

ABSTRACT

In the world of financial markets each day bullish and bearish investors determine with their trading activity in opposite directions the equilibrium price at which a security is exchanged. This dynamic interaction generally determines a fair market price under the standard law of supply and demand.

However, the presence of concerted trading activity can significantly distort this equilibrium, leading traders to experience higher than optimal execution costs.

A practical example of this phenomenon is the case of an institutional investor executing orders of large volume for a single security. Due to the substantial size of these trades, the price of the security may exert a direct and noticeable impact, especially in illiquid markets, where the market impact on the security price tends to persist over time.

The effectiveness of a trading strategy hinges on its adaptability to competitive market conditions, particularly distinguishing between liquid and illiquid market environments. In this context, the need for a tailored order execution strategy for each scenario becomes essential to reduce market impact.

Therefore, the primary objective of this thesis is to address the problem of optimal trading strategies under competition, aiming to achieve the most cost-efficient trading strategy for each market condition.

In the first chapter a mathematical framework developed by mathematician and hedge fund manager Neil A. Chriss will be presented and analyzed in detail.

The second chapter will delve into standard trading approaches such as moving averages and the volume weighted average price VWAP, the latter is commonly used as a benchmark by institutional investors for order execution.

The third chapter, the core of this thesis, will analyze and evaluate for cost efficiency the optimal trading strategies obtained by the Chriss model and by the VWAP approach in both liquid and illiquid scenarios.

The VWAP will be performed both by Monte Carlo simulation method and by using historical data, on daily trading volumes and prices for US equity securities, obtained from the Bloomberg Terminal.

All the research part, necessary to obtain these results, that are pivotal of this work, was done using the software R-Studio.

(The coding and programming .R-files can be provided, upon request)

TABLE OF CONTENTS

1. Introduction of the Chriss model.....	3
1.1 Preliminaries.....	3
1.2 Market impact.....	5
1.3 Total cost functions & Sensitivity.....	9
1.4 <i>K-regimes</i>	14
1.5 Passive vs Active optimization.....	16
1.6 Chriss model with volatility.....	32
1.7 Summary of main theoretical results.....	35
 2. Standard trading approaches & Benchmarks.....	 38
2.1 Simple moving averages.....	38
2.2 VWAP approach.....	43
2.3 Summary of main differences.....	47
2.4 VWAP Total cost function for <i>k-regimes</i>	48
 3. Performance analysis for cost efficiency.....	 52
3.1 Cost functions time discretization.....	53
3.2 Liquid vs Illiquid scenarios with Monte Carlo method.....	55
3.3 Limit case study as N traders $\rightarrow +\infty$	62
3.4 Liquid vs Illiquid scenarios with Historical data.....	64
 Conclusions.....	 67
Annex.....	68
References.....	69

1. Introduction of the Chriss model

1.1 Preliminaries

Notation: Traders are indicated with capital letters A,B,C.....etc. instead trading strategies are indicated with lowercase letters a,b,c.....etc.

$\dot{a} = \frac{da}{dt}$ the first derivative of the trading strategy w.r.t time

$\ddot{a} = \frac{d^2a}{dt^2}$ the second derivative of the trading strategy w.r.t time

Definition of a Trading strategy:

A trading strategy $x(t)$ is a C^2 (continuous and twice differentiable) function of time, that describes the units of stocks held by a trader at each time t , between the starting time $t = 0$ to the end time $t = 1$.

There are several important types of trading strategies and here is non-exhaustive list for a strategy $x(t)$:

- **Liquidation:** Strategies for which $x(0) > 0$ and $x(1) = 0$, in other words a strategy that starts with a positive quantity of stock and ends with none.
- **Position-building:** Strategies for which $x(0) = 0$ and $x(1) > 0$.
- **Unit:** When a trader seeks to acquire a single unit of stock, $\lambda = 1$.
- **λ -Scaled:** Position-building strategies whose rate of trading is scaled by a constant $\lambda > 1$ at each time $t \in [0,1]$.




Figure 1: Example of a trading strategy $y(t) = t^2$

Note that scaled trading strategies represent strategies that have the "shape" of unit position strategies, but which are scaled at each time t by a fixed constant $\lambda > 1$.

$$y(t) = \lambda \frac{y(t)}{\lambda} = \lambda b(t)$$

This means that for a λ -scaled strategy $b(t)$, its trajectory is $\lambda \cdot b(t)$, while its shape is given by $b(t)$.

Definition of a Trading Trajectory:

A trading trajectory is the path that a trading strategy takes from the starting time $x(0)$ to the end time $x(1)$. It can be graphed on a two-dimensional space and can assume different shapes.

There are three basic shapes all of which have constant sign of second derivative:

- ***Risk-neutral:*** Strategies for which $\ddot{x}(t) = 0$ for all $t \in [0,1]$ this represents a straight line of the form $x(t) = x(0) + \lambda t$. Given the starting condition $x(0) = 0$, then a risk neutral λ -Scaled trading strategy is of the form $x(t) = 0 + \lambda t$ for some $\lambda > 0$. This can be referred as a passive order execution benchmark.
- ***Risk-averse:*** Strategies for which $\ddot{x}(t) > 0$ for all $t \in [0,1]$ this is a convex curve. For example, the function $x(t) = \lambda t^2$ passing through the interval $t \in [0,1]$
- ***Eager:*** Strategies for which $\ddot{x}(t) < 0$ for all $t \in [0,1]$ this is a concave curve. For example, the function $x(t) = \lambda \sqrt{t}$ passing through the interval $t \in [0,1]$




Figure 2: A visual representation of the three trading trajectories for risk preferences. The blue line is a 45° sloped line $x(t) = t$, with $\lambda = 1$ it is considered the risk neutral passive execution benchmark. Both eager (red curve) and risk-averse (green curve) belong to the family of functions $x(t) = t^\alpha$ with $0 < \alpha < 1$ and par for the eager (concave curve); while $\alpha > 1$ and par for the risk-averse (convex curve).

- **Bucket:** bucket strategies acquire more than their target quantity immediately after the start time and then sell down to their target quantity as close to the completion time as possible.
- **Barbell:** a barbell strategy buys a portion of its target quantity at the very start of trading and the remaining amount very close to the end of trading.




Figure 3: Example of bucket and barbell strategies.

1.2 Market impact

In financial markets, asset prices are influenced by the continuous interaction of market participants placing trading orders. Formally, the trading activity can be divided into two broad categories:

- **Noise Trading:** This represents background trading activity by a large number of participants. It is assumed to occur persistently and randomly, with no net impact on the price of the asset.

- **Concerted Trading:** This refers to trading by one or more agents in a deliberate and directional manner (e.g., continuously buying or selling over a short time window), which is expected to move the price in the direction of trading.

Building on this clear differentiation between noise trading and concerted trading, the core challenge faced when implementing an optimal trading strategy in a competitive setting can be introduced.

Definition of Market impact:

The price deviation directly caused by trading activity that influences both **execution costs** and **overall market dynamics**. This phenomenon arises not only in illiquid environments but also in highly liquid and competitive markets, where factors such as trading speed, order size, and informational asymmetries can significantly amplify its effects.

The main components of the market impact:

Definition of Temporary Market Impact:

It is the immediate price deviation caused by a trader's execution rate at a given point in time. It affects the execution price only **momentarily** and **does not persist** beyond the instant of the trade. Therefore, it is only a short run phenomenon, the asset's price will revert back to its long run mean equilibrium level.

Economic Interpretation: Temporary impact reflects the premium paid for immediacy in execution, the liquidity premium.

Without Competition (Trader A only): Let $\dot{a}(t)$ be the trading rate of trader A at time t . Then the instantaneous temporary cost is proportional to $\dot{a}(t)$ and

Temporary Cost for trader A = $\dot{a}^2(t)$

With Competition (Trader A and Trader B trading concurrently): Let $\dot{b}(t)$ be the trading rate of trader B. Then the instantaneous temporary cost for trader A becomes: **Temporary Cost for trader A = $[\dot{a}(t) + \dot{b}(t)] \dot{a}(t)$**

Definition of Permanent Market Impact:

It is the lasting long run price movement caused by the cumulative trading activity over time. It reflects the sustained pressure that trading places on the market, and it persists as long as the position is being built.

Economic Interpretation: Permanent impact represents the market's adjustment to persistent demand or supply changes.

Without Competition: Let $a(t)$ be the cumulative position of trader A at time t . Then the instantaneous permanent cost is proportional to $a(t)$ and

Permanent Cost for trader A = $a(t) \dot{a}(t)$

With Competition: Let $b(t)$ be the cumulative position of trader B. Then the instantaneous permanent cost for trader A becomes:

Permanent Cost for trader A = $[a(t) + b(t)] \dot{a}(t)$

After defining the precise mathematical expressions for the market impact cost components, both in competition and in absence of it; it is important to understand from a practical perspective, what are the main real world scenarios in which market impact could play a significant role. Two distinct scenarios come to light; yet, despite their contrasting nature, both are characterized by a pronounced market impact on the asset's price driven by trading activity.

Scenario: 1 Institutional Investor Trading

An institutional investor trading a large quantity of a given asset may significantly affect its price. This can occur in an illiquid market, such as certain segments of the blockchain space where large trades may encounter wide bid-ask spreads and limited volume on the order book. For example, executing a large position in a lesser-known cryptocurrency can significantly and persistently impact its price, irrespective of underlying fundamentals.

Alternatively, even within a liquid market, an institutional trader may target an illiquid asset, such as a thinly traded corporate bond or an equity with low float. The attempt to exit a large position rapidly as in the case of a dumping can have a disproportionate and lasting impact on the asset market price, distorting it far from any fair value benchmark.

In both cases, the concentration of trading activity whether due to market structure or supply dominance produces both permanent and temporary market impact, necessitating sophisticated trading strategies to minimize execution costs and slippage.

Scenario: 2 Coordinated Retail Trading and Behavioral Dynamics

In contrast to institutional dominance, **large-scale market impact** can also result from the **coordinated actions of numerous small traders**. These behaviors are often studied within the framework of **Behavioral Finance**, which examines how psychological and social forces influence trading decisions. Two primary examples of coordinated retail behavior include:

- **Flight to Safety:** During periods of macroeconomic uncertainty or crisis, investors collectively abandon risky assets in favor of perceived "safe havens" such as U.S. Treasury bonds or gold. A clear example occurred in March 2020, as the outbreak of COVID-19 triggered panic selling in equity markets and a simultaneous surge in demand for Treasuries. Though no single trader moved the market, the aggregated reaction of many small players caused significant asset price shifts.
- **Sentiment Trading:** This occurs when investor behavior is driven by emotion, crowd psychology, or viral social media narratives rather than fundamental analysis. A notable case is the **GameStop (GME)** short squeeze in early 2021. Coordinated through social platforms like Twitter-WallStreetBets, retail investors collectively bought shares and call options, driving the price from under \$20 to nearly \$500 in a matter of weeks. Causing short-positioned hedge fund Melvin Capital to lose over \$6.8 billion and require a \$2.75 billion bailout credit from Citadel. This episode outstandingly remarks how **sentiment driven coordination can generate extreme market impact, rivaling or even exceeding that of large institutional trades**.

These examples underscore that market impact is not solely a function of single trade size, but also of timing, coordination, market structure, and investor psychology. Understanding in detail both temporary and permanent components of the market impact is therefore critical when designing optimal trading strategies, particularly when trading in highly competitive or unstable environments. In such scenarios, each trader must navigate not only the effects of their own trades but also the interactive impact caused by the actions of others. **The goal is to minimize the total cost of trading compared to a hypothetical unperturbed price path i.e., the price that would have prevailed in the absence of concerted trading.**

1.3 Total cost functions & Sensitivity

Recall the mathematical expressions introduced in section 1.2 for the market impact components for trader A, competing with trader B we have:

$$\text{Temporary Cost for trader A} = [\dot{a}(t) + \dot{b}(t)] \dot{a}(t)$$

$$\text{Permanent Cost for trader A} = [a(t) + b(t)] \dot{a}(t)$$

Starting from these expressions, it is finally possible to derive a unique mathematical expression for the total cost of trading which incorporates both the temporary and permanent market impact components, under the influence of the predominant *k-regime*, in competition.

Instantaneous Total Cost of Trading for trader A in competition with trader B:

$$\propto \underbrace{[\dot{a}(t) + \dot{b}(t)] \dot{a}(t)}_{\text{Temporary component}} + k \underbrace{[a(t) + b(t)] \dot{a}(t)}_{\text{Permanent component}}$$

With k called market impact coefficient: $k > 0$ always, more details on this key coefficient will be explained in the next section: 1.4 *K-regimes*.

