick that corresponds to the updated price crosses a boundary of a liquidity interval, then the protocol must account for the shift into the new interval by applying the change in liquidity ΔL that is associated with this new interval. Furthermore, the tick value is reset to the boundary value of the new liquidity interval. Any unfilled portion of the swap proceeds to the next trade, executed at the updated tick, with the updated liquidity. This process repeats until the entire input is consumed or slippage constraints are met.

3.1.6 Liquidity provision within a specified price range

Liquidity providers in Uniswap v3 specify a price range $[i_l, i_u]$, denoted in ticks, over which their liquidity is active. The change in liquidity is represented by ΔL , and the required token amounts depend on the current square root price $\sqrt{P}$ and the tick boundaries $\sqrt{p(i_l)}$ and $\sqrt{p(i_u)}$.

The following piecewise functions ensure that LPs supply the correct token proportions, based on whether the current price is below, within, or above their specified range.

The amount of `token1` required:

$$\Delta Y = \begin{cases} 0 & i_c < i_l \\ \Delta L \cdot (\sqrt{P} - \sqrt{p(i_l)}) & i_l \leq i_c < i_u \\ \Delta L \cdot (\sqrt{p(i_u)} - \sqrt{p(i_l)}) & i_c \geq i_u \end{cases} \quad (13)$$

The amount of `token0` required:

$$\Delta X = \begin{cases} \Delta L \cdot \left(\frac{1}{\sqrt{p(i_l)}} - \frac{1}{\sqrt{p(i_u)}} \right) & i_c < i_l \\ \Delta L \cdot \left(\frac{1}{\sqrt{P}} - \frac{1}{\sqrt{p(i_u)}} \right) & i_l \leq i_c < i_u \\ 0 & i_c \geq i_u \end{cases} \quad (14)$$

3.2 Agent-based model

We construct an Agent-Based Model (ABM) in `Julia` in order to implement the Uniswap v3 algorithm and investigate its performance. An ABM is a computational model which simulates the behaviour and interactions of autonomous agents. In the context of our research, the agents are liquidity providers and traders of various behavioural characteristics. The structure of the model and its behavioural rules are crucial in emulating realistic, applicable market dynamics. This computational framework provides the tools to simulate an automated market maker (AMM) such as Uniswap v3.

The ABM is implemented using the `Agents.jl` package in `Julia`, where the model is defined using the `StandardABM` framework. This allows for the agent behaviour described in Section 3.3 to be modelled using swap functions, which are then executed at each time step, for every agent. The swap functions implement the various equations and mechanisms discussed in Section 3.1. The order of action for the agents in the model is set randomly to ensure that the simulation is unbiased, thus following realistic market conditions.

3.3 Ecology of agents

3.3.1 Noise trader

In ABMs, noise traders are typically characterised by incorrect expectations of future prices (DeLong et al., 1990). Although irrational, they may still have a persistent and systematic impact on asset prices — occasionally exceeding that of rational traders under specific conditions. In this study, noise traders are implemented as random agents who either swap `token0` (x) for `token1` (y), or vice versa. Their

trading actions are driven purely by chance, without any informational or strategic foundation.

3.3.2 Trend follower

The trend follower is a speculative agent in our model. Representing momentum-driven traders, this agent makes trading decisions based solely on recent price movements, without reference to fundamental values or external signals (Bottazzi et al., 2005).

Behavioral logic and trading strategy

At each time step, the trend follower evaluates recent price dynamics to guide its trading behavior. In our model, trend followers base their trading decisions on the recent trend observed over the past 20 price changes. Specifically, they count the number of upward versus downward movements, under the assumption that the last 20 price changes provide a reasonable window to detect momentum. This count serves as a simple momentum signal that captures the prevailing market trend.

The initial amount of tokens used in a swap is selected at random, between the values of 0 and 10. The agent applies the following rule-based strategy in order to make a swap decision:

- **Strong upward trend:** If 14 or more of the last 20 price changes are positive, the agent detects upward momentum and swaps `token1` for `token0`.
- **Strong downward trend:** If 14 or more of the last 20 price changes are negative, the agent detects downward momentum and swaps `token0` for `token1`.
- **No clear trend:** If neither condition is met, the agent remains inactive during that trading round.

This rule-based strategy captures the essence of momentum trading and introduces speculative dynamics into the decentralised trading environment of the model.

3.3.3 Fundamentalist

Fundamentalist agents form expectations about future asset prices based on their perception of intrinsic value. In real-world markets, such valuations are typically guided by economic indicators and macroeconomic fundamentals, under the assumption that prices eventually revert to their fundamental or equilibrium levels.

In our agent-based Uniswap v3 model, no external economic data is available. As a simplifying assumption, fundamentalists treat the initial pool price of 1.0, as the

fundamental value. Therefore, we observe that they follow a mean-reversion strategy, responding to deviations between the current market price and the fundamental value.

Behavioural logic and trading strategy

At each time step, the fundamentalist keeps track of the current market price. The initial amount of tokens swapped is uniformly randomised between 0 and 10, and the agent applies the following rule-based strategy before executing a swap:

- **Overpriced condition:** If the market price exceeds 1.05 (i.e. 5% above the fundamental value), then the agent considers the asset overvalued and swaps `token1` for `token0`. This swap is tantamount to a selling an overpriced asset.
- **Underpriced condition:** If the market price falls below 0.95 (i.e., 5% below the fundamental value), then the agent considers the asset undervalued and swaps `token1` for `token0`, effectively buying the underpriced asset.
- **Near-fair-value condition:** If the price lies between 0.95 and 1.05, the agent abstains from trading during that step.

4 Results & Discussion

In Section 4 we present and investigate the results produced by the Uniswap v3 model across a range of market dynamics and ecologies. We begin by visualizing our results in the form of plots, provided in Section 4.1, before considering tables of performance metrics, provided in Section 4.2.

4.1 Visual comparison of market dynamics

We make use of four plots to compare the market conditions present when specific market dynamics are at play. The first plot presents the price at each time step, which also represents the exchange rate between the two tokens (`token0` and `token1`) at the given time. The second plot presents the current active liquidity at each time step, with the liquidity value on the y-axis indicating the amount of active liquidity in the current trading range. The third plot depicts a spread of the volumes of `token0` and `token1` traded at each time step, while the fourth plot illustrates the liquidity level at each tick where an agent action has occurred. (Please refer to Section 3.1 for the relationship between ticks and prices.)

4.1.1 Market dynamics with two liquidity providers

The following figures illustrate market dynamics when keeping our LPs constant, while varying our trader behaviours. Throughout this section, we have one LP supplying 1000 units of liquidity across the full range of ticks (-10000 to 10000), and another LP which provides 100 units of liquidity in a more concentrated range, between ticks -1000 and 1000. Both LPs provide liquidity at the beginning of the simulation.

Figure 2 presents the simplest case, in which one active noise trader acts on the Uniswap v3 market. In the top two subplots, (subplots numbered 1 through 4 henceforth, moving downwards), we observe how the price fluctuates over time, and how the active liquidity jumps between two different liquidity regions as the trader executes their trades. The fact that the price series follows no trend comes as no surprise, since noise traders act randomly. By comparing plots 1 and 3, we notice that the featured dynamics are rough inverses of each other; this is simply due to how their values are calculated. A decrease in plot 3 indicates an exchange of `token1` for `token0`, which according to Equation (3), leads to an increase in price. The effect of concentrated liquidity is evident in the final plot, which shows that liquidity provided within the $[-1000, 1000]$ tick range increases the available liquidity for traders operating within this tick interval.




Figure 2: Market behaviour with the configuration of one noise trader.

To analyse the impact of a trend-following trader on market conditions, we consider two scenarios. In the first experiment, the trend follower's initial trade quantity ranges from 1 to 2 units, whereas in the second scenario, it ranges from 1 to 10 units. This approach allows us to compare how market dynamics differ when the trend follower exerts a relatively small influence on the market, versus when it has a much larger impact.

We start by looking at the former case. In Figure 3, we observe in the first plot that the trend follower does not have a large impact on price, due to their relatively small trades. However, we notice three instances where there is a prolonged trend in the price, such as the upward trends between time points 175 and 375 and 800 and 1000. Plots 2 and 4 illustrate how the ticks jump between the different liquidity regions due to the traders' activities. For example, when comparing plot 2 from Figure 2 to its counterpart in Figure 3, we remark that the tick jumps between different liquidity regions are more frequent. We hypothesize that this could be due to the fact that the trend follower exacerbates any existing trends, thereby increasing (or decreasing) the price further than in a market with just a single noise trader.

In Figure 4, we consider the second trend follower scenario, in which their initial quantity of tokens traded is chosen randomly between 1 and 10 units. We observe




Figure 3: Market behaviour with the configuration of one noise trader and one trend follower. The trend follower trades either one or two tokens.

a substantial difference in market behaviour, with the price ranging from below 0.75 up to approximately 1.60. These price fluctuations are larger than any of the cases we have encountered thus far. We also observe smoother, more cyclical behaviour in plots 1 and 3 compared to previous experiments. These dynamics occur due to the effect of trend followers prolonging the presence of an existing trend in the market. Since trend followers are now trading larger quantities of tokens, the prolonging of trends is stronger, which thus results in smoother dynamics.

Looking at the first plot in Figure 4, we observe how the trend follower quickly identifies a market trend and then acts accordingly, thereby causing a shift in market dynamics. The moment that the trend reverses in direction, the trend follower identifies the new trend and causes the market to adjust accordingly. Plot 2 further illustrates the size of price movements within the market due to the presence of a trend follower. Due to the large fluctuations in price, the ticks jump between different liquidity regions with greater regularity. It is evident that when we only have a trend follower and a noise trader, the trend follower disrupts the market completely.




Figure 4: Market behaviour with the configuration of one noise trader and one trend follower. The trend follower trades a random amount of tokens in the $[1, 10]$ interval.

This is where the necessity for a fundamentalist trader arises. Looking at the first plot in Figure 5, we see that the price does not fluctuate drastically from 1.00 as it does in Figure 2, when we only consider one noise trader. A fundamentalist believes that an asset is either over- or under-priced when its price diverges from an equilibrium level or expected range, and subsequently trades in the opposite direction. The decision to act in the opposite direction is a counterbalance in the model, constraining the price within the fundamentalist's preferred bounds. In Figure 5, the upper and lower bounds for the fundamentalist trader are 1.05 and 0.95, respectively. The resulting price dynamics can be observed in the uppermost plot, with the price oscillating within the bounds of 1.05 and 0.95, and reversing course whenever the fundamentalist trades against any developing price bubbles.

When considering only one noise trader and one fundamentalist, the resulting market is small and simplistic, with the fundamentalist's beliefs (or bounds, equivalently) having a notable impact on market conditions. Plots 2 and 4 further establish the fact that the price does not fluctuate as drastically, compared to the initial experiment with just one noise trader. This is due to the traders' activities never causing the ticks to move between different liquidity regions (the tick always re-




Figure 5: Market behaviour with the configuration of one noise trader and one fundamentalist.

mains between -1000 and 1000). Lastly, Plot 3 reiterates the inverse relationship that exists between the price and the volume spread.

In Figure 6 we observe the market conditions when we have all three trader ecologies acting on the market at every time step. Notably, the presence of a fundamentalist counteracts the effects of the trend follower, even when the trend follower is allowed to trade the same quantities as noise traders (1 to 10 token units). From Figure 6, we may study the effects of the trend follower and fundamentalist when the latter widens their perceived price boundary from $[0.95, 1.05]$ to $[0.9, 1.1]$. Firstly, there are increased upward and downward trends in the price dynamics due to the trend follower intensifying existing market trends, especially around time steps 250 and 400. Secondly, the fundamentalist bounds the price between 0.9 and 1.1, acting against the market when the price moves outside of these bounds.

In addition, we note that the impact of the trend follower is weakened due to the behaviour of the fundamentalist, who bursts any bubbles created by the trend follower. While the trend follower effectively drives the price away from equilibrium, the fundamentalist pulls it back. As a result, price fluctuations are smaller than




Figure 6: Market behaviour with the configuration of one noise trader, one trend follower, and one fundamentalist.

when only trend followers are present in the model. The stabilizing effects of the fundamentalist highlight the necessity of a diverse trader ecology in the Uniswap v3 framework.

Figure 7 presents the market behaviour after increasing the number of agents in our ABM simulation to mimic a more realistic market under which Uniswap v3 would operate. We observe similar market conditions when compared to previous figures. Plot 2 seems a bit erratic, but is still sensible, as a market with 100 noise traders, 50 trend followers, and 50 fundamentalists inevitably results in frequent jumps between liquidity regions. Moreover, we observe the concentration of liquidity across the tick range of $[-1000, 1000]$ in plot 4. The ability to swap within and across liquidity intervals confirms the key characteristic of concentrated liquidity in Uniswap v3, and how it alters the overall landscape of the market.

4.1.2 Market dynamics with three liquidity providers

In Section 4.1.2, we consider a scenario in which the dynamics of the market are very similar to that of Figure 2 but with an increased number of agents that are supposed to represent a more realistic Uniswap v3 market. We now introduce a




Figure 7: Market behaviour with 100 noise traders, 50 trend followers, and 50 fundamentalists, all with initial traded amounts drawn randomly from a $\mathcal{U}[1, 10]$ distribution. Note that the fundamentalists have bounds of 0.9 and 1.1).

scenario with three LPs, each with different concentrated liquidity amounts, and 100 noise traders. The first LP supplies a liquidity of 1000 units across the entire tick range (from -10000 to 10000), the second LP introduces 100 units of liquidity between ticks -1000 and 1000, while the third supplies 50 units of liquidity within the narrowest range of -500 to 500.

Analysing the first plot in Figure 8, we see that the price moves erratically due to the large amount of noise traders in the model. Comparing plots 1 and 3, we note that the inverse relationship between price and the volume spread cannot be observed as clearly as in the earlier scenarios. This is due to the increased number of traders and subsequent increased price volatility. Plots 2 and 4 are where we see the difference between having 2 LPs and 3 LPs.

While it is difficult to see the jumps between different liquidity regions in Figure 8, we can conclude that the ticks jump within seven different liquidity levels, instead of the aforementioned three regions. The creation of new, incorrect liquidity levels occurs due to the logic of the `swap_x_for_y` function in our code. Trading y for x is easier to program, as tick crossings occur going rightward. When initializing the




Figure 8: Market behaviour with the configuration of 100 noise traders and three LPs.

positions of LPs, the leftmost tick of their specified interval serves as the liquidity addition point; liquidity is then removed at the rightmost tick. The logic cannot be reversed cleanly in the x for y trading case, as the program requires a specific set of conditions that account for when and where the tick lands.

There is a very specific issue that occurs when the tick “trickles” around the end-point of a liquidity interval. Secondary trades with residual, untraded amounts of tokens can occur near an active tick before actually landing on that tick, which thus causes the associated liquidity regions to be activated multiple times in a single swap. Consequently, liquidity is added to or removed from the pool *twice* in one interval. While the case for two LPs has been handled for the most part (see Figure 7 in which there are two LPs with two distinct levels of trading), the addition of further LPs causes errors, as the code has not been updated to account for the increased number of “trickling” cases. The extension of the program to account for three or more LPs is left as an avenue for future research.

4.1.3 Market dynamics with an active liquidity provider

In this scenario, the simulation starts with two LPs: one supplying 1000 units of liquidity across the full tick range (from -10000 to 10000), and the other supplying 100 units between -1000 and 1000. At time step 250, a third LP adds 100 units of

liquidity across $[-2000, 2000]$, which is later removed at time step 750.

Figure 9 illustrates the market dynamics when a third LP mints and burns liquidity during the simulation. All traders are included, and from Figure 9 we observe the typical behaviour of price and volume spread, but with the addition and subsequent removal of concentrated liquidity.




Figure 9: Market behaviour with the configuration of one noise trader, one trend follower, one fundamentalist, two static LPs and one active LP.

In the second plot, we observe that the active liquidity increases from 1100 to 1200 at time step 250. This upward shift reflects the addition of 100 units of concentrated liquidity by the third LP, which becomes active immediately. The immediate activation of the third liquidity range occurs because the current tick lies within the concentrated range of $[-1000, 1000]$, in which all three LPs are active. This behaviour is further validated by the fourth plot, which highlights the feature of concentrated liquidity. Notably, in the fourth plot, traders make swaps at liquidity levels below 1100 outside the tick range of $[-1000, 1000]$, due to the fact that the concentrated liquidity between -2000 and 2000 is not active across the entire time frame.

