Essays on Computational Finance

Evangelos Litos
School of Business, Economics Division
University of Leicester

A thesis submitted for the degree of
Doctor of Philosophy at the University of Leicester.

January 2018

To my mother, my father and to Becky

Essays on Computational Finance

by

Evangelos Litos

Abstract

This thesis explores the behaviour of different types of individuals (*i.e.*, traders, experimental subjects) in financial markets under two different environments. It consists of three chapters/papers. Chapters 2 and 3 focus on an electronic market environment where traders submit orders through an electronic platform (*i.e.*, electronic (limit) order book). Chapter 4 develops a portfolio choice problem - implemented in a laboratory environment - where subjects choose under risk, the allocation of wealth between two assets.

Chapters 2 and 3 develop a dynamic continuous-time model of trade in a single financial asset under two different setups. In chapter 2 under the Single-Market scenario, the set of preferences of all market participants is aggregated on a limit order book related to a 'lit' market. Under the Multi-Markets scenario, trading takes place on two limit order markets that operate in parallel: an open (transparent) 'lit' market and a 'dark' (opaque) market. The aim of this chapter is to study the effects of the latter market on the 'lit' market's quality measures, trading behaviour and welfare of traders. The results show that the presence of the 'dark' market harms the quality of the incumbent exchange (*i.e.*, reduced liquidity, quoted spread increase, the effective spreads of all market participants increase). All market participants across the two scenarios make higher profits when they participate in the 'lit' market under the Multi-Markets scenario.

Chapter 3 looks into the recent regulatory interests on 'dark' trading (*e.g.*, European Commission, 2010; Securities and Exchange Commission, 2010). To address the policy question a competitive setting is considered - with two different groups of agents - where two limit order markets operate simultaneously under two different scenarios. This chapter looks into whether the different price mechanisms applied to the 'dark' market - midpoint dark pool ('mid-market' scenario) and Volume-Weighted Average Price (VWAP) dark pool ('VWAP-market' scenario) - affect traders' welfare (and behaviour) and the 'lit' market's quality. The results suggest that across the two scenarios the 'VWAP-market' scenario benefits the 'lit' market's quality while only a certain group of agents (speculators) make less profits in this environment as compared to the alternative scenario.

The aim of chapter 4 is to study - in a laboratory environment - subjects' decision making under risk in a financial asset setting, with contingent claims fixed across frames. By testing the fundamental and untested hypothesis imposed on the majority of the asset demand tests (*i.e.*, Arrow-Debreu contingent claim setup); the study focuses on the application of revealed preference to the model of choice under risk and under uncertainty. Applying Afriat's Theorem (Afriat, 1967) we check if the experimental data is consistent with the maximization hypothesis. The results suggest that the power of the test is close to perfect; while the test's predictive success is within the positive range.

Acknowledgements

What a journey this has been. I would like to acknowledge the Department of Economics / School of Business (Economics Division) for accepting me and financing my studies. Sam thank you very much for your support especially during the probation and the extension periods. I have met great people in my academic and personal life that made this journey special to me. Without them I could have never made it.

First and foremost, I will be forever indebted to my supervisors and co-authors Prof. Dan Ladley, Dr. David Rojo-Arjona and my ex-supervisor Dr. Matthew Polisson. They have been always there for me encouraging me and supporting me during all these years. I am extremely grateful to Prof. Dan Ladley for introducing me to dark pools literature and providing me the opportunity to work with him. He took me in from the maths department and gave me the best training and guidance any PhD student could have ever asked for. David (Dr. David Rojo-Arjona) you are an inspiration for me. Without your continuous support especially over the last two years, this thesis would not have been possible. I am honoured to be counting David among my friends. Last but not least I would like to thank Matt (Dr. Matthew Polisson) for introducing me to revealed preference theory and for always being available to discuss any issues. I will never forget the hours that we spend in his office when he first introduced me to Afriat's Theorem. I am honoured to have been their student.

I also want to thank Subir Bose for accepting me to this program and for providing great advise and support. A big thank you to Andre Stenzel, Aris Boukouras and Dimitris Varvarigos for the their insightful comments and for the time we spent discussing.

Special thanks to my academic siblings: Ioli and Pasqua (soon you will be free - all the best for your theses), my two Italian legends Matteo and Alessandro, Alex (I love you so much), Stelios (the bee keeper), Nikos, Stamatoula, Aggelos and of course the new additions Natassa and Arpita (what more to ask from you - extremely grateful for all your help).

Thank you Giacomo, Gianfranco, Sneha, Narges, Sebastian, Laura, Valeria, Katerina, Maria, Rui, Robin, Sonia, Barbie and Maria for all the great moments that we shared together.

I have these people that I love with all my heart: Sofia, Simon, Kostas, Kostis, Katerina, Nadia and Dimitris; I feel so lucky having you in my life.

Becky, the words to describe my feelings for you are not enough; thank you for your continuous support; you are always there for me - I could not have asked for anything more. I am so proud for the things that you accomplished; and even more proud that you are part of my life. I love you so much (thank you soarpoint!). Linda I hope one day you will find what you are looking for, you deserve the best.

Last but not least, I would like to thank my parents Maria and Christos; I owe everything to you. Thank you for all your support, your love and your dedication. I miss you so much and I cannot wait to be close to you again. Georgia no matter what you do, no matter how you behave I will never stop loving you.

Declaration

Chapter 2 “Dynamic Equilibrium in Limit Order Markets: Analysis of Microstructure Characteristics, Trading Behaviour and Welfare” and chapter 3 “Dynamic Equilibrium in Limit Order Markets: Midpoint Dark Pool Vs VWAP Dark Pool, Analysis of Market Quality, Trading Behaviour and Welfare” are joint work with Dr. Alejandro Bernales, Prof. Dan Ladley, Rodrigo Orellana and Dr. Marcela Valenzuela. The author names are listed alphabetically. An early version of chapter 2 has been presented at the 4th Young Finance Scholars’ Conference 2017 (University of Sussex, Brighton) under the title “Dynamic Equilibrium in Limit Order Markets: Analysis of Depth Disclosure and Lit Fragmentation”.

Chapter 4 is a joint work with Dr. David Rojo-Arjona.

Contents

List of Tables

List of Figures

Chapter 1

Introduction

This thesis examines the behaviour of different types of individuals (*i.e.*, traders, experimental subjects) in financial markets under two different settings. That is, in chapters 2 and 3 we model the behaviour of different groups of traders in dynamic (limit) order book based financial markets. In contrast, in chapter 4 we develop and implement - in a laboratory setting a financial asset demand test where subjects choose under risk the allocation of wealth between two assets.

Without publicly displaying orders, dark pools have increased their market share - in recent years - in global equity markets. The increased percentages of 'dark' trading (*e.g.*, Banks, 2014; Securities and Exchange Commission, 2015) coupled with the decline in order sizes and depths on the incumbent exchanges (*e.g.*, Chordia, Roll, and Subrahmanyam, 2011) have triggered regulatory and academic concerns on the effect of these Alternative Trading Systems (ATS) on the information fragmentation, liquidity, market competitiveness and fairness. Another concern related with the order migration away from the traditional exchanges towards dark pools is the degree that these venues have on the distribution of welfare among market participants (retail and institutional investors); since these venues were used primarily by institutions in order to avoid market impact or adverse price movements. Yet the effects of 'dark' trading are still debatable. These trading platforms are important - high market share - and there is limited and contradictory empirical or theoretical work on dark pool activity (*e.g.*, Ye, 2010; Zhu, 2013, (price discovery); or Degryse, De Jong, and Kervel, 2015; Foley and Putniņš, 2016, (market quality)).

In chapters 2 and 3, motivated by the aforementioned, we consider dynamic continuous-time models of trade in a single financial asset, in order to capture the concerns related to 'dark' trading. Given the multidimensionality of the state space, we employ in each model the numerical approach of Pakes and McGuire (2001) and Goettler, Parlour, and Rajan (2009) in order to derive the optimal strategies adopted by the market participants. Optimal strategies are state dependent and each trader submits an order in order to maximize his expected payoff.

Chapter 2, sheds light on the effects that 'dark' trading has on traders (*i.e.*, trading behaviour and welfare) and on the quality of the 'lit' market. To be able to answer the above, we test two scenarios; under scenario I (Single-Market) trading takes place on a single limit order book that offers at any instant full pre-trade transparency. Under scenario II (Multi-Markets) two marketplaces exist and operate simultaneously, a 'lit' market and a 'dark' market. To be consistent with how real financial markets operate, the 'dark' market at any instant offers no pre-trade transparency. All traders can trade on share of the asset and each of them has an individual private value for the asset, allocated before they take any trading decision. Analysis of the model shows that the 'dark' market harms the 'lit' market's quality. Moreover under the Multi-Markets scenario, liquidity suppliers (speculators) under the risk of being picked off submit less limit orders to the 'lit' market as compared to the alternative scenario. Due to the different trading strategies adopted by the market participants under the Multi-Markets scenario, all trades irrespective of valuation make higher profits when they participate in the 'lit' market as compared to the Single-Market scenario.

Chapter 3 focusses on the regulatory concerns related to 'dark' trading. This chapter addresses the policy question related to whether or not 'dark' trading increases the traders' profits - by adopting different trading strategies - and the market quality of the incumbent exchange. For our purposes we test two scenarios: (i) a 'mid-market' scenario where a 'lit' market competes with dark pool that executes orders at any instant at the midpoint of the incumbent exchange; and (ii) a 'VWAP-market' scenario where the 'dark' market matches traders orders at any instant at the Volume-Weighted Average Price (VWAP). Irrespective of the scenario, the market contains two groups of agents: large traders who can trade $z > 1$ shares of the asset and small traders who can trade one share of the asset. Our results indicate that across the two scenarios there is an increase in liquidity supply in the 'lit' market under the 'VWAP-market' scenario. Moreover, across the two scenarios, the trading cost of submitting a market order that executes against an outstanding limit order under the 'mid-market' scenario, is higher. Traders with $\beta = 0$ (speculators) are the only market participants who make less profits when they participate in any market under the 'VWAP-market' scenario as compared to the alternative scenario.

Chapter 4 focuses on the recent strand of papers that have been developed and implemented - in a laboratory setting - related to nonparametric asset demand tests (*e.g.*, Varian, 1983, 1988; Green and Srivastava, 1986; Kübler, Selden, and Wei, 2014; Choi et al., 2007, and; Polisson, Quah, and Renou, 2017). In most of these models, asset demand tests make use of the Arrow-Debreu contingent claim setting; that is an individual posses preferences over state contingent consumption. Markets are assumed to be complete (i.e., same number of assets and states of the world) so observations can be transformed to contingent claim demands and prices. The focus of this chapter is to test this fundamental and untested hypothesis used in financial asset demand tests. To do that, we design an experiment in which each subject is asked to purchase securities under four different scenarios in the asset space with contingent claims *fixed* across frames (*i.e.*, scenario I: Arrow-Debreu securities, scenarios II, III, IV: ordinary securities). Subjects are facing risk where the probabilities of the two available states are objectively known. Under certain assumptions, we are able to generate - for each of the four scenarios - an equivalent environment in the asset space were subjects are asked to make asset demand choices. Using Afriat's Theorem (Afriat, 1967) we are able to apply a *nonparametric* test of utility maxmization within and across the available scenarios. In our analysis, we observe that the test's predictive success across and within scenarios is within the positive range. That is, the test has explanatory power against the alternative hypothesis of random choice. The power of the model is close to perfect given that the probability of a random dataset to be consistent with the utility maximization model across all scenarios at an efficiency index of 0.85 is effectively zero.

Chapter 2

Dynamic Equilibrium in Limit Order Markets: Analysis of Microstructure Characteristics, Trading Behaviour and Welfare

2.1 Introduction

The organization of trading in the financial markets has encountered enormous adjustments at the end of the 20th century. Quote-driven markets (dealer markets) have been substituted progressively by order-driven markets. While in quote-driven markets trades are only executed by dealers, who quote the bid and ask prices, in order-driven markets there are no designated intermediaries.[1] In these markets, trades are matched through computers, clerks, or brokers while traders interact directly with each other by submitting orders that are sorted according to certain trading rules.[2]

To be in line with the empirical findings of Jain (2005) the focus of this study is based on an electronic market environment (order-driven computerized system) where traders submit orders through an electronic trading platform (*i.e.*, electronic limit order book). Jain (2005) by examining the impact of market microstructure on asset pricing, observes that the leading stock exchange in 101 out of the 120 sample

[1] Some markets (hybrid) often combine features from both quote and order-driven protocols. Some typical examples of these markets are the NYSE, AMEX (organised by specialists who manage the order book and guarantee liquidity), NASDAQ and LSE.

[2] In all exchanges and electronic platforms traders that offer the most competitive prices are queued first. The next rule that most of the systems follow is the time rule, where priority is given to traders that improved first the prevailing quotes. Finally some markets give priority to displayed against non-displayed orders or obey the size precedence rule.

countries has introduced screen-based electronic trading.

Equity trading has seen a rapid increase in the number of new trading venues (*e.g.*, Alternative trading systems (ATSs) in the US and multilateral trading facilities (MTFs) in Europe). Trading systems structured as electronic limit order books (*i.e.*, electronic communication networks (ECNs)), generally offer pre- and post-trade transparency ('lit' venues).[3] Trading interests and prices are publicly available to the market participants prior to execution, while the transaction details of executed orders are announced in real time. Dark pools in contrast with ECNs, are alternative equity trading systems that operate with limited or no pre-trade transparency ('dark' venues). Details of the orders (*i.e.*, price and quantity) submitted in these off-exchange trading platforms, are not displayed to the other market participants. Any order executed in these venues is reported in a consolidated tape (in the U.S. equity markets), a mechanism that offers post-trade transparency.[4]

The motivation of this study lies on (i) the significant high percentage of 'dark' trading, in terms of average daily trading volume, in the major global equity markets (i.e. U.S., Europe, Canada, Australia, Asia) over the past few years. According to Banks (2014) at the largest U.S. platforms, the average daily trading volume absorbed by dark pools in 2013 ranged from 100 to 300 million shares; (ii) the empirical findings of Chordia, Roll, and Subrahmanyam (2011), that point clear evidence on the continuous decrease of the order sizes and depths on exchanges; and (iii) the fundamental regulatory concerns linked with 'dark' trading that include the potential impact on the price discovery process, on the fragmentation of information, liquidity and on market competition and fairness. Yet, the effect of dark trading on quoted spreads and on market quality as a whole is still debatable (see for example empirical findings of Degryse, De Jong, and Kervel, 2015; Foley and Putniņš, 2016). Another concern related with the order migration away from the traditional exchanges towards dark pools is the degree that these venues have on the distribution of welfare among market participants (retail and institutional investors); since dark pools primary target was to 'protect' institutions from trading large blocks of shares.

[3] Some global stock exchanges, such as the NASDAQ OMX, Deutsche Boerse/Xetra and LSE, feature non-displayed orders. These orders are either 'controlled' by specialists within their order books or held by the exchange.

[4] For further discussion on the ATSs/MTFs and primarily dark pools see Subsection 2.1.1.

Thus, this study will focus on answering the following question: "How does the microstructure of a dark pool affect the 'lit' market's quality measures, trading behaviour and welfare of market participants ?"

The effects of this market mechanism are extremely interesting. Minute effects on these markets have massive effects on the rest of the economy; by absorbing huge amounts of trade from the traditional exchanges, they generate huge turnover in the financial markets. These venues are important (*i.e.*, high market share) and there is limited empirical or theoretical work on dark pool activity and how these markets behave.[5] Dark pools nowadays add structural complexity in markets since the design and the trading rules they obey are completely dependent on whoever operates them (i.e. banks, agency brokers, exchanges).[6] As a result, there is limited and contradictory evidence of whether their design is sensible, if it encourages predatory trading or generates efficiency within the market.[7]

Given the aforementioned, this study tests two scenarios: (i) a Single-Market scenario where the set of preferences of all market participants is aggregated on a limit order book related to a 'lit' market; and (ii) a Multi-Markets scenario where liquidity is scattered between two marketplaces - a 'lit' market and a 'dark' market - that operate simultaneously. Under the latter scenario (Multi-Markets) each market has a respective limit order book and different levels of transparency. The 'lit' market offers at any instant full pre-trade transparency; that is the limit order book related to the 'lit' market reveals all the relevant information to the market participants concerning the available trading opportunities for any available price. In contrast the limit order book related to the 'dark' market offers no pre-trade transparency (*i.e.*, trading opportunities are not observable to the agents).

We consider a dynamic continuous-time model of trade in a single financial asset. Traders are risk neutral, observe the fundamental value of the asset at any instant with a time lag and have different motives to trade. An agent's private value, denoted by β, is drawn independently across traders from a discrete distribution $\mathscr{F}_\beta$ with finite support. Agents with the most extreme private values are keen to trade quickly since their main benefits arise from their exogenous private values.

[5] (*e.g.*, Ye, 2010; Zhu, 2013, (price discovery); or Degryse, De Jong, and Kervel, 2015; Foley and Putniņš, 2016, (market quality)).

[6] For further discussion on the different types of dark pools see Subsection 2.1.1.

[7] Brunnermeier and Pedersen (2005) define predatory trading as "trading that induces and/or exploits the need of other investors to reduce their positions".

In contrast agents with zero private values may be willing to trade only when the fundamental is mispriced. Market participants can trade one share of the asset and under the Multi-Markets scenario they can choose either the 'lit' or the 'dark' market in order to submit their order. If their order does not execute, traders re-enter the market and revise their order. Traders can re-enter the market multiple times until execution.