The same reasoning holds for trader B, the Instantaneous Total Cost of Trading for trader B is proportional to:

$$[\dot{a}(t) + \dot{b}(t)] \dot{b}(t) + k [a(t) + b(t)] \dot{b}(t)$$

Focusing on the temporary market impact component, it is useful to understand the impact that the trading direction of A and B, sign of first derivatives w.r.t. time $\dot{a}(t)$ and $\dot{b}(t)$ have on the temporary component of the total cost function of A:

A action	B action	A sign	A + B sign	A Impact Sign	A Impact
A buy	B buy	$\dot{a} > 0$	$\dot{a} + \dot{b} > 0$	$(\dot{a} + \dot{b})\dot{a} > 0$	Cost
A sell	B sell	$\dot{a} < 0$	$\dot{a} + \dot{b} < 0$	$(\dot{a} + \dot{b})\dot{a} < 0$	Profit
A sell	B buy	$\dot{a} < 0$	$\dot{a} + \dot{b} > 0$	$(\dot{a} + \dot{b})\dot{a} > 0$	Cost
A buy	B sell	$\dot{a} > 0$	$\dot{a} + \dot{b} < 0$	$(\dot{a} + \dot{b})\dot{a} < 0$	Profit

Impact of Actions on Sign

Then, the average cost of trading for the trading window that starts at $t = 0$ and ends at $t = 1$, for trader A in competition with trader B, will be the integral of the instantaneous total cost of trading expression. Assuming both traders use a unit trading strategy:

$$\text{Average Total cost} = \int_0^1 [\dot{a}(t) + \dot{b}(t)] \dot{a}(t) + k [a(t) + b(t)] \dot{a}(t) dt$$

If instead trader B uses a λ -scaled strategy, having $\lambda > 0$ then the expression is modified as in the next expression:

Average Total cost for trader A: C_A

$$C(a_\lambda; b_\lambda; k; \lambda) = \int_0^1 [\dot{a}_\lambda(t) + \lambda \dot{b}_\lambda(t)] \dot{a}_\lambda(t) + k [a_\lambda(t) + \lambda b_\lambda(t)] \dot{a}_\lambda(t) dt$$

The same reasoning holds for trader B, *Average Total cost for trader B: C_B*

$$C(b_\lambda; a_\lambda; k; \lambda) = \int_0^1 [\dot{a}_\lambda(t) + \lambda \dot{b}_\lambda(t)] \dot{b}_\lambda(t) + k [a_\lambda(t) + \lambda b_\lambda(t)] \dot{b}_\lambda(t) dt$$

Considering now S as the set of all trading strategies which follow the mathematical properties described before.

$$S = \{a: [0,1] \rightarrow R \mid a(0) = 0, a(1) = 1, \ddot{a} \text{ exist}\} \subseteq C^2 \{[0,1]\}$$

Now assuming to assign the set S with a probability measure m , then the expected cost of trading for trader A with respect to the measure m as:

$$\begin{aligned} E[C|S] &= \int_{b_\lambda \in S} C(a_\lambda; b_\lambda; k; \lambda) dm(\varpi) \\ &= \int_{b_\lambda \in S} dm \int_0^1 [\dot{a}_\lambda(t) + \lambda \dot{b}_\lambda(t)] \dot{a}_\lambda(t) + k [a_\lambda(t) + \lambda b_\lambda(t)] \dot{a}_\lambda(t) dt \end{aligned}$$

With: $b_\lambda = \varpi$ then $dm(\varpi) = dm(b_\lambda)$

And the corresponding variance:

$$\text{Var}[C|S] = E[C^2|S] - E[C|S]^2$$

Therefore, since total costs can be modeled as a random variable, it can be useful to perform a sensitivity analysis on the average total cost expressions. This will give insights on how structurally total costs reacts to changes in k (market impact) and λ (size)

$$\text{Sensitivity of trader A total cost of trading for } \lambda = \frac{dC(a; b, \lambda, k)}{d\lambda}$$

$$\text{Sensitivity of trader A total cost of trading for } k = \frac{dC(a; b, \lambda, k)}{dk}$$

$$\text{Sensitivity of trader B total cost of trading for } \lambda = \frac{dC(b; a, \lambda, k)}{d\lambda}$$

$$\text{Sensitivity of trader B total cost of trading for } k = \frac{dC(b; a, \lambda, k)}{dk}$$

The following results, summarized on tables 1 and 2, are taken from Neil A. Chriss' Research Paper titled "Optimal Position-Building Strategies in Competition"

λ	κ	Temp Impact	Perm Impact	Total Impact	Total Diff
5	0.1	2.01	0.10	2.10	
5	0.075	2.00	0.07	2.08	-1.28%
5	0.5	2.17	0.43	2.61	
5	0.375	2.10	0.34	2.44	-6.60%
5	1	2.69	0.74	3.43	
5	0.75	2.39	0.60	2.99	-12.68%
5	5	14.52	0.28	14.79	
5	3.75	9.97	0.73	10.70	-27.68%
5	25	84.99	-4.99	79.99	
5	18.75	63.66	-3.71	59.96	-25.05%
5	100	340.00	-20.00	320.00	
5	75	255.00	-15.00	240.00	-25.00%

Table 1: Pairs of market impact costs for the linear unit strategy $a(t)$, in two-trader equilibrium strategies for $\lambda = 5$ and various κ values. Each pair of rows shows the temporary, permanent and total market impact costs for strategy $a(t)$ that differ only in a 25% shift in the value of κ .

λ	κ	Temp Impact	Perm Impact	Total Impact	Total Diff
25	0.1	2.22	0.09	2.30	
25	0.075	2.12	0.07	2.19	-4.96%
25	0.5	7.43	0.15	7.58	
25	0.375	5.06	0.18	5.24	-30.90%
25	1	23.42	-0.36	23.06	
25	0.75	14.14	-0.02	14.12	-38.75%
25	5	388.63	-19.57	369.06	
25	3.75	248.94	-11.94	236.99	-35.78%
25	25	2459.70	-130.96	2328.80	
25	18.75	1841.70	-98.02	1743.70	-25.12%
25	100	9841.30	-524.00	9317.30	
25	75	7381.00	-393.00	6988.00	-25.00%

Table 2: Pairs of market impact costs for the linear unit strategy $a(t)$, in two-trader equilibrium strategies for $\lambda = 25$ and various κ values. Each pair of rows shows the temporary, permanent and total market impact costs for strategy $a(t)$ that differ only in a 25% shift in the value of κ .

These results lead to several important conclusions. In both tables, as predicted by theory, an increase in the market impact coefficient k results in higher total costs. However, the total difference in each groups tends to stabilize to a -25% for large levels of k . This is a significant empirical finding: it indicates that when market impact becomes more pronounced (i.e., at high k values), the Chriss optimization model is able to effectively stabilize average total trading costs arising from market impact and order size.

In other words, **the Chriss model's total cost structure exhibits lower sensitivity to estimation errors in k when operating in a high k -regime. This suggests a form of robustness in more illiquid markets or in other high market impact scenarios.**

Moreover, when analyzing subsets within the same k level for example for $k = (0.375, 0.5)$ but with different trade sizes λ (i.e., $\lambda = 5$ in table 1 versus $\lambda = 25$ in table 2), the percentage difference in total costs becomes significantly larger for the higher trading size. Specifically, the cost difference is -30.90% for $\lambda = 25$, compared to just -6.60% for $\lambda = 5$.

This confirms that the model's sensitivity to errors in the market impact parameter k increases with order size λ . However, this increased sensitivity is primarily observed in low k -regimes; as previously discussed, the Chriss model demonstrates greater cost stability and reduced sensitivity under high k -regimes.

1.4 *K-regimes*

Introduction

A significant portion of the analysis in this thesis focuses on the structure of the cost functions introduced earlier and their influence on the form of optimal trading strategies under competitive conditions. Central to this investigation is the parameter k , which appears in the Chriss model total trading cost expressions and governs the relative weight of permanent versus temporary market impact.

The k coefficient encapsulates the proportion of execution cost attributable to permanent market impact, as opposed to temporary impact.

It emerges as a critical parameter in the optimization framework for trading in competition. Notably, in the context of optimal execution problems against a passive risk neutral trader, in other words those without competition, permanent impact is considered negligible or irrelevant and thus omitted from the optimization process altogether. A mathematical proof of this assertion will be presented later on in section 1.5 Passive vs Active optimization.

However, when multiple active traders are simultaneously building positions in the same asset, the role of permanent market impact becomes not only relevant but central to strategic decision-making.

In order to understand its influence more clearly, it is useful to introduce the concept of *k-regimes*, which classify the market environment based on the dominance of temporary or permanent impact in the total cost function.

These regimes provide an intuitive and formal lens through which to interpret how variations in k , shape optimal trading behavior.

Definition of k-regimes

K-regimes describe how the cumulative trading cost is distributed between its temporary and permanent components:

- When $k < 1$, particularly in the limit where $0 < k \ll 1$, the total trading cost is **dominated by the temporary market impact component**. In such settings, traders prioritize minimizing the costs associated with aggressive execution, the instantaneous price pressure caused by rapid trading. As a result, optimal strategies tend to resemble a **risk-neutral execution paths, spreading the volume of trades more evenly over the trading window**, and paying less attention to the future price trajectory caused by accumulated position.
- Conversely, when $k > 1$, and especially when $k \gg 1$, the cost structure becomes **dominated by permanent market impact**. In this regime, the trader's concern shifts toward the long-term consequences of their own and their competitors' cumulative order flow. Since each additional unit traded contributes to a persistent long-lasting shift in price, **it becomes strategically advantageous to trade ahead of others, "First mover advantage"** executing larger volumes earlier to avoid adverse price moves caused by aggregate demand. The resulting trading profiles are characteristically **eager, front-loaded**, and exhibit a more pronounced deviation from the risk-neutral baseline.

This crucial distinction between *k-regimes* not only completes the Chriss framework modeling capability for optimal trading execution in competitive contexts; but also emphasizes how the balance between temporary and permanent impact fundamentally alters the incentives driving a trader behavior. The shape and efficiency of a trading strategy are tightly linked to the prevailing *k-regime*, and optimizing accordingly can yield substantial cost savings over naïve or benchmark-based approaches, especially under the previously described scenarios number 1 and 2

1.5 Passive vs Active optimization

Introduction

In this section the mechanism of the optimization process under the Chriss model will be presented. Moreover, this section will start with the simple two-trader equilibrium without competition, then it will arrive at the most complex case which is the multi-trader equilibrium in competition.

The main idea behind the Chriss optimization model is to find the strategy which minimizes the total cost of trading, driven by both temporary and permanent market impact components.

Two-trader Chriss model optimization without competition

In this situation two traders A and B are trading without considering the impact that the other trader's orders will have on the asset price. Therefore, both traders assume that the market impact is only caused by their own trading activity and that the opponent party trading activity is negligible in terms of market impact.

This situation can also be called passive optimization problem, since the optimal trading solution is obtained against risk neutral trading strategy (unitary $\lambda = 1$ or λ -scaled $\lambda > 1$), a passive benchmark.

Moreover, with its trading activity the large institutional investor is expected to impact the price of the asset mainly through the temporary cost component, in a liquid market.

Recalling the mathematical expressions for the market impact components, under no competition, from section 1.2

Permanent Cost for trader A = $a(t) \dot{a}(t)$

Temporary Cost for trader A = $\dot{a}^2(t)$

It will now be proven mathematically that, under the Chriss model optimization process applied to this passive scenario, the permanent cost component of the market impact vanishes.

Statement:

Given the set of all eligible trading strategies, set of possible solutions:

$$S = \{a: [0,1] \rightarrow R | a(0) = 0, a(1) = 1, \ddot{a} \text{ exist}\}$$

$$\begin{cases} \min_{a \in S} \text{ of } \{L_A\} = \min \{ \dot{a}^2(t) + k a(t) \dot{a}(t) + \lambda \sigma^2 \dot{a}^2(t) \} \\ \text{s.t. } a(0) = 0 \text{ and } a(1) = 1 \end{cases}$$

The Chriss model optimization applied to a passive scenario has as solution $a \in S$, the trading strategies of the form:

$$a(t) = \frac{\sinh(\sigma t)}{\sinh(\sigma)}; \text{ with } \sigma = \sqrt{\lambda} \sigma$$

Proof:

Instantaneous total cost of trading for trader A, called also Loss function:

$$L_A = \dot{a}^2(t) + k a(t) \dot{a}(t) + \lambda \sigma^2 \dot{a}^2(t)$$

Where: $\lambda \sigma^2 \dot{a}^2(t)$ is accounting for the risk aversion preference, it is a component depending directly on the volatility of the asset and on the size component λ .

From an economical perspective it represents the risk of holding the asset, due to price movements, during the trading window.

At this point the target of the Chriss optimization model will be to minimize the loss function L_A with boundary conditions $a(0) = 0$ and $a(1) = 1$, which ensure that trading must occur.

When averaging using the integral expression of the instantaneous total cost of trading it becomes:

$$\begin{cases} \min_{a \in S} \text{ of } \{ \int_0^1 L_A \} = \min \{ \int_0^1 \dot{a}^2(t) + k a(t) \dot{a}(t) + \lambda \sigma^2 \dot{a}^2(t) dt \} \\ \text{s.t } a(0) = 0 \text{ and } a(1) = 1 \end{cases}$$

In order to efficiently solve this, instead of computing directly the integral expression, it is useful to use a fundamental tool in the calculus of variations for constrained optimization problems:

The Euler-Lagrange formula:

$$\frac{dL_A}{da} - \frac{d}{dt} \left(\frac{dL_A}{d\dot{a}} \right) = 0$$

Then by computing the necessary derivatives:

$$\frac{dL_A}{da} = k\dot{a}(t) + 2\lambda \sigma^2 a(t)$$

$$\frac{dL_A}{d\dot{a}} = 2\dot{a}(t) + ka(t)$$

$$\frac{d}{dt} \left(\frac{dL_A}{d\dot{a}} \right) = 2\ddot{a}(t) + k\dot{a}(t)$$

Applying Euler-Lagrange equation:

$$k\dot{a}(t) + 2\lambda \sigma^2 a(t) - [2\ddot{a}(t) + k\dot{a}(t)] = 0$$

$$\cancel{k\dot{a}(t)} + 2\lambda \sigma^2 a(t) - 2\ddot{a}(t) - \cancel{k\dot{a}(t)} = 0$$

$$\cancel{2}\lambda \sigma^2 a(t) - \cancel{2}\ddot{a}(t) = 0$$

$$\ddot{a}(t) - \lambda \sigma^2 a(t) = 0$$

The permanent cost component of the market impact has disappeared, proving that in a passive optimization process under the Chriss model, or equivalently under no competition, a trader must focus only on the temporary cost component of the market impact.

$$\begin{cases} \ddot{a}(t) - \lambda \sigma^2 a(t) = 0 \\ s.t. \ a(0) = 0 \text{ and } a(1) = 1 \end{cases}$$

Optimal trading strategy under Chriss model for a two-trader passive scenario:

$$a(t) = \frac{\sinh(\sigma t)}{\sinh(\sigma)} ; \text{ with } \sigma = \sqrt{\lambda} \sigma ; \text{ where: } \sinh(\text{hyperbolic sin}) = \frac{e^{a(t)} - e^{-a(t)}}{2}$$

Since this optimization problem starts with the assumption of no competition, each trader thinks that the other trader's orders will not have an effect on price, therefore market impact is only due to the temporary component of their own trading activity, depending highly on the total size λ .