4.2 Comparison of performance metrics

Section 4.2 presents the performance metrics of the Uniswap v3 model across different configurations, each of which correspond to different figures in Section 4.1.

Metric	Figure 2	Figure 3	Figure 4	Figure 5
Volatility of Price	0.025148	0.03757	0.11627	0.01262
Average of Price	0.99079	1.00324	0.98243	1.00368
Skewness of Price	-0.14169	0.45213	0.15076	-0.35437
Volatility of Volume Traded	51.83629	50.52490	165.97832	20.11628
Average of Volume Traded	-63.91345	-22.03097	-109.01045	-6.77913
Skewness of Volume Traded	0.30984	-0.78189	-0.43030	0.044875

Table 1: Metrics of Uniswap v3 market across different scenarios, each with two liquidity providers.

Table 1 illustrates the impact of the different trader ecologies on the model when considering two liquidity providers. At the beginning of the model simulation, the first LP supplies liquidity over the $[-10000, 10000]$ interval, while the second provides liquidity across the $[-1000, 1000]$ interval.

The average price remains similar across the four different configurations. Notably, the simulation displayed in Figure 4 (featuring one noise trader and one trend follower, with a larger amount traded) exhibits both the greatest deviation in average price and the highest price volatility. This outcome is expected, as the trend follower in this setup is able to trade quantities comparable to the noise trader, thereby exerting a significant influence on market conditions.

Similarly, the volatility of the volume traded shows comparable patterns. The skewness of the price and the traded volume indicates that the price distribution is relatively symmetric, oscillating around 1. This behaviour stems from the fact that the initial tick is set to 0 and serves as the reference point, as shown in each of the figures. In contrast, the skewness of the volume traded suggests that the delta volume spread is less symmetric than the price distribution, reflecting the random nature of the trade sizes executed by the agents.

By broadening the ecology of the types of traders in the ABM, we obtain the results presented in Table 2. The volatility and mean of the price indicate that the market behaves as expected, exhibiting low volatility and maintaining an average price close to 1. Additionally, the skewness of the price confirms the relative symmetry of price fluctuations around this central value.

Metric	Figure 6	Figure 7
Volatility of Price	0.01887	0.02733
Average of Price	1.00019	1.00037
Skewness of Price	0.14473	0.09971
Volatility of Volume Traded	27.06547	397.17210
Average of Volume Traded	-19.73281	-702.49706
Skewness of Volume Traded	-0.06269	-0.74988

Table 2: Metrics of Uniswap v3 market across varying scenarios with two liquidity providers.

The metrics for the delta volume spread reveal noticeable skewness, particularly evident in Figure 7. Here, the average volume traded is negative, indicating that the market traded more of `token0` than `token1` (due to random selection of amount of tokens traded). This is further supported by the substantial volatility of 397.1720 and an average volume traded of -702.49706, highlighting both the variability and the directional bias in trading activity.

Metric	Figure 8	Figure 9
Volatility of Price	0.02869	0.03642
Average of Price	1.00114	1.00712
Skewness of Price	0.10919	-0.35404
Volatility of Volume Traded	413.21110	48.32451
Average of Volume Traded	229.82505	11.46620
Skewness of Volume Traded	-0.74304	0.16399

Table 3: Metrics of Uniswap v3 market across varying scenarios with three liquidity providers.

Increasing the number of LPs with varying concentrated liquidity levels yields the results shown in Table 3. Overall, the price behaves as expected, exhibiting low volatility and a mean value close to 1, while remaining approximately symmetric. However, in Figure 9, there is a noticeable increase in price skewness. The negative skewness likely results from the addition and removal of liquidity during the simulation, which introduces asymmetry in price movements.

Examining the volume spread metrics in Figure 8, we observe that an increase in the number of noise traders and LPs affects the quantity traded. Specifically, this configuration leads to higher volatility in traded volumes and a negative skewness, indicating that, on average, `token1` was exchanged for `token0` with greater magnitude and regularity. This outcome reflects the combined effects of increased trading activity and the uneven distribution of trade sizes.

5 Tree-based sensitivity analysis

Classification and regression trees were first introduced in Breiman et al. (1984). In Hastie et al. (2009), the authors provide a detailed summary of the foundational theory of decision trees, much of which originates from Breiman et al. (1984). We will provide a summary of the section on regression trees, random forests, and gradient boosted models in Hastie et al. (2009).

First, let us review how regression trees and ensemble methods can be used to conduct a preliminary sensitivity analysis. Sobol indices and other variance-based methods (as proposed in Saltelli et al. (2008)) are generally considered to be the most comprehensive forms of sensitivity analysis. They are often computationally intensive, however, and as a result, machine learning techniques have been proposed as an alternative form of global sensitivity analysis (Harper et al., 2011; Jaxa-Rozen and Kwakkel, 2018). These tree-based methods require the creation of a data set that represents the “reasonable” parameter space of the model. We construct this space as such:

1. For each model, all parameters are varied within “reasonable” intervals. As we do not know the underlying distributions of many key model parameters, we make the assumption that all input values are drawn from different uniform distributions.
2. The model is then simulated for each combination of parameters.
3. Several key performance metrics are computed at the end of each simulation, and will serve as target variables for the regression machine learning techniques. For the sake of brevity, we only consider price volatility and mean price as target variables in this report.

Note that in our sensitivity analysis, we treat model parameters as predictors. The goal is not to predict final performance metric values from a given parameter combination, but rather to use decision trees to visualize cuts in the parameter space, and ensemble methods to understand variable importance in the model. Predictive ability matters, however, as poor predictive power will give rise to erroneous visualizations (and a poor understanding) of the parameter space.

5.1 Theory of regression trees

For all theoretical sections henceforth, the notation and conventions will follow those from the original source (Hastie et al., 2009).

As a general idea of how the regression tree algorithm works at its most basic level, let us assume that we have a set of p continuous predictors $\{x_1, \dots, x_p\}$, as well as

a continuous target variable Y . We begin by initially splitting the space into two regions, and designating the response as the mean of Y in each of the two regions. We then choose the feature and associated split-value that produce the optimal fit. The two regions are then split using the technique explained above. This process continues recursively until a stopping condition is satisfied in each of the regions (Hastie et al., 2009).

Now, consider a data set that contains N observations, with one target variable $Y = \{y_1, \dots, y_N\}$ and p continuous features/predictors $\{x_{i1}, \dots, x_{ip}\}$, for observations $i = \{1, \dots, N\}$. As explained in (Hastie et al., 2009), the algorithm partitions the feature space into M regions, denoted by $R_1, \dots, R_M$. If we employ a mean squared error loss function, then the response will be the mean of the target variable y_i in each region, for each observation.

In general, let us model the response $f(x)$ in some region by a constant c_m , for each of the regions $m = \{1, \dots, M\}$. Formally, we obtain

$$f(x) = \sum_{m=1}^M c_m I(x \in R_m), \quad (15)$$

where $I(x \in R_m)$ is an indicator function, dependent on the location of x within any of the M regions. When minimizing a least-squares loss function, it follows that the optimal response value $\hat{c}_m$ will simply be the mean of the target variable values y_i in region R_m , which is to say

$$\hat{c}_m = \text{mean}(y_i \mid x_i \in R_m).$$

To reduce computational complexity, the algorithm uses a greedy, top-down approach to determine the optimal splitting variable and split-point. From (Hastie et al., 2009), consider some splitting variable j (i.e. one of the p predictors in our data set) and a split-point s . Two regions are then defined: the first region will contain *all* values of all predictors, such that each predictor value X_j is less than the split value s , while the second region will contain all values X_j greater than or equal to s . In other words,

$$R_1(j, s) = \{X \mid X_j < s\} \quad \text{and} \quad R_2(j, s) = \{X \mid X_j \geq s\}. \quad (16)$$

The splitting variable j and the split-point s that minimize the least-squares loss in each of the two formulated regions are then chosen:

$$\min_{j,s} \left[\min_{c_1} \sum_{x_i \in R_1(j,s)} (y_i - c_1)^2 + \min_{c_2} \sum_{x_i \in R_2(j,s)} (y_i - c_2)^2 \right],$$

with both of the interior minimizations solved by

$$\hat{c}_1 = \text{mean}(y_i | x_i \in R_1(j, s)) \quad \text{and} \quad \hat{c}_2 = \text{mean}(y_i | x_i \in R_2(j, s)),$$

for any combination of j and s , provided that least-squares loss is to be minimized. This process is then repeated on the two split regions, as well as all other regions, until the desired size of the tree has been attained.

The size of the tree is initially dependent on the minimum number of observations permitted to exist in a given region (or node). The minimum depth of a region is set as a *hyperparameter* in the algorithm; we will tune it to ensure that the structure of the tree can be understood without overfitting the data. The preferred strategy is to grow a large tree with a relatively lower minimum node size, before employing *cost-complexity pruning* [Hastie et al. \(2009\)](#). Given the larger, initial tree T_0 , a generic subtree T , and the number of terminal nodes $|T|$, we may define the cost complexity criterion:

$$C_\alpha(T) = \sum_{m=1}^{|T|} N_m Q_m(T) + \alpha |T|. \quad (17)$$

Here, m denotes a specific terminal node of the subtree T and N_m represents the number of observations in the region R_m (i.e. the number of $\{x_i \in R_m\}$, while

$$Q_m(T) = \frac{1}{N_m} \sum_{x_i \in R_m} (y_i - \hat{c}_m)^2,$$

$$\text{where } \hat{c}_m = \frac{1}{N_m} \sum_{x_i \in R_m} y_i$$

The goal is to find the value of α that minimizes $C_\alpha(T)$ and thus prunes the tree to an ideal size. The hyperparameter α is estimated using five- or tenfold cross-validation; we will refrain from explaining those processes in this report, as we simply choose a sensible α in order to visualize the parameter space.

5.2 Random forests and gradient boosting

The hierarchical nature of tree-building means that small changes in the data can result in wildly different splits. Consequently, an individual regression tree is unstable with high variance [\(Hastie et al., 2009\)](#). Ensemble methods such as random forests and gradient boosting (featured in the next subsection) can reduce this instability and provide us with greater insight as to the potential drivers of sensitivity and instability in our model. While decision trees will be used to visualize cuts in the parameter space, random forests and gradient boosting will be employed to understand variable importance in each model.

5.2.1 Theory of random forests

Random forests are not computationally intensive to build and train. As a result, they are a suitable alternative to Sobol indices for global sensitivity analysis. The method was first introduced in Breiman (2001) as a modification of bagging, in which a large collection of *de-correlated* trees is built and then averaged. The algorithm for random forests for regression trees (as presented in Hastie et al. (2009)) is displayed below:

1. For $b = 1$ to B trees:
 - (a) Given a training set, draw a bootstrap sample of the training data.
 - (b) Build an individual tree T_b using the bootstrapped predictors and target variable, by recursively repeating the following steps:
 - (i) Randomly select m variables of the set of p predictors. For regression problems, m is often chosen to be $p/3$, whereas for classification problems, m is chosen to be an approximation of the $\sqrt{p}$ (Hastie et al. 2009). Note that m is a hyperparameter that will be tuned.
 - (ii) Of the m chosen predictors, pick the variable c that minimizes the loss function over the two split nodes.
 - (iii) Use the optimal splitting variable c and its corresponding split value to split the node into two daughter nodes.
 - (iv) Repeat steps (i) to (iii) until the minimum node size n_{min} has been reached (Hastie et al. 2009). The minimum node size is a hyperparameter that will be tuned.
2. Once Step 1 has been completed for all B trees, we can make a prediction for some observation point x . In the case of a regression random forest, this prediction is given as the average predicted value across all trees B :

$$\hat{f}_{rf}^B(x) = \frac{1}{B} \sum_{b=1}^B T_b(x).$$

When considering the size of random forests, we may first note that since each tree is identically distributed, the bias of a forest of trees will be equal to the bias of each individual tree. Considering variance, however, we may find that if we have B identically distributed random variables (trees, in this case), each with variance σ^2 and pairwise correlation ρ , then the variance of the average of those variables is given by

$$\rho\sigma^2 + \frac{1-\rho}{B}\sigma^2.$$

Since our random variables are trees, we may thus reduce the variance of the random forest by not only increasing the number of trees B , but also by decreasing

the pairwise correlation ρ through random predictor variable selection during the tree-growing process [Hastie et al. \(2009\)](#).

5.2.2 Theory of gradient boosting

Consider a tree of the form

$$T(x; \Theta) = \sum_{j=1}^J \gamma_j I(x \in R_j), \quad (18)$$

where R_j represents the disjoint, partitioned regions of the parameter space, $j = 1, 2, \dots, J$ are the terminal nodes of the tree, and γ_j is the predicted value of the response function in some node j . We can condense the regions R_j and predicted response values γ_j into a parameter set $\Theta = \{R_j, \gamma_j\}_1^J$. The values of Θ are found by minimizing a loss function (mean squared error, for example) over each of the J nodes, for all of the x_i observations in the corresponding R_j region ([Hastie et al. \(2009\)](#)).

According to [Hastie et al. \(2009\)](#), a boosted tree model is expressed as a sum of M trees

$$f_M(x) = \sum_{m=1}^M T(x; \Theta_m), \quad (19)$$

where the risk $\hat{\Theta}$ must be minimized for the regions and response values $\Theta = \{R_{jm}, \gamma_{jm}\}_1^{J_m}$ of the *next* tree, given the current model $f_{m-1}(x)$:

$$\hat{\Theta}_m = \arg \min_{\Theta_m} \sum_{i=1}^N L(y_i, f_{m-1}(x_i) + T(x_i; \Theta_m)). \quad (20)$$

If the regions are known, then the response values for some node j in a given tree m are given by

$$\hat{\gamma}_{jm} = \arg \min_{\gamma_{jm}} \sum_{x_i \in R_{jm}} L(y_i, f_{m-1}(x_i) + \gamma_{jm}). \quad (21)$$

[Hastie et al. \(2009\)](#) state that for regression trees with a mean squared error loss function, the solution to Equation (20) is the regression tree that is best at predicting the residuals for the current tree (the difference between actual and predicted values; i.e. $y_i - f_{m-1}(x_i)$ for some observation i), while the predicted response value $\hat{\gamma}_{jm}$ is the mean of the residuals in the region j .

The gradient boosting algorithm for regression is presented in [Hastie et al. \(2009\)](#), and can be summarized as such:

1. Initialize the model with some constant value, which corresponds to the response of a single-node tree. This initial model is given by

$$f_0(x) = \arg \min_{\gamma} \sum_{i=1}^N L(y_i, \gamma).$$

2. Then, for trees $m = 1$ to M :

- (a) For observations $i = 1, 2, \dots, N$ of the current model, compute the gradient of the loss function:

$$r_{im} = - \left[\frac{\partial L(y_i, f(x_i))}{\partial f(x_i)} \right]_{f=f_{m-1}}.$$

- (b) Treating the gradient values as targets, fit a regression tree to the targets r_{im} . Terminal regions R_{jm} follow from the tree, for nodes $j = 1, 2, \dots, J_m$.
- (c) For each of the j nodes, compute the optimal response

$$\gamma_{jm} = \arg \min_{\gamma} \sum_{x_i \in R_{jm}} L(y_i, f_{m-1}(x_i) + \gamma).$$

- (d) Update the model according to

$$f_m(x) = f_{m-1}(x) + \sum_{j=1}^{J_m} \gamma_{jm} I(x \in R_{jm}),$$

and proceed to the next tree, continuing the loop at step (a) again.

3. Once all trees have been estimated, the output $\hat{f}(x)$ of the gradient boosting algorithm will be equal to the final tree $f_M(x)$.

5.2.3 Variable importance

Variable importance is derived in the exact same way for both random forests and gradient-boosted models (Hastie et al., 2009). For each split in each tree, the “importance” of a predictor variable is calculated as the improvement/difference in the split-criterion, from the parent node to the daughter node. These difference values are then summed for all nodes in all trees to obtain a ranking of the most important variables. In a more technical sense, the variable importance ranking lists the predictors in order of contribution to improved prediction accuracy.