Given the properties of the 'dark' market - under the Multi-Markets scenario - an agent's trading decision depends on the prevailing market conditions of the 'lit' market (state-dependent action). Conditional on the state of the trading game each trader submits an order in order to maximize his expected discounted payoff. Since our model does not have a closed-form solution we follow the numerical approach of Pakes and McGuire (2001) - originally used by Goettler, Parlour, and Rajan (2005) - and Goettler, Parlour, and Rajan (2009) so to be able to find a Markov-perfect equilibrium in which a trader irrespective of valuation acts optimally. Once the equilibrium is found we simulate the model for some billion order submissions were the dynamics of the model are examined.

Our key findings include: (i) (market quality) the presence of the 'dark' market under the Multi-Markets scenario harms the 'lit' market's quality. Liquidity is reduced; the quoted spread increases and the number of limit orders at the best quotes and at both sides of the limit order book decrease. Moreover the effective spreads of all market participants are greater and the 'lit' market's informational inefficiency increases; (ii) (trading behaviour) agents with zero private values decrease their liquidity provision in the LM under the Multi-Markets scenario. The 'lit' market's wider spread changes the order placement strategies adopted by the agents with nonzero private values (and as a result reduces the 'lit' market's microstructure noise). The above coupled with the properties of the DM (no pre-trade transparency) under the Multi-Markets scenario indicate that the placement of non-aggressive limit orders (*i.e.*, at the best quotes) to the LM involve more risk and are associated with a high probability of being picked off; (iii) (trading behaviour) the non-display properties of the 'dark' market drive agents with no intrinsic motives to trade; to place more aggressive limit orders in that venue; (iv) (trading behaviour) returning agents with zero private values prefer the 'dark' market for the placement of their orders. For returning agents with nonzero private values this is less of an issue given their exogenous motives to trade and the 'lit' market's reduced microstructure noise; (v) (trading behaviour) across the two scenarios, although the agents with nonzero private values tend to be a little more conservative in the 'lit' market - in terms of mean price of submitted limit orders - they reduce their picked off probabilities, increase their percentages of limit orders executing and obtain better terms of trade in that venue; and (vi) (welfare) agents irrespective of valuation across the two scenarios make higher profits when they participate in the LM under the Multi-Markets scenario.

Our results stand in line with the theoretical findings of Zhu (2013) and Buti, Rindi, and Werner (2017) and the empirical findings of Nimalendran and Ray (2014); Degryse, De Jong, and Kervel (2015) and Kwan, Masulis, and McInish (2015). Zhu (2013) examines the impact the 'midpoint' dark pool has on price discovery and shows that the incumbent exchange (dark pool) is more attractive to informed (uniformed) traders, and as a result harms the 'lit' market's liquidity. Buti, Rindi, and Werner (2017) show that the introduction of a dark pool that competes with a 'lit' market widens the spread and lowers the depth of the latter. Nimalendran and Ray (2014) analyse data from one US dark pool and conclude that trading in the 'dark' market increases the spread and price impact on the 'lit' market. Degryse, De Jong, and Kervel (2015) using data from 52 Dutch stocks in the Amsterdam Exchange (AEX) with large and mid-cap indices conclude that dark trading has a negative effect on the order books' liquidity. Kwan, Masulis, and McInish (2015) use a dataset of US off-exchange trading classified into five dark venues and show that the ability to queue jump in some dark venues discourages market participants from providing liquidity to the incumbent exchanges and hence widens the spread and lowers the depth of these markets.

Our results stand against the theoretical findings of Boulatov and George (2013) and the empirical findings of Foley and Putniņš (2016). Boulatov and George (2013) by examining the impact of hidden liquidity on market quality they show that opacity in limit order markets enhances liquidity and the market's informational efficiency. Foley and Putniņš (2016) using the introduction of the minimum price improvement regulation in Canada and Australia examine the effects that 'two-sided' dark trading has on market quality. Their results suggest that 'dark' markets benefit market quality by reducing the quoted and effective spreads and increasing the informational efficiency of the market.

2.1.1 Diving into Dark Pools

In 2005, the market share of dark pools in equity markets was low. Dark pools' primary purpose was to minimize information leakage by allowing certain investors, *i.e.*, institutions, to anonymously trade large blocks of shares (see Zhu, 2013). Thus, the benefits associated with the choice of these venues was the protection they offered to large institutional investors against predatory trading and unfavourable price movements. The adoption of new regulations, the evolution of communication networks together with the rise in high frequency trading, and decimalization brought changes in the market structure of equity trading systems (see Banks, 2014). In the U.S., the adoption in 2005 and full implementation in 2007 of the Regulation National Market System (Reg NMS) with its main target to enhance competition with incumbent exchanges, led to the introduction of new alternative electronic trading venues (ATSs).[8] Similarly in Europe, in force since November 2007, the Markets in Financial Instruments Directive (MiFID) enabled fast expansion to multilateral trading facilities (MTFs).[9]

According to Securities and Exchange Commission (2010), as of September 2009, liquidity in the U.S. equity markets was scattered among 5 electronic communication networks (ECNs), almost 10 registered exchanges, 32 dark pools and 200 broker-dealers. Registered exchanges and ECNs ('lit' venues) offer pre- and post-trade transparency since the consolidated quotation data is publicly available to market participants. These venues executed approximately 74.6% of the share volume in NMS stocks. Dark pools and broker-dealer internalization (dark venues) by matching orders anonymously offer limited or no pre-trade transparency. These venues executed 25.4% of the share volume.

In recent years, not only the market share but also the number of ATSs, and particularly that of dark pools, has increased even more. From the third quarter of 2009 until the second quarter of 2015, the number of dark pools has increased from 32 to 50. Moreover, the market share of these dark venues approximately doubled from 7.9% to 14.9% (Securities and Exchange Commission, 2015). As of September 2015, based on data available from BATS Chi-X Europe, the five days average of total con-

[8] For further discussion on rules adopted by the Securities and Exchange Commission (SEC) under Regulation National Market System (Reg NMS), see Securities and Exchange Commission (2005).

[9] For further discussion on Markets in Financial Instruments Directive (MiFID) and the developments in market structures in Europe, see European Commission (2010).

solidated volume traded in dark pools accounted for almost 10% up from approximately 2.7% five years ago.

Apart from the aforementioned regulations that enhanced competition, a confluence of factors led to the rapid increase in market power of dark pools. The continuous decrease in consolidated equity volume (see for example Chordia, Roll, and Subrahmanyam, 2011) and the introduction of maker-taker or taker-maker pricing schemes on equity exchanges forced certain market participants (*i.e.*, institutions) to shift their attention towards off-exchange trading platforms.[10] The matching and pricing mechanisms the majority of dark pools follow offer not only more attractive fee structures but potential price improvements relative to the national best bid and offer (NBBO) of exchanges.

According to Zhu (2013) and Buti, Rindi, and Werner (2017) dark pools are subdivided into four main categories: (i) bank/broker pools, operated by banks, hold the majority of the market share in the U.S., Canada and Europe. By offering limited price discovery, the execution price these continuous nondisplayed limit order books follow, is bounded between the national best bid and offer (NBBO). Examples involve Goldman Sachs SIGMA X (SIGMA X MTF in Europe), Barclays LX and UBS PIN ATS (UBS MTF in Europe). (ii) Independent/agency pools, operated by agency brokers, match orders at the midpoint of NBBO or the Volume-weighted average price (VWAP). Examples involve, ITG Posit (Posit in Europe), Liquidnet and Knight Match. (iii) Exchange-Based pools, which are owned by exchanges and provide no direct price discovery since execution prices are derived form the NBBO; and (iv) market maker pools where liquidity is supplied by the manager of these venues. For our purposes in this model, we follow the trading mechanism of bank/broker dark pools.

This study is organised as follows. Sections 2.2 and 2.3 present the details of the model; Section 2.4 depicts the results and Section 2.5 concludes.

2.2 Preliminaries

The trading game involves traders who can trade one share of the asset. We denote the variable that controls the number of available shares a particular trader has

[10] See Harris (2013) for further information on the effects of maker-taker or taker-maker pricing schemes.

to trade by z, where $z \in \{0,1\}$. A factor that gives heterogeneity to the market participants, is their intrinsic private value to trade the asset.[11] We denote an agent's private value by β; this value - allocated to an agent before taking any action - is drawn independently across traders from a discrete distribution $\mathscr{F}_\beta$ with finite support. Traders with greater absolute private values in β are impatient to trade since the main benefits arise from their private values. Traders with a private value of zero (no intrinsic motivation to trade) may only be willing to trade profitably when the fundamental value of the asset is mispriced.

We consider a dynamic continuous-time model of trade in a single financial asset. At any time t, the asset has a common (fundamental) value, denoted as v_t, that evolves over time following a random walk.[12] The time between innovations in v_t follow a Poisson process with parameter λ_v. When an innovation occurs, the fundamental value equiprobably increases or decreases by σ ticks. Any change in the fundamental value reveals new information about the asset. Assume at time $t > 0$, the fundamental value is drawn to decrease by σ ticks and the next event drawn is the arrival of a new trader at time $t' > t$. This trader arriving at time t', is better informed compared to all active traders who submitted limit orders in the time interval $[0, t)$.

Trading takes place on two limit order markets that operate simultaneously, an open (transparent) 'lit' market and a 'dark' (opaque) market.[13] While the 'lit' market offers at any instant full pre-trade transparency the 'dark' market represents a typical form of a dark pool, in terms of the trading mechanism applied, that provides limited price discovery. The dark pool in our model is opaque, that is an order submitted to the 'dark' market is not observable to any other active market participant but the order submitter (no pre-trade transparency). At any instant this continuous nondisplayed 'two-sided' limit order book (*i.e.*, 'dark' limit order market) can execute buy and sell orders at different prices as long as resting limit orders exist on

[11] Private value for the asset is also present in the models of Biais, Hillion, and Spatt (1995); Parlour (1998); Theissen (2000); Handa, Schwartz, and Tiwari (2003); Goettler, Parlour, and Rajan (2005, 2009); Hoffmann (2014); Chiarella and Ladley (2016) and Buti, Rindi, and Werner (2017); among many others.

[12] The fundamental value is considered to be the expectation of the future cash flows on the stock.

[13] From this point and onwards (until Section 2.4); given that the Single-Market scenario is a subdivision of the Multi-Markets scenario (*i.e.*, only one limit order book), our focus lies only on the latter scenario.

the other side of the market.[14] Due to these properties a trader's decision whether or not to submit an order to the 'dark' market is endogenously determined at any instant by the prevailing market conditions of the 'lit' market.[15]

New traders arrive at the market according to a Poisson process with parameter λ_N. Upon first arrival, traders can submit an order of one share to either market or wait until the market conditions change (submit no order). Traders, after submitting an order which has not been executed, re-enter the market and revise their previous trading strategies after a random amount of time (for further analysis, see Subsection 2.3.1). The re-entry time is drawn from a Poisson process with parameter λ_R.

The model discourages agents from infinitely postponing trade by introducing a discount rate. The payoffs of order executions, are discounted back to the order's submission time, at rate ρ. This discount rate which is the same for all agents irrespective of valuation, captures costs related to delaying transactions or even potential lost trading opportunities.[16] That is the discount rate captures the agents' desire to trade sooner rather than later.

2.3 Model Description

Let $m \in \{m_L, m_D\}$ be a binary variable where $m = m_L$ if we are referring to the 'lit' market (LM) characteristics and $m = m_D$ if we are referring to the 'dark' market (DM) characteristics. As in real limit order markets, the limit order book related to the market m is described by a discrete set of prices such that $\mathscr{P}_m = \{p^i\}_{i=-\infty}^{+\infty}$. We denote the tick size, which is constant and defines the distance between any two consecutive prices, by $d_m^{\Delta p}$. At any instant $t > 0$, with each price $p^i \in \mathscr{P}_m$ there are backlogs of outstanding (unexecuted) limit orders ($\ell^i_{m,t}$), obeying time priority, to buy ($\ell^i_{m,t} > 0$) or to sell ($\ell^i_{m,t} < 0$) the asset. That is, the limit order book related to the market m at time t denoted by $\mathscr{L}_{m,t}$, is the vector of outstanding orders such

[14] For further discussion about the two types of dark trading (*i.e.*, 'one-sided', 'two-sided') and the effects that these have on market quality measures see Foley and Putniņš (2016).

[15] As we will discuss later, the market conditions at any instant are influenced by a range of factors such as optimal decisions adopted by the other active market participants (returning or new) or changes in the fundamental value of the asset.

[16] In our model the discount rate ρ does not represent the opportunity cost of capital related to an investment.

that $\mathscr{L}_{m,t} = \{\ell^i_{m,t}\}^{+\infty}_{i=-\infty}$.[17]

The quotes, at each instant of time, are defined by the orders in the limit order book related to the market m. The best bid price is the highest price at which there is a buy limit order on the book while the best ask price is the lowest price at which there is a sell limit order. That is, given a limit order book $\mathscr{L}_m$, the best bid and the best ask are defined as $B(\mathscr{L}_m) = \max\{p^i | \ell^i_m > 0\}$ and $A(\mathscr{L}_m) = \min\{p^i | \ell^i_m < 0\}$ respectively.[18]

In this model, all orders submitted are considered to be *marketable*. A buy order submitted at price p^i to market m and time t' is a *market* buy order if there exists an outstanding limit order at the same price on the sell side of the limit order book $\mathscr{L}_{m,t'}$. In contrast if there is no matching order on the sell side, the buy order submitted at time t' is a *limit* buy order and is placed on the book at the submitted price level and quantity (*i.e.*, for this particular example the order will be placed at the positive values in $\ell^i_{m,t'}$).[19] Continuing with the example above, if the price of the limit buy order submitted is greater than the best bid price such that $p^i > p^{i-1}$, we set $B(\mathscr{L}_{m,t'}) = \max\{p^i | \ell^i_{m,t'} > 0\}$. Finally an order submitted to sell (buy) the asset at price below (above) the bid (ask) price is executed immediately at the bid (ask) quote.

All agents that take part in the trading game, are risk neutral and they observe the fundamental value of the asset at time t with a time lag Δ_t ($\mathrm{v}_{t-\Delta_t}$). That is, traders have to estimate the actual fundamental value by observing the current (at time t) market conditions of the LM.[20] In this model an agent's trading decisions at any instant are state-dependent. An agent upon first arrival at time t, conditional on the state of the trading game, may choose to submit an order or not to submit any order. In the first case, the agent has to choose between $\mathscr{L}_{m_L,t}$ (LM) or $\mathscr{L}_{m_D,t}$ (DM), the submitted quantity (buy or sell one unit) and the price of the submitted order. The price of the order submitted to either market (at or above/below quotes) implies

[17] We adopt similar notation to Goettler, Parlour, and Rajan (2005, 2009); Bernales (2014); and Chiarella and Ladley (2016) to describe the microstructure features of the model for the dynamic order book market.

[18] If the order book $\mathscr{L}_m$ is empty on the buy side or on the sell side we set $B(\mathscr{L}_m) = -\infty$ or $A(\mathscr{L}_m) = +\infty$ respectively.

[19] Similar order characterizations are used if we focus on the buy side of the market m when a sell order is submitted.

[20] For further analysis on the learning process applied for the estimation of the fundamental value see Subsection 2.3.3.

the agent's trading decision to submit a market or a limit order.

After submitting an order which has not been executed, agents re-enter the market and monitor their previous unfilled orders. Upon re-entry, agents can either keep or cancel their existing limit order. In the latter case, agents can immediately submit a new order or wait for different market conditions to apply, in order to submit a new one. Cancellation and re-submission of an unfilled limit order requires from agents to choose the market (LM or DM), the quantity (buy or sell) and the price of the new order, given the new state of the trading game.

The benefit of keeping an existing order in the LM or in the DM, is that the order's time priority in the queue is reserved. Conversely, the cost of keeping an order is that the value of the asset might have moved in a direction before the trader's re-entry that will affect the expected payoff from that order. For example, suppose at time $t > 0$, a trader submits an order to the LM to buy the asset at price p^{buy} when the fundamental value is v_t. Suppose further that the depth on the other side of the market m is such that the order represents a limit buy order. The trader's net common benefit for submitting an order of one share at time t is $\mathrm{v}_t - p^{buy}$. If at time $t' > t$, before the agent's re-entry at time t_{re}, the following events occur:

$$\begin{cases} \mathrm{v}_{t'} = \mathrm{v}_t - \sigma, & \mathrm{v}_t \textit{ has fallen to } \mathrm{v}_{t'}, \quad (2.1) \\ \mathrm{v}_{t'} = \mathrm{v}_t + \sigma, & \mathrm{v}_t \textit{ has risen to } \mathrm{v}_{t'}. \quad (2.2) \end{cases}$$

In the first event, the agent's net common benefit reduces since $\mathrm{v}_{t'} - p^{buy} \overset{(2.1)}{=} (\mathrm{v}_t - \sigma) - p^{buy} = \mathrm{v}_t - (p^{buy} + \sigma)$. In the second event, the agent's net common benefit increases and there is smaller chance of the order to be filled, since $\mathrm{v}_{t'} - p^{buy} \overset{(2.2)}{=} (\mathrm{v}_t + \sigma) - p^{buy}$. The former event, captures the agent's implicit transaction cost of being 'picked off' when the fundamental value ($\mathrm{v}_{t'}$) has moved in an unfavourable direction after a limit order has been submitted.