In this situation the optimal trading strategy is to deploy a λ -scaled risk-averse trading trajectory which has $\ddot{a}(t) > 0$ for all t . It is a convex set of solutions depending on the parameter of risk aversion: volatility σ .




Figure 4: A graphical representation of the λ -scaled risk-averse trading trajectories. As σ (volatility) increases, the trajectories become more convex moving away to the right of the red dashed line, which is representing the benchmark risk neutral trading trajectory.

Two-trader Chriss model optimization with asymmetric competition

In this situation, both trader A and trader B are executing trades on the same stock within the same trading window, with trader B employing a λ -scaled strategy. However, there exists a competitive asymmetry: trader B is a passive market participant who can adopt any form of λ -scaled trading strategy whether risk-averse, risk-neutral, or risk-seeking.

Moreover, trader B lacks any information regarding trader A's execution activity. Once a specific λ -scaled strategy is selected, it remains fixed for the entire trading window without any adaptive adjustments.

Consequently, B is trading passively with respect to A.

Conversely, trader A is an active participant who, upon acquiring information about B's trading activity, adjusts their strategy accordingly by employing a best-response optimization approach.

From a financial perspective, this scenario mirrors a competitive dynamic between institutional market participants, where one trader possesses an informational advantage over the other.

Setting aside the illegal practice of insider trading, a comparable real-world parallel emerges in the distinction between mutual funds and hedge funds.

Mutual funds and hedge funds differ significantly in terms of the flexibility of their investment mandates and more broadly, the regulatory frameworks governing their disclosure requirements, particularly concerning their investment strategies.

Mutual funds, designed to cater to retail investors, are subject to stringent regulations aimed at protecting this category of investors. On the other hand, hedge funds are structured for sophisticated investors and, as a result, are permitted to implement high-return speculative strategies without being legally required to disclose them publicly.

This discrepancy in mandate flexibility and consequently in the ability to rapidly adjust investment strategies, can create a competitive advantage for hedge funds over mutual funds when both entities are competing for the same asset at the same time.

The mathematical optimization problem of the Chriss model becomes in this case:

$$\begin{cases} \min_{a \in \mathcal{S}} \{ [\dot{a}(t) + \lambda \dot{b}(t)] \dot{a}(t) + k [a(t) + \lambda b(t)] \dot{a}(t) \} \\ s.t. \ a(0) = 0, b(0) = 0, a(1) = 1 \end{cases}$$

Applying the Euler-Lagrange equation:

$$2\ddot{a}(t) + \lambda \ddot{b}(t) + k\dot{a}(t) + k\lambda \dot{b}(t) - k\dot{a}(t) = 0$$

$$\ddot{a}(t) = -\frac{\lambda}{2} [\ddot{b}(t) + k\dot{b}(t)]$$

At this point of the optimization, trader A will use the information advantage and substitute for $b(t)$, $\dot{b}(t)$ and $\ddot{b}(t)$ the appropriate expression.

There could be three possible cases:

1. B employs a λ -scaled risk averse trading strategy, which has $\ddot{b}(t) > 0$
2. B employs a λ -scaled risk neutral trading strategy, which has $\ddot{b}(t) = 0$
3. B employs a λ -scaled risk seeking trading strategy, which has $\ddot{b}(t) < 0$

Let's see the final expressions once all the substitutions are done:

- For case 1 B risk averse, the optimal best response trading strategy of A is:

$$q(t) = \frac{\sinh(\sigma t)}{\sinh(\sigma)} + \frac{k}{\sigma} \frac{\cosh(\sigma t)}{\cosh(\sigma)}$$

$$a(t) = -\frac{\lambda}{2} q(t) + \left\{ 1 + \frac{\lambda}{2} [q(1) - q(0)] \right\} t + \frac{\lambda}{2} q(0)$$

- For case 2 B risk neutral, the optimal best response trading strategy of A is:

$$a(t) = \left(1 + \frac{\lambda k}{4} \right) t^2 + \frac{\lambda k}{4} t$$

- For case 3 B risk seeking, the optimal best response trading strategy of A is:

$$a(t) = \frac{e^{-\lambda t} \{ e^{\lambda} (\lambda - k) - t e^{\lambda t} (\lambda - k + 2) + e^{\lambda(t+1)} [k - \lambda + t(\lambda - k + 2)] \}}{2(e^{-\lambda} - 1)}$$

$$b(t) = \frac{e^{\lambda t} - 1}{e^{-\lambda} - 1} \text{ } \lambda\text{-scaled eager trading strategy of B}$$

It is useful to analyze graphically each case:




Figure 5: These figures represent the trading trajectories of case 1 in a low k -regime, k ranging from 0.05 to 0.25, for different increasing levels of λ , $y = t$ is the risk neutral passive benchmark.





Figure 6: These figures represent the trading trajectories of case 1 in a high k -regime, k ranging from 2 to 10, for different increasing levels of λ , $y = t$ is the risk neutral passive benchmark.

From both figures clearly the best response trading strategy of trader A, optimized using the Chriss model, has the trajectory of an eager strategy for which $\ddot{a}(t) < 0$ for all t , mathematically $a(t)$ it's always a concave function of time.

From a financial perspective A knows that B is trading using a risk averse trading strategies and anticipates it. The larger k , the more significant will be the market impact of both traders' orders on the asset's price, therefore the more eager will be the trading strategy of A, buying large volumes to get ahead of B before the price of the asset increases too much.

Figure 7: The following figures represent the trading trajectories of case 2. for k ranging from 0.05 to 25, while keeping λ constant. Where trader B is always using a risk neutral trading strategy, in fact the dashed line representing the passive benchmark $y = t$ coincides with the trading trajectories $b(t)$.




Also in this case, trader A best response to trader B is to trade aggressively using an eager trading trajectory, which is increasingly eager as the scenario enters in a high k -regime.

Moreover, for both cases:

- B trading using a risk averse trading strategy
- B trading using a risk neutral trading strategy

A best response under the Chriss model for the two-trader optimization with asymmetric competition can be generalized to deploy always an eager trading strategy which must be directly proportional to k .

Figure 8: The following figures represent the trading trajectories of case 3, k ranging from 0.1 to 25 while λ ranging from 1 to 25, $y = t$ is the risk neutral passive benchmark.



When B uses a $b(t)$ eager trading strategy the best response of A can't be generalized as in the previous two cases but rather depends on highly on the single scenarios simulated. For scenarios in which there is a small market impact coefficient $k < 1$ and where the impact of the eager trader B is low $\lambda = 1$, trader A optimal trading trajectory is to be risk averse. As $\lambda \uparrow \Rightarrow b(t)$ will mainly influence the instantaneous total cost function, then the best response of A will be to short sell initially and buy once the price has dropped, toward the ending of the trading window. When instead A faces an extremely high k -regime, with $k > 2.5$; $\lambda < 10$, the optimal strategy is to be more eager than trader B, buying more volume than necessary and selling it once the price has increased, primarily due to the impact of k on the instantaneous total cost function.

As also $\lambda \uparrow$, it counterbalances the effect of an high k -regime and A in some cases, (when $k = \lambda$) deploys as best response, a trading strategy which is risk neutral, for example in scenarios: $k = \lambda = 10$ and $k = \lambda = 25$

Two-traders Chriss model optimization under perfect competition

In this situation, both trader A and trader B are executing trades on the same stock within the same trading window and both traders are actively optimizing their trading strategies with respect to each other. Trader B uses a λ -scaled trading strategy while A is using a unitary one.

Moreover, there are no information asymmetries among them. In this scenario an important concept from Game Theory can be introduced, the so called Nash equilibrium:

Definition of Two traders joint equilibrium (Nash equilibrium)

Let A and B traders with strategies $a(t)$ and $b(t)$.

Let $\hat{a}(t)$ be the best response to $b(t)$ and $\hat{b}(t)$ be the best response to $a(t)$.

Then A and B are in equilibrium if and only if:

$$\hat{a}(t) = a(t) \text{ and } \hat{b}(t) = b(t)$$

Therefore, the Chriss model must be used to solve a joint optimization problem, to arrive at the Nash Equilibrium.

$$\begin{aligned} & \min \{L_A \text{ and } L_B\} \\ & s.t. \ a(0) = 0, b(0) = 0, a(1) = 1, b(1) = 1 \end{aligned}$$

which becomes:

$$\begin{aligned} \min \quad & \begin{cases} \{[\dot{a}(t) + \lambda \dot{b}(t)] \dot{a}(t) + k [a(t) + \lambda b(t)] \dot{a}(t)\} \\ \{[\dot{a}(t) + \lambda \dot{b}(t)] \lambda \dot{b}(t) + k [a(t) + \lambda b(t)] \lambda \dot{b}(t)\} \end{cases} \\ & s.t. \ a(0) = 0, b(0) = 0, a(1) = 1, b(1) = 1 \end{aligned}$$

By applying the Euler-Lagrange equation, it is obtained the following ODE system:

$$\begin{cases} \ddot{a}(t) = -\frac{\lambda}{2} (\ddot{b}(t) + k \dot{b}(t)) \\ \ddot{b}(t) = -\frac{\lambda}{2\lambda} (\ddot{a}(t) + k \dot{a}(t)) \end{cases}$$

With boundary conditions: $a(0) = 0, b(0) = 0, a(1) = 1, b(1) = 1$

Before deriving the final closed formula equilibrium solutions for the ODE system, under the Chriss model, it is useful to analyze two extreme cases:

1) Market impact is driven only by the temporary cost component:

ODE becomes:

$$\begin{cases} \ddot{a}(t) = -\frac{\lambda}{2} \ddot{b}(t) \\ \ddot{b}(t) = -\frac{\lambda}{2\lambda} \end{cases}$$

By applying the boundary conditions, the following equilibrium optimal solution are obtained:

$$a(t) = t$$

$$b(t) = t$$

Therefore, in the absence of the permanent market impact component, both traders will focus on minimizing total cost of trading at each pint in time. The best way to achieve this is to spread the trading volume equally across the whole duration of the trading window.

2) Market impact is driven only by the permanent cost component:

$$\begin{cases} L_A(t) = k[(a(t) + \lambda b(t))\dot{a}(t)] \\ L_B(t) = k[(a(t) + \lambda b(t))\lambda \dot{b}(t)] \end{cases}$$

Applying the Euler-Lagrange formula:

$$\begin{cases} k\dot{a}(t) + k\lambda \dot{b}(t) = 0 \\ k\lambda \dot{a}(t) + k\lambda^2 \dot{b}(t) = 0 \end{cases}$$

Dividing by k the first equation and by $k\lambda$ the second one, the system reduces to the single equation:

$$\dot{a}(t) + \lambda \dot{b}(t) = 0$$

By integrating both sides of the equation for some integration constant C , it becomes:

$$a(t) = -\lambda b(t) + C$$

Given the boundary conditions: $a(0) = 0$, $b(0) = 0$, $a(1) = 1$, $b(1) = 1$

substituting for $t = 0$

$$a(0) = -\lambda b(0) + C \Rightarrow 0 = 0 + C \Rightarrow C = 0$$


substituting for $t = 1$, knowing the integrating constant $C = 0$

$$a(1) = -\lambda b(1) \Rightarrow 1 = -\lambda(1) \Rightarrow \lambda = -1$$

which is not possible by construction of the Chriss model, λ must always be positive $\lambda > 0$! Therefore, the Chriss model optimization process in this case leads to no solution.

By studying the limit as $k \rightarrow +\infty$ of the original, it is possible to comprehend the behavior of the Chriss model in extremely high k -regimes.

Figure 9: Representing the optimal trading trajectories for $a(t)$ and $b(t)$, for a joint optimization problem under high k -regimes. k ranging from 10 to 500, while λ is kept constant at 5.



In this extreme high k scenarios, under a perfectly symmetric competition, trader A deploys a bucket shaped trading strategies while trader B, as best response to A, uses a barbel shaped trading strategies.


Intuitively the unit trader A prefers to get ahead of trader B, buying a volume bigger than the one needed mostly at the beginning of the trading window and then selling the excessive volume at the end of it. Instead, trader B that can trade for a λ -scaled factor, prefers to partially buy a certain volume of its position at the beginning and then waiting to the very end of the trading window to cover the remaining volume.