The improved accuracy of decision trees and ensemble methods is essential in our research. Variance importance plots provide us with a more realistic understanding of the parameter space, as well as the predictors that are most responsible for final price volatility and mean price values.

5.3 Tree-based sensitivity results

Tree-based sensitivity analysis methods assume that the underlying model is unknown. The primary goal of these techniques is to identify the predictors that are most responsible for driving the model to the observed price volatility and mean price values. Consequently, we must construct a parameter space that is representative of the wide range of possible price movements in the model.

We begin our sensitivity analysis by constructing a representative data set, with the model inputs drawn from the distributions specified in Table 4. Future research is required to understand the actual underlying distributions of some of these inputs; as a result, we make the simplifying assumption that all parameters are uniformly distributed within sensible intervals. Once all parameters have been sampled and arranged into a dataframe of inputs, we run our agent-based model for each combination of predictors (corresponding to each row in the dataframe), assuming an ecology where all types of traders are included, with two LPs. In total, we consider 1000 different simulations, each for 1000 time steps.

Code name	Parameter description	Sampling distribution
num_NT	Number of noise traders	$\mathcal{U}(100, 200)$ (integer)
num_TF	Number of trend followers	$\mathcal{U}(50, 100)$ (integer)
num_fund	Number of fundamentalists	$\mathcal{U}(50, 100)$ (integer)
pa	Lower tick for second liquidity provider interval	$\mathcal{U}(-2000, 2000)$ (integer)
pb	Upper tick for second liquidity provider interval	$pa + \mathcal{U}(200, 1000)$ (integer)
k	Second LP's delta liquidity	$\mathcal{U}(50.0, 200.0)$ (float)
y_in	Amount of Y traded for noise traders	$\mathcal{U}(50.0, 100.0)$ (float)
x_in	Amount of X traded for noise traders	$\mathcal{U}(50.0, 100.0)$ (float)
TF_in	Amount traded for trend followers	$\mathcal{U}(25.0, 50.0)$ (float)
fundamount_in	Amount traded for fundamentalists	$\mathcal{U}(25.0, 50.0)$ (float)
fundupper_bnd	Upper fundamentalist bound	$\mathcal{U}(1.05, 1.25)$ (float)
fundlower_bnd	Lower fundamentalist bound	$\mathcal{U}(0.75, 0.95)$ (float)
L_initial	First LP's delta liquidity	$\mathcal{U}(10000.0, 20000.0)$ (float)

Table 4: Predictors in the agent-based Uniswap v3 model. Note that the initial LP has a constant interval of $[-10000, 10000]$.

Our global sensitivity analysis enables us to study sub-ecologies of traders and LPs. For example, fundamentalist traders with boundaries that are further from 1 can be seen as more apprehensive to bet against developing bubbles, whereas fundamentalist traders with boundaries that are closer to 1 can be interpreted as being more aggressive in their actions. As for liquidity providers, LPs that set their position with wider ranges can be seen as relatively more risk-averse in their approaches, opting to collect consistent but smaller fee totals. On the other hand, LPs with relatively narrower positions are less conservative in their approaches, with higher potential fees collected, but a much greater risk of illiquidity.

We assume that the initial LP has a constant interval of $[-10000, 10000]$, and that the initial tick is always set at 0 (thus corresponding to an initial price of 1). Fur-




Figure 10: Regression tree using the `rpart` (Therneau and Atkinson, 2023) and `rpart.plot` (Milborrow, 2024) packages. The price volatility is selected as the target variable, with the initial cut made by the first LP’s delta liquidity L_{initial} . The upper value in each rectangle denotes the mean predicted value of the price volatility in that specific node, while the bottom value in each rectangle represents the number of observations in the node.

thermore, we include two dependent variables as columns in the input data frame: the width of the second LP, computed as $pb - pa$, and the midpoint of the interval, defined as $(pb - pa)/2$.

Our analysis begins by considering basic regression trees, which assist us in visualizing cuts in the parameter space. In Figure 10, we denote price volatility (the standard deviation of price) as the target variable, and observe that the initial cut in the tree is made by L_{initial} , the delta liquidity value of the initial LP. The tree provides intuition as to the most sensitive input in the model, as not only does L_{initial} form the first cut, but it also creates cuts further down the tree, as well.

In general, we observe that greater L_{initial} values produce lower price volatility values, as do smaller values of y_{in} . This confirms our prior understanding of the model, as from Equation (7), we observe that greater liquidity values produce relatively smaller price shifts. Similarly, we recall from Equation (11) that lower values of y_{in} produce smaller shifts in Δy , which thereby produce smaller decreases in price movement, as per Equation (7).

Turning our attention to Figure 11, we observe that the cuts in the tree are formed




Figure 11: Regression tree when the mean price level is selected as the target variable. The cuts in the tree are formed by the initial amounts of `token1` and `token0` used for trading, denoted by y_{in} and x_{in} , respectively.

by the initial amounts of the two tokens. Perhaps unsurprisingly, greater amounts of y_{in} and lower amounts of x_{in} correspond to higher average price levels, as `token1` becomes more valuable and is thus traded more often than `token0`. The opposite holds as well, as lower amounts of y_{in} and higher amounts of x_{in} produce lower average price levels, with `token0` traded more often than `token1`.

As mentioned earlier, an individual regression tree is unstable with high variance, and is more likely to overfit data than ensemble methods such as random forests and gradient boosted models (Hastie et al., 2009). In order to perform a more comprehensive tree-based sensitivity analysis, we use the random forest algorithm from the `randomForest` package (Liaw and Wiener, 2002) and the gradient boosting algorithm from the `xgboost` package (Chen et al., 2024) in R. Our tuning algorithm considers a 70%-30% training-testing split and varies hyperparameters according to the sets of values in Table 5.

In Figure 12, we confirm the validity of the results of the corresponding regression tree from Figure 10. The delta liquidity of the first LP is the most important feature in the model, while the initial amount of `token1` used for trading and the number of noise traders are second and third in the sensitivity rankings, respectively. While these results hold for both the random forest and gradient boosting methods, we observe that the difference between the three most important predictors is far more pronounced in the case of the former.

Name	Description	Values considered for tuning
mtry	Number of randomly sampled predictors considered for splitting (RF)	{4, 5, 6, 7}
ntree	Number of trees (RF)	{500, 1000, 2000, 3000}
nodesize	Minimum number of observations in terminal nodes (RF)	{3, 10, 30}
nrounds	Number of trees (GB)	{100, 500, 1000}
eta	Shrinkage/learning rate (GB)	{0.05, 0.15, 0.3}
max_depth	Maximum depth of a single tree (GB)	{3, 6, 9}
min_child_weight	Minimum number of observations required for each node (GB)	{3, 10, 30}

Table 5: Hyperparameters for the random forest (RF) and gradient boosting (GB) methods. Tree-based analysis is done in R.




Figure 12: Feature importance plot for parameters of the Uniswap v3 ABM, when using the random forest (RF) and gradient boosting (GB) methods. We designate the price volatility as the target variable for both techniques. The optimal RF hyperparameters are $mtry = 7$, $ntree = 3000$, $nodesize = 3$ and the optimal GB hyperparameters are $nrounds = 500$, $eta = 0.05$, $max_depth = 3$, $min_child_weight = 30$. The test set MSE values for the optimal RF and GB models are 3.7912×10^{-6} and 1.9842×10^{-6} , respectively. In both cases, we observe that the delta liquidity of the first LP is the most important predictor.

Similarly, the results presented in Figure 13 confirm the results of the regression tree presented in Figure 11. The average price level is almost solely determined by the initial amounts of tokens used for trading. Furthermore, we can conclude from our tree-based sensitivity analysis that the features in our model that contribute most to pricing dynamics (in terms of both volatility and average level) are the amounts of tokens used for trading (y_{in} and x_{in}), the amount of liquidity provided by the first LP ($L_{initial}$), and the number of noise traders in the model (num_NT).

We then consider price volatility and mean price level as target variables, and analyze parameter sensitivity versus parameter uncertainty. Each feature's con-




Figure 13: Feature importance plot for the predictors in the model when using the random forest (RF) and gradient boosting (GB) methods. We designate the mean price as the target variable for both techniques. The optimal RF hyperparameters are $mtry = 7$, $ntree = 2000$, $nodesize = 3$ and the optimal GB hyperparameters are $nrounds = 500$, $eta = 0.05$, $max_depth = 6$, $min_child_weight = 3$. The test set MSE values for the optimal RF and GB models are 3.0736×10^{-4} and 1.5966×10^{-4} , respectively. In both cases, we observe that the most sensitive feature is the initial amount of `token1`.

tribution to information gain in the gradient boosted model serves as a proxy for parameter sensitivity, while the features' normalized variance values are used as a proxy for uncertainty. In Figure 14, we observe that y_{in} and `num_NT` are quite sensitive in determining the price volatility, but less uncertain when compared to `L_initial`, which exhibits both high uncertainty *and* sensitivity. This result demands further investigation into the distribution of initial LP behaviour. When considering the mean price level as a target variable, we observe that there are no predictors that exhibit both high uncertainty and high sensitivity.

Lastly, we isolate the most important model parameters (as observed from our tree-based sensitivity analysis) and visualize their effects on the target variables. In Figure 15 (a), we confirm that the price volatility increases when there is less liquidity provided by the initial LP, and when traders put forth a greater amount of `token1` for swaps. In Figure 15 (b), we observe once again that greater amounts of y_{in} result in higher mean price levels, with the opposite holding true for x_{in} amounts.





Figure 14: Sensitivity versus uncertainty plots, for price volatility and mean price. Sensitivity is measured by each predictor's contribution to information gain in the gradient boosted model, while uncertainty is measured by normalized feature variance.



(a) Sensitivity of $L_initial$, y_{in} , and num_NT for price volatility. (b) Sensitivity of y_{in} and x_{in} for mean price.

Figure 15: Parameter sensitivity heat map for the most important model parameters, when considering the price volatility and mean price as target variables.

6 Conclusion

The popularity and efficiency of decentralised finance have driven the development of technologies such as decentralised exchanges, which rely on AMMs like Uniswap v3. This makes Uniswap v3 a powerful and influential platform in the evolving landscape of decentralised currency and finance.

The goal of this paper is to analyse how liquidity providers and different trader ecologies interact and affect each other within such an exchange. This is accomplished by implementing an agent-based model in `Julia`, which thus enables the exploration of these dynamics in a controlled, simulated environment. The study successfully achieves its objectives, providing insights into the mechanisms of Uniswap v3 and highlighting the unique role of concentrated liquidity. The results demonstrate that Uniswap v3 effectively facilitates cryptocurrency token swaps, with trader behaviors significantly shaping market conditions, liquidity, and price dynamics.

In particular, trend followers are found to have the most pronounced impact on price movements, while the presence of other agents, such as fundamentalists, are able to counteract and stabilise such effects. From our sensitivity analysis, we demonstrate that the initial LP has a considerable impact on price volatility through the amount of liquidity that it supplies to the market, and that the number of noise traders and initial amounts of tokens traded are also sensitive to changes. These four parameters (`L_initial`, `num_NT`, `yin`, and `xin`) are most responsible for the richness in price dynamics in the model, and the further study of their typical values and distributions is essential for future research.

The model exhibits some limitations, particularly when introducing multiple liquidity providers alongside a large number of traders. In these cases, swapping `token0` for `token1` occasionally results in liquidity being added or removed incorrectly, due to numerical issues arising near the endpoints of active liquidity ranges. Resolving this issue is a logical next step for our research, as it enables the investigation of a model that features more than two liquidity providers, and a variety of liquidity provider ecologies, as well. Furthermore, a more robust analysis could include a third currency, which would thus introduce arbitrageurs into the model framework.

Overall, this work illustrates that Uniswap v3 offers a robust and adaptable marketplace for decentralized finance, capable of accommodating diverse trader strategies and liquidity configurations. As the model is extended and refined, we will gain further insight into the diversity of outcomes, risks, and instabilities that are intrinsic to the framework.

Bibliography

- Adams, H., Zinsmeister, N., Keefer, R., Robinson, D., Salem, M., 2021. Uniswap v3 core. <https://uniswap.org/whitepaper-v3.pdf>. Whitepaper.
- Adams, H., Zinsmeister, N., Robinson, D., 2020. Uniswap v2 core. <https://uniswap.org/whitepaper-v2.pdf>. Uniswap Labs Technical Whitepaper.
- Aigner, A.A., Gurvinder Dhaliwal, 2021. Uniswap: Impermanent loss and risk profile of a liquidity provider URL: <http://rgdoi.net/10.13140/RG.2.2.32419.58400/6>, doi: [10.13140/RG.2.2.32419.58400/6](https://doi.org/10.13140/RG.2.2.32419.58400/6).
- Angeris, G., Kao, H.T., Chiang, R., Noyes, C., Chitra, T., 2021. An analysis of uniswap markets. URL: <https://arxiv.org/abs/1911.03380>, arXiv:1911.03380.
- Auer R., Haslhofer B., K.S.S.P., F., V., 2023. The Technology of Decentralized Finance (DeFi). BIS Working Paper 1066. Bank for International Settlements. URL: <https://www.bis.org/publ/work1066.htm>.
- Axtell, R.L., Farmer, J.D., 2025. Agent-based modeling in economics and finance: Past, present, and future. Journal of Economic Literature 63, 197–287. URL: <https://www.aeaweb.org/articles?id=10.1257/jel.20221319>, doi: [10.1257/jel.20221319](https://doi.org/10.1257/jel.20221319).
- Bar-On, Y., Mansour, Y., 2023. Uniswap liquidity provision: An on-line learning approach. URL: <https://arxiv.org/abs/2302.00610>, arXiv:2302.00610.
- Boswijk, H.P., Hommes, C.H., Manzan, S., 2007. Behavioral heterogeneity in stock prices. Journal of Economic Dynamics and Control 31, 1938–1970. URL: <https://www.sciencedirect.com/science/article/pii/S0165188907000061>, doi: <https://doi.org/10.1016/j.jedc.2007.01.001>. tenth Workshop on Economic Heterogeneous Interacting Agents.