Limit order submitters are active members of the trading game, by revising their orders until they execute. The model allows agents with unfilled orders to re-enter the market multiple times. However, agents stop being part of the trading game, and leave the market forever, after their order executes. Despite the fact that the (re-) entry rate - drawn from a Poisson process with parameter (λ_R) λ_N - is the same for all agents; trading decisions are endogenously determined by the prevailing market conditions of the LM and the number of active agents. As a result, there is a random number of active market participants at any instant, that monitor the market.

2.3.1 Traders Maximization Problem

Consider a trader who submits an order to either market at time $t > 0$. Assume further that the order is fully-matched with its respective limit order that stands on the other side of the book. That is, the order submitted represents a market order. The trader's instantaneous utility at time t is given by:

$$u_t = \begin{cases} \beta + \left(\mathrm{v}_t - p^i\right) & \text{if the trader's order is a buy order executed at price } p^i, \\ \left(p^i - \mathrm{v}_t\right) - \beta & \text{if the trader's order is a sell order executed at price } p^i. \end{cases} \tag{2.3}$$

Alternatively, if the order is placed in the book as a limit order (no matching order on the other side), then the trader's instantaneous utility is equal to zero. Recall that the fundamental value evolves over time following a random walk, and thus the price at which a trade can occur can be unbounded. Nevertheless, as it can be observed from Eq. (2.3) a trader conditional on the value of β, cares only about the price relative to the fundamental rather than the absolute price of a trade.

Each trader in the game conditional on the state of the LM he observes at any instant, submits an order to either market to maximize his expected discounted payoff. We define as $\alpha = \left(m, x, p\right)$ the action taken by an agent who arrives at the market at time t. Here, p denotes the price at which a trader submits an order.[21] Let $m \in \{m_L, m_D\}$ denote the limit order book choice, where $m = m_L$ if the trader chooses the LM ($\mathscr{L}_{m_L,t}$) and $m = m_D$ if the trader chooses the DM ($\mathscr{L}_{m_D,t}$). Finally, x denotes the quantity of the submitted order to buy or sell one unit of the asset such that

$$x = \begin{cases} +1 & \text{if a buy order of one unit is submitted,} \\ -1 & \text{if a sell order of one unit is submitted,} \\ 0 & \text{if no order is submitted.} \end{cases} \tag{2.4}$$

We denote the state observed by a trader with private valuation β on a particular

[21] The superscript for the price parameter p is removed for ease of notation.

entry to the market at time $t > 0$, by $s = (mc_{m_L,t}, \beta, \alpha, z_R) \in \mathscr{S}_\beta$. Here $mc_{m_L,t}$ are the market conditions of the LM at time t and α is the status of the trader's previous action.[22] Moreover the variable $z_R \in \{0, 1\}$ indicates the remaining number of shares a trader has to trade; while $\mathscr{S}_\beta$ is the set of possible states that particular trader may face. We define the market conditions at any instant $t > 0$ by the three-tuple

$$mc_{m_L,t} = \left\{ \overset{\star}{\mathscr{L}}_{m_L,t}, \mathrm{v}_{t-\Delta_t}, q_{m_L,t}(\cdot) \right\}. \tag{2.5}$$

Here, $\overset{\star}{\mathscr{L}}_{m_L,t}$ is a set of variables that characterize the 'open' limit order book, $\mathrm{v}_{t-\Delta_t}$ is the lagged fundamental value and $q_{m_L,t}(p, x) \geq 0$ is the trader's order priority among all its respective orders in terms of price and quantity.

Suppose at time $t > 0$, a new trader by observing state $s = (mc_{m_L,t}, \beta, \emptyset, z_R)$ takes an optimal action by submitting an order at price p to the LM. If the action does not represent a market order we define the order's priority as $q_{m_L,t}(p, x) = |\ell^p_{m_L,t} + x|$, where $\ell^p_{m_L,t}$ denotes the depth of the limit order book related to m_L.[23] Upon re-entry at time $t' > t$, if the trader's order has not been filled ($z_R \neq 0$) the trader by observing the state of the trading game, which includes the term $q_{m_L,t'}(\cdot) \leq q_{m_L,t}(\cdot)$, can take further actions in order to maximize his expected payoff. Thus the trader's order priority might evolve over time due to the trading decisions adopted by the rest active market participants before the agent's re-entry to the market at time t'.

On any entry, a trader conditional on the available information set, chooses an optimal action. For computational tractability if a trader's optimal action in state $s \in \mathscr{S}_\beta$ and time $t > 0$ does not represent a market order (*i.e.*, a limit order) the price of it is bounded within k ticks away from the agent's expectation of the fundamental value of the asset.[24] Denote the agent's expectation of the fundamental at time t by $\hat{\mathrm{v}}_t = E(\mathrm{v}_t \mid mc_{m_L,t})$. Thus, the set of possible actions a trader can take in state s, denoted as $A(s)$, is defined to be

[22] If the trader is entering the market for the first time we set the state variable $\alpha = \emptyset$.

[23] It is worth highlighting that although in Eq. (2.5) our focus lies on the prevailing market conditions of the LM, the priority of an order submitted at price $p \in \mathscr{P}_{m_D}$ to the DM is defined in the same way.

[24] To ensure that the fixed point is not affected we choose a large value for k (for further analysis see Subsection 2.3.4).

$$
\begin{aligned}
A(s) = \Big\{\emptyset, \big(m, x, p\big) \mid \ & p \in [\hat{\mathrm{v}}_t - k, \hat{\mathrm{v}}_t + k] \cap \mathscr{P}_m, \\
& m \in \{m_L, m_D\}, \\
& x \in \{-1, +1\} \Big\}.
\end{aligned} \tag{2.6}
$$

Consider a particular trader entering at the market at time $t > 0$ and taking an optimal action $\bar{\alpha} = (\bar{m}, \bar{x}, \bar{p}) \in A(s)$ by submitting an order to the DM. Assume further that the trader re-enters the market at some future random time $t' > t$ and observes state s'. Upon re-entry, if the order is still active (*e.g.*, $z_R \neq 0$), in the new state $s' = \big(mc_{m_L,t'}, \beta, \bar{\alpha}, z_R\big) \in \mathscr{S}_\beta$ the trader has the option not only to keep, but also to cancel the order and submit a new one of z shares to either market. That is, each action $\bar{\alpha} \in A(s)$ generates an expected payoff that consists of: (i) a payoff given that the order executes, *i.e.*, $z_R = 0$, at the time interval $[t, t']$ before the trader's re-entry to the market; and (ii) a value related to the trader's re-entry to the market when $z_R \neq 0$.

The probability of a limit order - submitted in state s - to be filled depends on the strategic decisions adopted by the active agents in other states. Each order submission alters the limit order book $\mathscr{L}_{m_L,t}$ ($\mathscr{L}_{m_D,t}$) and as a result the state of the trading game for the later traders. Assume at time $t = 0$ a trader's optimal action $\hat{\alpha} = (\hat{m}, \hat{x}, \hat{p})$ when facing state s is to submit an order at price $\hat{p}$ to market m.[25] The expected payoff of this order being filled at price $\hat{p}$ and time t_{ex} before the trader's re-entry to the market at time t_{re}, is given by

$$
\pi(t_{re}, \hat{\alpha}, s) = \int_0^{t_{\mathrm{re}}} \int_{-\infty}^{+\infty} \left(e^{-\rho\, t_{ex}} \big(\big(\beta + \mathrm{v}_{t_{ex}} - \hat{p} \big)\, \hat{x} \big)\, \eta(t_{ex} \mid \hat{\alpha}, s) \right) \gamma(\mathrm{v} \mid s, t_{\mathrm{ex}})\, d\mathrm{v}\, dt. \tag{2.7}
$$

The payoff from execution before re-entry at the time t_{ex}, is discounted back to the trader's first arrival to the market, at rate ρ. The term $\beta + \mathrm{v}_{t_{ex}} - \hat{p}$, captures the instantaneous payoff of the order at time t_{ex} (see Eq. (2.3)), where $\mathrm{v}_{t_{ex}}$ is the fundamental value of the asset at that time. As introduced in Eq. (2.4), $\hat{x}$ is the quantity of the submitted order and $\gamma(\mathrm{v} \mid s, t_{\mathrm{ex}})$ is the density function of the fundamental value of the asset at the execution time given state s. Additionally, $\eta(t_{ex} \mid \hat{\alpha}, s)$ denotes the prob-

[25] Normalizing the trader's entry time to zero.

ability that the trader's action $\hat{\alpha}$ in state $s = \left(mc_{m_L,t=0}, \beta, \emptyset, z\right)$ at time $t = 0$ results in execution at time t_{ex}, before his re-entry to the market at time t_{re}. If the trader's optimal action $\hat{\alpha}$ in state s and time $t = 0$, is to submit an order to either market that results in immediate execution we set $\eta(0 \mid \hat{\alpha}, s) = 1$.

Consider now the alternative case in which the trader's previous order submitted at price $\hat{p}$ to market m and time $t = 0$ is still unfilled. The trader by re-entering the market at time t_{re}, he can take further actions given the new state $s' = \left(mc_{m_L,t_{re}}, \beta, \hat{\alpha}, z_R \neq 0\right)$ of the LM. The trader can (i) cancel and resubmit a new order to the same market m at a different price $\overset{\star}{p} \neq \hat{p}$ and quantity $\overset{\star}{x} \pm 1$; (ii) cancel and resubmit an order to the alternative market at price $\overset{\star}{p} \neq \hat{p}$ or $\overset{\star}{p} = \hat{p}$ and quantity $\overset{\star}{x} \pm 1$; (iii) cancel and submit no order ($\overset{\star}{x} = 0$); or (iv) keep his unfilled order in market m at the same price such that $\overset{\star}{\alpha} = (\overset{\star}{m} = \hat{m}, \overset{\star}{x} = \hat{x}, \overset{\star}{p} = \hat{p})$. The value $\mathrm{V}(s)$ to the agent of being in state s is given by the Bellman Equation of the trader's maximization problem

$$\begin{aligned} \mathrm{V}(s) = \max_{\hat{\alpha} \in A(s)} \int_{t_{re}=0}^{+\infty} \Bigg[\pi(t_{re}, \hat{\alpha}, s) + e^{-\rho\, t_{re}} \int_{s' \in \mathscr{S}_\beta} & \mathrm{V}(s') \times \\ \times \psi\left(s' \mid \hat{\alpha}, s, t_{re}\right) ds' & \Bigg] dG(t_{re}). \end{aligned} \tag{2.8}$$

Here, $\mathscr{S}_\beta$ is the set of possible states a trader with private valuation β may face on re-entry. $G(t_{re})$ is the exogenous probability distribution of the re-entry that follows an exponential distribution with parameter λ_R (see Section 2.2). In addition, $\psi(s' \mid \hat{\alpha}, s, t_{re})$ is the probability that at time t_{re} on re-entry the trader faces the new state s', given previous action $\hat{\alpha}$ and previous state s. The first term is defined in Eq. (2.7); while the second term reflects the subsequent payoff on re-entry at time t_{re}.

2.3.2 Existence of Equilibria

On each entry to the market, a particular trader facing state $s = \left(mc_{m_L,t}, \beta, \alpha, z_R\right) \in \mathscr{S}_\beta$, takes an action $\hat{\alpha} = (\hat{m}, \hat{x}, \hat{p})$ that maximizes his expected discounted utility, as introduced in Eq. (2.8). In this model, the optimal action $\hat{\alpha} \in A(s)$, depends on the prevailing market conditions of the LM (state-dependent action). That is the state of the game at any instant contains all the relevant information required for the trader in order to choose his optimal action.

Given the set of possible actions a trader can take in a given state s, all the elements of b and x are integers (see Eq. (2.6)). Moreover, by bounding the prices of the submitted limit orders k ticks away from the agent's expectation of the fundamental value of the asset (*i.e.*, finite set of prices); we obtain a discrete and finite set of decisions. The state space of the trading game at any instant can change either due to fundamental value changes or optimal trading decisions adopted by the active market participants (new or returning). Since these events can occur a countable number of times within a finite simulation, it follows that the state space is countable. Based on the above it follows that the game has Markov-perfect equilibria (see for example Rieder (1979); Fudenberg and Tirole (1991)). We focus on stationary symmetric equilibria, where traders with the same type choose the same strategy, and this strategy does not depend on the arrival time of traders to the market.[26]

2.3.3 Model Solution of the Trading Game

Ideally, we would like traders to take optimal actions conditional on the entire information set of the trading game. In practice, this makes the model analytically intractable due to the size of the state space. Therefore, to obtain a numerical approximation of the equilibrium, we follow the numerical approach of Pakes and McGuire (2001). Similar to Goettler, Parlour, and Rajan (2005, 2009) and Chiarella and Ladley (2016) the equilibrium concept of the game is an approximation of the true Markov-perfect equilibrium, in which traders' optimal strategies are conditional on the available information set of the reduced state space.[27]

Agents in the trading game start with initial expected payoffs of different action and state pairs. The learning process allows us firstly to dynamically update traders' expected payoffs of different actions and states by observing the realized payoffs of their actions. The equilibrium is reached when traders' expected payoffs of each action and state pair match the realized outcomes. Secondly, given that the full state space of the trading game is too big, the Pakes and McGuire (2001) algorithm allows us to reduce the dimensionality of the state space size by updating the agents' expected payoffs only for states that are actually visited. That is, states that belong to

[26] The optimal strategy for each trader at time $t > 0$ when facing state s is the same when facing the same state at some future time $t' > t$.

[27] For further analysis on the reduced state space see Subsection 2.3.4.

the recurrent class $\mathscr{R}_\beta$, such that $R_\beta \subset \mathscr{S}_\beta$.[28]

The simulation starts at time $t = 0$, with both limit order books $\mathscr{L}_{m_L,t=0}$ (LM) and $\mathscr{L}_{m_D,t=0}$ (DM) empty, and is driven by three exogenous events. Recall that the arrival of new traders, the re-entry of traders with $z_R \neq 0$ (*i.e.*, more shares to trade) and the fundamental value changes, follow the Poisson process with parameters λ_N, λ_R and λ_{v} respectively. The inter-event times in a Poisson process with parameter λ are i.i.d. exponentially distributed with mean λ.

Suppose a new trader enters the market at time $t > 0$. The trader, conditional on the state $s = \left(mc_{m_L,t}, \beta, \emptyset, z_R\right)$ which incorporates fundamental value changes and entry (re-entry) rates of new (old) traders up to time t, takes an action. At any instant t, each action $\alpha = \left(m, x, p\right) \in A(s)$ in each state $s \in \mathscr{R}_\beta$ has an expected payoff. We denote the expected payoff by $\mathrm{U}_t(\alpha \mid s) \in \mathbb{R}$. Let $\alpha^\star$ be the agent's payoff-maximal action in state s, such that $\alpha^\star \in \underset{\alpha \in A(s)}{\operatorname{argmax}} \mathrm{U}_t(\alpha \mid s)$.[29] That is, the value of state s is determined as $\mathrm{V}(s) = \mathrm{U}_t(\alpha^\star \mid s)$.

To ensure that traders explore the whole space for each new pair (α, s) in the simulation we assign to them optimistic initial expected payoffs. The first time a state s is visited, the initial expected payoff - denoted as $\mathrm{U}_0(\alpha \mid s)$ - of taking action α in state s that might involve submitting a limit buy (sell) order at price p to either market, is the payoff $e^{-\rho t_y}\left(\beta + \mathrm{v} - p\right)\left(e^{-\rho t_y}\left(p - \mathrm{v} - \beta\right)\right)$. Here, t_y is the expected time to the arrival of a new trader who submits a sell (buy) market order to the respective market at price p while facing state $s \in \mathscr{R}_\beta$. The trader by submitting this limit order to market m presumes that the new trader will be willing to be a counterparty by submitting a market order and not a limit order at price $\hat{p} \neq p$. Given any new pair (α, s) encountered in the simulation, the initial action α taken by a trader, involves submitting an order with high expected payoff but with very small probability of execution. Through a learning process the term $\mathrm{U}_0(\alpha \mid s)$ is updated and the trader moves to orders with smaller expected payoffs but with higher probability of execution.[30] If the trader was pessimistic (*i.e.*, by submitting orders with higher execution

[28] A state i is recurrent iff the expected total number of visits to i is infinite, that is $\sum_{k=0}^{\infty} P_{ii}^k$. Two different states i and j communicate $i \longleftrightarrow j$ if $i \to j$ and $j \to i$. All the states that communicate with each other form a class. A recurrent class is absorbing *i.e.*, onces the system enters to the class it never leaves.

[29] Since on any entry the action set is finite, the maximum over all other actions exists and it is well defined.

[30] For details about the learning process see Subsection 2.3.3.

probability but lower expected payoffs) he would not have been able to identify if more profitable orders will execute.