The closed formula expression, solution of the two-traders Chriss model optimization under perfect competition are the following:

$$a(t) = - \frac{\left(1 - e^{-\frac{\kappa t}{3}}\right) \left(-e^{\kappa/3} (e^{\kappa/3} + e^{2\kappa/3} + 1) (\lambda + 1) + (\lambda - 1)e^{\frac{\kappa t}{3}} + (\lambda - 1)e^{\frac{2\kappa t}{3}} + (\lambda - 1)e^{\kappa t}\right)}{2(e^{\kappa} - 1)}$$

$$b(t) = \frac{\left(1 - e^{-\frac{\kappa t}{3}}\right) \left(e^{\kappa/3} (e^{\kappa/3} + e^{2\kappa/3} + 1) (\lambda + 1) + (\lambda - 1)e^{\frac{\kappa t}{3}} + (\lambda - 1)e^{\frac{2\kappa t}{3}} + (\lambda - 1)e^{\kappa t}\right)}{2(e^{\kappa} - 1) \lambda}$$

Figure 10: Plots of the optimal trading trajectories for the two-traders Chriss model optimization under perfect competition, κ ranging from 0.05 to 75, while λ is kept constant to 5. $y = t$ is the risk neutral passive benchmark.



As $k \uparrow$ trader A uses an increasingly eager strategy, that at $k = 75$ resembles a bucket one. Instead, B starts with a risk neutral strategy for low k -regime scenarios that for $k = 1$ becomes risk averse, as k continues to increase the risk averse strategy becomes a barbel shaped trading strategy, spreading the position building activity at the extremes of the trading window.

Multi-traders Chriss model optimization under perfect competition

If there are $n+1$ unit traders $A_1, \dots, A_{n+1}$ all trading in perfect competition, the instantaneous total cost function or loss function for trader A_i becomes:

$$L_i(t) = \dot{a}_i(t) \sum_j^n \dot{a}_j(t) + k a_i(t) \sum_j^n \dot{a}_j(t)$$

With the following minimization problem:

$$\begin{cases} \min \{L_i(t)\} \Rightarrow \min \{ \dot{a}_i(t) \sum_j^n \dot{a}_j(t) + k a_i(t) \sum_j^n \dot{a}_j(t) \} \\ s.t. \ a_i(0) = 0, a_i(1) = 1, a_i = a_j \text{ for all } i \end{cases}$$

Applying the Euler-Lagrange equation:

$$\ddot{a}_i(t) = -\frac{1}{2} (\sum_{j \neq i}^n \ddot{a}_j + k \sum_{j \neq i}^n \dot{a}_j) ; a_i(0) = 0, a_i(1) = 1, a_i = a_j \text{ for all } i$$


This is the multi-trader symmetric equilibrium equation. If all $n+1$ traders will trade identical optimal strategies in equilibrium by solving this equation simultaneously for all traders by setting $\lambda = n$:

$$a(t) = \frac{e^{\frac{\lambda k}{\lambda+2}} \left(1 - e^{-\frac{\lambda k}{\lambda+2} t} \right)}{e^{\frac{\lambda k}{\lambda+2}} - 1}$$

By studying the limit as $n \rightarrow +\infty$ it is equivalent to let the size parameter of the Chriss model $\lambda \rightarrow +\infty$ in that case $\frac{\lambda k}{\lambda + 2} \rightarrow k$ then the limit becomes:

$$a_i(t) = \frac{e^k - e^{k(1-t)}}{e^k - 1}$$

Figure 11: Plots of the optimal trading trajectories for the multi-traders Chriss model optimization under perfect competition, k ranging from 0.05 to 500, for different levels of λ . The red dashed line represents the risk neutral passive benchmark.



From the plots it can be seen clearly that for low k -regimes, where the temporary market impact cost component dominates, the optimal trading strategy for trader A_i is to follow a risk neutral trading trajectory. As k increases entering inside an high k -regime, where the permanent market impact component dominates the temporary one, the optimal strategy of trader A_i is to deploy an eager trading strategy that increases in a direct proportionality relation with the parameters k and λ .

However, when comparing the limit case scenario with both high k and λ , the shape of the trading strategy doesn't resemble the bucket of the limit case seen previously under the two-traders Chriss model optimization in symmetric competition.

This result is key to understand that treating a two-traders equilibrium scenario with a large λ , for the λ -scaled trader (B), as equal to the scenario with multiple unit traders, thinking that the size parameter will be split among them, is incorrect and will lead to different optimal trading strategies.

Moreover, trader A_i before choosing between a two-traders or a multi-traders perfect symmetric optimization process, must know if in the market there is one large adversary or if there are many smaller one.

1.6 Chriss model with volatility

In this section an extension of the two-traders Chriss model optimization under perfect competition will be presented. From section 1.3 under the two-trader passive optimization the instantaneous cost function was modified to account also for the volatility of the asset σ . As in that case the volatility represents the risk of holding the asset, due to price movements, during the trading window.

Therefore, the extended the two-traders Chriss model optimization under perfect competition will be the following problem:

$$\begin{aligned} \min \quad & \begin{cases} \{[\dot{a}(t) + \lambda \dot{b}(t)] \dot{a}(t) + k [a(t) + \lambda b(t)] \dot{a}(t) + \xi_a \sigma^2 a^2\} \\ \{[\dot{a}(t) + \lambda \dot{b}(t)] \lambda \dot{b}(t) + k [a(t) + \lambda b(t)] \lambda \dot{b}(t) + \xi_b \sigma^2 b^2\} \end{cases} \\ \text{s.t.} \quad & a(0) = 0, b(0) = 0, a(1) = 1, b(1) = 1 \end{aligned}$$

Where:

$\xi_a \sigma^2 a^2$ and $\xi_b \sigma^2 b^2$ are the mathematical expressions accounting for the asset volatility inside the loss function also called instantaneous total cost function. In particular ξ_a and ξ_b are the trader-specific risk preference coefficients.

By applying the Euler-Lagrange formula:


$$\begin{cases} 2\ddot{a}(t) + \lambda \ddot{b}(t) + k \lambda \dot{b} - 2\xi_a \sigma^2 a(t) = 0 \\ \ddot{a}(t) + 2\lambda \ddot{b}(t) + k \dot{a} - 2\frac{\xi_b}{\lambda} \sigma^2 b(t) = 0 \end{cases}$$

Which reduces to the second order ODE system:

$$\begin{cases} \ddot{a}(t) = -\frac{\lambda}{2} (\ddot{b}(t) + k \dot{b}(t)) + \xi_a \sigma^2 a(t) \\ \ddot{b}(t) = -\frac{\lambda}{2\lambda} (\ddot{a}(t) + k \dot{a}(t)) + \frac{\xi_b}{\lambda^2} \sigma^2 b(t) \end{cases}$$

With boundary conditions: $a(0) = 0, b(0) = 0, a(1) = 1, b(1) = 1$

Figure 12: Plots of the optimal trading trajectories for the two-traders Chriss model optimization under perfect competition, including risk preferences and σ . k and λ are kept constant at 5, σ is kept constant at 2 while ξ_a, ξ_b are changing.



From the first row: $\xi_a = \xi_b$ trader A and trader B have the same level of risk-aversion and are increasing from 1.5 to 50, all the other parameters are kept at a constant level ($\sigma = 2$ and $k = \lambda = 5$). Trader A deploys an eager trading strategy until ξ_a grows too much.

From the second row: $\xi_a > \xi_b$, trader A has greater risk aversion than trader B with ξ_a is increasing from 10 to 200 while ξ_b is kept constant at 1.5 all the other parameters are kept at a constant level ($\sigma = 2$ and $k = \lambda = 5$). In this case trader A is more conservative and as optimal trading strategy deploys a risk-averse trajectory, as can be seen in the central and last plots (2ND row; 2ND and 3RD columns).

From the third row: $\xi_a < \xi_b$, trader B has greater risk aversion than trader A; ξ_b is increasing from 10 to 200 while ξ_a is kept constant at 1.5, all the other parameters

are kept at a constant level ($\sigma = 2$ and $k = \lambda = 5$) In this case trader B uses a roughly risk-neutral strategy while A trades a very eager strategy.

The main result is that when including inside the Chriss model the trader personal risk preferences parameter ξ_i and the asset volatility σ , a trader does not engage in significant overbuying (meaning deploying an eager shaped trading strategy) unless the opponent party is significantly more risk-averse than him.

1.7 Summary of main theoretical results

In this first chapter of this master's degree in finance thesis, the Chriss model developed in the research paper titled: Optimal Position-Building Strategies in Competition by Neil A. Chriss was rigorously presented and extensively analyzed in each subcase; exploring its results both under a mathematical perspective but also from a financial one.

An important concept to keep in mind as this work proceeds is the structure of the instantaneous cost function, which is made by both the temporary and permanent market impacts components governed by the *k-parameter*, which tells in what *k-regime* the model is optimizing.

The goal of each optimization problem is always the same: minimize the instantaneous total cost function, at least theoretically. Perhaps, as this work proceeds, it will be explained as in practice, it is only possible to minimize the average total cost function (the integral of the instantaneous total cost function, presented in details in section 1.3). It is not possible to work and compute integrals with continuous time, time needs to be discretized in very small time steps Δt .

As $\Delta t \rightarrow 0$, the discrete time steps approach an infinitesimal variation dt , i.e.,

$$\lim_{\Delta t \rightarrow 0} \Delta t = dt$$

However, this does not mean that the theory is unfeasible to use in real life, it just needs to be approximated using the correct procedure in full compliance with the theoretical framework of Chriss.

From the Chriss model theoretical model, some “axioms” can be derived, the most important ones are:

- From the sensitivity analysis carried out in section 1.3 it is clear that the **Chriss model's total cost structure exhibits lower estimation errors for the parameter k when operating in a high k -regime. This suggests a form of robustness of the model when used in more illiquid markets or in all the other high market impact scenarios, high k -regime.**
- **In a case in which the temporary market impact component dominates the permanent one, (if $0 < k \ll 1$) the optimal trading strategy will have**

the shape of a risk neutral trading trajectory. Intuitively a trader wants to minimize the total cost function that in this case depends only on the temporary market impact, therefore it must be optimized at each time t because there is no persistence in the asset price shift.

Then in general, the optimal choice would be to spread the trading volume as even as possible during the entire trading window.

- In a case in which the permanent market impact component is predominant, (if $k \gg 1$) the optimal trading strategy will have in general the shape of a risk seeking eager trading trajectory.

However, for this case some distinction must be made:

- a) **In a passive two-traders optimization framework the permanent market impact disappears** (as mathematically proven at the beginning of section 1.5). Moreover, only the personal temporary market impact component, which depends directly on size parameter λ , and the personal risk aversion preferences matter in this scenario.
- b) **When the optimization problem is derived for a two-traders asymmetric competition scenario, the trader which has more information regarding the opponent party trading strategy, will always have an advantage when minimizing its total cost function.** In particular: when trader B is trading using a risk neutral or risk averse strategy, the best response of trader A will be to deploy an eager strategy directly proportional to k . If instead trader B uses an eager strategy, trader A best response can't be generalized and depends on the specific scenario.
- c) **When the optimal problem is derived from the perfect competition, the solution of the problem represents the Nash equilibrium. The limit case in which $k \rightarrow +\infty$ (high k -regime) leads to the following optimal trading strategies: Trader B (λ -**

scaled) deploys as best response to trader A a barbel shaped trading strategy, while trader A best response to trader B will be to deploy a bucket shaped trading trajectory.

- When comparing this limit case scenario (with both high k and λ) between the two-trader and the multi-traders equilibriums, the shape of the multi-traders optimal trading strategy doesn't resemble the bucket of the two-traders under perfect competition. This result is key to understand that treating a two-traders equilibrium scenario with a large λ , for the λ -scaled trader (B), as equal to the scenario with multiple unit traders, thinking that the size parameter will be split among them, is incorrect and will lead to different optimal trading strategies.

Therefore, a trader must know if in the market there is one large adversary or if there are many smaller one, before choosing to carry out a two-traders or a multi-traders perfect symmetric optimization process.

- When including inside the Chriss optimization process, for a two-trader symmetric competition scenario, the trader personal risk preferences parameter ξ_i and the asset volatility σ . A trader does not engage in significant overbuying (meaning deploying an eager shaped trading strategy) unless the opponent party is significantly more risk-averse than him:

A_i uses an eager trading strategy for which $\ddot{a}_i(t) < 0$ for all t

If and only if $\Rightarrow \xi_i < \xi_j$

2. Standard trading approaches & Benchmarks

Introduction

The principal aim of this thesis is to find the most efficient order execution methodology (the one that minimizes the total cost function), tailored to the specific scenario (liquid or illiquid) in which trading occurs.

This chapter will focus on analyzing standard trading approaches, that rely on a simpler mathematical framework, when compared to the sophisticated Chriss optimization model. Nevertheless, standard trading approaches like moving averages and the volume weighted average price VWAP, have been a corner stone for disciplines like technical analysis and algorithmic trading.

Having said that, it is outside the scope of this thesis to discuss whether technical analysis signals can be used to make profitable investment decisions. Moreover, arguments like technical vs fundamental analysis will not be touched in this work. Again, the main target here is to find the most cost-efficient execution approaches when trading in different scenarios.

However, tools from technical analysis and algorithmic trading can be used in order to reach this target and possibly provide a benchmark for the performance evaluation of the Chriss model that will be the focus of the next chapter.

2.1 Simple moving averages

In the field of financial econometrics and quantitative trading, the moving average (MA) represents one of the most fundamental and widely utilized tools in the analysis of time series data, particularly in the study of security prices. Its simplicity, combined with its dynamical interpretive power, makes it an essential component in both academic research and practical trading strategies.