- Bottazzi, G., Dosi, G., Rebesco, I., 2005. Institutional architectures and behavioral ecologies in the dynamics of financial markets. *Journal of Mathematical Economics* 41, 197–228.
- Breiman, L., 2001. Random forests. *Machine learning* 45, 5–32.
- Breiman, L., Friedman, J., Stone, C.J., Olshen, R.A., 1984. Classification and regression trees. CRC press.
- Álvaro Cartea, Drissi, F., Monga, M., 2024. Decentralised finance and automated market making: Predictable loss and optimal liquidity provision. URL: <https://arxiv.org/abs/2309.08431>, arXiv:2309.08431.
- Chen, T., He, T., Benesty, M., Khotilovich, V., Tang, Y., Cho, H., Chen, K., Mitchell, R., Cano, I., Zhou, T., Li, M., Xie, J., Lin, M., Geng, Y., Li, Y., Yuan, J., 2024. xgboost: Extreme Gradient Boosting. URL: <https://CRAN.R-project.org/package=xgboost>. r package version 1.7.8.1.
- Cong, L., Dong, Y., Lu, Y., Ruan, Q., Wang, G., 2024. Agent-Based Modeling for DAOs and DeFi. URL: <https://ssrn.com/abstract=5171559>.
- Datseris, G., Vahdati, A.R., DuBois, T.C., 2024. Agents.jl: a performant and feature-full agent-based modeling software of minimal code complexity. *SIMULATION* 100, 1019–1031. URL: <https://doi.org/10.1177/00375497211068820>, doi:10.1177/00375497211068820, arXiv:<https://doi.org/10.1177/00375497211068820>.
- DeLong, J.B., Shleifer, A., Summers, L.H., Waldmann, R.J., 1990. Noise trader risk in financial markets. *Journal of Political Economy* 98, 703–738. doi:10.1086/261703.
- Fagiolo, G., Guerini, M., Lamperti, F., Moneta, A., Roventini, A., 2019. Validation of Agent-Based Models in Economics and Finance. Springer International Publishing, Cham. pp. 763–787. URL: https://doi.org/10.1007/978-3-319-70766-2_31, doi:10.1007/978-3-319-70766-2_31.
- Fan, Z., Marmolejo-Cossío, F., Altschuler, B., Sun, H., Wang, X., Parkes, D.C., 2022. Differential liquidity provision in uniswap v3 and implications for contract design. URL: <https://arxiv.org/abs/2204.00464>, arXiv:2204.00464.
- Fan, Z., Marmolejo-Cossío, F., Moroz, D.J., Neuder, M., Rao, R., Parkes, D.C., 2024. Strategic liquidity provision in uniswap v3. URL: <https://arxiv.org/abs/2106.12033>, arXiv:2106.12033.
- Harper, E.B., Stella, J.C., Fremier, A.K., 2011. Global sensitivity analysis for complex ecological models: a case study of riparian cottonwood

- population dynamics. *Ecological Applications* 21, 1225–1240. URL: <https://esajournals.onlinelibrary.wiley.com/doi/abs/10.1890/10-0506.1>, doi:<https://doi.org/10.1890/10-0506.1>, arXiv:<https://esajournals.onlinelibrary.wiley.com/doi/pdf/10.1890/10-0506.1>.
- Hastie, T., Tibshirani, R., Friedman, J., 2009. *The elements of statistical learning: data mining, inference, and prediction*. Springer Science & Business Media.
- Heimbach, L., Schertenleib, E., Wattenhofer, R., 2022. Risks and returns of uniswap v3 liquidity providers, in: *Proceedings of the 4th ACM Conference on Advances in Financial Technologies*, ACM. p. 89–101. URL: <http://dx.doi.org/10.1145/3558535.3559772>, doi:[10.1145/3558535.3559772](https://doi.org/10.1145/3558535.3559772).
- Hommes, C.H., 2005. *Heterogeneous Agent Models in Economics and Finance*. Tinbergen Institute Discussion Paper 05-056/1. Amsterdam and Rotterdam. URL: <https://hdl.handle.net/10419/86441>.
- Jaxa-Rozen, M., Kwakkel, J., 2018. Tree-based ensemble methods for sensitivity analysis of environmental models: A performance comparison with sobol and morris techniques. *Environmental Modelling & Software* 107, 245–266. URL: <https://www.sciencedirect.com/science/article/pii/S1364815217311581>, doi:<https://doi.org/10.1016/j.envsoft.2018.06.011>.
- Lehar, A., Parlour, C.A., 2023. Decentralized exchange: The uniswap automated market maker. *Journal of Finance* URL: <https://ssrn.com/abstract=3905316>. forthcoming.
- Liaw, A., Wiener, M., 2002. Classification and regression by randomforest. *R News* 2, 18–22. URL: <https://CRAN.R-project.org/doc/Rnews/>.
- Milborrow, S., 2024. *rpart.plot: Plot 'rpart' Models: An Enhanced Version of 'plot.rpart'*. URL: <https://CRAN.R-project.org/package=rpart.plot>. r package version 3.1.2.
- Othman, A.M., 2018. *Automated Market Making: Theory and Practice* URL: https://kilthub.cmu.edu/articles/thesis/Automated_Market_Making_Theory_and_Practice/6714920, doi:[10.1184/R1/6714920.v1](https://doi.org/10.1184/R1/6714920.v1).
- Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., Tarantola, S., 2008. *Global Sensitivity Analysis. The Primer*. John Wiley & Sons, Ltd.
- Therneau, T., Atkinson, B., 2023. *rpart: Recursive Partitioning and Regression Trees*. URL: <https://CRAN.R-project.org/package=rpart>. r package version 4.1.23.

Wang, M., Kampakis, S., 2024. Modeling speculative trading patterns in token markets: An agent-based analysis with tokenlab. URL: <https://arxiv.org/abs/2412.07512>, [arXiv:2412.07512](https://arxiv.org/abs/2412.07512).

xVA Greeks using the Likelihood Ratio Method within an American Monte Carlo Framework

TEAM 4

ADRIEN MATHIEU, University of Oxford
GAVRIEL BURGIN, University of Cape Town
GATHUKU MATHERI, University of Cape Town
LEHLOHONOLO MATLONYA, University of Cape Town

Supervisor:
OBEID MAHOMED, University of Cape Town

Contents

1 Introduction	4
1 Background and Motivation	4
2 Objectives	5
3 Scope & Limitations	5
4 Paper Outline	5
2 Theory Development	6
1 Definition of the Probability Space	6
2 American Monte-Carlo Estimation	7
2.1 Regression with Laguerre Polynomials	7
3 XVA Adjustments: FCVA	8
3.1 FCVA: Funding Credit Valuation Adjustment	8
4 Vasicek Interest Rate Model	9
5 Risk Sensitivities - The Greeks	9
5.1 Bump and Revalue Method	10
5.2 Pathwise Method	11
5.3 Likelihood Ratio Estimation Method	11
3 Methodology	13
1 Simulation of Paths	13
2 Implementation of AMC on FCVA	13
2.1 European style derivative with multiple assets and stochastic interest rate	13
3 Bump and Revalue Method for Greek Evaluation at $t=0$	15
3.1 Impracticality of bump and revalue method for Greeks through time	16
4 Pathwise Method for Greek Evaluation	18
5 Likelihood Ratio Estimation for Greek Evaluation	19
5.1 Case (i): One asset price with constant interest rates model	19
5.2 Case (ii): d asset prices with constant interest rates model	20
5.3 Case (iii): One asset price with stochastic interest rates model	21
5.4 Using the likelihood ratio estimator to compute FCVA Greeks	24

6	Method Comparison Metrics	29
4	Results and Analysis	33
5	Conclusion	36
1	Opportunities for Further Research	36
A	Covariance Matrix	39
1	Joint Covariance Matrix of stochastic processes	39

Chapter 1

Introduction

1 Background and Motivation

In modern financial markets, especially in the aftermath of the 2008 Global Financial Crisis, the calculation of valuation adjustments (xVAs) and their corresponding sensitivities has become an essential component in derivative pricing. Financial institutions could no longer overlook the impact of counterparty credit risk and its associated funding and hedging costs. xVAs represent a collection of adjustments that are applied to the theoretical risk-free price of over-the-counter derivatives. Their primary purpose is to incorporate a variety of real-world costs and risks that an associated financial institution encounters in these derivative transactions. Prior to 2008, existing valuation and risk management frameworks significantly underestimated counterparty credit risk, which led to the collapse of major financial institutions. Furthermore, there was a growing credit spread between the risk-free rate and funding benchmarks such as LIBOR and JIBAR, which required more realistic valuations (Gregory, 2018).

To address the shortcoming of the risk-free valuation approach, xVAs were created. We will focus on a particular xVA being the Funding Credit Valuation Adjustment (FCVA). The FCVA adjusts the value of a derivative to account for the cost to shareholders of funding potential future credit losses arising from counterparty default, thereby capturing the true economic cost of the position (Conti et al., 2022).

In addition to xVAs, their sensitivities, commonly known as Greeks measure how the value of an xVA responds to shifts in market factors such as asset prices and volatility. Understanding these sensitivities is vital to managing risk and implementing effective hedging strategies. This research paper will focus on two key Greeks: Delta and Gamma. Gregory (2018) describes Delta as a measure of how the value of a derivative changes in response to small changes in the underlying

asset price, while Gamma measures how sensitive Delta itself is to movements in the underlying, capturing the curvature of the price relationship.

The industry standard for determining these sensitivities is to use Algorithmic Automatic Differentiation (AAD), a computation technique that uses a reverse mode of the chain rule to propagate through sensitivities in a numerical model (Capriotti et al., 2015). A challenge arises when existing proprietary quantitative libraries are not equipped to integrate AAD. Thus, there is a need to use more traditional methods to determine the sensitivities of the xVAs. These traditional approaches include the bump and revalue method, the pathwise method, and, of particular interest, the likelihood ratio estimation technique. Although there is literature on calculating xVA Greeks (Conti et al., 2022) using the likelihood ratio method and using American Monte Carlo (AMC) simulation (Longstaff and Schwartz, 2001), the combined application of these procedures is not well documented in existing literature. This paper aims to explore the practicality of an integrated approach.

2 Objectives

This research paper seeks to compute the risk sensitivities for xVA metrics, in particular FCVA, using the Likelihood Ratio Method (LRM) within an AMC framework, with a focus on computational feasibility. The primary objective is to evaluate the efficacy of the LRM approach in comparison to other techniques, such as the bump and revalue method and the pathwise method for determining xVA Greeks.

3 Scope & Limitations

The project focuses specifically on the computation of the FCVA metric under the general xVA framework as detailed by (Kjaer, 2017). Furthermore, the interest rate model will only utilise the single factor Vasicek model and the sensitivities will focus on Delta and Gamma.

4 Paper Outline

In Chapter 2, we will focus on providing a background on xVAs, the American Monte-Carlo simulation, and the estimation techniques to determine the sensitivities. This will detail the theoretical and mathematical framework for each of these areas along with detailing of the model chosen for interest rates. In Chapter 3, the simulation procedure and regression framework is then further detailed. In Chapter 4, the results are presented across different scenarios using different estimation techniques. In Chapter 5, we conclude the research paper by highlighting our key insights and areas for future research.

Chapter 2

Theory Development

1 Definition of the Probability Space

Throughout this study, we consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}, \mathbb{P})$, where the filtration $\mathbb{F}$ is complete. We consider a d-dimensional asset process $(S_t)_{t \in [0, T]}$ with dynamics:

$$dS_t = \mu^S(S_t, r_t) dt + \sigma^S(S_t) dW_t^S, \quad S_0 = s_0, \quad (2.1)$$

where the process $(r_t)_{t \in [0, T]}$ is a one-dimensional interest rate process, with dynamics given by:

$$dr_t = \mu^r(r_t) dt + \sigma^r(r_t) dW_t^r. \quad (2.2)$$

The functions $\mu^r, \mu^S, \sigma^S, \sigma^r$ are assumed to satisfy the usual conditions such that the processes $(S_t)_{t \in [0, T]}$ and $(r_t)_{t \in [0, T]}$ are well defined. The filtration is the canonical filtration generated by the correlated Brownian motions $(W_t^S, W_t^r)_{t \in [0, T]}$, i.e., $\mathcal{F}_t = \sigma((W_u^S, W_u^r)_{u \leq t})$.

We remark that we also have

$$\mathcal{F}_t = \sigma((S_u, r_u)_{u \leq t}).$$

Let $(X_t)_{t \in [0, T]}$ be the solution to:

$$dX_t = a(X_t) dt + b^S(X_t) dW_t^S + b^r(X_t) dW_t^r, \quad (2.3)$$

where a, b satisfy the usual conditions. (see [Carmona \(2016\)](#) for more details). Furthermore, we then have the following Markov property:

$$\mathbb{E}_t[X_T] := \mathbb{E}[X_T | \mathcal{F}_t] = \mathbb{E}[X_T | S_t, r_t]. \quad (2.4)$$

2 American Monte-Carlo Estimation

American Monte Carlo (AMC) is a numerical approach that involves stepping through time and deriving Monte Carlo estimates at each time point through linear regression to determine the conditional expectation with respect to a current value compared to a future value.

To estimate the conditional expectation, we resort to the use of a special type of polynomial (termed Laguerre polynomials) that forms a basis and is dense in the space of L^2 functions (Longstaff and Schwartz (2001) discuss the use of other polynomials basis functions). In essence, the conditional expectation is a linear combination of a subset of Laguerre polynomials with the goal to estimate the required coefficients which will then determine the conditional expectation.

For example, we want to estimate a conditional expectation under the following form:

$$V_{t,u} = \mathbb{E}[G(X_u)|Y_t], \quad \text{where } u \geq t, \quad (2.5)$$

and $(X_t)_{t \in [0,T]}$ and $(Y_t)_{t \in [0,T]}$ solve two stochastic differential equations (SDEs). Then, we estimate V_t by

$$\sum_{j \leq N} \beta_{t,u}^j L_j(Y_t), \quad (2.6)$$

where $(L_j(x))_{j \leq n}$ are the polynomial basis functions we choose, and $(\beta_{t,u}^j)_{j \leq N}$ are solved by projecting $G(X_u)$ onto the polynomial basis functions. Longstaff and Schwartz (2001) discuss the convergence of this estimate to the true function V .

2.1 Regression with Laguerre Polynomials

The Laguerre polynomials can be defined in several equivalent ways with one of the simpler formulations being that they can be defined recursively as:

$$\begin{aligned} L_0(x) &= 1 \\ L_1(x) &= 1 - x \\ L_{k+1}(x) &= \frac{(2k+1-x)L_k(x) - kL_{k-1}(x)}{k+1}, \quad \forall k \geq 1 \end{aligned}$$

These polynomials are dense in the Hilbert space $L^2(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ and can thus be employed to estimate an L^2 function. Given that L^2 is a Hilbert space, a natural and reasonable approximation of a given function g in terms of Laguerre polynomials is the projection of g into the subspace generated by the countable collection of Laguerre polynomials.

Since linear regression involves minimising a squared distance, which is an L^2 norm, projection into an L^2 subspace is equivalent to linear regression and therefore we can estimate the required conditional expectation by regressing into a set of chosen regressors as detailed by Longstaff and Schwartz (2001).

3 XVA Adjustments: FCVA

Gregory (2018) refers to XVAs as a group of valuation adjustments that are applied typically on the risk-free price of OTC (Over-The-Counter) derivatives to account for a variety of real-world costs and risks that a financial institution may face in these transactions, such as funding cost spreads and counterparty default risk. This paper will focus specifically on the (FCVA) as defined in Kjaer (2017).

3.1 FCVA: Funding Credit Valuation Adjustment

The purpose of FCVA is to adjust a derivative's value so that it represents the cost to shareholders of funding potential future credit losses due to counterparty default risk. The adjustment arises when the bank must actively hedge this risk by estimating the expected loss from a counterparty. This hedging requires funding at the bank's funding rate which becomes a liability to shareholders. The equation below captures the expected funding cost to hedge these losses

$$\text{FCVA}_t = -\mathbb{E}_t \left[\int_t^T \lambda_C(u) D_{r+\lambda_C}(t, u) (V(u) - g_C(u)) du \right], \quad (16)$$

where

- $(r_t)_{t \in [0, T]}$ is the stochastic interest rate,
- $\lambda_C(u)$ is the spread of a counterparty zero recovery bond and is given by $\lambda_C = \frac{r_C - \gamma_C}{1 - R_C}$, chosen constant in our framework,
- $r_C(u)$ and $\gamma_C(u)$ are the counterparty overnight bond rate and the repo rate secured against C 's bonds respectively,
- R_C is the recovery rate in the event of counterparty default
- $D_{r+\lambda_C}(t, u) = \exp(-\int_t^u (r_s + \lambda_C(s)) ds)$,
- $V(u)$ is the value of the derivative portfolio at time u ,
- $g_C(u)$ represents the present value of the portfolio immediately after default and can represent contractual features such as standard ISDA closeouts, netting (or absence thereof) and collateral. We only consider uncollateralised transactions in this paper. In line with Kjaer (2017), we consider $g_C(u) =$

$R_C V(t)^+ + V(t)^-$ (since there's no collateral in our implementations, $\phi = \psi = 0$). This implies that $V(u) - g_C(u) = (1 - R_C)V(t)^+ = V(u)^+$ (since we assume zero-recovery). We remark that for a derivative with a non-negative payoff (such as options), $g_C(u) = 0$ ($R_C = V(t)^- = 0$) and since we only consider options in our implementation, we set $g_C(u) = 0$.

This expectation incorporates the bank's hazard rate (λ_C) over its funding rate which represents the additional cost of unsecured borrowing. An appropriate discount factor is applied with respect to the bank's actual funding rate where all cash flows are discounted using this rate to characterise the cost of financing uncollateralised exposure. For the purposes of this project, the recovered value of the derivative in the event of default $g_C(u)$ is taken to be 0 partly because the derivative is considered uncollateralised. In our framework, where we consider V to be the price of an option, the Feynmann-Kac formula allows us to represent the value of the portfolio as the expected discounted terminal payoff of the portfolio. From now, we define it as

$$V_t := \mathbb{E}_t \left[e^{-\int_t^T r_s ds} H(S_T, r_T) \right]. \quad (2.7)$$

4 Vasicek Interest Rate Model

Asset prices and their associated pricing are driven by interest rate dynamics which are characterised as stochastic in nature. The Vasicek model is a widely-used Ornstein-Uhlenbeck mean-reverting process for the short-rate that is intended to capture the basic features of interest rate dynamics that has the following form:

$$dr_t = \kappa (\theta - r_t) dt + \sigma dW_t^r,$$

where κ represents the rate of mean-reversion, θ is the long-term mean, σ is the volatility, and $(W_t^r)_{t \in [0, T]}$ is a standard Brownian motion. The model maintains the single factor Gaussian distribution for the short-rate while allowing for a drift term that is incorporated in to the short-rate dynamics.