Learning Process

Over repeated visits to a state $s \in \mathscr{R}_\beta$, the updating process described below calculates each trader's expected payoff for each action taken that results in an order being executed or cancelled so that the trader can pick the optimal action. Consider a new trader entering the market at time $t > 0$ while facing state $s = \left(mc_{m_L,t}, \beta, \emptyset, z_R\right) \in \mathscr{R}_\beta$. Assume the trader's optimal action does not represent a market order, *e.g.*, a limit order submitted to the LM or to the DM at price $p \in \mathscr{P}_m$ or no order. This action $\alpha^\star$ provides the highest expected payoff among all the other possible actions such that $\alpha^\star \in \underset{\alpha \in A(s)}{\operatorname{argmax}} \mathrm{U}_t(\alpha \mid s)$. Consider now two alternative events (i) at time $t'' < t'$ and state $s'' \in R_\beta$, the trader's order is executed, that is $z_R = 0$ (the trader has no more shares to trade). This is due to the presence of another trader with private valuation $\hat{\beta}$ who takes an action $\hat{\alpha}^\star \in A(s'')$ which involves submitting a market order at price $\hat{p} = p$ and time t'' to the respective market m; or (ii) the trader's order is still unexecuted and remains on the book as a limit order before the traders re-entry. Suppose at some future time $t' > t$, the trader re-enters and observes state s'. Under the aforementioned event (i) since the trader's order is filled (*i.e.*, $z_R = 0$) we update his expected payoff for taking action $\alpha^\star$ in state s and time t, as follows

$$\mathrm{U}_{t''}(\alpha^\star \mid s) = \frac{n_{\alpha^\star,s}}{n_{\alpha^\star,s}+1}\mathrm{U}_t(\alpha^\star \mid s) + \frac{1}{n_{\alpha^\star,s}+1} e^{-\rho(t''-t)}(\beta + \mathrm{v}_{t''} - p)x. \tag{2.9}$$

The term $n_{\alpha^\star,s}$ is a positive integer that records the number of times the action $\alpha^\star$ is chosen in state s. For the trader with private valuation $\hat{\beta}$ who submitted the market order at price $\hat{p} = p$ in state s'' and time t'', we update his expected payoff as follows

$$\mathrm{U}_{t''}(\hat{\alpha}^\star \mid s'') = \frac{n_{\hat{\alpha}^\star,s''}}{n_{\hat{\alpha}^\star,s''}+1}\mathrm{U}_{t''}(\hat{\alpha}^\star \mid s'') + \frac{1}{n_{\hat{\alpha}^\star,s''}+1}(\hat{\beta} + \mathrm{v}_{t''} - \hat{p})\hat{x}. \tag{2.10}$$

That is, the expected payoff is updated so that it is equal to the average payoff attained when the optimal action $\hat{\alpha}^\star$ is taken in state s''. For each pair (α, s) in the simulation process we assign an initial positive integer $n_0 = n_{a,s}$. As $n_{a,s}$ gets very large, this integer reduces the convergence speed of the algorithm. To achieve faster convergence, after 3×10^8 actions, we reset the counter to its initial value n_0. Moreover, by assigning optimistic initial expected payoffs to all agents, states in R_β that are only visited at the start of the optimization, once execution probabilities are updated, stop being part of the recurrent class $\mathscr{R}_\beta$.

If the order submitted at state s and time t is still active (see event (ii)), then given the new state $s' = \left(mc_{m_L,t'}, \beta, \alpha^\star, z_R\right)$ at time t' the trader will choose an optimal action with value $\mathrm{V}(s') = \mathrm{U}_{t'}(\alpha^{\star\star} \mid s')$ since $z_R \neq 0$. We update the agent's expected payoff for taking action $\alpha^\star$ in state s and time t, as follows

$$\mathrm{U}_{t'}(\alpha^\star \mid s) = \frac{n_{\alpha^\star,s}}{n_{\alpha^\star,s}+1}\mathrm{U}_t(\alpha^\star \mid s) + \frac{1}{n_{\alpha^\star,s}+1}e^{-\rho(t'-t)}\mathrm{V}(s'). \tag{2.11}$$

The continuation value of state s', denoted as $\mathrm{V}(s')$, allows the agent in the trading game to profitably execute the order in the future. Being consistent with the notation used above, as introduced in Eq (2.8), the value $\mathrm{V}(s')$ of the agent in state s' at time t' is decomposed into the payoff from execution before the trader's next re-entry, at time *i.e.*, t'' and the continuation payoff on re-entry at time *i.e.*, $t''' > t''$. In the event where the trader's order executes at time *i.e.*, t'' we set the continuation payoff $\mathrm{V}(s') = 0, \forall s'$, since z_R (*i.e.*, the trader has no more shares to trade). If the trader's order is still unfilled at time t''', the agent while facing state $s'' = \left(mc_{m_L,t'''}, \beta, \alpha^{\star\star}, z_R\right)$, will choose an optimal action such that $\alpha^{\star\star\star} \in \underset{\alpha \in A(s'')}{\operatorname{argmax}} \mathrm{U}_{t'''}(\alpha \mid s'')$. That is, the agent's future optimal actions at different states and order execution, are used to determine the expected payoff from taking action $\alpha^\star$ in state s and time t. Over repeated visits to a state, the learning process applied computes the trader's expected payoff and continuation values so that the trader can pick the action that will maximize his payoff.

All agents in the trading game observe the fundamental value at any instant t with a time lag Δ_t. That is, conditional on the prevailing market conditions, denoted $mc_{m_L,t}$, agents have to estimate the fundamental value before taking any action. Recall that an agent's estimate of the fundamental value at time t is denoted as $\hat{\mathrm{v}}_t = E(\mathrm{v}_t \mid mc_{m_L,t})$. We set $\hat{\mathrm{v}}_t = \mathrm{v}_{t-\Delta_t} + \zeta_n(mc_{m_L,t})$, where $\zeta_n(mc_{m_L,t}) = E(\mathrm{v}_t \mid mc_{m_L,t}) - \mathrm{v}_{t-\Delta_t}$ is the extent by which an agent with private valuation β at time t, revises the term $E(\mathrm{v}_t \mid mc_{m_L,t})$ given lagged value $\mathrm{v}_{t-\Delta_t}$. We define

$$\zeta_{n+1}(mc_{m_L}) = \frac{\hat{n}}{\hat{n}+1}\zeta_n(mc_{m_L}) + \frac{1}{\hat{n}+1}(\mathrm{v}_t - \mathrm{v}_{t-\Delta_t}) \tag{2.12}$$

where $\hat{n}$ counts the number of times market conditions m are encountered during the simulation process.[31] Thus, using the estimate $\hat{\mathrm{v}} = \mathrm{v}_{t-\Delta_t} + \zeta_n(mc_{m_L})$ the action set $A(s)$ is defined as in Eq. (2.6).

Traders use the term $\hat{\mathrm{v}}_t = \mathrm{v}_{t-\Delta_t} + \zeta_n(mc_{m_L,t})$ to estimate what the current state at time t is and choose an action so as to maximize their expected discounted payoffs. However during the learning process we described above, we record the payoffs in the real state of the trading game (*i.e.*, without the estimate) that the action was taken and not the perceived state. In equilibrium these estimates for any given (α, s) pair are correct given that we require - as we will see in the next Subsection - the difference between the expected and the realized payoffs to be small.

Convergence Criteria

To check for convergence we follow the same approach proposed by Bernales (2014) and Chiarella and Ladley (2016). We initially run the learning process (see Subsection 2.3.3) for several billion order submissions, by updating agents' expected payoffs for each pair (α, s). After this point, we fix these payoffs denoted as $\mathrm{U}^{\star}(\alpha, s)$ (*i.e.*, the learning process is not applied) and we simulate the model for a further billion more order submissions (for new and returning agents). During this period the realized payoffs for all executed orders are recorded. Denote the realized payoff for an order executed in state $s \in R_{\beta}$, by $\hat{\mathrm{V}}(s)$. Suppose for example a trader with private valuation β enters to the market at time t and submits and order to the DM at price $p \in \mathscr{P}_{m_D}$. If this order fully executes at some future time $t' > t$ without the trader modifying it, his realized payoff is given by $\hat{\mathrm{V}}(s) = e^{-\rho(t'-t)}(\beta + \mathrm{v}_{t'} - p)x$.[32] The convergence of the algorithm is achieved if difference between the expected payoffs and the realized ones is small (i) the correlation between beliefs $\mathrm{U}^{\star}(\cdot)$ and realized payoffs $\hat{\mathrm{V}}(\cdot)$ is greater than 0.99; and (ii) the mean absolute error in beliefs, that is $|\mathrm{U}^{\star}(\cdot) - \hat{\mathrm{V}}(\cdot)|$, weighted by the numbers of times the pair $(\alpha,\ s)$ is observed, is less than 0.04. If both of these criteria are satisfied, we simulate a market session for a further billion order submissions in which the trading data is collected. Else, we repeat the learning process through equations (2.9) to (2.11) until the equilibrium is reached.

[31] Since our focus lies on stationary equilibria, the time subscript on market conditions is omitted.
[32] The realized payoff $\hat{\mathrm{V}}(\cdot)$ can be thought as a direct measure of benefits to trade.

2.3.4 Numerical Parametrization

In the model, to ensure numerical tractability: (i) all agents given their private valuation β, care about the relative price of a trade, that is the difference between the execution price and the fundamental value v (see discussion after Eq. (2.3)); and (ii) given all the possible states a trader with private valuation β might face, we consider only states that belong to the recurrent class such that $s \in R_\beta \subset \mathscr{S}_\beta$. All the above, allow us by restricting the state space of the trading game for each agent, to compute the fixed point. Hence, the state at any instant $t > 0$ is defined by the four-tuple

$$s = (\underbrace{\{\overset{\star}{\mathscr{L}}_{m_L,t}, \mathrm{v}_{t-\Delta_t}, q_{m_L,t}(p,x)\}}_{= mc_{m_L,t}}, \beta, \alpha, z_R) \tag{2.13}$$

where $mc_{m_L,t}$ are the market conditions of the LM observed by the agent with private valuation β (see Eq. (2.5)). If the returning trader's order, submitted at time $t' < t$ and state s', has not been executed before his re-entry, α is the trader's previous action taken in state s' and z_R denotes the remaining shares the trader can trade. The variables in $\overset{\star}{\mathscr{L}}_{m_L,t}$ that characterize $\mathscr{L}_{m_L,t}$ at time t, are: (i) the bid price and the ask price; (ii) the depth at the bid quote and the ask quote (*i.e.*, $\ell^B_{m_L,t}$ and $\ell^A_{m_L,t}$); and (iii) the buy and the sell quantities in the $\mathscr{L}_{m_L,t}$ (*i.e.*, $\sum_{i=0}^N \ell^i_{m_L,t} > 0$ and $\sum_{i=0}^N \ell^i_{m_L,t} < 0$).

We set the support in ticks of the private value distribution $\mathscr{F}_\beta$ to be $\{-8,-4,0,4,8\}$ with cumulative distribution $\{0.15, 0.35, 0.65, 0.85, 1.0\}$. Speculators, agents with private value equal to zero may only trade profitably when the fundamental is mispriced. Agents with the most extreme private values (*i.e.*, $\beta = |8|$) are keen to trade quickly given that their main benefits arise from their private values.

The values assigned to the private value distribution $\mathscr{F}_\beta$ follow approximately the findings of Hollifield et al. (2006). By splitting the private value into five different intervals, they estimate the distributions of private values and the optimal order submission strategies for three sample stocks on the Vancouver exchange.[33] They observe that across the three stocks the mean probability of private value in the interval $(-2.5\%, +2.5\%)$ is 44%; in the intervals $(-5\%, -2.5\%]$ and $[-2.5\%, +5\%)$ is 26%;

[33] The three sample stocks are Barkhor Resources (BHO), Eurus Resources (ERR), and War Eagle Mining Company (WEM).

and in the intervals (–5%, –2.5%] and [–2.5%, +5%) is 30%.

New agents arrive randomly to the market following a Poisson process with intensity λ_N. To support the value assigned to λ_N, we use the empirical findings of Conrad, Wahal, and Xiang (2015). By sampling all common stocks in the US with a stock price and market cap greater than \$1 and \$100 million respectively, Conrad, Wahal, and Xiang (2015) examine the relation between high frequency quotations and the behaviour of stock prices for the period $2009-2011$. By splitting firms into different size quantiles and by averaging over the entire time series, they observe that the average number of trades for stocks in the largest size quantile is almost one trade per second.[34] Given the above and conditional that each trade involves two traders we set $\lambda_N = 2$. That is, a trader enters to the market on average every 0.5 seconds.

In the model, an agent with $z_R \neq 0$ (more shares to trade) re-enters the market and revises his previous order according to a Poisson process with intensity λ_R. The agent conditional on the new state of the game, can take further actions so as to maximize his expected discounted payoff. We set $\lambda_R = 0.1$, that is traders re-enter the market and revise their orders on average every 10 seconds. This is approximately in line with the empirical findings of Hautsch and Huang (2011), where the median cancellation time of one asset (WCRX) among the 200 biggest stocks listed on NASDAQ is 10 seconds.[35]

The fundamental value of the asset at any t, denoted as v_t, follows a Poisson process with rate λ_{v}. This parameter denotes the expected number of innovations of the fundamental value in one second. Similar to Goettler, Parlour, and Rajan (2009) we set $\lambda_{\mathrm{v}} = 0.125$, that is the fundamental value increases or decreases by σ ticks every 8 seconds on average. We set $\sigma = 1$, so when an innovation occurs the fundamental value increases or decreases with equal probability by one tick.

At any instant, agents of any valuation, can place their limit orders to either the LM or the DM at any feasible price that lies in the interval $p \in [\hat{\mathrm{v}}_t - k, \hat{\mathrm{v}}_t + k] \cap \mathscr{P}_m$. We set the value of the parameter $k = 13$; that is, the price of a limit order at any t may be submitted in a range between thirteen ticks above and below $\hat{\mathrm{v}}_t$.[36] Recall that $\hat{\mathrm{v}}_t = E(\mathrm{v}_t \mid mc_{m_L,t})$ is the trader's best estimate of the fundamental value since

[34] Stocks in the two largest size quantiles accounted for over 80% of the market cap.

[35] In terms of market capitalization on October 1, 2010.

[36] Similar to Goettler, Parlour, and Rajan (2009); Bernales (2014) and Chiarella and Ladley (2016) we simulate scenarios with $k > 13$ but no effect has been observed on the market outcomes since market participants were not submitting limit orders so far away from the fundamental value.

all agents in the trading game observe v with a time lag Δ_t. The lag with which all traders observe the fundamental value is set to ten seconds, that is $\Delta_t = 10$.

In the model we set the discount rate, denoted as ρ, to 0.03. This parameter is essential to compute the fixed point since it restricts agents from infinitely postponing trade, *e.g.*, submit only limit orders with very small probability of execution, denoted as $\eta(\cdot)$. Consider the case in which $\rho = 0$ and a trader enters to the market with z shares to trade. The agent, by observing any state $s \in R_\beta$ at any time t, will always pick an action that involves submitting a limit buy (sell) order to the OM at price $p \ll \mathrm{v}_t$ ($p \gg \mathrm{v}_t$) with $\eta(\cdot)$ converging asymptotically to zero. That is, given $\rho = 0$ and any state, an agent's strategic profile will include taking an optimal action by submitting orders at prices far away from the fundamental value (not an aggressive order).

2.4 Results

This section displays the results derived from the simulations for 9×10^8 order submissions given the numerical parametrization as introduced in Subsection 2.3.4. To ensure that the model converges to the same equilibrium we run several optimizations under different random seeds. For our analysis we simulate two scenarios. Under scenario I (Single-Market) trading takes place on a single limit order book that offers at any instant full pre-trade transparency (LM). Under scenario II (Multi-Markets) two marketplaces exist and operate simultaneously, a 'lit' market (LM) and a 'dark' market (DM).

2.4.1 Market Quality

We first examine how the introduction of the DM affects the market quality of the LM. To do this we focus on standard market quality measures: the bid-ask spread, the effective spread, the number of limit orders at the ask and at the sell side (the sell side is symmetric to the buy side), the effectively traded limit orders at the sell side of the book related to market m, the tracking error (*i.e.*, the average absolute difference between the market's midprice and the fundamental value of the asset) and the microstructure noise (*i.e.*, the mean absolute difference between the execution price and the fundamental value).[37] Table (2.1) depicts the above measures. The presence of the DM under the Multi-Markets scenario harms LM's market quality. For example across the two scenarios, the quoted spread of the LM widens from 1.135 to 4.780 ticks. Liquidity is reduced as the number of limit orders at the ask and at the sell side of the book related to the LM decline from 3.711 to 1.344 and from 46.302 to 4.783 ticks respectively and the tracking error increases. In contrast the microstructure noise decreases from 0.620 to 0.483 ticks.

Single-Market		Multi-Markets	
LM		LM	DM
	Spread		
1.135		4.780	1.701
	Effective Spread		
1.028		3.212	-0.346
	N. of limit orders at the ask		
3.711		1.344	6.025
	N. of limit orders at the sell side of the LOB		
46.302		4.783	15.300
	N. of limit orders at the sell side of the LOB (effectively traded)		
4.527		1.626	2.891
	Tracking Error		
0.716		1.496	0.925
	Microstructure noise: $\|p - \mathrm{v}_t\|$		
0.620		0.483	0.856

Table 2.1: **Market quality and liquidity**. The sell side is symmetric to the buy side. The spread is the distance in ticks, between the best ask and the best bid. The number of effectively traded limit orders at the sell side of the book are the limit orders submitted and executed without any modification. The effective spread is defined to be $2(p-mp)x$ where p is the execution price, mp is the mid-price when the order is submitted and x is the quantity of the order to buy or to sell one unit of the asset, with $x = 1$ ($x = -1$) if the transaction involves a market buy (sell) order. The tracking error is the average absolute distance between the midprice and the fundamental value. The microstrucutre noise is defined to be the absolute mean difference between the execution price and the fundamental value of the asset. All the statistics are calculated by observing the market every 10 simulation minutes. Number of observations: 45516 (Single-Market), 80183 (Multi-Markets). The fixed point for each of the available scenarios is obtained independently.