The origins of moving averages in finance can be traced back to the early 20th century, when analysts and traders began using average price levels to discern market trends. One of the earliest adopters of the method was Charles H. Dow, whose writings laid the foundation for modern technical analysis. However, it was not until the mid-20th century, with the advent of computational tools and the formalization of technical analysis frameworks, that moving averages gained

widespread academic and professional adoption. With the development of modern portfolio theory and the efficient market hypothesis (EMH) in the 1960s and 1970s, moving averages initially faced criticism from academics who questioned the predictive power of past price data. Nevertheless, empirical studies particularly those exploring momentum effects, mean reversion, and algorithmic trading strategies have continued to demonstrate that moving averages can, under certain market conditions, offer valuable insights into price dynamics and investor behavior.

Definition of moving average (MA)

A moving average is a statistical calculation that smooths a time series by averaging a subset of data points typically closing prices over a defined period. **The average “moves” because it continuously recalculates as new data becomes available, thereby adjusting dynamically to the latest information.**

There are several types of moving averages, with the most common being the **Simple Moving Average (SMA)**, which assigns equal weight to each observation within the period, and the **Exponential Moving Average (EMA)**, which can be set to give more weight to recent data or to past observations in order to respond more dynamically to price changes.

Mathematically, for a time series of prices P_t for $t = 1, \dots, n$ the simple moving average over n periods is defined as:

$$SMA_t = \frac{1}{n} \sum_{i=0}^{n-1} P_{t-i}$$

In contrast, the exponential moving average incorporates a smoothing factor α , such that:

$$EMA_t = \alpha P_t + (1 - \alpha) EMA_{t-1} ; \text{ where: } \alpha = \frac{2}{n+1}$$

The factor α is the weighting factor between the last period EMA_{t-1} and the current price P_t . As $n \rightarrow +\infty \Rightarrow \lim_{n \rightarrow +\infty} \alpha = \lim_{n \rightarrow +\infty} \frac{2}{n+1} = 0$

The larger the period considered, the smaller the smoothing factor, this means that less weight is given to recent prices P_t and more weight is retained by the last period EMA_{t-1} . Moreover, the EMA becomes less sensitive to recent price changes, behaving more like a long-term moving average, reacting slowly to new data. This

is almost equivalent to a very long Simple Moving Average (SMA) but with a key difference: the EMA never fully "forgets" the past, whereas the SMA drops older data entirely once it falls out of the fixed window. This property allows EMA to have a “memory” of past prices that decays exponentially, this is very useful to detect only major, persistent trends, avoiding noise.

On the Bloomberg Terminal, the three commonly used moving averages that appear by default on charts and in technical analysis functions are:

1. Simple Moving Average (SMA) Short-term, usually: SMA (50)
Period (trading window): 50 trading days (~2.5 month)
Purpose: Captures short-term price trends; often used by traders for recent momentum.
2. Simple Moving Average (SMA) Medium-term, usually: SMA (100)
Period (trading window): 100 trading days (~5 months)
Purpose: Filters out short-term noise; used to identify intermediate trend direction.
3. Simple Moving Average (SMA) Long-term, usually: SMA (200)
Period (trading window): 200 trading days (~10 months)
Purpose: Identifies long-term trends and major support/resistance levels.



Figure 13: A screenshot of the three standard moving averages of the Bloomberg terminal for the Bitcoin (XBT) against the US Dollar (USD).

Moving averages are a cornerstone of technical analysis, a discipline predicated on the notion that historical price patterns and trading volumes can provide predictive insight into future price movements. MAs can be used as trend-following indicators and support/resistance levels, making them a dual-purpose tool for market participants. **A moving average effectively acts as a smoothing function, filtering out short-term price fluctuations or “noise” to reveal the underlying trend. When prices remain consistently above a moving average, it typically indicates an upward trend, while sustained movement below the average suggests a downward trend.** Furthermore, the slope of the moving average itself provides information about the momentum and strength of the prevailing price trend.

In practice, moving averages generate actionable trading signals through various mechanisms, the most common of which include:

- **Crossovers:** One of the most well-known strategies is the “moving average crossover,” where a shorter-term MA (e.g., 21-day) crosses above or below a longer-term MA (e.g., 50-day or 200-day). A bullish crossover (also known as a "golden cross") occurs when the short-term MA crosses above the long-term MA, suggesting upward momentum. Conversely, a bearish crossover ("death cross") signals potential downward pressure.
- **Price-to-MA Crosses:** When the market price crosses above a moving average, it may signal a buy opportunity, while a move below the average can indicate a sell opportunity or weakening support.
- **Support and Resistance:** Moving averages often act as dynamic support or resistance levels, particularly in trending markets. Institutional traders frequently monitor key MAs (such as the 50-day and 200-day) to anticipate potential reversal zones or breakout points.

These signals, while simplistic in structure, are widely used due to their intuitive appeal and ability to be enhanced through algorithmic frameworks or combined with other technical indicators such as Relative Strength Index (RSI), Bollinger Bands, or MACD.

In summary, moving averages offer a foundational yet powerful mechanism for smoothing price data and identifying underlying trends, making them widely used in both discretionary and algorithmic trading contexts. While inherently simple in construction, their adaptive nature and ability to produce trading signals through crossovers or trend identification have made them enduring tools in market practice. However, despite their widespread use, moving averages are not without limitations. They are inherently lagging indicators, meaning they reflect past price action and may respond slowly to sudden market reversals or volatility shocks. The choice of period length also significantly influences their effectiveness; shorter periods offer quicker responses but are prone to false signals, while longer periods are more stable but may react too slowly for active trading.

For the aim of this thesis, moving averages are not evaluated for their predictive capability or investment profitability per se, but rather as a baseline methodology for trade execution under cost efficiency. While moving averages effectively capture price dynamics and are widely adopted in liquid market environments, their use as a benchmark for evaluating optimization-based execution strategies, such as the Chriss model, presents certain limitations. At first glance, moving averages may appear as suitable candidates for such benchmarking, given that the total cost function in the Chriss framework is inherently sensitive to price evolution, which is affected by both the temporary and permanent market impact components. **However, moving averages do not incorporate a critical component required for accurate cost assessment in optimal execution models: the traded volume. Since the effectiveness of a trading strategy under the Chriss model depends not only on price trajectories but also on the size and timing of executed orders, any benchmark must reflect both price and volume. Consequently, while moving averages can serve as a preliminary reference for price aligned strategies, they are insufficient as standalone benchmarks for evaluating the trading execution cost efficiency of the Chriss model. Therefore, a more complex approach is needed to pursue this task.**

2.2 VWAP approach

In modern market microstructure and algorithmic execution, the Volume Weighted Average Price (VWAP) stands as one of the most widely implemented benchmarks for institutional trading. Unlike moving averages, which rely solely on price information, **VWAP integrates both price and volume data**, thereby offering a more comprehensive and exhaustive representation of the market activity. This dual dependency enables VWAP to better approximate the real average transaction price over a given period, thus allowing for more refined execution cost evaluation. The evolution of VWAP as an execution benchmark can be traced back to the increased automation of trading during the 1980s and 1990s. As institutional investors sought to minimize the market impact of large orders and align execution performance with average market prices, VWAP emerged as a preferred strategy. It quickly became a standard for post-trade cost analysis and subsequently evolved into a basis for algorithmic execution models, especially for passive or participation-based trading strategies. Its widespread adoption was further catalyzed by the increased availability of high-frequency market data and the demand for execution transparency under regulatory frameworks such as MiFID II in Europe and Reg NMS in the United States.

Definition of the volume weighted average price VWAP

The VWAP approach, captures the average price at which a security is traded throughout a given time interval, weighted by the total traded volume. Mathematically, the VWAP for a trading interval is defined as:

$$VWAP_t = \frac{\sum_{i=1}^t P_i V_i}{\sum_{i=1}^t V_i}$$

Where:

- P_i is the trade price of the security at time interval (i)
- V_i is the volume traded at that price P_i during the interval (i)
- t is the current time interval for the sum, starting at (i = 1) ending at (i = t)
- If the VWAP is computed for the entire trading session, which is composed of a total of n intervals, then the sum is executed setting $t = n$

This formulation ensures that trades executed at higher volumes exert proportionally greater influence on the final value of VWAP, making it inherently sensitive to both price dynamics and liquidity conditions. Unlike simple or exponential moving averages, which treat all prices (or recent prices) with equal or decaying weight, **VWAP dynamically reflects where the majority of trading activity is concentrated throughout the session.**

From a practical perspective, VWAP is often used in two distinct but complementary ways: as a benchmark and as a trading strategy, known as the VWAP execution algorithm approach.

As a benchmark, institutional traders compare the average execution price of their orders against the VWAP to evaluate whether they traded “better” or “worse” than the market average. **Executing below the VWAP for a buy order (or above for a sell order) typically implies favorable execution performance.**

As a trading strategy, the **VWAP execution algorithm approach aims to distribute order flow proportionally throughout the entire trading session, matching the expected market volume curve (also known as the volume profile).** Specifically, the algorithm aims to execute buy or sell orders at prices that are close to the n interval $VWAP_t$, for each time interval in the entire trading session. The main idea is that the total execution price should be as close as possible to the average price weighted by volume during the trading session.

In order to implement this strategy, the algorithm breaks down the order into smaller, time-based portions and then executes these portions according to the expected trading volume at each time point. **This "expected market volume curve" is essentially the volume profile, which shows how the volume of trades is typically distributed throughout the day. In practice, the VWAP algorithm tries to execute orders in proportion to this volume distribution, executing larger portions of the order when the market volume is expected to be high, and smaller portions when volume is low.**

This approach is especially effective in liquid markets, where there is ample volume throughout the trading session. In such markets, the algorithm can smoothly match the pace of the market, without significantly moving the price

or encountering slippage. The liquidity ensures that the smaller trades executed by the algorithm can be absorbed without disrupting the market or drastically changing the price.

In summary, the VWAP execution strategy is about passively participating in the market by aligning with its natural volume flow, which helps to avoid drawing attention to large orders or causing unnecessary price fluctuations. It works well when there is sufficient trading volume to absorb the passive participation throughout the day.

On the Bloomberg Terminal, the VWAP can be visualized in multiple ways. It appears as a time-series line overlaid on price charts (function: VWAP<GO>) and is available for both real-time and historical analysis. Additionally, the Execution Quality Analysis (EQA) tool enables traders to benchmark individual or grouped trades against VWAP to quantify slippage and assess trading desk performance. Bloomberg computes intraday VWAP using tick-by-tick trade data, ensuring accuracy across various timeframes.

From a signal generation perspective, the VWAP line serves multiple roles:

- **Price Relative to VWAP:** Traders often consider prices above the VWAP as a sign of bullish sentiment (overvaluation in the short term), while prices below VWAP may suggest bearishness or undervaluation.
- **Mean Reversion:** Since VWAP reflects the average traded price, short-term deviations from it are often interpreted as temporary inefficiencies, leading to reversion trades.
- **Support and Resistance:** Similar to moving averages, VWAP may act as a dynamic support or resistance level, particularly during intraday trading.




Figure 14: An execution performance evaluation carried out with the Bloomberg terminal, the VWAP (green vertical line) represents the mean of the VWAP_t distribution for each time interval t during the entire trading session.

Despite its versatility, VWAP is not devoid of limitations. It assumes that past intraday volume patterns are a good predictor of current liquidity, which may not hold true during volatile or news-driven sessions. Moreover, in illiquid markets or for large block orders, VWAP strategies may result in adverse selection or increased opportunity costs due to slow execution.

Nonetheless, VWAP's integration of both price levels and execution volume renders it especially suitable as a benchmark for optimal execution frameworks, such as the Chriss model. Given that the Chriss model seeks to minimize a total cost function dependent on market impact (a function of both traded volume and price trajectory), VWAP offers a natural benchmark against which cost efficiency execution can be evaluated. Unlike moving averages, which abstract away order size, **VWAP mirrors the execution environment more faithfully, capturing both the intensity of trading and prevailing market prices.**

2.3 Summary of main differences

$$SMA_t = \frac{1}{n} \sum_{i=0}^{n-1} P_{t-i} \text{ vs } VWAP_t = \frac{\sum_{i=1}^t P_i V_i}{\sum_{i=1}^t V_i}$$

From a financial perspective, the key distinction between Moving Averages (MAs) and the Volume Weighted Average Price (VWAP) lies in their treatment of price and volume dynamics. MAs are primarily used for identifying price trends but do not incorporate volume, which is a crucial factor in determining execution costs. This limitation is especially problematic in illiquid markets, where volume and trade frequency can significantly influence price trajectories. Furthermore, MAs are inherently lagging indicators, relying on a fixed window of historical prices that may not fully capture sudden market shifts.

In contrast, VWAP integrates both price and volume, offering a more accurate representation of the average price at which a security is traded over a specified period. Since VWAP is tied to transaction volumes, it provides a more precise picture of how price dynamics unfold within the context of market liquidity. This makes VWAP particularly valuable for evaluating trade execution under both liquid and illiquid scenarios.