The model is analytically tractable due to its Gaussian structure which allows for closed-form solutions for zero-coupon bond prices.

5 Risk Sensitivities - The Greeks

Risk sensitivities measure the sensitivity of a financial position to factors such as the underlying asset price. (Gottesman, 2016).

Calculating xVAs includes estimating how exposed a derivative portfolio is to potential counterparty defaults, which depends on how underlying market risk factors evolve. Understanding how this exposure changes when these factors move, as measured by the Greeks, is essential for effective risk management and hedging.

In this paper, we will focus on the following Greeks:

- **Delta:** Measures the change in valuation based on changes in the price of the asset.
- **Gamma:** Measures the impact of convexity of changes in asset price.

5.1 Bump and Revalue Method

This method is a rudimentary approach to calculating sensitivities by implementing small shocks to each of the state variables. This method is independent of the actual simulation procedure, thus making it highly implementable. However, Gaussian noise during simulation can lead to inaccuracies. These challenges may be mitigated by utilising the same set of random numbers in the implementation of the state variable bump [Glasserman \(2004\)](#). //

The generalised formula for a sensitivity calculation $\hat{\Delta}$ with respect to a state variable (θ) is defined as:

$$\hat{\Delta}(\theta) \approx \frac{\bar{V}(\theta + h) - \bar{V}(\theta)}{h},$$

where:

- θ is the parameter/state variable of interest,
- $\bar{V}$ is the (estimated) derivative value function whose arguments include θ ,
- h is some suitably small positive real number.

The use of finite difference methods within the context of xVAs requires considering the future expectation of exposures on simulated paths with respect to simulated risk factors. The computation of sensitivities requires the entire Monte Carlo simulation to be rerun with the bumped parameters and the recalculation of regression polynomials which are then used to calculate risk sensitivities via the finite difference methods.

While this method serves as the benchmark due to its practicality and simplicity to implement, the requirement to run multiple simulations increases the computational cost which is not ideal [\(Buis, 2023\)](#). Nonetheless, this method can serve as the benchmark to determine the effectiveness of using the LRM.

5.2 Pathwise Method

The pathwise derivative method implements sensitivity calculations utilising the direct differentiation of a simulation path. Let $Y(\theta)$ be the discounted payoff or exposure computed under a target parameter and define $\alpha(\theta) = \mathbb{E}[Y(\theta)]$ represents a valuation. The pathwise method would thus estimate the derivative as follows:

$$\frac{d\alpha}{d\theta} = \frac{d}{d\theta} \mathbb{E}Y(\theta) = \mathbb{E} \left[\frac{dY}{d\theta}(\theta) \right],$$

where the interchange of the differentiation and expectation operators are valid when the function is continuous and smooth. These continuous and smooth functions are required for pathwise estimation to work and leads to the following estimator:

$$\frac{d\alpha}{d\theta} \approx \frac{1}{N} \sum_{i=1}^N \frac{dY^{(i)}}{d\theta}.$$

In contrast to the bump and revalue method, the advantage of the pathwise method is its ability to reuse simulation paths to calculate the derivative with limited additional computation in contrast to the bump and revalue method (Glasserman, 2004).

Our framework simplifies the method. We have

$$V_t := F(S_t), \quad (2.8)$$

where F is a C^n function. Then, if we want access to the n -th order sensitivity with respect to the process S_t of V_t , we can compute it directly with

$$\partial_{S_t}^n V_t = \partial_{S_t}^n F(S_t). \quad (2.9)$$

Moreover, if we want access to the m -th sensitivity with respect to a parameter x of V_t , and if we assume that $x \rightarrow S_t$ is at least C^m , then we obtain

$$\partial_x^m V_t = \partial_{S_t}^m F(S_t) \partial_x^m S_t. \quad (2.10)$$

5.3 Likelihood Ratio Estimation Method

The LRM estimates the sensitivities by deriving the underlying probability density function with respect to the target variable, denoted here by θ .

Consider the standard estimation of a function under a given probability density:

$$\mathbb{E}_\theta[Y] = \int f(x) g_\theta(x) dx.$$

According to [Glasserman \(2004\)](#), the LRM derives a score function from the probability density function:

$$\frac{d}{d\theta} \mathbb{E}_\theta[Y] = \int f(x) \frac{\dot{g}_\theta(x)}{g_\theta(x)} g_\theta(x) dx = \mathbb{E}_\theta \left[f(X) \frac{\dot{g}_\theta(X)}{g_\theta(X)} \right],$$

where $\dot{g}_\theta(x)$ is the partial derivative with respect to θ . In our framework, it allows us to obtain the following results. We want to compute the sensitivities of a function $(V_t)_{t \in [0, T]}$ given by

$$V_t = \mathbb{E}[F(X_T) | \mathcal{F}_t], \quad (2.11)$$

where the process $(X_t)_{t \in [0, T]}$ is a d -dimensional process that follows the SDE

$$dX_t = a(X_t) dt + \sigma(X_t) dW_t, \quad X_0 = x_0, \quad (2.12)$$

and where $(\mathcal{F}_t)_{t \in [0, T]}$ is the filtration generated by the stochastic process $(Y_t)_{t \in [0, T]}$, that satisfies

$$dY_t = \bar{a}(Y_t) dt + \bar{\sigma}(Y_t) d\bar{W}_t, \quad Y_0 = y_0. \quad (2.13)$$

We assume that $\langle W, \bar{W} \rangle_t = \rho t$. By the Markov property of V_t , we can rewrite it as

$$V_t = \mathbb{E}[F(X_T) | Y_t]. \quad (2.14)$$

The crux of the likelihood estimation method is the fact that we can find the conditional law of $X_T | Y_t$, given by $f_{X_T | Y_t}$, so that V_t becomes

$$V_t = \int F(x) f_{X_T | Y_t}(x) dx. \quad (2.15)$$

Then, computing the sensitivity to S_t is equivalent to computing

$$\partial_{S_t} V_t = \int F(x) \partial_{S_t} f_{X_T | Y_t}(x) dx = \int F(x) \partial_{S_t} \log(f_{X_T | Y_t}(x)) f_{X_T | Y_t}(x) dx, \quad (2.16)$$

and this is rewritten as

$$\partial_{S_t} V_t = \mathbb{E} \left[F(X_T) \log(f_{X_T | Y_t}(X_T) | Y_t) \right]. \quad (2.17)$$

The second order derivative is computed the same way, using the fact that

$$\partial_{x,x}^2 f_x = \partial_x (\partial_x \log(f_x) f_x) = \left(\partial_{x,x}^2 \log(f_x) + (\partial_x \log(f_x))^2 \right) f_x. \quad (2.18)$$

We can also perform sensitivity with respect to parameters in a similar way.

Chapter 3

Methodology

1 Simulation of Paths

In order to price xVAs using an AMC simulation, the stochastic processes need to be simulated. Multiple assets and the OIS rate are simulated where asset prices are simulated using a geometric Brownian motion, where the log-return of the asset is driven by a stochastic short rate and constant volatility. Furthermore, interest rates evolve also as a stochastic process under the Vasicek model with constant mean reversion and constant volatility. Given the link between the short rate and evolving asset prices, these models are linked with a fixed correlation structure to represent these dependencies.

2 Implementation of AMC on FCVA

2.1 European style derivative with multiple assets and stochastic interest rate

We consider a d-dimensional asset price process:

$$dS_t = S_t (r_t dt + \sigma_s) dW_t^S, \quad (3.1)$$

where the stochastic interest rate follows

$$dr_t = \kappa (\theta - r_t) dt + \sigma_r dW_t^r, \quad (3.2)$$

where the Brownian motions $(W_t^S)_{t \in [0, T]}$ and $(W_t^r)_{t \in [0, T]}$ are correlated. We wish to obtain:

$$\text{FCVA}_t = -\mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-\int_t^u r_s ds} V_u du \right].$$

First we consider V_u as non-negative for each $u \in [t, T]$. Recall that derivative prices are conditional expectations of future discounted payoffs, i.e., $V_u = \mathbb{E}_u [D_r(T, u) \cdot H(S_T, r_T)]$. Using Laguerre polynomials of order 5, we project the product $D_r(T, u) \cdot H(S_T, r_T)$ onto the stochastic processes $X_u = [S_u^1, S_u^2, \dots, S_u^d, r_u]$ like so:

$$\begin{pmatrix} D_r(T, u) \cdot H(S_T^1, S_T^2, \dots, S_T^d, r_T)(\omega_1) \\ D_r(T, u) \cdot H(S_T^1, S_T^2, \dots, S_T^d, r_T)(\omega_2) \\ D_r(T, u) \cdot H(S_T^1, S_T^2, \dots, S_T^d, r_T)(\omega_3) \\ \vdots \\ D_r(T, u) \cdot H(S_T^1, S_T^2, \dots, S_T^d, r_T)(\omega_N) \end{pmatrix} = \sum_{i \in \mathcal{I}} \alpha_i^u \prod_{k=1}^{d+1} L_i \begin{pmatrix} X_u^k(\omega_1) \\ X_u^k(\omega_2) \\ X_u^k(\omega_3) \\ \vdots \\ X_u^k(\omega_N) \end{pmatrix},$$

where $\mathcal{I} = \{0, 1, 2\} \otimes \{0, 1, 2\} \otimes \dots \{0, 1, 2\} : (d+1) \text{ times. i.e., } |\mathcal{I}| = 3^{(d+1)}$. We now use the α 's to obtain $V_\phi(u, \omega)$ for each $\omega \in \Omega$ (length N) and each time step $u \in [t, T]$, i.e.,

$$V_u^{\text{Laguerre}}(\omega) = \sum_{i \in \mathcal{I}} \alpha_i^u \prod_{k=1}^{d+1} L_i(X_u^k(\omega)).$$

Next we wish to calculate

$$\int_t^T \underbrace{\lambda_c(u) e^{-\int_t^u \lambda_c(s) ds}}_{\text{PD}} \underbrace{e^{-\int_t^u r_s(\omega) ds} V_u^{\text{Laguerre}}(\omega)}_{\text{EAD}} du.$$

For each $\omega \in \Omega$, we have that

$$FCVA(u) \approx \sum_{u=t}^T \text{PD (deterministic)} \cdot \text{EAD (stochastic)} \cdot \Delta u.$$

Once again we wish to circumvent the conditional expectation by utilising projection. This time we project on $X_t = [S_t^1, S_t^2, \dots, S_t^d, r_t]$ like so:

$$\begin{pmatrix} \sum_{u=t}^T \text{PD} \cdot \text{EAD}(\omega_1) \cdot \Delta u \\ \sum_{u=t}^T \text{PD} \cdot \text{EAD}(\omega_2) \cdot \Delta u \\ \vdots \\ \sum_{u=t}^T \text{PD} \cdot \text{EAD}(\omega_N) \cdot \Delta u \end{pmatrix} = \sum_{i \in \mathcal{I}} \beta_i^t \prod_{k=1}^{d+1} L_i \begin{pmatrix} X_t^k(\omega_1) \\ X_t^k(\omega_2) \\ X_t^k(\omega_3) \\ \vdots \\ X_t^k(\omega_N) \end{pmatrix}.$$

Through Laguerre regression, we have projected the conditional expectation of discounted exposure onto the current state variable S_t and r_t , allowing us to express FCVA at any time t without requiring nested simulations.

$$\text{FCVA}_t = -\mathbb{E}_t \left[\int_t^T \underbrace{\lambda_c(u) e^{-\int_t^u \lambda_c(s) ds}}_{\text{PD}} \underbrace{e^{-\int_t^u r_s ds} V_u^{\text{Laguerre}}}_{\text{EAD}} du \right] \approx - \sum_{i \in \mathcal{I}} \beta_i^t \prod_{k=1}^{d+1} L_i(X_t^k).$$

Given the coefficients β_i^t from the projection onto X_t , we can now compute FCVA_t for any $t \in [0, T]$ directly as a function of the current state, bypassing the need for evaluating conditional expectations.

3 Bump and Revalue Method for Greek Evaluation at t=0

We consider the most simple example with a single underlying asset with dynamics governed by constant parameters:

$$dS_t = r S_t dt + \sigma S_t dW_t.$$

Our objective is to analyse the time evolution of the Δ_t of FCVA for a European call option written on this asset.

First, we consider Δ_0 of an FCVA. We simulate a set of representative sample paths for the underlying asset price S_t as well as its bumped counterpart $\tilde{S}_t = S_t + \Delta S$.




Figure 3.1: Toy demonstration of S_t and $\tilde{S}_t$ paths

For each of these paths, we compute the exposure at default. We regress the exposures at default on the paths of S_0 to obtain coefficients β_i^0 , and separately regress on the bumped paths $\tilde{S}_0$ to obtain $\tilde{\beta}_i^0$.

- We define the estimated FCVA under the original paths as:

$$\text{FCVA}_0 = \sum \beta_i^0 L_i(S_0),$$

- and the FCVA under the bumped paths as:

$$\widetilde{\text{FCVA}}_0 = \sum \tilde{\beta}_i^0 L_i(S_0 + \Delta S).$$

Then the finite difference approximation of the sensitivity is:

$$\Delta \text{FCVA}_0 = \frac{\widetilde{\text{FCVA}}_0 - \text{FCVA}_0}{\Delta S}.$$

3.1 Impracticality of bump and revalue method for Greeks through time

Next we consider ΔFCVA_1 , the change in FCVA due to variations in the underlying asset at $t = 1$. Unlike the initial FCVA estimation at time $t = 0$, computing ΔFCVA_1 requires re-evaluating the conditional expectations at time $t = 1$ for each ω . This means that for every realization $S_1(\omega)$, a fresh set of inner Monte Carlo simulations must be run.

We are essentially sitting here and wish to know what the delta is given that we made it to this time point.




Figure 3.2: Toy demonstration of $S_t(\omega)$ paths at $t = 1$.

For each $\omega \in \Omega$ we simulate paths of length $n\text{-step} - 1$. For example consider the path of $S_t(\omega_1)$ below.




Figure 3.3: Toy demonstration of $S_t(\omega_1)$ and $\tilde{S}_t(\omega_1)$

Again we compute the exposure at default along each new path of $S_t(\omega_1)$ and $\tilde{S}_t(\omega_1)$. We then regress on $S_1(\omega_1)$ to obtain coefficients $\beta_i^1(\omega_1)$, and separately regress on $\tilde{S}_1$ to obtain $\tilde{\beta}_i^1(\omega_1)$.

- We define the conditional FCVA_1 as:

$$\text{FCVA}_1(\omega_1) = \sum \beta_i^1(\omega_1) L_i(S_1(\omega_1)),$$

- and the conditional FCVA under the bumped paths as:

$$\widetilde{\text{FCVA}}_1 = \sum \tilde{\beta}_i^1(\omega_1) L_i(S_1(\omega_1) + \Delta S).$$

Then the finite difference approximation of the sensitivity is:

$$\Delta \text{FCVA}_1(\omega_1) = \frac{\widetilde{\text{FCVA}}_1(\omega_1) - \text{FCVA}_1(\omega_1)}{\Delta S}.$$

Finally we obtain:

$$\Delta \text{FCVA}_1 = \frac{1}{N} \sum_{\omega \in \Omega} \Delta \text{FCVA}_1(\omega)$$

Summary

- At each outer path $\omega \in \{1, \dots, N\}$ and each time step t_i , we perform an inner simulation of N paths to estimate the conditional expectation from t_i to T . Each of these inner paths has a remaining time horizon of $n_{\text{steps}} - i$ time steps.
- The total number of inner simulation time steps across all outer paths and time steps is therefore:

$$\sum_{\omega=1}^N \sum_{i=0}^{n_{\text{steps}}-1} N \cdot (n_{\text{steps}} - i) = N^2 \sum_{i=0}^{n_{\text{steps}}-1} (n_{\text{steps}} - i) = N^2 \sum_{j=1}^{n_{\text{steps}}} j = N^2 \cdot \frac{n_{\text{steps}}(n_{\text{steps}} + 1)}{2}.$$

- This implies a total computational complexity of order:

$$\mathcal{O}\left(N^2 \cdot n_{\text{steps}}^2\right).$$

Table 3.1: Total Simulated Time Steps in Nested Monte Carlo

N	n_{steps}	Total Simulated Steps $N^2 \cdot \frac{n_{\text{steps}}(n_{\text{steps}}+1)}{2}$
1,000	100	5.05×10^9
5,000	100	1.26×10^{11}
10,000	100	5.05×10^{11}
10,000	252	3.18×10^{12}
100,000	252	3.18×10^{14}

4 Pathwise Method for Greek Evaluation

In the context of AMC estimation, the Laguerre polynomials provide a continuous and differentiable set of functions which eliminate the primary limitation of the pathwise method. Thus, the pathwise method offers a useful solution when using the AMC estimate.