The decrease in the quality of the LM is amplified further by observing the effective spread. Similar to Chiarella and Ladley (2016) we define the effective spread as $2(p - mp)x$ where p is the execution price, mp is the mid-price when the order is submitted and x is the quantity of the order to buy or to sell one unit of the as-

[37] Given the bid and ask quotes as defined in Section 2.3 the mid-price (midpoint) of the LM at any instant t is defined as $mp_t = \frac{A(\mathscr{L}_{m_L,t})+B(\mathscr{L}_{m_L,t})}{2}$.

set, with $x = 1$ ($x = -1$) if the transaction involves a market buy (sell) order. While the market spread illustrates posted positions and given the fact that a trade may occur within the bid and ask quotes the effective spread is a direct measure of execution costs related to a trade. That is the effective spread which is based on trade price (when an incoming order at price p (market order) matches with a limit order on the other side of the market) captures the microstructure features of the market (*e.g.*, stale quotes). Comparing the two scenarios the effective spread of the LM increases from 1.028 to 3.212 ticks. Thus, under the Multi-Markets scenario market order submitters that choose the LM have to pay relatively higher transaction costs. The increase of the effective spread in the LM is twofold. Firstly, the likelihood of finding a trader who is willing to be a counterpart is reduced (see for example depth at the sell side of the book related to the LM). Secondly, the quotes are less informative due to the increase of the tracking error and less competitive as the number of effectively traded limit orders is reduced (see for example row five in Table (2.1)).

In order to shed more light into the execution costs of the market participants we introduce Table (2.2). Table (2.2) depicts the effective spreads that market participants face in the LM across the two scenarios. As can be seen from the table the increment in liquidity provision in the LM coupled with the lower spread and the smaller tracking error under the Single-Market scenario signifies that market order submitters have to pay low cost to transact. In contrast, the effective spreads for agents irrespective of valuation under the Multi-markets scenario are higher. For example the effective spreads for agents with extreme private values (*i.e.*, with a β of 4 and 8) are increasing from approximately 1.02 to 3.17 ticks. However the effective spreads of agents with zero private values are greater. As we will see later (Subsection 2.4.2) these agents follow a more conservative approach in the LM and may only place market orders when significant mispricing is present. Thus the speculators' increased effective spread signifies the relative deficiency of trading opportunities of these traders when they participate in the LM by posting market orders under the Multi-Markets scenario.

Focussing now on the Multi-Markets scenario the DM, as compared to the LM, offers a more liquid and competitive environment to the market participants.[38] The

[38] Although our focus in this paragraph lies on the Multi-Markets scenario; as can be seen from Table (2.1) the LM under the Single-Market scenario, as compared to the DM under the Multi-Markets scenario, is more competitive and more liquid (see rows one, four and five).

Scenario	Venue	Effective Spread		
		Private Value $\lvert\beta\rvert$		
		0	4	8
Single-Market	LM	1.049	1.024	1.023
Multi-Markets	LM	3.481	3.161	3.192

Table 2.2: Table showing the effective spreads by each trader valuation under the two available scenarios. The effective spread is defined to be $2(p-mp)x$ where p is the execution price, mp is the mid-price when the order is submitted and x is the quantity of the order to buy or to sell one unit of the asset, with $x=1$ ($x=-1$) if the transaction involves a market buy (sell) order. All the statistics are calculated by observing the market every 10 simulation minutes.

spread decreases and the quantities available at the ask quote and at the sell side of the limit order book related to the DM increase (see rows one, three and four in Table (2.1)). A result to the above is the DM's reduced tracking error from 1.496 to 0.925 ticks. Although the DM reduces informational inefficiency the properties related to that market (at any instant the DM's depth at the best quotes and throughout the book are not observable to the traders) do not allow market participants to benefit from the informational advantage this venue offers considering the fundamental value of the asset.

The above is justified even more by observing the increase in the DM's microstructure noise. Similar to Goettler, Parlour, and Rajan (2009) and Bernales (2014) to be able to measure the trading frictions present in our model (*e.g.*, agents with different exogenous motives to trade observing the fundamental value with the same lag, discrete set of prices and cognition limits among others) that make execution prices deviate from the fundamental value we introduce the microstructure noise.[39] As can be seen from Table (2.1) under the Multi-Markets scenario the microstructure noise is greater in the DM (*i.e.*, 0.856 ticks) as compared to the LM (*i.e.*, 0.483 ticks). As we will see later the optimal trading strategies adopted by the agents in the LM - as a result primarily of the increased spread and the lower quantities available at the best quotes and throughout the book related to the LM - will make quotes in that venue be more often set and executed at prices closer to the fundamental.

Given the properties of the DM, market order submitters participating in that venue cannot directly infer the costs associated to these orders. That is in our calculations we define the DM's effective spread relative to the LM's midprice (*i.e.*, ef-

[39] For further analysis about the decomposition of transaction price see Hasbrouck (2002).

fective spread = $2(p - mp)x$ where p is the DM's execution price, mp is the LM's mid-price when the order is submitted and x is the quantity of the order to buy or to sell one unit of the asset, with $x = 1$ ($x = -1$) if the transaction involves a market buy (sell) order.). However the results in Table (2.1) show an aggregate picture of the transaction cost related to submitting a market order to the DM (*i.e.*, market sell (buy) orders submitted to the DM execute on average at prices 0.173 ticks above (below) the LM's midprice). Table (2.3) decomposes the agents effective spreads in the DM under the Multi-Markets scenario.

Scenario	Venue	Effective Spread $[\lvert p_{m_L} - p_{m_D}\rvert]$		
		Private Value $\lvert\beta\rvert$		
		0	4	8
Multi-Markets	DM	0.449	-0.912	-0.215
		[1.516]	[2.036]	[1.704]

Table 2.3: Table showing the DM's effective spreads by each trader valuation under the Multi-Markets scenario. The effective spread is defined to be $2(p - mp)x$ where p is the execution price, mp is the LM's mid-price when the order is submitted and x is the quantity of the order to buy or to sell one unit of the asset, with $x = 1$ ($x = -1$) if the transaction involves a market buy (sell) order. Given the definition above, row two measures the average absolute difference between the LM's market orders' execution price and the DM's market orders' execution price. All the statistics are calculated by observing the market every 10 simulation minutes.

As can be seen from Table (2.3) agents with nonzero private values conditional on the LM's midprice have to pay relatively lower transaction costs when a market order matches a resting limit order on the other side of the book in the DM. For example agents with a β of $|4|$ execute market sell (buy) orders in the DM at prices 0.912 ticks above (below) the LM's midprice. In contrast given the LM's midprice, agents with zero private values have to pay higher transaction costs. The above is justified even further by looking at the average absolute distance between the LM's market orders' execution price and the DM's market orders' execution price for these traders. As can be seen from the table an agent with a β of zero executes market sell (buy) orders to the DM at prices 1.51 ticks above (below) the LM's market orders' execution price.

Our results stand in line with the theoretical findings of Buti, Rindi, and Werner (2017) which show that the introduction of a dark pool that competes with a transparent limit order book (LOB) affects negatively the market quality of the latter (*i.e.*, wider spread and lower depth). Although Foley and Putniņš (2016) find that 'two-sided' dark trading drives prices close to the random walk; they also observe that it benefits liquidity (low quoted, effective and realized spreads) and informational efficiency.

2.4.2 Trading Behaviour

Conditional on the state of the trading game at any instant t, agents submit orders to maximise their expected discounted payoff. Recall that in our model all orders submitted are *marketable*. That is an order at price p executes if there exists a limit order at the same price standing on the other side of the respective book that the order is submitted. If no matching order is present the order is placed on the book as a limit order. While market orders offer immediate execution, the cost associated with them is that traders might end up overpaying for the asset. On the other hand, limit orders offer potentially better terms for trade but the cost associated with them is that traders might have to wait longer for their trade to execute. Moreover, limit order submitters face the risk of being picked off, by submitting a limit sell (buy) order and the price of the fundamental value rises (falls) in the time interval between the order's submission and execution.

An agent, by submitting a market order, acts as liquidity demander, conversely an agent who posts a limit order provides liquidity to the market (*i.e.*, liquidity supplier). Agents with higher positive or negative intrinsic private values to trade are impatient to trade since their main benefits arise from their private values. Speculators, agents with $\beta = 0$ have no exogenous reasons to trade. Table (2.4) depicts the percentage of limit and market orders submitted by each trader type under the two available scenarios. As can be seen from the table, speculators act as liquidity suppliers under the 'Single-Market' scenario by providing approximately 92% of the limit orders in the market. Agents with an absolute private value of 4 or 8 act as liquidity demanders as they prefer to trade by submitting market orders and relatively few limit orders. Speculators under the Multi-Markets scenario decrease their supply of liquidity in both the LM and the DM by posting fewer limit orders. More interestingly comparing the speculators' liquidity provision strategy across the two available scenarios the percentage of limit orders submitted by them in the LM when competing with the DM is reduced significantly from 91.99% to 58.35%. Agents with nonzero private values (*i.e.*, |4| and |8|) increase their liquidity provision in the LM approximately by 33% compared to the Single-Market case, follow the same strategy in both venues by submitting relatively fewer limit orders and more market orders.

Scenario	Venue	Order's Type	Private Value $\|\beta\|$		
			0	4	8
Single-Market	LM	Limit orders submitted	91.99%	5.81%	2.20%
		Market orders	14.55%	41.35%	44.10%
Multi-Markets	LM	Limit orders submitted	58.35%	24.29%	17.36%
		Market orders	12.51%	52.25%	35.24%
	DM	Limit orders submitted	87.13%	10.37%	2.50%
		Market orders	19.93%	37.79%	42.28%

Table 2.4: Table showing proportion of limit and market order submitted by each trader valuation (new and returning) under the two available scenarios, as a percentage of all limit and market orders. Statistics are calculated from 9×10^8 order submissions. The fixed point for each of the available scenarios is obtained independently.

The reasons for the decrease in the percentage of the limit orders submitted in the LM by the speculators under the Multi-Markets scenario are (i) the switch in the strategies adopted primarily by the agents with a β of |4| in the LM (*e.g.*, increase in liquidity demand to 52.25%) together with the properties of the DM (no pre-trade transparency) under the Multi-Markets scenario indicate that the placement of non-aggressive limit orders (*i.e.*, at the best quotes) to the LM involve more risk and are associated with high probability of being picked off. Recall that agents with no intrinsic motive to trade are vulnerable to this effect (for further discussion see analysis of Table (2.7)); and (ii) the increased market spread of the LM that supports the placement of limit orders from traders with nonzero private values in order to avoid overpaying for the asset. Thus due to the increased competition and the low quantities available, quotes in the LM are more often set by traders with higher private valuations that are willing to trade at prices - as opposed to the speculators - closer to the fundamental value of the asset (*i.e.*, more aggressive orders).

Tables (2.5), (2.6) and (2.7) support further the aforementioned observation as well as our findings in Subsection 2.4.1. As can be seen from Table (2.5) under the Multi-Markets scenario the percentage of traders with nonzero private values (*i.e.*, |4| and |8|) who have not yet executed and switch to the DM (LM) after the random re-entry is on average approximately 52% (48%). Recall that these traders are keen to trade quickly by submitting aggressive orders (*i.e.*, market or price improving limit orders) to either market so as to capture their exogenous private values. In contrast returning speculators seem to prefer the DM and its non-displayed properties for the placement of their orders. For example, comparing the speculators strategies related to switching/no switching between markets, the percentage of these agents that move to the DM more than triples, from 22.40% to 77.60%.

Scenario	Venue	Switching/no switching between markets	Private Value $\|\beta\|$ 0	4	8
Multi-Markets	LM	Remain to LM	13.94%	19.31%	35.05%
		From DM to LM	8.46%	24.23%	17.52%
		Total	22.40%	43.54%	52.57%
	DM	Remain to DM	68.51%	48.93%	34.00%
		From LM to DM	9.09%	7.54%	13.43%
		Total	77.60%	56.46%	47.43%

Table 2.5: Table showing the proportion of returning agents who switch from the DM to the LM and vise-versa (rows two and five respectively) split by each trader valuation. Rows one and four show the proportion or returning agents who remain to the same venue. Statistics are calculated from 9×10^8 order submissions.

Given that traders irrespective of valuation can re-enter the market multiple times before execution we next examine the LM's (LM's and DM's) order flow after the random re-entry to the market under the Single-Market (Multi-Markets) scenario. Table (2.6) depicts the percentage of orders submitted by each trader valuation with $z_R = 1$ (*i.e.*, more shares to trade) after re-entering the market under the two available scenarios.[40] Not surprisingly under both scenarios the LM's - DM's order flow is predominately determined by the order submissions strategies adopted by the returning market participants with zero private values. For example after the random re-entry under the Single-Market scenario, 83.41% of the orders submitted at the LM are coming from the agents with zero private values, 10.00% from the agents with a β of $|4|$ and 6.59% from the agents with the most extreme private values ($\beta = |8|$). What is worth noticing is the switch in the traders' order placement strategies under the Multi-Markets scenario. Comparing the two scenarios, despite the transparent characteristics of the LM, the percentage of orders submitted at this venue by the returning traders with $\beta = 0$ drops to 52.90%. In contrast the smaller effective spreads of agents with $\beta \neq 0$ in the LM (see Table (2.2)) along with the LM's wider spread

[40] That is we consider traders with shares available to trade who re-enter the market and either retain their previous limit order, cancel and resubmit a new order to the LM under the Single-Market scenario (or to either the LM or the DM under the Multi-Markets scenario) or submit an order (market or limit) for the first time.

enhance trading activity of these agents in the LM. For example the percentage of orders submitted in the LM by the returning agents with nonzero private values almost doubles from 14.00% to 27.99% and from 8.16% to 19.11%.

Scenario	Venue	Order flow returning agents (%)		
		Private Value $\|\beta\|$		
		0	4	8
Single-Market	LM	83.41%	10.00%	6.59%
Multi-Markets	LM	52.90%	27.99%	19.11%
	DM	77.85%	14.00%	8.16%

Table 2.6: Table showing the proportion of market and limit orders submitted by each returning trader valuation as a percentage of all orders under the two scenarios (LM's - DM's order flow after re-entry). For each active trader type with $z_R = 1$ (more shares to trade) we find the percentage of orders submitted after re-entry to the market. These orders involve returning traders that choose either to retain, cancel and resubmit or submit for the first time. Statistics are calculated from 9×10^8 order submissions.

It is worth highlighting, that although speculators submit relative fewer limit orders in the LM under the Multi-Markets scenario (see row 3, Table (2.4)) they do not switch to demand liquidity by posting more market orders. In fact the percentage of market orders submitted by these traders in the LM has its lowest value (*i.e.*, 12.51%). In contrast to the traders with nonzero private values, agents with $\beta = 0$ follow a more conservative approach by keeping their limit orders in the book related to market m for longer period and waiting for opportunities to arise (see analysis of Table (2.7)). That is, they submit market orders only when significant mispricing is observed. Although the presence of the DM harms the market quality of the LM under the Multi-Markets scenario - by shifting the trading behaviour of the liquidity suppliers - as we will see later this competition seems to be beneficial with respect to obtaining better terms of trade for the traders that participate in the LM.

Consider a new trader irrespective of valuation, entering the market at time t and state s when the fundamental value is v_t. Assume further that the trader's optimal action in state s involves submitting a market sell order to the LM at price p.[41] Since a trade occurs, the new agent's net common benefit for submitting the sell order at price p when the fundamental value is v_t is given by $p - \mathrm{v}_t$ (see Eq. (2.3)). Agents with $\beta = 0$ have no intrinsic motives to trade and may only be willing to sell

[41] Here we consider the case in which there exists an outstanding limit order in the LM.

the asset at prices above the fundamental value. Conversely, traders with nonzero private values (*i.e.*, $\beta = -8$) should be willing to sell the asset at lower prices since their main benefits arise from their exogenous private values.[42]

Table (2.7) justifies the aforementioned discussion in terms of aggressiveness of quotes. As can be seen from the table speculators submit more conservative limit orders in comparison to the traders with nonzero private values (see 'Price of sell limit orders'). More specifically, under the Single-Market scenario speculators submit limit sell orders on average at prices 3.68 ticks above the fundamental value, while traders with a β of |4| and |8| are more aggressive submitting sell orders at prices 0.03 and 0.01 ticks respectively above the fundamental value. Under the Multi-Markets scenario the decrease in the market quality measures of the LM make traders - and primarily traders with a β of 0 and |4| - more cautious by submitting sell orders further away from the fundamental. Speculators submit sell limit orders on average at prices 4.55 ticks above the fundamental value; agents with a β of |4| tend to be a little more conservative, as compared to the Single-Market scenario, by submitting sell limit orders on average at prices 0.33 above the fundamental. The speculators' increased probability of being picked off in the LM and the properties of the DM (*i.e.*, no pre-trade transparency) switch the order submissions strategies (in terms of aggressiveness) of the aforementioned traders. Specifically, when speculators participate in the DM they submit limit sell orders on average at prices 2.88 (2.01) ticks lower as compared to the LM under the Multi-Markets scenario (Single-Market scenario).