The inclusion of volume in VWAP further enhances its ability to reflect the true cost of execution. In real-world trading, larger orders tend to create greater market impact, and VWAP adjusts for this by giving more weight to trades with higher volumes. **This feature makes VWAP a more realistic benchmark for assessing execution strategies, particularly in the context of large block orders or strategies aimed at minimizing market impact.**

The Chriss model, which seeks to minimize total execution costs through optimized order flow, is sensitive to both price changes and trade volume. It explicitly accounts for the temporary and permanent market impact of order execution, highlighting the importance of liquidity and timing in minimizing overall trading costs. **In this context, MAs, while valuable for understanding general price trends, fail to capture the volume component that directly impacts execution costs, particularly in the case of large trades.** Therefore, using MAs as a benchmark for the Chriss model could lead to misleading conclusions about the true cost of trade execution. **Conversely, VWAP, which incorporates both**

price and volume, offers a comprehensive view of market dynamics. Given that the Chriss model aims to minimize execution costs by considering the interaction between price and market impact, VWAP serves as a more appropriate benchmark. It reflects both price fluctuations and trade size, making it an ideal reference for evaluating strategies designed to minimize market impact, especially in liquid or semi-liquid markets.

In conclusion, while Moving Averages provide a useful tool for analyzing price trends, they do not capture the critical volume aspect essential to the Chriss model's cost function. **VWAP's dual reliance on price and volume makes it a far more suitable benchmark for assessing optimal execution strategies that aim to minimize market impact.**

2.4 VWAP Total cost function for *k*-regimes

For the aim of this thesis in finding the most efficient order execution methodology, fitted for the specific market scenario (illiquid and liquid). A total cost performance comparative evaluation between the Chriss optimization model and the VWAP execution algorithm approach must be carried out. In order to pursue this final objective, the VWAP total cost function must be formulated in a way that can be effectively comparable with the mathematical expression conceptualized by Neil A. Chriss for his optimization model. In particular, the VWAP instantaneous total cost function must also depend on the *k* parameter of the Chriss model.

The generalized VWAP instantaneous total cost function, restructured to incorporate the Chriss model's market impact parameter κ , is defined as:

$$C_t^{VWAP} = (\dot{v}(t))^2 + kv(t)\dot{v}(t)$$

Where:

- $(\dot{v}(t))^2$ is the temporary market impact at time *t*
- $v(t)\dot{v}(t)$ is the permanent market impact component at time *t*
- $\dot{v}(t) = \frac{dVWAP_t}{dt}$ is the VWAP execution rate at time *t*

As already done in Chapter 1 section 1.3 for the Chriss model, the instantaneous total cost function is integrated from 0 to 1, then the average total cost function is:

$$C_{VWAP} = \frac{1}{T} \sum_{t=1}^T \left[\left(\frac{dVWAP_t}{dt} \right)^2 + k \sum_{s=1}^t dVWAP_s \cdot dVWAP_t \right]$$

$$C_{VWAP} = \int_0^1 [(\dot{v}(t))^2 + k \int_0^t \dot{v}(s) ds \dot{v}(t)] dt$$

The term $\int_0^t \dot{v}(s) ds$ is the cumulative volume up to time t . The double summation in the second term reflects how prior trades influence the cost of future executions, aligning this VWAP formulation with the theoretical structure of dynamic market impact present in the Chriss model which is equivalent to $v(t)$. Therefore, the integral expression can be simplified and becomes:

$$C_{VWAP} = \int_0^1 [\dot{v}(t)^2 + k v(t) \cdot \dot{v}(t)] dt$$

This mathematical expression effectively models the cumulative market impact of executed volume over time, while at the same time accounting for the fitted leading k -regime of the scenario. In particular, the cost function of the VWAP approach has a double dependence on both the temporary market impact component (first term) and the permanent market impact component (second term) weighted by the market impact coefficient κ .

Recalling multi-traders optimization scenario:

If there are $n+1$ unit traders $A_1, \dots, A_{n+1}$ all trading in perfect competition, the instantaneous total cost function or loss function for trader A_i becomes:

$$L_i(t) = \dot{a}_i(t) \sum_j^n \dot{a}_j(t) + k a_i(t) \sum_j^n \dot{a}_j(t)$$

Here, $\dot{a}_i(t)$ represents the trading rate of trader i , while $a_i(t)$ is the cumulative quantity traded inside the market during the trading window. The first term reflects

the temporary cost due to simultaneous executions (liquidity competition), while the second term, scaled by κ , models the permanent impact due to accumulated execution pressure on the market.

Since this scenario assumes a perfect symmetric competition among all the n traders, then the Chriss structurally assumes that each trader uses a symmetric unit trading strategy:

$$\int_0^1 \dot{a}_j(t) dt = 1 \text{ for all } j$$

if all use the same strategy $a(t)$, since it's under perfect competition then:

$$\sum_{j=1}^{n+1} \dot{a}_j(t) = (n+1)\dot{a}(t)$$

$$\sum_{j=1}^{n+1} a_j(t) = (n+1)a(t)$$

Then the average total cost function (or loss function) for trader A_i will be:

$$L_1(t) = \dot{a}(t) \cdot (n+1)\dot{a}(t) + \kappa a(t) \cdot (n+1)\dot{a}(t)$$

That becomes:

$$L_1(t) = (n+1)[\dot{a}(t)^2 + \kappa a(t)\dot{a}(t)]$$

Then by taking the integral from 0 to 1, the average total cost function is obtained:

$$C_{Chriss} = \int_0^1 L_1(t) dt = (n+1) \int_0^1 [\dot{a}(t)^2 + \kappa a(t) \cdot \dot{a}(t)] dt$$

And by normalizing dividing by $(n+1)$, the following final expression is obtained:

$$C_{Chriss} = \int_0^1 [\dot{a}(t)^2 + \kappa a(t) \cdot \dot{a}(t)] dt$$

$$C_{Chriss} = \int_0^1 [\dot{a}(t)^2 + \kappa a(t) \cdot \dot{a}(t)] dt$$

$$C_{VWAP} = \int_0^1 [\dot{v}(t)^2 + \kappa v(t) \cdot \dot{v}(t)] dt$$

Comparing the VWAP average total cost function and the one of the Chriss model, It is evident that the VWAP approach is simpler and passive, relying on historical average volumes, therefore lacking the sophisticated scenario adaptability of the Chriss minimization cost procedure. Moreover, the VWAP is not derived under an equilibrium condition like the result of the Chriss multi-traders optimization, in perfect symmetric competition. Nevertheless, by embedding the κ parameter into the VWAP cost function, it is possible to achieve the target of deriving a meaningful benchmarking method, that can be used to effectively evaluate the cost efficiency of the Chriss model under different simulated scenarios.

From a theoretical perspective in liquid markets, where price impact is minimal and competition is less of a concern, the VWAP trading approach may be proven to be the most cost-efficient due to its simplicity and alignment with the market volume profile. In contrast, in illiquid or high-volatility environments, the Chriss model should exhibit a superior performance, thanks to its natural adaptability by explicitly optimizing trading in real-time, minimizing the total cost function under both temporary and permanent market impact components, through a dynamical mathematical framework, that accounts for size and leading *k-regime*.

In the following chapter a comparative cost efficiency evaluation of these two trading methodologies, that will be implemented in practice using the R-Studio software, will provide the empirical evidence to support this master's degree in Finance thesis' main objective: identifying the most efficient trading strategy (the one that minimizes the total cost function) tailored to the specific scenario (liquid or illiquid) in which trading occurs.

3. Performance analysis for cost efficiency

The following chapter constitutes the core of this master's degree in Finance thesis. Indeed, it finally addresses the primary objective of this work: identifying the most cost-efficient trading strategies tailored to specific market conditions.

In the previous chapter, a valid benchmark for cost efficiency comparison and performance analysis was identified through the VWAP trading algorithm.

In line with this objective, the VWAP approach and the Chriss model will be implemented and compared with a focus on cost efficiency, in full compliance with the theoretical framework established in the earlier sections of this thesis.

The tool used to conduct this analysis is R-Studio, an integrated development environment (IDE) specifically designed for the R programming language, widely adopted in statistics, data analysis, and quantitative finance.

It is important to recall a key insight already highlighted by their theoretical frameworks: namely, given the different structures of the cost functions and more broadly the fundamental distinct nature between the Chriss optimization model and the VWAP approach, their performance is expected to vary significantly depending on the specific market scenario in which they are applied.

This theoretical distinction is further reinforced by earlier analysis in the thesis.

As shown in the sensitivity analysis in Section 1.3, the Chriss model demonstrates greater robustness in illiquid market scenarios those characterized by a high value of the parameter k . Its cost structure is less sensitive to estimation errors in such environments, making it more reliable when the market impact is significant.

On the other hand, the VWAP algorithm, which spreads execution in line with volume distribution, performs best in liquid markets. In these conditions, abundant market volume allows the algorithm to execute smoothly with minimal price disruption and slippage.

In summary, these theoretical insights suggest that the Chriss model is more effective in illiquid scenarios, whereas the VWAP approach performs better in highly liquid ones. The central aim of this chapter is to investigate whether these theoretical expectations are confirmed by empirical evidence, and to assess the statistical significance of the observed results.

3.1 Cost functions time discretization

The instantaneous total cost functions and their integrated expressions, which represents the average total cost functions, all belong to the family of C^2 functions (C^2 means continuous and twice differentiable). This level of smoothness is essential for deriving analytical results within the continuous time theoretical framework.

However, practical implementation of these models, particularly in computational environments such as R-Studio, requires a shift from continuous to discrete time. Since such software environments do not work directly with continuous variables or functions, time discretization becomes necessary to simulate and evaluate trading strategies numerically.

The discretization method employed in this thesis involves dividing the total trading horizon T into N equally spaced intervals of length where:

$$\Delta t = \frac{T}{N} ; t_0 = 0, t_1 = \Delta t, t_N = N\Delta t = T;$$

in general, for an i -interval $t_i = i \Delta t$ for $i = 0, 1, \dots, N$

This transforms the continuous cost function into a stepwise approximation, allowing cumulative costs to be estimated using discrete values.

As $\Delta t \rightarrow 0$, the discrete time steps approach an infinitesimal increment dt , i.e.,

$$\lim_{\Delta t \rightarrow 0} \Delta t = dt$$

bridging the discrete approximation with the continuous-time formulation used in the theoretical frameworks of optimal trading strategies.

A useful way to visualize this process is through the use of Riemann sums, where the area under the curve of a continuous cost function is approximated by the sum of the areas of rectangles.

For the purpose of this thesis, each rectangle represents the average cost of a discrete time interval and mathematically speaking is referred to as a Riemann rectangle. The smaller the time step, the smaller the base of the Riemann rectangle, leading to a more precise approximation of the continuous time theoretical framework.

The software R-Studio is able to perform these summations calculations, using a very small discretization time step $\Delta t \approx dt$, leading to consistent and precise results.




Figure 15: Both figures visualize the approximation of continuous time thanks to its discretization. The first figure shows a quite large time step Δt for the Riemann sums, therefore it is not a very good approximation. Instead, the second figure uses a very small time step Δt , leading to a more precise approximation of the continuous time.

This discretized framework not only makes numerical summation feasible for computational purposes but also aligns the theoretical models with real-world trading, where as a matter of fact, decisions are made at discrete time intervals.

3.2 Liquid vs Illiquid scenarios with Monte Carlo method

The first step in a comparative analysis between the Chriss model's optimization process and the VWAP algorithmic approach for trading execution is to test whether their respective trading trajectories differ with statistical significance under both liquid and illiquid market scenarios.

Therefore, using the software R-Studio the Chriss model optimization procedure was implemented for the case of multi-traders under perfect competition; for the liquid market scenario the model parameters were set to $k = 0.1$ and $\lambda = 1$ while for the illiquid market scenario parameters were set to $k = 5$ and $\lambda = 10$.

Instead, the VWAP algorithm approach was simulated using the Monte Carlo method, using a standard normal distribution with standard deviation $\sigma = 0.02$, needed to introduce a random noise factor, typical for any Monte Carlo Simulation. Monte Carlo simulations were set to run 100 times.

The following figure shows the trading trajectories obtained for both frameworks.




Figure 16: R-Plots of trading trajectories for both liquid and illiquid market environments. The Chriss model trading trajectories (red lines) were built following the theoretical optimization process for the case of multi-traders under perfect competition while the VWAP (black dashed lines) were simulated using the Monte Carlo method.

It is clear from the figures that the Chriss optimization process is more flexible and adapts better to different market conditions when compared to the VWAP, which instead appears to be more stable showing rather risk averse trading trajectories for both simulated scenarios.

This is not surprising, in fact as expected from the theory, the Chriss optimization model is an active execution approach that highly depends on the environment in which it is operating in. Moreover, both trading strategies are of the eager type especially in the case of illiquid environment or equivalently (institutional trader trading large quantities).

The final step in this preliminary comparison of the two methodologies trading trajectories is to seek for statistical significance. A standard t-test would work if both trading outcomes were assumed to be normally distributed. While this is true for the VWAP, which has its trading trajectories distributed as a standard normal with standard deviation $\sigma = 0.02$, it could not be assumed that also the Chriss model trading trajectories follow a standard normal distribution.

Therefore, to pursue a more robust comparison, that does not rely on the assumption of normality, the Kolmogorov–Smirnov (K-S) test is employed. This non-parametric test is particularly well suited for this analysis, as it compares the entire empirical distribution functions of the two datasets without requiring any assumptions about their underlying distributions. Given that the Chriss model generates trading trajectories that are shaped by dynamic optimization in response to varying market impact and liquidity conditions, its output distribution may exhibit skewness, kurtosis, or other non-Gaussian features. The K-S test is thus preferable to the standard t-test, which is limited to comparing means under the assumption of normality and equal variances.

Therefore, by applying the Kolmogorov–Smirnov test, it is possible to rigorously determine whether the differences in trading trajectories between the Chriss model and VWAP are statistically significant across different market scenarios, capturing both the magnitude and distributional shape of the outcomes.