Consider the following problem: We want to compute the Δ of a product of the form

$$V_t = \mathbb{E}_t \left[\int_t^T F(S_u, r_u) du \right], \quad (3.3)$$

where we have

$$dS_t = S_t (r_t dt + \sigma^S dW_t^S), \quad (3.4)$$

and

$$dr_t = \kappa (\theta - r_t) dt + \sigma_r dW_t^r, \quad (3.5)$$

where $(W_t^S)_{t \in [0, T]}$ is a d -dimensional Brownian motion, and $(W_t^r)_{t \in [0, T]}$ a one dimensional Brownian motion, correlated with $(W_t^S)_{t \in [0, T]}$. First, we obtain, by the American Monte Carlo estimation, that

$$\mathbb{E}_t \left[\int_t^T F(S_u, r_u) du \right] \approx \sum_{j \leq N} \beta_t^j L_j(S_t, r_t), \quad (3.6)$$

where the $L_j(x)$ are the Laguerre polynomials, and the β^j the optimal regressors. Formally, we want to obtain

$$\partial_{S_t} \mathbb{E}_t \left[\int_t^T F(S_u, r_u) du \right] \approx \sum_{j \leq N} \beta_t^j \partial_{S_t} L_j(S_t, r_t). \quad (3.7)$$

If we define the Hessian matrix $\nabla_{S_t}^2$, we obtain

$$\nabla_{S_t}^2 V_t = \sum_{j \leq N} \beta_t^j \nabla_{S_t}^2 L_j(S_t, r_t). \quad (3.8)$$

As the Laguerre polynomials are C^∞ , we can perform as many derivatives of them as we want. Moreover, at fixed t , as S_t is C^∞ in its parameters $x \in \{r, \sigma^S\}$ we can perform sensitivities with respect to those by

$$\partial_x V_t = \sum_{j \leq N} \beta_t^j \partial_{S_t} L_j(S_t, r_t) \partial_x S_t. \quad (3.9)$$

5 Likelihood Ratio Estimation for Greek Evaluation

We compute the sensitivities of a product using the LRM in three different settings:
 (i) in the case of a one dimensional asset with a constant short term interest rate,
 (ii) in the case of a d-dimensional asset with a constant short term interest rate, and
 (iii) in the case of a one dimensional asset with stochastic short term interest rate.

5.1 Case (i): One asset price with constant interest rates model

We consider the process $(S_t)_{t \in [0, T]}$ with the dynamic

$$dS_t = r S_t dt + \sigma^S S_t dW_t^S. \quad (3.10)$$

We define

$$V_t = \mathbb{E}_t \left[e^{-r(T-t)} H(S_T) \right], \quad (3.11)$$

and we want to use the likelihood ratio estimator to estimate Δ_t of this option. To do so, we recall that

$$S_T = S_t e^{-r(T-t)} e^{(r - \frac{\sigma^2}{2})(T-t) + \sigma(W_T - W_t)}. \quad (3.12)$$

We define $Y_t := \log(S_t)$, hence

$$Y_t = \log(S_t) - \frac{\sigma^2}{2}(T-t) + \sigma(W_T - W_t). \quad (3.13)$$

Then Y_t is Gaussian, meaning that $S_T|S_t$ is log-normal. Given that we know the law of $S_T|S_t$, given by $f_{S_T|S_t}$, we can rewrite V_t as

$$V_t = \int_{\mathbb{R}_+} f_{S_T|S_t}(x) H(x) dx. \quad (3.14)$$

Then, if we want to differentiate V_t with respect to S_t , we obtain

$$\partial_{S_t} V_t = \int H(x) \partial_{S_t} f_{S_T|S_t}(x) dx, \quad (3.15)$$

which is equal to

$$\partial_{S_t} f_{S_T|S_t}(x) = \partial_{S_t} \log(f_{S_T|S_t}(x)) f_{S_t}(x). \quad (3.16)$$

Thus, we obtain

$$\partial_{S_t} V_t = \mathbb{E}_t \left[H(S_T) \partial_{S_t} \log(f_{S_T|S_t}(S_T)) e^{-r(T-t)} \right] = \mathbb{E}_t \left[H(S_T) \partial_{S_t} \log(f_{S_T|S_t}(S_T)) e^{-r(T-t)} \right],$$

and so

$$\partial_{S_t} V_t = \mathbb{E}_t \left[H(S_T) \left(\frac{\log(S_T) - \log(S_t) - (r - \frac{\sigma^2}{2})(T-t)}{S_t \sigma^2 (T-t)} \right) e^{-r(T-t)} \right]. \quad (3.17)$$

Finally, we estimate this value using the American Monte Carlo method. We obtain

$$\partial_{S_t} V_t \approx \sum_{j \leq N} \beta_t^j L_j(S_t). \quad (3.18)$$

Now, if we want to compute the Γ_t , we need to differentiate twice with respect to S_t , and we obtain

$$\partial_{\{S_t, S_t\}}^2 V_t = \mathbb{E}_t \left[H(S_T) \left(\partial_{\{S_t, S_t\}}^2 \log(f_{S_T|S_t}(S_T)) + (\partial_{S_t} \log(f_{S_T|S_t}(S_T)))^2 \right) e^{-r(T-t)} \right].$$

We have

$$\partial_{\{S_t, S_t\}}^2 \log(f_{S_T|S_t}(S_T)) = - \frac{\bar{\sigma}_{T,t} + (\log(S_T) - \log(S_t) - \bar{\mu}_{T,t}) \bar{\sigma}_{T,t}}{S_t^2 \bar{\sigma}_{T,t}^2}, \quad (3.19)$$

where $\bar{\sigma}_{T,t} = \sigma^2 (T - t)$ and $\bar{\mu}_{T,t} = (r - \frac{\sigma^2}{2})(T - t)$. Finally, we obtain

$$\Gamma_t = \mathbb{E}_t \left[H(S_T) \left(- \frac{1}{S_t^2 \bar{\sigma}_{T,t}} + \left(\frac{\log(S_T) - \log(S_t) - \bar{\mu}_{T,t}}{S_t \bar{\sigma}_{T,t}} \right)^2 - \left(\frac{\log(S_T) - \log(S_t) - \bar{\mu}_{T,t}}{S_t^2 \bar{\sigma}_{T,t}} \right) \right) e^{-r(T-t)} \right],$$

and once again, the estimation of the Γ_t is obtain with the American Monte Carlo method.

5.2 Case (ii): d asset prices with constant interest rates model

We consider now a d -dimensional asset price $(S_t)_{t \in [0, T]}$ given by

$$dS_t^i = S_t^i (r dt + \sigma dW_t^i), \quad S_0^i = S_0, \quad (3.20)$$

where $\langle W^i, W^j \rangle_t = \rho^{i,j} t$. Following the same approach as in the one dimensional case, we obtain that $(S_T^1, \dots, S_T^d | S_t^1, \dots, S_t^d)$ is log normal, with mean:

$$\mu_{T,t} = \begin{pmatrix} (r - \frac{\sigma^2}{2})(T - t) + \log(S_t^1) \\ \dots \\ (r - \frac{\sigma^2}{2})(T - t) + \log(S_t^d) \end{pmatrix} \quad (3.21)$$

and where the covariance is given by

$$\Sigma_{T,t} = (\Sigma_{T,t}^{i,j})_{0 \leq i, j \leq d}, \quad \text{where } \Sigma_{T,t}^{i,j} = \sigma^2 \rho^{i,j} (T - t), \quad (3.22)$$

where $\rho^{i,i} = 1$. A similar calculation gives

$$\nabla_{S_t} V_t = \mathbb{E}_t \left[e^{-r(T-t)} H(S_T) \left[D \left(\frac{1}{S_t^i} \right)_{i \leq d} \right] \Sigma_{T,t}^{-1} \begin{pmatrix} \log(S_T^1) - \mu_{T,t}^1 \\ \dots \\ \log(S_T^d) - \mu_{T,t}^d \end{pmatrix} \right], \quad (3.23)$$

where

$$\left[D \left(\frac{1}{S_t^i} \right)_{i \leq d} \right] = \begin{pmatrix} \frac{1}{S_t^1} & 0 & \dots & 0 \\ 0 & \frac{1}{S_t^2} & \dots & 0 \\ 0 & 0 & \dots & 0 \\ 0 & 0 & \dots & \frac{1}{S_t^d} \end{pmatrix}. \quad (3.24)$$

Remark 5.1. When $d = 1$, we recover the first case, given that $\Sigma_{T,t}^{-1}$ becomes $1/\sigma^2 (T - t)$.

5.3 Case (iii): One asset price with stochastic interest rates model

We consider the model where

$$dS_t = S_t (r_t dt + \sigma dW_t^S), \quad S_0 = s, \quad (3.25)$$

$$dr_t = \kappa (\theta - r_t) dt + \sigma_r dW_t^r, \quad r_0 = r, \quad (3.26)$$

where $\langle W^r, W^S \rangle_t = \rho t$. We want to estimate

$$\partial_{S_t} V_t := \partial_{S_t} \mathbb{E}[e^{-\int_t^T r_s ds} H(S_T, r_T) | r_t, S_t], \quad (3.27)$$

with the likelihood ratio estimator. To do so, we define

$$I_{u,t} = \int_t^u r_s ds. \quad (3.28)$$

We obtain by the use of the stochastic Fubini theorem (see Theorem 12.4.18 in [Cohen and Elliott \(2015\)](#)), for $u \geq t$, we have

$$I_{u,t} = \frac{1}{\kappa} (\bar{r}_t - \bar{r}_u + \sigma_r W_u^r - \sigma_r W_t^r) + \theta (u - t) + (r - \theta) \frac{(e^{-\kappa t} - e^{-\kappa u})}{\kappa}, \quad (3.29)$$

where $\bar{r}_t = \int_0^t \sigma_r e^{-\kappa(t-u)} dW_u^r$. Moreover, we define the stochastic process $\tilde{S}_t := \log(S_t)$. Then, $\tilde{S}$ is a Gaussian process. We rewrite $\tilde{S}_t$ as

$$\tilde{S}_t = \tilde{S}_0 + I_{t,0} - \frac{\sigma_s^2}{2} t + \sigma_s W_t^s, \quad (3.30)$$

where

$$I_{t,0} = \frac{1}{\kappa} (-\bar{r}_t + \sigma_r W_t^r) + \theta t + (r - \theta) \frac{(1 - e^{-\kappa t})}{\kappa}, \quad (3.31)$$

so that, we obtain

$$\tilde{S}_t = \underbrace{\tilde{S}_0 + \theta t + (r - \theta) \frac{(1 - e^{-\kappa t})}{\kappa}}_{:=c(t)} + \frac{\sigma_s^2 t}{2} + \frac{\sigma_r W_t^r - \bar{r}_t}{\kappa} + \sigma_s W_t^s. \quad (3.32)$$

We rewrite V_t as

$$V_t = \exp \left(-\frac{\bar{r}_t}{\kappa} - \theta(u-t) - (r-\theta) \frac{(e^{-\kappa t} - e^{-\kappa u})}{\kappa} \right) \mathbb{E} \left[e^{-\frac{1}{\kappa}(-\bar{r}_T + \sigma_r W_T^r - \sigma_r W_t^r)} H(e^{\tilde{S}_T}, r_T) | \tilde{S}_t, r_t \right].$$

Let

$$f_{(\bar{r}_T, W_t^r, W_T^r, \tilde{S}_T, r_T) | (\tilde{S}_t, r_t)}(\bar{r}^T, W^t, W^T, \tilde{S}^T, r^T)$$

denote the conditional joint density of $(r_T, W_t^r, W_T^r, \tilde{S}_T)$ given $(\tilde{S}, r) = (\tilde{S}_t, r_t)$. For readability, we will henceforth represent this conditional density using compact notation as

$$f_{X_{T,t} | Y_t = \underline{y}_t}(\underline{x}_{T,t}),$$

where

$$\begin{aligned} X_{T,t} &= (\bar{r}_T, W_t^r, W_T^r, \tilde{S}_T, r_T), \\ Y_t &= (\tilde{S}_t, r_t), \\ \underline{x}_{T,t} &= (\bar{r}^T, W^t, W^T, \tilde{S}^T, r^T), \\ \underline{y}_t &= (\tilde{S}^t, r^t). \end{aligned}$$

We define

$$H_t := \exp \left(-\frac{\bar{r}_t}{\kappa} - \theta(u-t) - (r-\theta) \frac{(e^{-\kappa t} - e^{-\kappa u})}{\kappa} \right), \quad (3.33)$$

then we have

$$V_t = H_t \int e^{-\frac{1}{\kappa}(-\bar{r}^T + \sigma_r W^T - \sigma_r W^t)} H(e^{\tilde{S}^T}, r^T) f_{X_{T,t} | Y_t = \underline{y}_t}(\underline{x}_{T,t}) d(\underline{x}_{T,t}).$$

We use $f_{X_{T,t} | Y_t} = \frac{f_{X_{T,t}, Y_t}}{f_{Y_t}}$ to compute the conditional density. Given that the process $(\bar{r}_t, r_t, W_t^r, \tilde{S}_t)_{t \in [0, T]}$ is Gaussian, then $(\bar{r}_T, W_t^r, W_T^r, \tilde{S}_T, r_T, \tilde{S}_t, r_t) = (X_{T,t}, Y_t)$ is a multivariate Gaussian random variable, with mean $\mu_{X_{T,t}, Y_t}$ and covariance $\Sigma_{X_{T,t}, Y_t}$. After tedious calculations, we obtain for the mean:

$$\mu_{X_{T,t}, Y_t} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ c(T) \\ r e^{-\kappa T} + \theta(1 - e^{-\kappa T}) \\ c(t) \\ r e^{-\kappa t} + \theta(1 - e^{-\kappa t}) \end{bmatrix}, \quad (3.34)$$

and the covariance $\Sigma_{X_{T,t}, Y_t}$ is given by $\Sigma_{X_{T,t}, Y_t} = \left(\Sigma_{X_{T,t}, Y_t}^{(i,j)} \right)_{1 \leq i, j \leq 7}$ where $\Sigma_{X_{T,t}, Y_t}$ is given by a table found in Appendix A.