The limit order submission strategies adopted by the different market participants (*i.e.*, the level of aggressiveness, percentages of limit orders submitted) determine the likelihood of a limit order executing, the probability of a limit order being picked off and the average time to execution. As we can see from Table (2.7) under all available scenarios speculators have the lowest percentages of executed limit orders followed by the agents with private values of |4| and |8|. For example under the Single-Market scenario, the percentage of executed limit orders by agents with private values 0, |4| and |8| increases on from 5.84% to 78.69% to 85.38%. The high percentages of the agents with nonzero private values are supported by their exoge-

[42] It is worth mentioning that agents with the higher private values (*i.e.*, $\beta = \pm 8$) might submit sell (buy) orders at prices lower (higher) than the fundamental. Recall the term $\beta + \mathrm{v}_t - p$ in Eq. (2.7) that captures the order's instantaneous payoff.

		Limit orders executed (%)			Picked Off			Time to Execution (Limit orders)			Price of sell limit orders		
		Private Value $\lvert\beta\rvert$			Private Value $\lvert\beta\rvert$			Private Value $\lvert\beta\rvert$			Private Value $\lvert\beta\rvert$		
		0	4	8	0	4	8	0	4	8	0	4	8
Single-Market	LM	5.84%	78.69%	85.38%	47.83%	16.15%	11.55%	129.87	4.01	2.53	3.68	0.03	0.01
Multi-Markets	LM	2.78%	79.95%	84.45%	56.96%	8.86%	7.12%	68.23	3.75	2.05	4.55	0.33	0.03
	DM	13.98%	45.79%	56.91%	48.22%	25.09%	16.45%	63.37	7.39	4.27	1.67	0.31	-0.01

Table 2.7: Table showing the percentage of limit orders executed per trader valuation, the average time to execution between the trader's first arrival to the market and his limit order executing ('Time to Execution') and the mean price of submitted limit sell orders. The probability of being picked off is calculated when $x(\mathrm{v}_t - \mathrm{v}_s) < 0$, where v_s is the fundamental value at submission time, v_t is the fundamental value at execution time and x is the limit order's quantity when submitted. Due to the symmetry of the model we do not report the price of buy limit orders. Statistics are calculated from 9×10^8 order submissions.

nous - higher than the speculators - motives to trade. As we have mentioned these traders prefer to trade quickly by submitting limit orders close to the fundamental value rather than waiting for more conservative orders to execute. Under the Multi-Markets scenario although traders with nonzero private values are acting relatively more aggressively in the DM, as compared to the LM, their percentages of limit orders executing is reduced. The uncertainty the DM offers in order placement strategies (*i.e.*, no information is revealed to any market participant about the available quantities on either side of the DM) coupled with the decrease in the percentages of limit orders submitted to the DM, by the traders with nonzero private values, justify the above. In contrast the protection dark trading offers to speculators as opposed to the LM for the placement of more aggressive limit orders increases the speculators' percentage of limit orders executing in the DM.

The characteristics of the centralized limit order book under the Single-Market scenario as compared to the Multi-Markets scenario (*e.g.*, lower market and effective spread, increased number of limit orders at the ask side and at the book as a whole) make the post of a market order or an aggressive limit order in the LM for a speculator under the former scenario, to be rarely profitable.[43] By keeping limit orders in the book so to exploit any mispricing, speculators increase the amount of time spent in the LM under the Single-Market scenario. In contrast agents with nonzero private values have much shorter waiting times. Comparing now the average execution time of agents with $\beta = 0$ across the two available scenarios, we can observe

[43] For further analysis see Subsection 2.4.1.

for example that although these trades are more cautious in the LM by placing limit order further away from the fundamental they have shorter waiting times in the LM in comparison to the Single-Market scenario. For example the average time interval between a trader's entry to the market and his limit order executing in the LM is almost halved from 129.87 to 68.32 seconds. The reasons for the reduction in speculators execution time across the two scenarios are due to (i) the protection dark trading offers to these traders as opposed to the LM, for the placement of more aggressive limit orders with a higher chance of executing; and (ii) the increased spread of the LM under the Multi-Markets scenario that supports the strategies adopted by the agents with nonzero values to place more limit orders; and as a result the relative deficiency of trading opportunities the LM offers to traders with $\beta = 0$.

Cognition limits that prevent agents from continuously monitoring the market and immediately modifying their previous limit orders after a change in the market conditions together with the exogenous characteristics of the market participants (patient-impatient traders) increase significantly speculators' probability of being picked off. As expected agents with extreme values ($\beta = |8|$) have the lowest probability of being picked off followed by the traders with the intermediate private values under both scenarios (see Table (2.7)). The non-display liquidity properties of the DM and the speculators strategies in this venue increase the picked off probability from 8.80% to 25.09% for agents with $\beta = |4|$ and from 7.12% to 16.45% for agents with $\beta = |8|$. Under the Multi-Markets scenario, the higher risk associated with the placement of conservative limit orders to the LM (see discussion after Table (2.4)) and the increase in the percentage of market orders submitted to the same venue, by the traders with $\beta = |4|$ increase the speculators probability of being picked off. Namely from 47.83% under the Single-Market scenario to 56.96% under the Multi-markets scenario when a trader with $\beta = 0$ decides to submit a limit order to the LM.

Table (2.8) shows the difference between the execution price p_{ex} and the fundamental value of the asset $\mathrm{v}_{t_{ex}}$ for all executed sell orders (market and limit). Given a trader with valuation β, an increase in the term $p_{ex} - \mathrm{v}_{t_{ex}}$ signifies an improvement in terms of trade. As can be seen from the table, under the Single-Market scenario agents with zero private value execute sell orders on average at prices 0.35 ticks above the fundamental while agents with nonzero values execute at worse terms of trade. Namely traders with $\beta = |4|$ execute their sell orders at prices 0.11 ticks

Scenario	Venue	Price of executed sell orders (benefits from trade)		
		Private Value $\|\beta\|$		
		0	4	8
Single-Market	LM	0.35	-0.11	-0.30
Multi-Markets	LM	0.72	-0.04	-0.12
	DM	0.38	-0.12	-0.49

Table 2.8: Table showing the mean price in ticks relative to the fundamental value for all executed sell orders (market and limit orders). Given all traders with valuation β we obtain the mean price in ticks of all sell orders executed realtive to the fundamental value by traders with the given valuation. Statistics are calculated from 9×10^8 order submissions.

and traders with $\beta = |8|$ at prices 0.30 ticks below the fundamental value. Across the two scenarios, agents' - irrespective of valuation - order submission strategies adopted in the LM lead then to attain better terms of trade in that venue under the Multi-Markets scenario. The increment in mispricing risk impatient agents face in the DM leads to worse terms of trade for these agents (see rows two and three in 'Picked Off' section in Table (2.7)). For example, agents with the extreme private values (*i.e.*, $\beta = |8|$) execute their sell orders on average at prices 0.49 ticks below the fundamental.

2.4.3 The Gains From Trade

In order to test the effects the dark market (DM) has on the profits of the market participants, we calculate the average payoff per trader valuation (*i.e.*, the surplus) under the two available scenarios. To do that we apply the analysis of gains from trade and definitions used by Hollifield et al. (2006).

Suppose a returning trader with private value β enters the market under the Multi-Markets scenario at time t with $z_R = 1$ (*i.e.*, available shares to trade). Assume further that the trader's optimal action $\alpha = \left(m, x, p\right)$ while facing state $s = (mc_{m_L,t}, \beta, \hat{\alpha}, z_R)$ represents a market sell order submitted to the DM. Recall from Eq. (2.3) the term $\left(p - \mathrm{v}_t\right) - \beta$ that captures the instantaneous payoff of an order submitted and executed at time t. This payoff is discounted back to the traders first arrival to the market (*i.e.*, $t' < t$) at rate ρ such that $e^{-\rho(t'-t)}(\left(p - \mathrm{v}_t\right) - \beta)$. Given the above we can decompose the agent's - with private value β - profit as follows

$$
\begin{aligned}
\mathrm{GT} &= -\beta + \beta + e^{-\rho(t'-t)}(\big(p - \mathrm{v}_t\big) - \beta) \\
&= \beta + e^{-\rho(t'-t)}\big(p - \mathrm{v}_t\big) + \beta - \beta e^{-\rho(t'-t)} \\
&= \beta + \beta(1 - e^{-\rho(t'-t)}) + e^{-\rho(t'-t)}\big(p - \mathrm{v}_t\big)
\end{aligned}
\tag{2.14}
$$

where the term $\beta(1-e^{-\rho(t'-t)})$ captures the trader's with private value β cost associated with delaying the trade (waiting cost) and $e^{-\rho(t'-t)}\big(p-\mathrm{v}_t\big)$ captures the gains or losses associated with a trade excluding any gains from the agent's private value (money transfer).[44]

Table (2.9) shows the average payoff, the waiting cost and the money transfer per trader valuation under the two available scenarios. As we can see from the table across the two scenarios traders with the extreme private values (*i.e.*, $\beta = |8|$) in the LM (Single-Market scenario) and in the DM (Multi-Markets scenario) have the lowest waiting costs of -0.180 and -0.247 respectively. In contrast under the Multi-Markets scenario traders with intermediate private values (*i.e.*, $\beta = |4|$) trade faster in the LM than the traders with a β of $|8|$. For example the waiting cost in the LM for the former traders is -0.277 and for the later is -0.290. The above is explained by the switch in the trading strategies adopted by the agents with a β of $|4|$ (*i.e.*, submit 52% of the market orders in the LM, see Table (2.4)). Under the Multi-Markets scenario, the more aggressive strategies adopted by the agents with nonzero private values explain the DM's total lower waiting costs as compared to the LM.

Scenario	Venue	Average Payoff				Waiting Cost				Money Transfer			
		Private Value $\|\beta\|$				Private Value $\|\beta\|$				Private Value $\|\beta\|$			
		0	4	8	Total	0	4	8	Total	0	4	8	Total
Single-Market	LM	0.111	3.667	7.530	3.759	0.000	-0.242	-0.180	-0.151	0.111	-0.091	-0.291	-0.090
Multi-Markets	LM	0.328	3.690	7.603	4.897	0.000	-0.277	-0.290	-0.258	0.328	-0.033	-0.107	-0.031
	DM	0.181	3.486	7.274	3.007	0.000	-0.404	-0.247	-0.190	0.181	-0.110	-0.479	-0.077

Table 2.9: Table showing the average payoff, the waiting cost and the money transfer per trader valuation under the two available scenarios. For full details on how profits are defined see text. Statistics are calculated from 9×10^8 order submissions. The fixed point for each of the available scenarios is obtained independently.

Across the two scenarios irrespective of the venue agents with $\beta = |8|$ incur the higher money transfer. Driven by their exogenous motives to trade, these agents are

[44] In our model we do not pre-assume that a trader with a negative (positive) private valuation will sell (buy) the asset. All decisions are state-dependent and thus they depend on the prevailing market conditions of the LM at any instant.

keen to trade quickly even when the asset is not mispriced. This loss in the money transfer is primarily absorbed by the speculators. For example in the DM under the Multi-Markets scenario and in the LM under the Single-Market scenario agents with extreme values of β incur high money transfer (*i.e.*, -0.479 and -0.291) while speculators obtain money transfer equal to 0.181 and 0.111 ticks respectively. The lower quantities available at the best quotes and throughout the book related to the LM under the Multi-Markets scenario that led to the adoption of more conservative strategies resulted in better money transfer and terms to trade for all agents irrespective of valuation.

As can be seen from the table, across the two available scenarios the average payoffs of all agents irrespective of valuation are the highest in the LM under the Multi-Markets scenario. The average payoff for agents with $\beta = 0$ is equal to 0.328 ticks, while agents with nonzero private values obtain gains equal to 3.960 and 7.603 ticks. Comparing LM's payoffs across scenarios the speculators obtain higher average payoff under the Multi-Markets scenario. In the DM an agent's gains from trade irrespective of valuation are reduced. The properties related to the DM seem to do more harm to the agents with the extreme private values (*i.e.*, $\beta = |8|$).

2.5 Conclusion

In recent years, dark pools' market share in equity markets has increased substantially. In order to study the effects that a 'dark' market has on the 'lit' market we build a dynamic continuous-time model of trade in a single financial asset. For our purposes we test two scenarios where risk neutral traders can choose either to buy or to sell one share of the asset. Under the Single-Market scenario trading takes place in a single limit order book with full pre-trade transparency. Under the Multi-Markets scenario liquidity is scattered between two limit order markets that operate simultaneously. An open 'lit' market (LM) and a 'dark' market (DM). To be consistent with how real financial markets operate, the DM offer no pre-trade transparency to the market participants. This continuous non-displayed 'two-sided' limit order book matches buy (sell) orders at different prices as long as resting limit orders exist on the sell (buy) side of the DM.

To derive our results we use the numerical approach of Pakes and McGuire (2001) and Goettler, Parlour, and Rajan (2009) where market participants at any instant choose optimal actions given the state of the trading game. Trading decisions - irrespective of scenario - are endogenously determined at any instant by the prevailing market conditions of the LM (state-dependent actions). Analysis of the model shows that under the Multi-Markets scenario the DM harms the LM's quality. Across the two scenarios the LM's quoted and effective spread increase; the number of limit orders at the best quotes and at the book as a whole substantially decrease. Moreover the tracking error (*i.e.*, market's inefficiency) increase. In contrast, the order placement strategies adopted by the agents with nonzero private values move transactions in the LM closer to the fundamental value of the asset (*i.e.*, decrease in the LM's microstructure noise).

Under the Multi-Markets scenario, the decrease in the LM's quality (i) increase the picked off probabilities of speculators when they participate in the LM; returning agents with zero private values prefer the DM and its non-displayed properties for the placement of their orders. For returning agents with nonzero private values this is less of an issue, the LM's reduced mictostructure noise coupled with their exogenously desire to trade quickly reduces their picked off probabilities in the LM; (ii) the uncertainty the DM offers in order placement strategies has a strong effect (in terms of higher picked off probabilities, increased time to execution of limit orders and lower percentages of limit orders executing) to agents with nonzero private values; and (iii) all agents irrespective of type to move backwards - in terms of aggressiveness - when they participate in the LM by submitting limit orders. The above strategies drive all agents to obtain better terms of trade when they participate in the LM as compared to the DM.

Agents irrespective of valuation across the two scenarios make higher profits when they participate in the LM under the Multi-Markets scenario. The different strategies adopted by all agents irrespective of valuation in the LM, the reduced LM's microstructure noise coupled with the pre-trade transparency this venue offers benefit all market participants in terms of money transfer and average payoffs.

Chapter 3

Dynamic Equilibrium in Limit Order Markets: Midpoint Dark Pool Vs VWAP Dark Pool, Analysis of Market Quality, Trading Behaviour and Welfare

3.1 Introduction

Dark pools are equity trading systems that offer limited or no pre-trade transparency. In recent years, these Alternative Trading systems (ATSs) - used primarily by institutions so as to avoid market impact or adverse price movements - have increased their market share in global equity markets (*e.g.*, U.S., Europe, Australia and Canada). According to Foley and Putniņš (2016), as of 2013, in the U.S. the most active dark pools - in terms of trading equity volume - accounted for approximately 15% of the consolidated volume; 'dark' trading in Europe represented 10% of the volume, in Australia 14% and in Canada 10%.

Dark pools are subdivided into different groups in terms of the matching mechanism applied to them. The most active dark pools in the U.S. equity trading markets - in descending order in terms of percentage of consolidated equity volume (see Figure 1 in Buti, Rindi, and Werner (2017)) - are the bank/broker pools, the Independent/agency pools and the market maker pools. Bank/broker pools, offer limited price discovery since they match traders' orders at execution prices that are bounded between the National Best Bid and Offer (NBBO). Independent/agency pools offer a single price mechanism to the market participants; that is at any instant these systems match orders at the midpoint of the NBBO or the Volume-Weighted Average Price (VWAP). For the purposes of this study we follow the trading mechanism of independent/agency pools.

By absorbing liquidity from the incumbent exchanges, 'dark' trading has raised concerns - especially among market regulators - on the effects these venues have on market quality and welfare. Numerous proposals and public consultations have been conducted by many regulators and policymakers (European Commission, 2010; Securities and Exchange Commission, 2010). However the inconclusive effects of 'dark' trading have resulted in a hesitant approach towards the development of new regulations related to 'dark' trading.[1] For example, the Securities and Exchange Commission (SEC) Chairman Mary Jo in her prepared speech said that "One of the most difficult tasks in developing the right regulatory response to such potentially disruptive trading strategies is the need to avoid undue interference with practices that benefit investors and market efficiency."[2]

Given the aforementioned (*i.e.*, rising market share of dark pools and the regulatory interest in 'dark' trading) by applying two different price mechanisms to the 'dark' market (*i.e.*, a midpoint dark pool and a VWAP dark pool) we address the following policy question: "How, in a competitive setting where two order markets operate simultaneously, does the different price setting rule of the 'dark' market affect the 'lit' market's quality, trading behaviour and welfare of market participants?"