The following lines show the R-Studio printouts of the K-S results for both liquid and illiquid market scenarios:

```
--- Trading Trajectories Comparison for Common Traders (Liquid Market) ---  
Kolmogorov-Smirnov Test (Trader A vs VWAP): p-value = 0.0158141  
Kolmogorov-Smirnov Test (Trader B vs VWAP): p-value = 0.0158141  
Linear Regression (Trader A ~ VWAP) Beta: 0.975422  
Linear Regression (Trader B ~ VWAP) Beta: 0.975422  
Trader A significantly diverges from VWAP.  
Trader B significantly diverges from VWAP.
```

```
--- Trading Trajectories Comparison for Institutional Traders (Illiquid Market) ---  
Kolmogorov-Smirnov Test (Trader A vs VWAP): p-value = 2.68708e-33  
Kolmogorov-Smirnov Test (Trader B vs VWAP): p-value = 0.006302223  
Linear Regression (Trader A ~ VWAP) Beta: 0.9298428  
Linear Regression (Trader B ~ VWAP) Beta: 0.9656463  
Trader A significantly diverges from VWAP.  
Trader B significantly diverges from VWAP.
```

As shown by the results of the Kolmogorov Smirnov test, the trading trajectories produced by the Chriss model and the VWAP algorithm differ with statistical significance across both market scenarios. Particularly for the illiquid environment, where the test p-values are considerably lower than those observed in the liquid case, the K-S test result is confirmed at all confidence levels. This outcome confirms that the choice between the Chriss model and VWAP for execution purposes leads to materially different trading paths, as expected from their distinct theoretical foundations and underlying assumptions.

Given this divergence, attention naturally shifts to a more fundamental issue namely, the question of which methodology offers greater cost efficiency under different market conditions. Specifically, it becomes necessary to evaluate which approach is more advantageous in liquid markets and which proves superior in illiquid environments. In order to address this, a comparative performance analysis based on the cost functions embedded in each model must be undertaken.

For this purpose, as already explained in sections: 1.3 (for the Chriss model cost functions), 2.4 (for the VWAP) and 3.1 (Time discretization)

The procedure is the following:

- 1) The instantaneous total cost functions also known as loss functions are derived, which reflect both the temporary and permanent market impact components. Both components are governed by the k parameter.
- 2) By integrating these expressions, for the trading interval from 0 (starting time) to 1 (closing time), the average total cost functions are obtained for both trading execution approaches.
- 3) Finally, to perform in practice the integral summation, time will be discretized in very small discrete time steps $\Delta t \rightarrow 0$, ensuring good accuracy (fine Riemann summation methodology).

The following figures show the comparative dynamic cost evolution, between the Chriss model optimization, for the case of multi-traders symmetric equilibrium, and the Monte Carlo simulated VWAP algorithm, during a trading window.

For the liquid scenario parameters were set to $k = 0.1$ and $\lambda = 1$ while for the illiquid market scenario parameters were set to $k = 5$ and $\lambda = 10$.

Monte Carlo simulations were set to 100 times. The number of traders was fixed.

It is important to notice that both Chriss model and VWAP cost functions are deeply depending on the parameter k , which is responsible to determine the equilibrium between the temporary and permanent market impact components as time evolves.

$$C_{Chriss} = \int_0^1 [\dot{a}(t)^2 + ka(t) \cdot \dot{a}(t)] dt$$

$$C_{VWAP} = \int_0^1 [\dot{v}(t)^2 + kv(t) \cdot \dot{v}(t)] dt$$




Figure 17: R-Plots of the average cost functions dynamic time evolution under both liquid and illiquid market scenarios. The Chriss model average total cost functions (red lines) were derived following the theoretical optimization process for the case of multi-traders under perfect competition while the VWAP average total cost functions (black dashed lines) were simulated using the Monte Carlo method, 100 times. N (number of traders) = 10

From the figure on the left, it is evident that when operating in a liquid market environment, the VWAP algorithm proves to be the most cost-efficient execution strategy. Indeed, as illustrated in the figure, the dashed black line depicting the average total cost evolution of the VWAP over time, is always below the red line representing the Chriss model average total cost time evolution.

Therefore, it is possible to conclude that in a liquid scenario the VWAP execution approach leads to lower costs in terms of both temporary and permanent market impact.

This picture changes significantly when analyzing the figure on the right.

Under an illiquid market environment, both cost functions initially rise at a similar rate. Nevertheless, around time $t = 0.25$, the advantage of the Chriss model's flexibility begins to emerge: its average total cost curve starts to decline, whereas the VWAP cost continues to rise until the end of the trading window.

From this empirical evidence, it is possible to conclude that when operating in an illiquid scenario the Chriss model optimization execution methodology has lower costs in terms of both temporary and permanent market impact.

The following R-table shows the mean values of both Chriss model and VWAP total costs under both liquid and illiquid market environments.

Market Scenario	Mean Chriss Cost	Mean VWAP Cost
Common Traders (Liquid Market)	84.813175	66.26729
Institutional Traders (Illiquid Market)	8.491506	66.26729

Looking at the means, the difference between average total cost of the two approaches is noticeable, especially for the illiquid market scenario, where the Chriss optimization procedure clearly dominates the VWAP in terms of cost efficiency.

For completeness the results obtained using the same methodology but for N (number of traders) = 100 are also reported below.




Figure 18: R-Plots of the average cost functions dynamic time evolution under both liquid and illiquid market scenarios. The Chriss model average total cost functions (red lines) were derived following the theoretical optimization process for the case of multi-traders under perfect competition while the VWAP average total cost functions (black dashed lines) were simulated using the Monte Carlo method, 100 times. N (number of traders) = 100

Figure 17 ($N = 10$) and Figure 18 ($N = 100$) display nearly identical visual patterns, with both showing similar shapes in the average total cost functions for the two methods, despite the expected differences in cost magnitudes, given by the difference in total traders and therefore total traded volumes.

Market Scenario	Mean Chriss Cost	Mean VWAP Cost
Common Traders (Liquid Market)	4240.6588	3336.164
Institutional Traders (Illiquid Market)	424.5753	3336.164

The two tables show different values from the $N = 10$ and the $N = 100$ cases, as expected by the difference in total traded volumes. **However, when analyzing the difference between the mean of the total cost function of the Chriss model and the one of the VWAP, the conclusion is exactly the same: The Chriss model is the best at cost efficiency when deployed in an illiquid scenario.**

In this analysis the number of traders was kept constant, in the first case $N = 10$ and in the second case to $N = 100$ with a dynamical comparative cost evaluation for each time step inside the interval starting at $t = 0$ and ending at $t = 1$.

At this point to have a more robust and complete analysis, it could be useful to simulate both liquid and illiquid market scenarios, keeping time at a fixed time interval and computing the limit as the number of traders increases, approaching infinity.

3.3 Limit case study as N traders $\rightarrow +\infty$

The following figure shows the limit case scenarios for both the Chriss model optimization and the VWAP approach. In this case N represents the number of traders and goes to infinity. The time is fixed to a closed interval, i.e. $(0, 1)$ therefore, this is not a dynamic, as in the previous section, but rather a static cost efficiency performance evaluation. Here the average total cost functions are analyzed at a given time horizon across different market densities.

The following figure shows a comparative cost efficiency performance evaluation between the Chriss model optimization, for the case of multi-traders symmetric equilibrium, and the Monte Carlo simulated VWAP algorithm, focusing on different market sizes, given by the increasing $N \uparrow \Rightarrow N \rightarrow +\infty$

For the liquid scenario parameters were set to $k = 0.1$ and $\lambda = 1$ while for the illiquid market scenario parameters were set to $k = 5$ and $\lambda = 10$.

Monte Carlo simulations were set to 100 times. In this case the time interval is fixed, the number of traders was varying approaching infinite.




Figure 18: R-Plots of the average cost functions under both liquid and illiquid market scenarios. The Chriss model average total cost functions (red lines) were derived following the theoretical optimization process for the case of multi-traders under perfect competition while the VWAP average total cost functions (black dashed lines) were simulated using the Monte Carlo method, $N \rightarrow +\infty$ where N is the number of traders.

When analyzing these limit case scenarios, the result does not change. Chriss model average total costs are always above the ones of the VWAP in the case of liquid market environment (left figure) and are always below the ones of the VWAP in the case of illiquid market environment (right figure). **Confirming the findings of the previous section's comparative dynamical cost efficiency performance evaluation:**

- **Chriss model is better in terms of cost efficiency when deployed in an illiquid market environment.**
- **VWAP algorithmic approach is better in terms of cost efficiency when deployed in a liquid market environment.**

At this stage, the primary objective set forth in this thesis has been successfully achieved through the application of advanced methodologies, fully aligned with the theoretical fundamentals of both trading execution strategies.

Nonetheless, while not strictly essential to the core analysis, a valuable extension of this work could involve the comparison of the average total cost functions based on historical data rather than relying solely on Monte Carlo simulations for the VWAP approach. Such an addition would enrich the empirical robustness of the findings and offer further insights into the practical applicability of VWAP under real-market conditions.

The next section focuses on this using real life trading data obtained from the Bloomberg terminal to construct the VWAP execution methodology.

3.4 Liquid vs Illiquid scenarios with Historical data

In this section an R-Studio code was developed to firstly compute the VWAP approach total costs using real life trading data, obtained from the Bloomberg terminal. Then carry out a comparative dynamic cost efficiency performance evaluation between the constructed VWAP based on historical data and the Chriss optimization model. These data include the most influent traded volumes in the Apple stock during the trading day of the 19TH February 2025 in the Nasdaq.

The Chriss model optimization was set to operate in a highly liquid market environment, as both the security Apple stock and market Nasdaq are very liquid ones. In particular, the model's parameters for the numerical optimization were set as the following: $k = 0.1$, $\lambda = 1$

The following figures represent the obtained average total cost functions, as time evolves, for both VWAP with historical data and Chriss model equilibrium.




Figure 19: R-Plot of the average cost function under a highly liquid market scenario. The Chriss model average total cost function (red line) was derived following the theoretical optimization process for the case of multi-traders under perfect competition. The green dashed line represents the sample mean of each total cost incurred during the trading window.



Figure 20: R-Plot of the average cost function for the daily volume trading data obtained from the Bloomberg terminal for the Apple stock on the 19TH February 2025 in the Nasdaq. The VWAP algorithm was applied to this data. The VWAP average total cost function dynamical evolution (black line). The green dashed line represents the sample mean of each total cost incurred during the trading window.

From both figures but also from the data shown in the table below, it could be seen that a trader, trading in the Apple stock on the Nasdaq market during the trading day of the 19TH February 2025, would incur on average lower trading costs, considering both the temporary and the permanent market impacts components, by using the VWAP algorithm as order execution method. This is not surprising as both the traded security and the market in which trading occurs are highly liquid ones. **Therefore, once again confirming the superiority in terms of cost efficiency of the VWAP algorithmic order execution approach in liquid market environments.**

Table: VWAP vs Chriss Cost Comparison Over Time

time	VWAP_Cost_Dynamic	Chriss_Cost
0.000000	0.0074889	0.000000
0.010101	0.0072591	3.963880
0.020202	0.0068493	7.852364
0.030303	0.0066201	11.666844
0.040404	0.0059327	15.408597
0.050505	0.0060276	19.079083
0.060606	0.0061229	22.679587
0.070707	0.0064434	26.211403
0.080808	0.0069806	29.675800
0.090909	0.0076875	33.074022
0.101010	0.0081726	36.407294
0.111111	0.0093515	39.676816
0.121212	0.0096735	42.883767
0.131313	0.0106110	46.029306
0.141414	0.0110581	49.114569
0.151515	0.0107770	52.140673
0.161616	0.0104272	55.108712
0.171717	0.0101004	58.019764
0.181818	0.0095038	60.874884
0.191919	0.0089797	63.675109
0.202020	0.0088274	66.421458
0.212121	0.0083847	69.114931
0.222222	0.0080240	71.756508
0.232323	0.0072429	74.347155
0.242424	0.0064247	76.887816
0.252525	0.0056357	79.379421
0.262626	0.0042696	81.822881
0.272727	0.0027712	84.219093
0.282828	0.0017314	86.568935

Kolmogorov-Smirnov Test Results for Grouped Time Intervals:

```
> print(ks_results, row.names = FALSE)
```

Group	Cost_Difference	Statistic	P_Value
1	41.27461	0.96	7.910729e-13
2	104.76409	1.00	1.582146e-14
3	143.54459	1.00	1.582146e-14
4	166.99460	1.00	1.582146e-14

The K-S test p-values confirm the statistical significance of the result founded in this comparative dynamic cost efficiency performance evaluation between the constructed VWAP based on historical data and the Chriss optimization model.

Conclusions

This master's thesis in Finance addresses the challenge of identifying cost efficient trading execution strategies under competitive conditions in both liquid and illiquid market environments. The objective was to determine the most effective methodology for minimizing execution costs in each scenario.

In the first chapter a mathematical framework developed by mathematician and hedge fund manager Neil A. Chriss was presented and analyzed in detail.

The theoretical mathematical optimization approach presented by Chriss is sophisticated showing high robustness and great adaptability to different market conditions. Nevertheless, the model remains easy to understand and interpret, relying on a limited set of parameters and boundary conditions for reaching an optimal solution, that minimizes the total cost for order execution.

Chapter two introduced the Volume Weighted Average Price (VWAP), a widely used and relatively simple algorithmic trading benchmark. VWAP executes orders based on historical price and volume data and is frequently employed by institutional investors due to its transparency and ease of implementation.