Moreover, $(\tilde{S}_t, r_t) = Y_t$ is also a multivariate Gaussian process, with mean:

$$\mu_{Y_t} = \begin{bmatrix} c(t) \\ r e^{-\kappa t} + \theta(1 - e^{-\kappa t}) \end{bmatrix} \quad (3.35)$$

and with covariance matrix:

$$\Sigma_{Y_t} = \begin{bmatrix} \Sigma_{X_{T,t}, Y_t}^{(6,6)} & \Sigma_{X_{T,t}, Y_t}^{(7,6)} \\ \Sigma_{X_{T,t}, Y_t}^{(7,6)} & \Sigma_{X_{T,t}, Y_t}^{(7,7)} \end{bmatrix} \quad (3.36)$$

We obtain that $f_{X_{T,t}|Y_t}$ is a Gaussian density. Let us rewrite the matrix $\Sigma_{X_{T,t}, Y_t}$ as

$$\begin{bmatrix} \Sigma_{X_{T,t}} & \Sigma_{T,t}^\dagger \\ \Sigma_{T,t}^* & \Sigma_{Y_t} \end{bmatrix}. \quad (3.37)$$

The other matrices in $\Sigma_{X,Y}$ are defined as sub-matrices with:

- $\Sigma_{X_{T,t}}$: top left 5x5 sub-matrix;
- $\Sigma_{T,t}^\dagger$: top right 5x2 sub-matrix;
- $\Sigma_{T,t}^*$: bottom left 2x5 sub-matrix;

and we rewrite the vector $\mu_{X_{T,t}, Y_t}$ as

$$\begin{bmatrix} \mu_{X_{T,t}} \\ \mu_{Y_t} \end{bmatrix}. \quad (3.38)$$

Finally, $f_{X_{T,t}|Y_t}$ is a Gaussian density with mean:

$$\mu_{X_{T,t}|Y_t=\underline{\mathbf{y}}_t} = \mu_{X_{T,t}} + \Sigma_{T,t}^\dagger \Sigma_{Y_t}^{-1} (\underline{\mathbf{y}}_t - \mu_{Y_t}), \quad (3.39)$$

and with covariance:

$$\Sigma_{X_{T,t}|Y_t=\underline{\mathbf{y}}_t} = \Sigma_{X_{T,t}} - \Sigma_{T,t}^\dagger \Sigma_{Y_t}^{-1} \Sigma_{T,t}^*. \quad (3.40)$$

Finally, when writing $\partial_{S_t} V_t$ we obtain

$$\partial_{S_t} V_t = \partial_{\tilde{S}_t} V_t \partial_{S_t} \tilde{S}_t = \partial_{\tilde{S}_t} \mathbb{E}[e^{-\int_t^T r_s ds} H(S_T, r_T) | r_t, S_t] \frac{1}{S_t}, \quad (3.41)$$

which is equal to

$$\mathbb{E}\left[e^{-\int_t^T r_s ds} H(e^{\tilde{S}_T}, r_T) \langle \Sigma_{X_{T,t}|Y_t}^{-1} (\mathbb{W}_{T,t} - \mu_{X_{T,t}|Y_t}), \Sigma_{T,t}^\dagger \Sigma_{Y_t}^{-1} e_1^2 \rangle | r_t, \tilde{S}_t\right] \frac{1}{S_t},$$

where

$$\mathbb{W}_{T,t} = \begin{pmatrix} \bar{r}_T \\ W_t^r \\ W_T^r \\ \tilde{S}_T \\ r_T \end{pmatrix},$$

and where, the vector e_1^2 means the first basis vector in $\mathbb{R}^2$. Finally, we obtain that $\partial_{S_t} V_t$ equals

$$\mathbb{E}_t \left[e^{-\int_t^T r_s ds} \frac{H(S_T, r_T)}{S_t} \langle \Sigma_{X_{T,t}|Y_t}^{-1} (\mathbb{W}_{T,t} - \mu_{X_{T,t}|Y_t}), \Sigma_{T,t}^\dagger \Sigma_{Y_t}^{-1} e_1^2 \rangle | r_t, \tilde{S}_t \right].$$

The following remark sheds light of one limitation of the use of the Likelihood ratio method with stochastic interest rate.

Remark 5.2. While t is relatively small, the elements of $\Sigma_{Y_t} \rightarrow 0$ which causes $\Sigma_{Y_t}^{-1}$ to explode. This is briefly offset by the $\Sigma_{Y_t}^{-1}$ term inside of $\Sigma_{X_{T,t}|Y_t}$ however as $t \rightarrow T$ the whole Likelihood ratio adjustment factor explodes and surprisingly crashes to 0.

5.4 Using the likelihood ratio estimator to compute FCVA Greeks

So far we have presented a way to compute the sensitivities of an option to an asset (and more generally to a parameter of the asset). What we are interested in is to study the sensitivities of the FCVA, that is, by construction, path dependent. This change creates difficulties to apply directly the Likelihood ratio method, because we would need to define an infinite dimensional density function. Indeed, to apply the Likelihood ratio method, we need to compute the conditional density of

$$e^{-\int_t^u r_s ds} V_u, \quad \text{given } (S_t, r_t), \quad (3.42)$$

for all values of $t \in [0, T]$ and all values of $u \in [t, T]$. Indeed, in the general case, we can write V_u as a functional of $(S_x, r_x)_{x \leq u}$, i.e., we obtain

$$V_u = F((S_x, r_x)_{x \leq u}), \quad (3.43)$$

but writing the density of V_u is not possible because it would be infinite dimensional. However, given that $(r_t)_{t \in [0, T]}$ is an Ornstein-Uhlenbeck process, we can rewrite $\int_t^u r_s ds$ as a function of $(\bar{r}_u, \bar{r}_t, W_u^r, W_t^r)$. By the Markov property, we rewrite V_u as a function of (F_u, r_u) , and thus the problem reduces to a finite dimensional problem. Formally, we obtain

$$V_t = F(S_u, r_u), \quad (3.44)$$

and then the FCVA becomes

$$\text{FCVA}_t = -\mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-\int_t^u r_s ds} F(S_u, r_u) du \right]. \quad (3.45)$$

Using the results proved above in the cases (ii) and (iii), we find a conditional density:

$$f_{\int_t^u r_s ds, S_u, r_u | S_t, r_t}, \quad (3.46)$$

and we obtain

$$\text{FCVA}_t = - \int_t^T \lambda_c(s) e^{-\int_t^s \lambda_c(s) ds} \left(\int e^{-I_{u,t}} F(s, r) f_{\int_t^u r_s ds, S_u, r_u | S_t, r_t}(I_{u,t}, s, r) \right) du,$$

and then we obtain the sensitivities by differentiating the likelihood function. Finally, we obtain the terminal value by using the American Monte Carlo method. However, to use it, we need to know exactly what is the functional F , is while we do not. To bypass this issue, we will consider its estimate calculated with the American Monte Carlo estimator, i.e., we consider a proxy

$$\tilde{V}_u^N = \sum_{j \leq N} \beta_u^j L_j(S_u, r_u), \quad (3.47)$$

Now, we want to compute the Greeks related to the FCVA, defined by

$$\text{FCVA}_t = -\mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-\int_t^u r_s ds} V_u du \right], \quad (3.48)$$

where we rewrite as

$$-\mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-\int_t^u r_s ds} V_u du \right] = -\mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-I_{u,t}} V_u du \right], \quad (3.49)$$

where we work under a one-dimensional log-normal price process, and with a stochastic interest rate. To do so, we use the American Monte Carlo method to estimate by $\tilde{V}_u^N$. Then, we obtain

$$\text{FCVA}_t \approx - \sum_{j \leq N} \int_t^T \beta_u^j \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} \mathbb{E}_t \left[e^{-I_{u,t}} L_j(S_u, r_u) \right] du, \quad (3.50)$$

and so we obtain, for Δ_t

$$\partial_{S_t} \text{FCVA}_t \approx - \sum_{j \leq N} \int_t^T \beta_u^j \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} \partial_t \mathbb{E}_t \left[e^{-I_{u,t}} L_j(S_u, r_u) \right] du. \quad (3.51)$$

Thus, what remains to compute is

$$\partial_t \mathbb{E}_t [e^{-I_{u,t}} L_j(S_u, r_u)], \quad (3.52)$$

for all $u \in [t, T]$ and for all $t \in [0, T]$. We observe that this computation is very close to what has been done in the previous subsection with the computation for Δ_t of an option, but with u instead of T . More precisely, if we define the process $\mathbb{W}_{u,t}$ as

$$\mathbb{W}_{u,t} = \begin{pmatrix} \bar{r}_u \\ W_t^r \\ W_u^r \\ \tilde{S}_u \\ r_u \end{pmatrix},$$

where $\tilde{S}_t = \log(S_t)$, we obtain

$$\begin{aligned} & \partial_{S_t} \mathbb{E}_t [e^{-I_{u,t}} L_j(S_u, r_u)] = \\ & \mathbb{E}_t \left[e^{-I_{u,t}} L_j(S_u, r_u) \langle \Sigma_{X_{u,t}|Y_t}^{-1} (\mathbb{W}_{u,t} - \mu_{X_{u,t}|Y_t}), \Sigma_{u,t}^\dagger \Sigma_{Y_t}^{-1} e_1^2 \rangle \frac{1}{S_t} \right], \end{aligned}$$

where $\tilde{\Sigma}$ and $\bar{\mu}$ are defined above. Finally, we obtain $\partial_{S_t} \text{FCVA}_t \approx$

$$\begin{aligned} & - \sum_{j \leq N} \mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-I_{u,t}} \beta_u^j L_j(S_u, r_u) \times \right. \\ & \left. \langle \Sigma_{X_{u,t}|Y_t}^{-1} (\mathbb{W}_{u,t} - \mu_{X_{u,t}|Y_t}), \Sigma_{u,t}^\dagger \Sigma_{Y_t}^{-1} e_1^2 \rangle \frac{1}{S_t} \right]. \end{aligned}$$

We rewrite this equation to finally obtain:

$$\begin{aligned} & - \mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-\int_t^u r_s ds} \frac{V_u}{S_t} \times \right. \\ & \left. \langle \Sigma_{X_{u,t}|Y_t}^{-1} (\mathbb{W}_{u,t} - \mu_{X_{u,t}|Y_t}), \Sigma_{u,t}^\dagger \Sigma_{Y_t}^{-1} e_1^2 \rangle \right]. \end{aligned}$$

Now, if we consider a d -dimensional log-normal price process, with a constant interest rate, we can again follow the previous ideas to compute the Greeks of the FCVA. Indeed, we obtain:

$$\begin{aligned} & \nabla_t \text{FCVA}_t \approx \\ & - \mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-r(u-t)} V_u \left[D \left(\frac{1}{S_t^i} \right)_{i \leq d} \right] \Sigma_{u,t}^{-1} \begin{pmatrix} \log(S_u^1) - \mu_{u,t}^1 \\ \dots \\ \log(S_u^d) - \mu_{u,t}^d \end{pmatrix} du \right]. \end{aligned}$$

For a one-dimensional asset, we can easily obtain Γ_t of the FCVA, given by

$$\Gamma_t \text{FCVA}_t = -\mathbb{E}_t \left[\int_t^T \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-r(u-t)} V_u \left(-\frac{1 + (\log(S_u) - \bar{\mu}_{u,t})}{S_t^2 \bar{\sigma}_{u,t}} + \left(\frac{\log(S_u) - \bar{\mu}_{u,t}}{S_t \bar{\sigma}_{u,t}} \right)^2 \right) du \right],$$

where $\bar{\sigma}_{u,t}$ is the matrix $\Sigma_{u,t}$ in one dimension. The following Theorem presents a limitation to the use of the Likelihood ratio estimator.

Theorem 5.3. The Likelihood ratio method is inconsistent to estimate the Γ_t of the FCVA through time.

Proof. We observe that the Γ_t behaviour is erratic. The calculation of the terms inside the integral gives

$$e^{-r(u-t)} V_u \left(-\frac{1 + N(0,1) \sqrt{u-t}}{S_t^2 \bar{\sigma}_{u,t}} + \left(\frac{N(0,1) \sqrt{u-t}}{S_t \bar{\sigma}_{u,t}} \right)^2 \right). \quad (3.53)$$

We observe that the term

$$\frac{N(0,1)}{\sigma^2 \sqrt{u-t}}, \quad (3.54)$$

is integrable. Moreover, given that V_u is defined as the conditional expectation of $H(S_T)$ onto S_u , by definition of the conditional expectation, $\mathbb{E}[|V_u|] < \infty$, and then V_u is finite almost surely for all $u \in [0, T]$. By continuity of the projection onto S_u , and given that $[0, T]$ is separable, we obtain that almost surely, for all $u \in [0, T]$, V_u is finite and continuous. Given that V_u is explained through its Laguerre projection, namely

$$V_u \approx \sum_{j \leq N} \beta_u^j L_j(S_u), \quad (3.55)$$

we may assume that, for N large enough, the sum is almost surely continuous and finite. This comes from [Longstaff and Schwartz \(2001\)](#). This implies that

$$V_u \frac{N(0,1)}{\sigma^2 \sqrt{u-t}}, \quad (3.56)$$

is integrable on $[t, T]$ for all $t \in [0, T]$. The remaining terms are under the form

$$V_u \frac{N(0,1)^2}{\sigma^2 (u-t)}, \quad (3.57)$$

and

$$\frac{V_u}{\sigma^2 (u-t)}. \quad (3.58)$$

Numerically, it shows that these two terms are the ones making Γ_t inconsistent. Indeed, as u approaches t , the integral diverges. We observe this behaviour numerically, when we plot the Γ_t of the FCVA _{t} using the Likelihood ratio method. We obtain Figure 3.4, depicted below.




Figure 3.4: Γ_t of the FCVA _{t} : Likelihood Ratio method

□

This implies that the likelihood method is not useful to compute the second order derivative of the FCVA. This issue is not encountered when we consider Γ_t of an option, because the diverging terms are only whenever u approaches T , and since we know the theoretical value of Γ of the option at T , we can fix it numerically. Here, the path dependency of the FCVA prevents us to do this. To be able to compute Γ_t using the likelihood ratio, we need to consider a mixed estimator. Firstly, we compute the Δ_t of the FCVA using the likelihood ratio estimator, and then, we compute the second order derivative using the pathwise approach. More precisely, once we compute Δ_t for the FCVA, we obtain

$$\Delta_t(\text{FCVA}_t) \approx \sum_{j \leq N} \beta_t^j L_j(S_t), \quad (3.59)$$

then, we obtain Γ_t of the FCVA by taking the pathwise derivative of the Laguerre polynomials. We obtain

$$\Gamma_t(\text{FCVA}_t) \approx \sum_{j \leq N} \beta_t^j \partial_{S_t} L_j(S_t). \quad (3.60)$$

Theorem 5.3 provides another limitation of the Likelihood ratio estimator, due to the lack of numerical stability as we increase the dimension of the problem.

Theorem 5.4. Let consider a model with a stochastic interest rate and a d -dimensional price process, given by

$$dS_t^i = S_t^i (r_t dt + \sigma dW_t^i), \quad S_0^i = S_0, \quad \text{where } \langle W^i, W^j \rangle_t = \rho^{i,j} t. \quad (3.61)$$

Then, the calculation of the sensitivities becomes computationally expensive as we increase d .

Proof. We can compute Δ_t of the FCVA by following the same approach as done previously, so we obtain that $\partial_{S_t} \text{FCVA}_t$ is equal to

$$- \mathbb{E}_t \left[\int_t^T \Lambda(u) V_u \left(\left[\tilde{D} \left(\frac{1}{S_t^i} \right)_{i \leq d} \right] \Sigma_{Y_t}^{-1} \Sigma_{u,t}^* \Sigma_{X_{u,t}|Y_t}^{-1} (\mathbb{W}_{u,t} - \mu_{X_{u,t}|Y_t}) \right) | r_t, S_t \right],$$

where $\Lambda(u) := \lambda_c(u) e^{-\int_t^u \lambda_c(s) ds} e^{-\int_t^u r_s ds}$ and

$$\left[\tilde{D} \left(\frac{1}{S_t^i} \right)_{i \leq d} \right] = \begin{pmatrix} \frac{1}{S_t^1} & 0 & \dots & 0 & 0 \\ 0 & \frac{1}{S_t^2} & \dots & 0 & 0 \\ 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & \dots & \frac{1}{S_t^d} & 0 \end{pmatrix}, \quad (3.62)$$

and $\mathbb{W}_{u,t}, \mu_{X_{u,t}, Y_t}, \Sigma_{u,t}^\dagger, \Sigma_{Y_t}, \Sigma_{X_{u,t}|Y_t}$ are defined from the ones in the one-dimensional case, by increasing the dimensions properly, and taking in account the correlations between the assets. Indeed, $\Sigma_{X_{u,t}, Y_t}$ needs now to take into account the correlation between the assets and the increase of the number of assets, as does μ . That means that, if we want to increase the dimension of the asset price by one (going from dimension d to $d+1$), the covariance matrix Σ that lived in $\mathbb{R}^{(5+2d) \times (5+2d)}$, lives now in $\mathbb{R}^{(5+2(d+1)) \times (5+2(d+1))}$. $\square$

6 Method Comparison Metrics

To ensure the feasibility of the methods to determine the sensitivities using different estimation methods, diagnostic tests are performed on a standard European call option and are then compared. We are considering a single dimension asset with constant interest rate. The strike K is set to 100 as well as S_0 . The interest rate is 0.05 as well as $\sigma_s = 0.2$. We set $T = 1$, and we discretize the space $[0, T]$ into 100 timesteps. We perform the American Monte Carlo method simulating 100.000 paths using Laguerre polynomials up to degree $N = 5$.