To be able to answer the above, this study tests two scenarios (i) a 'mid-market' scenario where an open limit order book competes with dark pool that executes orders at any instant at the midpoint of the 'lit' market; and (ii) a 'VWAP-market' scenario where the 'dark' market matches traders orders at any instant at the Volume-Weighted Average Price (VWAP)[3] Understanding the dynamic order choices under the two available scenarios is crucial since of our rational risk neutral agents may use different trading strategies across the two environments. We consider a dynamic continuous-time model of trade in a single financial asset where agents irrespective of the scenario, choose at any instant an optimal action conditional on the prevailing market conditions of the trading game (*i.e.*, state of the game). That is all trading decisions - given the properties of the 'dark' market (non-displayed depth) - depend on the state of the open limit order book. Due to the size of the state-

[1] Limited and contradictory evidence of whether the dark pools' design is sensible, if it encourages predatory trading or generates efficiency within the market.

[2] For more details see Wall Street Journal, (https://www.wsj.com/articles/dark-pools-convince-sec-to-delay-transparency-rules-mary-jo-white-says-1473876535).

[3] For further details about the different price mechanisms imposed on the 'dark' markets, see Section 3.3.

space the equilibrium cannot be determined analytically, so we use the numerical approach of Pakes and McGuire (2001) and Goettler, Parlour, and Rajan (2005, 2009) to identify a Markov-perfect equilibrium.

In this model, we consider two groups of traders who can trade different number of shares of the financial asset. Large traders (*i.e.*, institutions) who can trade at most $\bar{z} > 1$ shares and small traders who can trade one share of the asset. Another factor that gives additional heterogeneity to the market participants is their exogenous motives to trade (*i.e.*, private value of the asset). An agent's private value - irrespective of type - is drawn from a discrete distribution $\mathscr{F}_\beta$ with finite support and is assigned to him before taking any action. We assign the most extreme private values (*i.e.*, $\beta = |8|$) only to large traders so to capture their desire to trade quickly and to avoid adverse price movements (impatient traders). While both groups can have intermediate private values (*i.e.*, $\beta = |4|$); we assign the zero private value (*i.e.*, $\beta = 0$) only to small traders (*i.e.*, no exogenous reasons to trade) so that these traders may only trade profitably when the fundamental value of the asset is mispriced (speculators).[4]

The main findings of this study include (policy implications): (i) across the two scenarios there is an increase in liquidity supply in the 'lit' market under the 'VWAP–market' scenario. The above is supported by a reduced quoted spread, and the increased number of limit orders at the best quotes and throughout the open limit order book as a whole. The 'lit' market's narrower quoted spread reduces the market's informational inefficiency (*i.e.*, reduction in the 'lit' market's tracking error); (ii) the 'lit' market's narrower spread (increased competition) and smaller tracking error, under the 'VWAP-market' scenario, move transactions closer to the fundamental value of the asset when a trader participate in the VWAP 'dark' market as compared to the midpoint 'dark' market; (iii) across the two scenarios, the trading cost of submitting a market order that executes against an outstanding limit order under the 'mid-market' scenario, is higher. That is the effective spreads of all agents under the alternative scenario ('VWAP-market' scenario) are reduced; (iv) across the two 'dark' markets the VWAP proved to be a faster exchange to trade given that the average time between the trader's entry and his order (market or limit) executing is reduced; (v) under both scenarios the demand for liquidity in the 'dark' market

[4] The values assigned to the private value distribution $\mathscr{F}_\beta$ are empirically supported, for further analysis see Subsection 3.3.4.

shifts across small and large traders. Specifically, small traders with zero private values significantly increase their percentage of market orders submitted. The single price mechanism of the 'dark' market coupled with the 'lit' markets' competitive environment, increase the likelihood for agents with zero private value of finding mispriced stale limit orders at the 'dark' market; (vi) the aforementioned strategy adopted by speculators will lead them to obtain worse terms of trade - compared to the 'mid-market' scenario - when they participate in any market under the 'VWAP–market' scenario. Large traders and small traders with the intermediate values are better off - in terms of benefiting from trade - when participate in any market under the 'VWAP-market' scenario; and (vii) traders with $\beta = 0$ (speculators) are the only market participants that make less profits when participate in any market under the 'VWAP-market' scenario as compared to the alternative scenario.

To the best of our knowledge this is the first study that focusses on answering the policy question related to the effects that two different price mechanisms have on the 'lit' market's quality, trading behaviour and welfare of market participants. Some seminar papers - although focussing on answering similar questions on a completely different setup - with contradictory evidence of the effects of 'dark' trading are presented below. Ye (2010) using the Kyle (1985) setup by allowing only large informed traders to freely choose trading venues shows that the addition of the dark pool harms the 'lit' market's price discovery. In contrast Zhu (2013) using the Glosten and Milgrom (1985) framework, shows that when a midpoint dark poll operates in parallel with a 'lit' exchange the former trading system improves the price discovery of the exchange. Degryse, De Jong, and Kervel (2015) analyse 52 Dutch stocks in the Amsterdam Exchange (AEX) with large and mid-cap indices and show that 'dark' trading has a negative effect on the order books' liquidity. Foley and Putniņš (2016) examine the minimum price improvement regulation imposed in Canada and Australia and analyse the effects of 'dark' trading. They show that the 'two-sided' dark trading benefits the 'lit' market's quality (*i.e.*, reduced spreads and increased informational efficiency); in contrast they do not provide clear evidence on the effects of 'one-sided' dark trading (*i.e.*, midpoint or VWAP dark pool) on the 'lit' market's quality. Buti, Rindi, and Werner (2017) using the Parlour (1998) benchmark show that the introduction of a midpoint dark pool that competes with a 'lit' market widens the spread and lower the depth of the latter. Moreover they show that the midpoint dark pool on average reduces the welfare of traders with and without access to the 'dark' market.

This study is organised as follows. Sections 3.2 and 3.3 present the details of the model; Section 3.4 depicts the results and Section 3.5 focusses on the conclusion and policy implications.

3.2 Preliminaries

The trading game involves two groups of traders: large (institutional) traders who can trade $z \in \{1, 2, \ldots, \bar{z}\}$ shares of the asset and small (retail) traders who trade one share of the asset (*i.e.*, $z = 1$). We denote the proportion of large traders by γ^{LT} and the proportion of small traders by γ^{ST}. A factor that gives additional heterogeneity to the market participants, is their intrinsic private value to trade the asset.[5] We denote an agent's private value by β; this value - allocated to an agent before taking any action - is drawn independently across traders from a discrete distribution $\mathscr{F}_\beta$ with finite support. Traders with greater absolute private values in β are impatient to trade since the main benefits arise from their private values. Traders with private value equal to zero may only trade profitably when the asset is mispriced depending on the observable market conditions. Thus each trader in the game has a type $\delta = (\mu, \beta)$ where μ is the trader's group indicator and β is the private value for the asset. We set $\mu = \mu_S$ if the market participant is a small (retail) trader and $\mu = \mu_L$ if the market participant is a large (institutional) trader.

We consider a dynamic continuous-time model of trade in a single financial asset. At any time t, the asset has a common (fundamental) value, denoted as v_t, that evolves over time following a random walk.[6] The time between innovations in v_t follow a Poisson process with parameter λ_{v}. When an innovation occurs, the fundamental value equiprobably increases or decreases by σ ticks.

Trading takes place on two limit order books that operate simultaneously, an open 'lit' market (LM) and a single price 'dark' market (DM) (*i.e.*, 'one-sided' dark market).[7] For our analysis we consider two scenarios depending on the DM's match-

[5] Private value for the asset is also present in the models of Biais, Hillion, and Spatt (1995); Parlour (1998); Theissen (2000); Handa, Schwartz, and Tiwari (2003); Goettler, Parlour, and Rajan (2005, 2009); Hoffmann (2014); Chiarella and Ladley (2016) and Buti, Rindi, and Werner (2017); among many others.

[6] The fundamental value is considered to be the expectation of the future cash flows on the stock.

[7] The definition of 'one-sided' dark trading is taken from Foley and Putniņš (2016).

ing mechanism applied. Although in both scenarios the dark trade executes at a single price ('one-sided' dark trading) under the 'mid-market' scenario the DM matches orders at any instant at the midpoint of the LM. In contrast under the 'VWAP-market' scenario dark trades execute at any instant at the Volume-Weighted Average Price (VWAP). Given the full pre-trade transparency properties of the LM a buy (sell) order of z shares submitted to the DM - irrespective of the scenario - executes at a single price if there is sufficient depth on the sell (buy) of the book.[8] We use the term 'dark' since an order submitted to the DM is not observable to any other active trader but the order submitter. That is the DM irrespective of the scenario offers no pre-trade transparency to the market participants.

New traders arrive at the market according to a Poisson process with parameter λ_N. Given any scenario upon first arrival, small (large) traders can submit an order of one share ($\bar{z}$ shares) to either market or wait until the market conditions change (submit no order). Small (large) traders, after submitting an order which has not been (fully) executed, re-enter the market and revise their orders after a random amount of time (for further analysis, see Subsection 3.3.1). The re-entry time is drawn from a Poisson process with parameter λ_R.

The model discourages agents from infinitely postponing trade by introducing a discount rate. The payoffs of order executions, are discounted back to the order's submission time, at rate ρ. This discount rate which is the same for both groups of agents, captures costs related to delaying transactions or even potential lost trading opportunities (*e.g.*, agents' desire to trade sooner rather than later).[9]

3.3 Model Description

We set $m \in \{m_L, m_D\}$ such that if $m = m_L$ we are referring to the 'lit' market (LM) characteristics and if $m = m_D$ we are referring to the 'dark' market (DM) characteristics. The limit order book related to market m_L, is described by a discrete set of prices such that $\mathscr{P}_{m_L} = \{p^i\}_{i=-\infty}^{+\infty}$. We denote the tick size, which is constant and defines the distance between any two consecutive prices, by $d_{m_L}^{\Delta p}$. At any instant

[8] The factor that determines the sufficiency of the depth depends on the number of shares a trader can trade.

[9] In our model the discount rate ρ does not represent the opportunity cost of capital related to an investment.

$t > 0$, with each price $p^i \in \mathscr{P}_{m_L}$ there are backlogs of unexecuted limit orders ($\ell^i_{m_L,t}$), obeying time priority. Here, positive values in $\ell^i_{m_L,t}$ (*i.e.*, $\ell^i_{m_L,t} > 0$) denote buy orders and negative values (*i.e.*, $\ell^i_{m_L,t} < 0$) sell orders. That is, the limit order book related to the LM at time t denoted by $\mathscr{L}_{m_L,t}$, is the vector of outstanding orders such that $\mathscr{L}_{m_L,t} = \{\ell^i_{m_L,t}\}^{+\infty}_{i=-\infty}$.[10]

The limit order book of the LM obeys price and time priority. Given the limit order book $\mathscr{L}_{m_L,t}$ the bid price and the ask price are defined as follows

$$\begin{aligned} B(\mathscr{L}_{m_L,t}) &= \max\{p^i | \ell^i_{m_L,t} > 0\} \\ A(\mathscr{L}_{m_L,t}) &= \min\{p^i | \ell^i_{m_L,t} < 0\}. \end{aligned} \tag{3.1}$$

The best bid price is the highest price at which there is a buy limit order on the book while the best ask price is the lowest price at which there is a sell limit order. If the order book related to the LM is empty on the buy side or on the sell side at any instant t' we set $B(\mathscr{L}_{m_L,t'}) = -\infty$ or $A(\mathscr{L}_{m_L,t'}) = +\infty$ respectively. Given the bid and ask quotes as defined above, the LM's mid-price (midpoint) at time $t > 0$ will be

$$\text{mp}_t = \frac{A(\mathscr{L}_{m_L,t}) + B(\mathscr{L}_{m_L,t})}{2}. \tag{3.2}$$

That is under the 'mid-market' scenario the backlog of unexecuted limit orders ($\ell^{\text{mp}_t}_{m_D,t}$) to buy ($\ell^{\text{mp}_t}_{m_D,t} > 0$) or to sell ($\ell^{\text{mp}_t}_{m_D,t} < 0$) the asset at price $\text{mp}_t \in \mathscr{P}_{m_L}$ define the depth of the book related to the DM such that $\mathscr{L}_{m_D,t} = \ell^{\text{mp}_t}_{m_D,t}$.[11] In the simulations if the buy (sell) side of the book related to the LM is empty, the bid (ask) price is set to be the last defined price, so as the midpoint to be properly defined at any instant.

Under the 'VWAP-market' scenario the execution price of the DM is set to be

$$\text{VWAP}_t = \lambda \text{VWAP}_{t-1} + (1-\lambda) p_{ex} \tag{3.3}$$

where p_{ex} is the price of the order executed at time t in the LM and $0 < \lambda < 1$ is

[10] We adopt similar notation to Goettler, Parlour, and Rajan (2005, 2009);Bernales (2014); and Chiarella and Ladley (2016) to describe the microstructure features of the model for the dynamic order book market.

[11] In this content under both scenarios we avoid the characterization of a pure 'dark' limit order book. Due to the DM's characteristics (*i.e.*, single price mechanism) agents that choose to participate in the DM cannot post limit orders away from the 'best' price. Unexecuted limit orders can be thought as resting orders that define at any instant the DM's depth (*i.e.*, sell (buy) imbalance).

the smoothing factor. It is worth noticing that VWAP_t is updated every time a trade occurs in the LM. Thus under the 'VWAP-market' scenario the depth of the book related to the DM is defined by the backlog of outstanding limit orders ($\ell_{m_D,t}^{\text{VWAP}_t}$) to buy ($\ell_{m_D,t}^{\text{VWAP}_t} > 0$) or to sell ($\ell_{m_D,t}^{\text{VWAP}_t} < 0$) the asset at price $\text{VWAP}_t \in \mathscr{P}_{m_L}$. Similar to the 'mid-market' scenario we denote the book related to the DM under the 'VWAP-market' scenario by $\mathscr{L}_{m_D,t}$ such that $\mathscr{L}_{m_D,t} = \ell_{m_D,t}^{\text{VWAP}_t}$.

All orders of z shares submitted in this model are *marketable*. Consider for example the 'mid-market' scenario and denote with slight abuse of notation the depth on the sell side of the DM at price p by $\ell_{m_D}^{p}$. A buy order of z shares submitted at price p to the DM is a *market* buy order if the DM's depth $\ell_{m_L}^{p}$ is such that $\ell_{m_D}^{p} \geq z$. That is, the submitted order of z shares results in a trade and the order is removed from the book. If $0 \leq \ell_{m_L}^{p} < z$ the buy order of z shares submitted to the DM at price p partially executes against the available $\ell_{m_L}^{p}$ shares and the remaining shares enter into the 'dark' order book $\mathscr{L}_{m_D}$ at the buy side as a *limit* buy order of $z - \ell_{m_L}^{p}$ shares.[12] Given that there exists enough depth on the other side of the LM, an order submitted to sell (buy) the asset at price below (above) the bid (ask) price is executed immediately at the LM's bid (ask) quote. Finally, under both scenarios due to the price mechanism applied to the DM (single price) unexecuted sell (buy) orders in that venue submitted earlier are ahead in the queue (*i.e.*, no price priority).

All agents that take part in the trading game, irrespective of the group and the scenario, are risk neutral and they observe the fundamental value of the asset at time t with a time lag Δ_t ($\text{v}_{t-\Delta_t}$). That is, both groups have to estimate the actual fundamental value by observing the current (at time t) market conditions.[13] In this model an agent's trading decisions at any instant are state-dependent. Given any scenario upon first arrival at time t, a new small (large) agent, conditional on the state of the trading game, may choose to submit an order of one share ($\bar{z}$ shares) to the LM or to the DM, or not to submit any order. In the former case the agent's choice between the LM or the DM imply choosing the submitted quantity - buy or sell one ($\bar{z}$) unit(s) - and if he decides to participate in the LM the price of the submitted order. The price of the order submitted to the LM (at or above/below quotes) coupled with

[12] For small traders this is not the case since they can trade only one share of the asset. Irrespective of scenario identical order characterizations are used if we focus on the depth at the buy (sell) side of the DM or on the depth at the buy (sell) side of the LM when an order of z shares is submitted.

[13] For further analysis on the learning process applied for the estimation of the fundamental value see Subsection 3.3.2.

the depth (at or above/below the given quote) implies the agent's decision to submit a market or a limit order. In contrast the endogenous decision between market or limit order in the DM depends on the unobservable to the trader depth, on the other side of the DM.

After submitting an order which has not been (fully) executed, returning small (large) agents with available shares to trade re-enter the market and monitor their previous unfilled orders. Given that market participants cannot instantly respond to market changes (*i.e.*, order submission strategies adopted by other new or returning agents or fundamental value changes) they face the risk of being picked off if the fundamental values of the asset moves in an unfavourable direction. Upon re-entry agents can keep or cancel their previous limit order. In the latter case, agents can immediately submit a new order to either market by picking a new set of choices as described in the previous paragraph or wait for different market conditions to apply in order to submit a new one.

The model allows agents with unfilled orders to re-enter the market multiple times. However, small (large) agents stop being part of the trading game after their order (fully) executes. Despite the fact that the (re-) entry rate - drawn from a Poisson process with parameter (λ_R) λ_N - is the same for all agents; trading decisions are endogenously determined by the prevailing market conditions of the LM and the number of active agents. As a result, there is a random number of active market participants at any instant, that monitor the market.