In the third chapter, the core of this thesis, the Chriss model and the VWAP were set to compete against each other, carrying out a comparative performance analysis focused on cost efficiency in each market scenario.

The comparative analysis accurately recreated the Chriss optimization model procedure and compared it to the VWAP approach, both dynamically (fixed market size and evolving time intervals) and statically (fixed time interval and increasing market sizes). The VWAP was performed both by Monte Carlo simulation method and by using historical data, on daily trading volumes and prices for US equity securities, obtained from the Bloomberg Terminal.

The results of this thesis research show that the Chriss model outperforms VWAP in terms of cost efficiency within illiquid markets, whereas VWAP proves more cost efficient in liquid market environments.

All simulations, empirical analyses and statistical tests were carried out using the R-Studio software.

Annex

The original code provided by Neil A. Chriss' in his Research Paper titled "Optimal Position-Building Strategies in Competition". In this case the following code in Wolfram Mathematica is used to solve the two-traders symmetric equilibrium in for the Chriss model including volatility and traders' risk preferences as in section 1.6.

```

1
2 (* Define constants *)
3 λ = 5;
4 κ = 5;
5 σ = 0.5; (* Increased value for effect *)
6
7 (* Define the system of differential equations *)
8 solveSystem[ξa_, ξb_] := Module[{eqns, bcs, sol, aSol, bSol},
9   eqns = {
10     a''[t] == -(λ/2) (b''[t] + κ b'[t]) + ξa σ2 a[t],
11     b''[t] == -(1/(2 λ)) (a''[t] + κ a'[t]) + (ξb/λ2) σ2 b[t]
12   };
13   bcs = {a[0] == 0, a[1] == 1, b[0] == 0, b[1] == 1};
14   sol = DSolve[{eqns, bcs}, {a[t], b[t]}, t];
15   {a[t] /. sol[[1]], b[t] /. sol[[1]]}
16 ];
17
18 (* Define the grid of ξa and ξb values, multiplied by 3 *)
19 xiValues = {
20   {{1.5, 1.5}, {10, 10}, {50, 50}},
21   {{10, 1.5}, {50, 1.5}, {200, 1.5}},
22   {{1.5, 10}, {1.5, 50}, {1.5, 200}}
23 };
24
25 (* Generate the 3x3 grid of plots *)
26 gridPlots = Grid[Table[
27   Module[{aSol, bSol, ξa = xiValues[[i, j, 1]], ξb = xiValues[[i, j, 2]]},
28     {aSol, bSol} = solveSystem[ξa, ξb];
29     Plot[{aSol, bSol}, {t, 0, 1},
30       PlotLegends -> {"a(t)", "b(t)"},
31       PlotStyle -> {Blue, Red},
32       Frame -> True,
33       FrameLabel -> {"a(t), b(t)", "κ=" <> ToString[κ] <>
34         ", λ=" <> ToString[λ]}, {"t", "\(\xi_a=\)" <> ToString[ξa] <>
35         ", \(\xi_b=\)" <> ToString[ξb] <> ", σ=" <> ToString[σ]}},
36       PlotRange -> {{0, 1}, {0, 2.5}}]
37   ], {i, 1, 3}, {j, 1, 3}], Frame -> All
38 ];
39
40 (* Display the grid of plots *)
41 gridPlots

```

References

1. Agarwal, V., & Naik, N. Y. (2004). Risk and portfolio decisions involving hedge funds. *The Review of Financial Studies*, 17(1), 63–98
2. Almgren, R., & Chriss, N. (1997). *Optimal liquidation*. SSRN. <https://doi.org/10.2139/ssrn.53501>
3. Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, 3(2), 5–39.
4. Almgren, R., Thum, C., Hauptmann, E., & Li, H. (2005). Direct estimation of equity market impact. *Risk*, 18(7), 58–62.
5. Ammann, M., & Buetler, M. (2009). Optimal trading algorithms and market microstructure. *Financial Markets and Portfolio Management*, 23(2), 125–147.
6. Avellaneda, M., & Stoikov, S. (2008). High-frequency trading in a limit order book. *Quantitative Finance*, 8(3), 217–224.
7. Bachmann, K., De Giorgi, E. G., & Hens, T. (2018). *Behavioral Finance for Private Banking* (2nd ed.). Wiley.
8. Bialkowski, J., & Jakubowski, J. (2012). Algorithmic trading strategies in the context of VWAP and implementation shortfall. *The Journal of Trading*, 7(1), 16–28.
9. Bloomberg Terminal. (2024). *VWAP and benchmark order types*. Bloomberg L.P.
10. Bloomberg Market Concepts (BMC). (2023). *Trading and execution analytics*. Bloomberg L.P.
11. Bouchaud, J. P., Farmer, J. D., & Lillo, F. (2009). How markets slowly digest changes in supply and demand. In T. Hens & K. R. Schenk-Hoppe

- (Eds.), *Handbook of Financial Markets: Dynamics and Evolution* (pp. 57–160). North-Holland.
12. Brock, W., Lakonishok, J., & LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance*, 47(5), 1731–1764.
 13. Brown, S. J., Goetzmann, W. N., & Park, J. (2001). Careers and survival: Competition and risk in the hedge fund and CTA industry. *The Journal of Finance*, 56(5), 1869–1886.
 14. Brunnermeier, M. K., & Pedersen, L. H. (2005). Predatory trading. *The Journal of Finance*, 60(4), 1825–1863.
 15. Busseti, E., & Boyd, S. (2015). Volume weighted average price optimal execution.
 16. Cartea, Á., Jaimungal, S., & Penalva, J. (2015). *Algorithmic and high-frequency trading*. Cambridge University Press.
 17. Chaddha, A., & Yadav, S. (2022). Examining the predictive power of moving averages in the stock market. *Journal of Student Research*, 11(3).
 18. Chen, J. (2012). The world of moving averages. In *Essentials of Technical Analysis for Financial Markets* (pp. 123–140).
 19. Chriss, N. A. (2024, September 5). *Optimal position-building strategies in competition* [Preprint].
 20. Chriss, N. A. (2024, September 23). *Position-building in competition with real-world constraints* [Preprint].
 21. Chriss, N. A. (2024, October 17). *Competitive equilibria in trading* [Preprint].
 22. Chriss, N. A. (2025, January 2). *Position building in competition is a game with incomplete information* [Preprint].

23. Chong, T. T.-L., Li, J., & Ma, K. (2014). Profitability of trading rules based on three moving averages. *Labuan Bulletin of International Business and Finance*, 10, 1–6.
24. Cont, R., & Kokholm, T. (2013). A consistent pricing model for index options and volatility derivatives. *Mathematical Finance*, 23(2), 248–274.
25. Crypto Liquidity Dries Up as Billions in Market Value Evaporates. (2022, February 24). Bloomberg.
26. De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98(4), 703–738.
27. Di Lorenzo, R. (2013). Advanced moving averages. In *Basic Technical Analysis of Financial Markets* (pp. 89–102).
28. Donnelly, R. (2022). Optimal execution: A review. *Applied Mathematical Finance*, 29(3), 181–212.
29. ‘Dumb Money’ Looks at GameStop Short Squeeze. *Bloomberg*.
30. Eddelbuettel, D. (2013). *Seamless R and C++ Integration with Rcpp*. Springer.
31. Ellis, C. A., & Parbery, S. A. (2005). Is smarter better? A comparison of adaptive and simple moving average trading strategies. *Research in International Business and Finance*, 19(3), 399–411.
32. Elton, E. J., Gruber, M. J., Brown, S. J., & Goetzmann, W. N. (2014). *Modern portfolio theory and investment analysis* (9th ed.). Wiley.
33. Engle, R. F., & Russell, J. R. (2010). *Analysis of high-frequency data*. Princeton University Press.
34. Fama, E. F., & Blume, M. E. (1966). Filter rules and stock-market trading. *The Journal of Business*, 39(1), 226–241.

35. Fama, E. F. (1970). Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2), 383–417.
36. Foucault, T., Pagano, M., & Röell, A. (2013). *Market liquidity: Theory, evidence, and policy*. Oxford University Press.
37. GameStop Is a Bubble in Its Purest Form. *The Wall Street Journal*.
38. GameStop Resurgence Reinforces New Reality for Hedge Funds. *The Wall Street Journal*.
39. Gatheral, J. (2010). No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7), 749–759.
40. Gatheral, J. (2010). Three models of market impact and data. In *Market Microstructure and High Frequency Data* (pp. 1–37).
41. Gatheral, J., & Schied, A. (2011). Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework. *International Journal of Theoretical and Applied Finance*, 14(3), 353–368.
42. Gatheral, J., & Schied, A. (2013). Dynamical models of market impact and algorithms for order execution. In J.-P. Fouque & J. A. Langsam (Eds.), *Handbook on Systemic Risk* (pp. 579–599). Cambridge University Press.
43. Gelfand, I. M., & Silverman, R. A. (2000). *Calculus of variations*. Courier Corporation.
44. Givens, G. H., & Hoeting, J. A. (2012). *Computational statistics* (2nd ed.). Wiley.
45. Glasserman, P. (2004). *Monte Carlo methods in financial engineering*. Springer.

46. Gunasekarage, A., & Power, D. M. (2001). The profitability of moving average trading rules in South Asian stock markets. *Emerging Markets Review*, 2(1), 17–33.
47. Harris, L. (2003). *Trading and exchanges: Market microstructure for practitioners*. Oxford University Press.
48. Hautsch, N., & Huang, R. (2012). The market impact of a limit order. *Journal of Economic Dynamics and Control*, 36(4), 501–522.
49. Hedge funds can handle GameStop short squeeze: Schonfeld CEO. *Bloomberg*
50. Hull, J. (2022). *Options, futures, and other derivatives* (11th ed.). Pearson.
51. Humphery-Jenner, M. L. (2011). Optimal VWAP trading under noisy conditions. *Journal of Banking & Finance*, 35(9), 2319–2329.
52. Isaacs, R. (1999). *Differential games: A mathematical theory with applications to warfare and pursuit, control and optimization*. Dover Publications.
53. Kato, T. (2014). VWAP execution as an optimal strategy. *arXiv preprint arXiv:1408.6118*
54. Kearns, M., Kulesza, A., & Nevmyvaka, Y. (2010). Empirical limitations of the VWAP strategy. *ACM Transactions on Intelligent Systems and Technology*, 1(2), 1–19.
55. Khorana, A., Servaes, H., & Tufano, P. (2009). Mutual fund fees around the world. *The Review of Financial Studies*, 22(3), 1279–1310.
56. Kissell, R. (2013). *The science of algorithmic trading and portfolio management*. Academic Press.
57. Kissell, R., & Glantz, M. (2003). *Optimal trading strategies: Quantitative approaches for managing market impact and trading risk*. AMACOM.

58. Kloeden, P. E., & Platen, E. (1999). *Numerical solution of stochastic differential equations*. Springer.
59. Kon, S. J., & Jen, F. C. (1978). The application of moving averages to investment timing. *Journal of Financial and Quantitative Analysis*, 13(1), 65–78.
60. Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica*, 53(6), 1315–1335.
61. Litterman, R., & Winkelmann, K. (1998). Estimating the market impact of large trades. *Goldman Sachs Risk Management Series*.
62. Lo, A. W., Mamaysky, H., & Wang, J. (2000). Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4), 1705–1765.
63. Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, 3(3), 205–258.
64. Marshall, B. R., Nguyen, N. H., & Visaltanachoti, N. (2017). Time series momentum and moving average trading rules. *Quantitative Finance*, 17(3), 405–421.
65. Matloff, N. (2011). *The art of R programming: A tour of statistical software design*. No Starch Press.
66. Murphy, J. J. (1999). *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. New York Institute of Finance.
67. Neely, C. J., Rapach, D. E., Tu, J., & Zhou, G. (2014). Forecasting the equity risk premium: The role of technical indicators. *Management Science*, 60(11), 1772–1791.

68. Obizhaeva, A. A., & Wang, J. (2013). Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1), 1–32.
69. Owen, A. B. (2013). *Monte Carlo Theory, Methods and Examples*.
70. Patterson, S. (2012). *Dark pools: High-speed traders, A.I. bandits, and the threat to the global financial system*. Crown Business.
71. Robert, C. P., & Casella, G. (2010). *Introducing Monte Carlo Methods with R*. Springer.
72. Schönbucher, P. J. (2003). *Credit derivatives pricing models*. Wiley.
73. Sharpe, W. F., Alexander, G. J., & Bailey, J. V. (1999). *Investments* (6th ed.). Prentice Hall.
74. Shiller, R. J. (2003). From efficient markets theory to behavioral finance. *Journal of Economic Perspectives*, 17(1), 83–104.
75. Thaler, R. H. (2005). *Advances in Behavioral Finance* (Vol. II). Princeton University Press.
76. The GameStop short squeeze shows an ugly side of the investing world. *The Wall Street Journal*.
77. Treynor, J. L. (1995). The principles of performance measurement. *The Journal of Portfolio Management*, 22(2), 13–22.
78. Webster, K. T. (2023). *Handbook of price impact modeling*. Chapman and Hall/CRC.
79. Wolfram Research, Inc. (2024). *Mathematica* (Version 14.1). Champaign, IL.
80. Zarinelli, E., Treccani, M., Farmer, J. D., & Lillo, F. (2015). Beyond the square root: Evidence for logarithmic dependence of market impact on size and participation rate. *Market Microstructure and Liquidity*, 1(2), 1550004.