Figure 3.5: Comparison of analytical price evolution versus the AMC estimate for varying laguerre polynomials and moneyness

Figure 3.5 presents the dynamics of the price of the European Call option, using the analytical formulation (in dotted line is depicted the mean of its value), and the estimate using the American Monte Carlo method, for different values of K and for different polynomial degrees. We observe that, as we increase the degree of the polynomials, we increase the accuracy of the American Monte Carlo method. Moreover, as we decrease the Strike and consequently increase the number of paths in the money, we are able to regress more non-zero values onto the Laguerre basis which yields more accurate results. However we would expect the divergence of the American Monte Carlo estimate to be equal over all time points but rather we observe that the beta coefficients are particularly biased as $t \rightarrow T$.




Figure 3.6: Comparison of analytical delta evolution versus the AMC estimate.

Figure 3.6 presents the evolution of Δ_t of the European Call option using two different methods (in dotted line: the analytical Δ_t of the European Call option). Left panel uses the pathwise approach, while the right panel uses the Likelihood ratio method. We observe that the Likelihood ratio method performs better than the pathwise approach, and that the pathwise jumps at $t = 0$ that the pathwise approach is volatile for small values of t .




Figure 3.7: Comparison of analytical gamma evolution versus the AMC estimate.

Figure 3.7 presents the evolution of Γ_t of the European Call option using (from left to right) the pathwise method, the Likelihood ratio method and the mixed method. We observe that all methods fit the analytical Γ_t of the option (in dotted lines), but the pathwise approach displays more volatility around the initial time. This artefact is also in the mixed method.

It must be noted that due to the artefacts at $t = 0$ caused by the ill-conditioned basis matrices in the evolution plots above, we have chosen to regress on to $t = \Delta t$ for the following experiments:




Figure 3.8: Comparison of delta at time zero for a European call: analytical versus AMC estimate.

Figure 3.8 presents the convergence plots of the Δ_0 of a European call option as a function of the number of paths simulated for (from left to right) the Bump and Revalue method, the Pathwise method and the Likelihood ratio method. We observe that both the Bump and Revalue and the Likelihood ratio method converge quickly toward the analytical solution as we increase the number of paths, while the pathwise approach still displays significant volatility.




Figure 3.9: Comparison of gamma at time zero for a European call: analytical versus AMC estimate.

Figure 3.9 presents the evolution of the Γ_0 of the European call option as a function of the number of paths simulated for (from left to right) the Bump and Reset method, the pathwise, the Likelihood ratio method and the mixed method. Similarly to the Figure 3.8, we observe a fast convergence of all the methods except the pathwise method. This explains why the mixed method seems more erratic than the Likelihood ratio method. From the above Figures, we observe that the Likelihood ratio method estimates efficiently price sensitivities of the option, while the pathwise method seems to create artefacts around the initial time.

Chapter 4

Results and Analysis

We are considering a single dimension asset with constant interest rate. The portfolio considered is composed of one standard European call option. The strike is set to 100 as well as S_0 . The interest rate is 0.05 as well as $\sigma_s = 0.2$. We set $T = 1$, and we discretize the space $[0, T]$ into 100 timesteps. We perform the American Monte Carlo method simulating 100.000 paths. We perform sensitivity estimations for the FCVA through time, using the Likelihood ratio method, the pathwise and the mixed method.

Before we begin the analysis of the sensitivities, we observe in Figure 4.1 that the FCVA estimate though time is negative as one would expect on a long position on a European Call Option, additionally it tends to 0 as $t \rightarrow T$ which is sensible considering the integral is being computed over increasingly smaller ranges.




Figure 4.1: Evolution of the $FCVA_t$




Figure 4.2: Δ_t of the FCVA_t : Pathwise method




Figure 4.3: Δ_t of the FCVA_t : Likelihood Ratio method

Figure 4.2 and 4.3 present the evolution of Δ_t for the FCVA_t through time. We observe that both methods fall in the same range which is promising as we do not have an analytical solution to compare them to. Figure 4.2 displays an artefact at $t = 0$ and then jumps to a more consistent value. We may be able to attribute this to the numerical instability caused from regressing onto a matrix with linearly dependent rows at $t = 0$.

As expected, Delta of the FCVA tends to 0 as $t \rightarrow T$ as the integral tends to 0.




Figure 4.4: Γ_t of the FCVA $_t$: Pathwise method




Figure 4.5: Γ_t of the FCVA $_t$: Mixed method

Figure 4.2 and 4.3 present the evolution of Γ_t for the FCVA $_t$ through time. Once again we observe that both methods fall in the same range. Both Figures display an artefact at $t = 0$. We can attribute this to the ill-conditioned regression matrix at $t = 0$. As expected, Gamma of the FCVA tends to 0 as $t \rightarrow T$ as the integral tends to 0.

Chapter 5

Conclusion

The main goal of this report was to present three methods (the bump and revalue, the pathwise, and the Likelihood ratio) for estimating the sensitivities for the FCVA through time, using the AMC estimation. Firstly, we observed that the American Monte Carlo method, despite its simple approach, was adequately efficient in computing the required metrics

Moreover, we observed that the Likelihood ratio method works efficiently for the first order sensitivities calculation.

Furthermore, we highlighted that this method depends strongly on the dimension of the problem: It becomes unusable as we increase the dimension of the asset price, because of the joint matrices calculations and inversion.

Moreover, we observed that the method is also unusable to compute the second order sensitivities, as singularities appear. On the other hand, even though the structure given by the use of the AMC method seemed favourable towards to the use of the pathwise approach, we observed that this method was unstable and leads to erratic behaviour around the initial time.

1 Opportunities for Further Research

As a further study, it could be interesting to expand the range of xVAs studied, such as FVA. Moreover, we could consider more general models, such as the Libor Market Model, using more latent factors.

Regarding the implementation, we could compare the AMC method with the use of Neural Networks.

To reduce the dimension dependency of the Likelihood ratio method, we could consider the process $I_{u,t}$ himself instead of rewriting him into a few processes. It boils down the dimension of the conditional law.

Bibliography

- Buis, K., 2023. Computational Challenges in Risk Calculations for XVA: Likelihood Ratio Method to Compute XVA Greeks for OTC Derivatives. Master's thesis. Delft University of Technology. Faculty of Electrical Engineering, Mathematics & Computer Science. Supervisors: Prof.dr.ir. C. Vuik, Dr. F. Fang, J. van der Linden, S. Borst.
- Capriotti, L., Jiang, Y., Macrina, A., 2015. Real-time risk management: An aad-pde approach. *International Journal of Financial Engineering* 2, 1550039–1550070. doi:[10.1142/S2424786315500395](https://doi.org/10.1142/S2424786315500395), preprint posted July 15, 2015 (SSRN 2630328); date of writing: June 22, 2015.
- Carmona, R., 2016. Lectures on BSDEs, stochastic control, and stochastic differential games with financial applications. SIAM.
- Cohen, S.N., Elliott, R.J., 2015. Stochastic calculus and applications. volume 2. Springer.
- Conti, F., Palazzi, R., Pioppi, M., 2022. Xva greeks via likelihood ratio method. URL: <https://www.linkedin.com/in/francescoconti82/>. working paper.
- Glasserman, P., 2004. Monte Carlo Methods in Financial Engineering. volume 53 of *Stochastic Modelling and Applied Probability*. Springer, New York.
- Gottesman, A., 2016. Derivatives Essentials: An Introduction to Forwards, Futures, Options and Swaps. 1st ed., Wiley, Hoboken, NJ.
- Gregory, J., 2018. The xVA Challenge: Counterparty Credit Risk, Funding, Collateral and Capital. 3rd ed., Wiley, Chichester, UK.
- Kjaer, M., 2017. Consistent xva metrics part i: Single-currency. arXiv preprint arXiv:2506.18521 URL: <https://arxiv.org/pdf/2506.18521>.
- Longstaff, F.A., Schwartz, E.S., 2001. Valuing american options by simulation: a simple least-squares approach. *The Review of Financial Studies* 14, 113–147. doi:[10.1093/rfs/14.1.113](https://doi.org/10.1093/rfs/14.1.113).

Appendix A

Covariance Matrix

1 Joint Covariance Matrix of stochastic processes

Element	Value	Element	Value
$\Sigma_{X,Y}^{(1,1)}$	$\frac{\sigma_r^2(1 - e^{-2\kappa T})}{2\kappa}$	$\Sigma_{X,Y}^{(4,4)}$	(1)
$\Sigma_{X,Y}^{(1,2)}$	$\frac{\sigma_r}{\kappa}(e^{-\kappa(T-t)} - e^{-\kappa T})$	$\Sigma_{X,Y}^{(4,5)}$	(8)
$\Sigma_{X,Y}^{(1,3)}$	$\frac{\sigma_r^2}{\kappa}(1 - e^{-\kappa T})$	$\Sigma_{X,Y}^{(4,6)}$	(9)
$\Sigma_{X,Y}^{(1,4)}$	(3)	$\Sigma_{X,Y}^{(4,7)}$	(10)
$\Sigma_{X,Y}^{(1,5)}$	$\frac{\sigma_r^2}{2\kappa}(1 - e^{-2\kappa T})$	$\Sigma_{X,Y}^{(5,5)}$	$\frac{\sigma_r^2(1 - e^{-2\kappa T})}{2\kappa}$
$\Sigma_{X,Y}^{(1,6)}$	(4)	$\Sigma_{X,Y}^{(5,6)}$	(11)
$\Sigma_{X,Y}^{(1,7)}$	$\frac{\sigma_r^2}{2\kappa}(1 - e^{-2\kappa t})$	$\Sigma_{X,Y}^{(5,7)}$	$\frac{\sigma_r^2}{2\kappa}e^{-\kappa(T+t)}(e^{2\kappa t} - 1)$
$\Sigma_{X,Y}^{(2,2)}$	T	$\Sigma_{X,Y}^{(6,6)}$	(2)
$\Sigma_{X,Y}^{(2,3)}$	t	$\Sigma_{X,Y}^{(6,7)}$	(12)
$\Sigma_{X,Y}^{(2,4)}$	(5)	$\Sigma_{X,Y}^{(7,7)}$	$\frac{\sigma_r^2(1 - e^{-2\kappa t})}{2\kappa}$
$\Sigma_{X,Y}^{(2,5)}$	$\frac{\sigma_r}{\kappa}(e^{-\kappa(T-t)} - e^{-\kappa T})$	$\Sigma_{X,Y}^{(2,6)}$	(6)
$\Sigma_{X,Y}^{(2,7)}$	$\frac{\sigma_r}{\kappa}(1 - e^{-\kappa t})$	$\Sigma_{X,Y}^{(3,3)}$	t
$\Sigma_{X,Y}^{(3,4)}$	$\frac{\sigma_r}{\kappa}T - \frac{\sigma_r}{2\kappa}(1 - e^{-\kappa T}) + \sigma_s \rho T$	$\Sigma_{X,Y}^{(3,5)}$	$\frac{\sigma_r}{\kappa}(1 - e^{-\kappa T})$
$\Sigma_{X,Y}^{(3,6)}$	(7)	$\Sigma_{X,Y}^{(3,7)}$	$\frac{\sigma_r}{\kappa}(1 - e^{-\kappa t})$

Formula:

$$\begin{aligned}
(1) \quad \Sigma_{X,Y}^{(4,4)} &= \frac{T\sigma_r^2}{\kappa^2} + \sigma_s^2 T + \frac{2\sigma_r\sigma_s\rho T}{\kappa} - \frac{2\rho\sigma_s\sigma_r(1-e^{-\kappa T})}{\kappa^2} - \frac{2\sigma_r^2(1-e^{-\kappa T})}{\kappa^3} + \frac{\sigma_r^2(1-e^{-2\kappa T})}{2\kappa^3} \\
(2) \quad \Sigma_{X,Y}^{(6,6)} &= \frac{t\sigma_r^2}{\kappa^2} + \sigma_s^2 t + \frac{2\sigma_r\sigma_s\rho t}{\kappa} - \frac{2\rho\sigma_s\sigma_r(1-e^{-\kappa t})}{\kappa^2} - \frac{2\sigma_r^2(1-e^{-\kappa t})}{\kappa^3} + \frac{\sigma_r^2(1-e^{-2\kappa t})}{2\kappa^3} \\
(3) \quad \Sigma_{X,Y}^{(1,4)} &= \frac{\sigma_r^2}{\kappa^2}(1-e^{-\kappa T}) - \frac{\sigma_r^2}{2\kappa^2}(1-e^{-2\kappa T}) + \frac{\sigma_r\rho\sigma_s}{\kappa}(1-e^{-\kappa T}) \\
(4) \quad \Sigma_{X,Y}^{(1,6)} &= \frac{\sigma_r^2}{\kappa^2}(e^{-\kappa(T-t)} - e^{-\kappa T}) - \frac{\sigma_r^2}{2\kappa^2}e^{-\kappa(T+t)}(e^{2\kappa t} - 1) + \frac{\sigma_r\sigma_s\rho}{\kappa}(e^{-\kappa(T-t)} - e^{-\kappa T}) \\
(5) \quad \Sigma_{X,Y}^{(2,4)} &= \frac{\sigma_r}{\kappa}t + \sigma_r\sigma_s\rho t - \frac{\sigma_r}{\kappa^2}(e^{-\kappa(T-t)} - e^{-\kappa T}) \\
(6) \quad \Sigma_{X,Y}^{(2,6)} &= \frac{\sigma_r}{\kappa}t + \frac{\sigma_r(1-e^{-\kappa t})}{\kappa^2} + \sigma_s\rho t \\
(7) \quad \Sigma_{X,Y}^{(3,6)} &= \frac{\sigma_r}{\kappa}t - \frac{\sigma_r}{\kappa^2}(1-e^{-\kappa t}) + \sigma_s\rho t \\
(8) \quad \Sigma_{X,Y}^{(4,5)} &= \frac{\sigma_r^2(1-e^{-\kappa T})}{\kappa^2} + \frac{\sigma_r\sigma_s\rho(1-e^{-\kappa T})}{\kappa} - \frac{\sigma_r^2(1-e^{-2\kappa T})}{2\kappa^2} \\
(9) \quad \Sigma_{X,Y}^{(4,6)} &= \left(\frac{\sigma_r}{\kappa}\right)^2 t - \frac{\sigma_r^2}{\kappa^3}(1-e^{-\kappa t}) + \frac{\sigma_r\sigma_s\rho t}{\kappa} + \frac{\sigma_r^2}{\kappa^3}(e^{-\kappa(T-t)} - e^{-\kappa T}) \\
&\quad - \frac{\sigma_r^2}{2\kappa^3}e^{-\kappa(T+t)}(e^{2\kappa t} - 1) - \frac{\sigma_r\sigma_s\rho}{\kappa^2}(e^{-\kappa(T-t)} - e^{-\kappa T}) + \frac{\sigma_r\sigma_s\rho t}{\kappa} - \frac{\sigma_r\sigma_s\rho}{\kappa^2}(1-e^{-\kappa t}) + \sigma_s^2 t \\
(10) \quad \Sigma_{X,Y}^{(4,7)} &= \frac{\sigma_r^2}{\kappa^2}(1-e^{-\kappa t}) - \frac{\sigma_r^2(e^{-\kappa(T-t)} - e^{-\kappa(T+t)})}{2\kappa^2} + \frac{\sigma_s\sigma_r\rho}{\kappa}(1-e^{-\kappa t}) \\
(11) \quad \Sigma_{X,Y}^{(5,6)} &= \frac{\sigma_r^2 e^{-\kappa T}(e^{\kappa t} - 1)}{\kappa^2} + \frac{\sigma_s\sigma_r\rho e^{-\kappa T}(e^{\kappa t} - 1)}{\kappa} - \frac{\sigma_r^2 e^{-\kappa(T+t)}(e^{2\kappa t} - 1)}{2\kappa} \\
(12) \quad \Sigma_{X,Y}^{(6,7)} &= \frac{\sigma_r^2(1-e^{-\kappa t})}{\kappa^2} - \frac{\sigma_r^2(1-e^{-2\kappa t})}{2\kappa^2} + \frac{\rho\sigma_r\sigma_s(1-e^{-\kappa t})}{\kappa}
\end{aligned}$$