3.3.1 Traders Maximization Problem

Each trader in the game conditional on his type $\delta = (\mu, \beta)$ and on the state of the observable market conditions he observes (*i.e.*, the state of the LM) at any instant, submits an order of z shares to either market so to maximize his expected discounted payoff. We define as $\alpha = (m, x_z, p)$ the action taken by an agent with type $\delta = (\mu, \beta)$ who arrives at the market at time t. Here, p denotes the price at which a trader submits an order.[14] Recall that under the 'mid-market' scenario if the order is submitted to the LM then $p \in \mathscr{P}_{m_L}$ or if it is submitted to the DM then $p = \mathrm{mp}_t \in \mathscr{P}_{m_L}$. Under the 'VWAP-market' scenario if the order is submitted to the DM then $p = \mathrm{VWAP}_t$. Recall that VWAP is updated every time a trade occurs in the

[14] The superscript for the price parameter p is removed for ease of notation.

LM (see Eq. (3.3)).[15] Let $m \in \{m_L, m_D\}$ denote the book choice, where $m = m_L$ if the trader chooses the limit order book related to the LM ($\mathscr{L}_{m_L,t}$) and $m = m_D$ if the trader chooses the book realted to the DM ($\mathscr{L}_{m_D,t}$). Finally, x_z denotes the quantity of the submitted order to buy or sell z units of the asset such that[16]

$$x_z = \begin{cases} +1 & \text{if a buy order of } z \text{ units is submitted,} \\ -1 & \text{if a sell order of } z \text{ units is submitted,} \\ 0 & \text{if no order is submitted.} \end{cases} \tag{3.4}$$

We denote the state observed by a trader with type δ on a particular entry to the market at time $t > 0$, by $s = (mc_{m_L,t}, \delta, \alpha, \phi_\delta) \in \mathscr{S}_\delta$; where $mc_{m_L,t}$ denotes the market conditions of the LM and $\mathscr{S}_\delta$ is the set of possible states a trader with type δ may face. The variable ϕ_δ indicates the remaining number of shares a trader with type δ has to trade.[17] If a returning trader had previously entered the market m by submitting a limit order of z shares (*i.e.*, $x_z \pm 1$) to either market that has not been (fully) executed; the state also contains the action α previously taken by the trader. We define the market conditions at any instant $t > 0$ by the three-tuple

$$mc_{m_L,t} = \left\{ \overset{\star}{\mathscr{L}}_{m_L,t}, \mathrm{v}_{t-\Delta_t}, q_{m_L,t}(\cdot) \right\}. \tag{3.5}$$

Here, $\overset{\star}{\mathscr{L}}_{m_L,t}$ is a set of variables that characterize the limit order book related to the LM, $\mathrm{v}_{t-\Delta_t}$ is the lagged fundamental value and $q_{m_L,t}(p, x_z) \geq 0$ is the trader's order priority among all its respective orders in terms of price and quantity (*i.e.*, $x_z \in -1, 0, +1$).

Suppose at time $t > 0$ a new trader with type $\delta = (\mu, \beta)$ by observing state $s = (mc_{m_L,t}, \beta, \emptyset, \phi_\delta)$ takes an optimal action by submitting an order of z shares at price p to the LM. If the action does not represent an order that fully executes at time t (*i.e.*, market order) we define the order's priority as $q_{m_L,t}(p, x_z) = |\ell^p_{m_L,t} + x_z|$, where $\ell^p_{m_L,t}$ denotes the depth of the limit order book related to the LM.[18] Upon re-entry

[15] From this point and onwards (until Section 3.4) for the sake of repetition - unless stated - given that decisions are state-dependent (*i.e.*, state of the LM) we consider a single price 'dark' market irrespective of the scenario.

[16] Large traders in this model cannot split their orders between either market. That is they cannot submit orders that are in part buy orders, and in part sell orders.

[17] We set $\phi_\delta \in \{0, 1\}$ if the agent is a small trader and $\phi_\delta \in \{0, \bar{z}\}$ if he is a large trader.

[18] It is worth highlighting that in Eq. (3.5) our focus lies on the prevailing market conditions of

at time $t' > t$, if the trader's order has not been fully filled ($\phi_\delta \neq 0$) the trader by observing the state of the trading game, which includes the term $q_{m_L,t'}(\cdot) \leq q_{m_L,t}(\cdot)$, can take further actions in order to maximize his expected payoff.

The state of the trading game at any instant determines the action a particular trader with type δ will choose (*i.e.*, state-dependent trading decisions). Denote the set of possible actions a trade can take in state s by $A(s)$. Formally we define the action space as

$$\begin{aligned} A(s) = \Big\{\emptyset, \big(m, x_z, p\big) \mid\ & p \in [\hat{\mathrm{v}}_t - k, \hat{\mathrm{v}}_t + k] \cap \mathscr{P}_{m_L}, \\ & m \in \{m_L, m_D\}, \\ & x_z \in \{-1, +1\} \Big\}. \end{aligned} \tag{3.6}$$

Recall that an agent irrespective of group observes the fundamental value at any t with a time lag Δ_t ($\mathrm{v}_{t-\Delta_t}$), thus conditional on the prevailing market conditions we denote his best estimate by $\hat{\mathrm{v}}_t = E(\mathrm{v}_t \mid mc_{m_L,t})$.[19] Since the fundamental value of the asset follows a random walk, the set of prices a trade can occur at the process of the game can be unbounded. That is any market participant in this model cares about the relative price of a trade (*i.e.*, all prices are expressed in terms relative to the fundamental value of the asset).

Assume at time $t = 0$ a trader's optimal action $\hat{\alpha} = (\hat{m}, \hat{x}_z, \hat{p})$ when facing state s is to submit an order of z shares at price $\hat{p}$ to market m.[20] We define the trader's expected payoff of the ith ($i \in \{1, 2, \ldots, z\}$) share executing at price $\hat{p}$ and time t_{ex} in the time interval [0, t_{re}] before his re-entry to the market at time t_{re} as

$$\pi(t_{re}, \hat{\alpha}, s, i) = \int_0^{t_{\mathrm{re}}} \int_{-\infty}^{+\infty} \left(e^{-\rho\, t_{ex}} \big((\beta + \mathrm{v}_{t_{ex}} - \hat{p})\, \hat{x}_z \big) \eta(t_{ex}, i \mid \hat{\alpha}, s) \right) \gamma(\mathrm{v} \mid s, t_{\mathrm{ex}}) d\mathrm{v} d t. \tag{3.7}$$

The term $\beta + \mathrm{v}_{t_{ex}} - \hat{p}$ captures the trader's with type $\delta = \big(\mu, \beta\big)$ instantaneous utility of the ith share executing at price $\hat{p}$. We denote the fundamental value of the as-

the LM. The priority of an order submitted at the single price $p \in \mathscr{P}_{m_D}$ of the DM follows only time priority.

[19] For further analysis, see discussion before Eq. (3.13) and Subsection 3.3.4 for the restriction imposed on the price placement of limit orders.

[20] Normalizing the trader's entry time to zero.

set at time t_{ex} by $\mathrm{v}_{t_{ex}}$. As introduced in Eq. (3.4), $\hat{x}_z$ is the quantity of the submitted order and $\gamma(\mathrm{v} \mid s, t_{\mathrm{ex}})$ is the density function of the fundamental value of the asset at time t_{ex} given state s. Additionally, $\eta(t_{ex}, i \mid \hat{\alpha}, s)$ denotes the probability that the trader's action $\hat{\alpha}$ in state $s = \left(mc_{m_L,t=0}, \delta, \emptyset, \phi_\delta\right)$ at time $t = 0$ results in the ith share executing at time t_{ex}, before his re-entry at time t_{re}. If there exist a matching order(s) on the other side of market m the order of z shares submitted at time $t = 0$ immediately executes at price $\hat{p}$ (*i.e.*, market order) and we set $\eta(t_{ex}, i \mid \hat{\alpha}, s) = 1$ for each $i \in \{1, 2, \ldots, z\}$.

Continuing with the above example, consider now the case where the trader re-enters the market m with available shares to trade (*i.e.*, $\phi_\delta \neq 0$). The trader by re-entering the market at time t_{re}, can take further actions given the new state $s' = \left(mc_{m_L,t_{re}}, \delta, \hat{\alpha}, \phi_\delta\right)$ of the LM. The trader can either (i) cancel and resubmit a new order of the available shares at a different price to the same market m. If $m = m_L$ at $p \in \left[\hat{\mathrm{v}}_{t_{re}} - k, \hat{\mathrm{v}}_{t_{re}} + k\right] \cap \mathscr{P}_{m_L} \neq \hat{p}$ else if $m = m_D$ at price $p = \mathrm{mp}_{t_{re}} \neq \hat{p}$ ('mid-market' scenario) or $p = \mathrm{VWAP}_{t_{re}} \neq \hat{p}$ ('VWAP-market' scenario) with $x_z \pm 1$;[21] (ii) cancel and resubmit an order to the alternative market. If the trader's optimal action is to jump from the DM to the LM after re-entry he can choose to submit an order at price $p = (\neq)\mathrm{mp}_t$ or $p = (\neq)\mathrm{VWAP}_t$ with $x_z \pm 1$. Alternatively if his optimal action is to jump from the LM to the DM after re-entry he can choose to submit an order at a single price p with $x_z \pm 1$; (iii) cancel and submit no order ($x_z = 0$); or (iv) keep his unfilled order in market m at the same price and same direction such that $\alpha = (b = \hat{b}, x_z = \hat{x}_z, p = \hat{p})$. The value $\mathrm{V}(s)$ to the agent of being in state s is given by the Bellman Equation of the trader's maximization problem

$$\begin{aligned} \mathrm{V}(s) = \max_{\hat{\alpha} \in A(s)} \int_{t_{re}=0}^{+\infty} \Bigg[\Bigg[\sum_{i=1}^{z} \pi(t_{re}, \hat{\alpha}, s, i) \Bigg] + e^{-\rho\, t_{re}} \int_{s' \in \mathscr{S}_\delta} \mathrm{V}(s') \times \\ \times \psi\left(s' \mid \hat{\alpha}, s, t_{re}, i\right) ds' \Bigg] dG(t_{re}). \end{aligned} \tag{3.8}$$

Here, $\mathscr{S}_\delta$ is the set of possible states a trader with type δ may face on re-entry. $G(t_{re})$ is the exogenous probability distribution of the re-entry that follows an exponential distribution with parameter λ_R (see Section 3.2). In addition, $\psi(s' \mid \hat{\alpha}, s, t_{re}, i)$ is the

[21] The execution price of the DM is determined at any instant by the order placements strategies adopted by the active agents in the LM before the trader's re-entry to the market.

probability that at time t_{re} on re-entry the trader faces state s', given previous action $\hat{\alpha}$, previous state s and i executed orders.[22] The first term - defined in Eq. (3.7) - captures the trader's payoff of the i ($i \in \{0, 1, \ldots, z\}$) shares executing before his re-entry at time t_{re}. The second term reflects the agent's subsequent payoff on re-entry at time t_{re}.

3.3.2 Model Solution of the Trading Game

Ideally, we would like traders to take optimal actions conditional on the entire information set of the trading game. In practice, this makes the model analytically intractable due to the size of the state space. Therefore, to obtain an approximation of the equilibrium, we follow the numerical approach of Pakes and McGuire (2001). Similar to Goettler, Parlour, and Rajan (2005, 2009) and Chiarella and Ladley (2016) the equilibrium concept of the game is an approximation of the true Markov-perfect equilibrium, in which traders' optimal strategies are conditional on the available information set of the reduced state space.[23]

Traders in the trading game - irrespective of type - start with initial overly optimistic expected payoffs of different action and state pairs (*i.e.*, (a, s)). The learning process which is described below allows us firstly to dynamically update traders' expected payoffs for each pair (a, s) by observing the realized payoffs of their actions. The equilibrium is reached when traders' expected payoffs of each action and state pair match the realized outcomes.[24] Additionally, given that the full state space of the trading game is too big, the Pakes and McGuire (2001) algorithm allows us to reduce the dimensionality of the state space size by updating the agents' expected payoffs only for states that are actually visited. That is, states that belong to the recurrent class $\mathscr{R}_\delta$, such that $R_\delta \subset \mathscr{S}_\delta$.[25]

The simulation starts at time $t = 0$, with both the limit order book $\mathscr{L}_{m_L, t=0}$ (LM) and the book $\mathscr{L}_{m_D, t=0}$ (DM) being empty, and is driven by three exogenous events.

[22] Recall that the state of a returning trader contains the remaining order(s) he has available to trade.

[23] For further analysis on the reduced state space see Subsection 3.3.4.

[24] For further details about the convergence criteria used in our simulations see Subsection 3.3.3.

[25] A state i is recurrent iff the expected total number of visits to i is infinite, that is $\sum_{k=0}^{\infty} P_{ii}^{k}$. Two different states i and j communicate $i \longleftrightarrow j$ if $i \rightarrow j$ and $j \rightarrow i$. All the states that communicate with each other form a class. A recurrent class is absorbing *i.e.*, onces the system enters to the class it never leaves.

Recall that the arrival of new traders irrespective of type, the re-entry of traders with $\phi_\delta \neq 0$ (*i.e.*, more shares to trade) and the fundamental value changes, follow the Poisson process with parameters λ_N, λ_R and λ_v respectively.

Suppose a new trader with type $\delta = (\mu, \beta)$ and z shares to trade, enters the market at time $t > 0$. The trader, conditional on the state $s = \left(mc_{m_L,t}, \delta, \emptyset, \phi_\delta\right)$ which incorporates fundamental values changes and entry (re-entry) rates of new (old) traders up to time t, takes an action.[26] At any instant t, each action $\alpha = (m, x_z, p) \in A(s)$ in each state $s \in \mathscr{S}_\delta$ has an expected payoff. We denote the expected payoff by $\mathrm{U}_t(\alpha \mid s) \in \mathbb{R}$. Let $\alpha^\star$ be the agent's payoff-maximizing action in state s, such that $\alpha^\star \in \underset{\alpha \in A(s)}{\operatorname{argmax}} \mathrm{U}_t(\alpha \mid s)$.[27] That is, the value of state s is determined as $\mathrm{V}(s) = \mathrm{U}_t(\alpha^\star \mid s)$.

Over repeated visits to a state $s \in \mathscr{R}_\delta$, the updating process described below calculates each trader's expected payoff for each action taken that results in an order been executed or cancelled so that the trader can pick the optimal action. Consider a new large trader with type $\delta = (\mu_L, \beta)$ entering the market at time $t > 0$ while facing state $s = \left(mc_{m_L,t}, \delta, \emptyset, \phi_\delta\right) \in \mathscr{R}_\delta$. Assume that the action (*i.e.*, $\alpha^\star \in \underset{\alpha \in A(s)}{\operatorname{argmax}} \mathrm{U}_t(\alpha \mid s)$) taken in state s from that particular trader represent a limit order of z shares submitted to the LM or the DM that fully executes at time t' before the trader's re-entry to market m at time t_{re} ($t' < t_{re}$). We update the agent's expected payoff for taking action $\alpha^\star$ in state s and time t, as follows

$$\mathrm{U}_{t'}(\alpha^\star \mid s) = \frac{n_{\alpha^\star,s}}{n_{\alpha^\star,s}+1} \mathrm{U}_t(\alpha^\star \mid s) + \frac{1}{n_{\alpha^\star,s}+1} e^{-\rho(t'-t)} z(\beta + \mathrm{v}_{t'} - p) x_z. \tag{3.9}$$

Here, p is the price of the order with submitted quantity x_z which executes at time t' prior to the trader's re-entry at time $t_{re} > t'$, when the fundamental value is $\mathrm{v}_{t'}$. The expected payoff is updated so that it is equal to the average payoff attained when the optimal action $\alpha^\star$ is taken in state s. The term $n_{\alpha^\star,s}$ is a positive integer that records the number of times the action $\alpha^\star$ is chosen in state s. Conversely, if the trader's order does not execute (*i.e.*, $\phi_\delta = z$) before his re-entry to the market at time t_{re}, we update his expected payoff as follows:

[26] Every order submission strategy adopted by an active participant changes the state of the limit order book related to market m_L (state variable) and consequently changes the states observed by the next agents.

[27] Since on any entry the action set is finite, the maximum over all other actions exists and it is well defined.

$$\mathrm{U}_{t_{re}}(\alpha^\star \mid s) = \frac{n_{\alpha^\star,s}}{n_{\alpha^\star,s}+1}\mathrm{U}_t(\alpha^\star \mid s) + \frac{1}{n_{\alpha^\star,s}+1}e^{-\rho(t_{re}-t)}\mathrm{V}(s_{re}). \tag{3.10}$$

Here $\mathrm{V}(s_{re})$ is the continuation value of state $s_{re} = \left(mc_{m_L,t_{re}}, \delta, \alpha^\star, \phi_\delta\right)$ when the trader re-enters the market at time t_{re}, given previous action $\alpha^\star$ and previous state s.

Suppose now that the action $\alpha^\star \in \underset{\alpha\in A(s)}{\operatorname{argmax}}\,\mathrm{U}_t(\alpha \mid s)$ chosen by the new large trade in state $s = \left(mc_{m_L,t}, \delta, \emptyset, \phi_\delta\right) \in \mathscr{R}_\delta$ and time $t > 0$ represents a limit order of $z = 2$ shares that partially executes in market m and time t'' (*i.e.*, $\phi_\delta = 1$) before the trader's re-entry at time t_{re}.[28] Since the order submitted in state s and time t is still active the trader given the new state $s' = \left(mc_{m_L,t_{re}}, \delta, \alpha^\star, \phi_\delta\right)$ at time t_{re} will choose an optimal action with value $\mathrm{V}(s') = \mathrm{U}_{t_{re}}(\alpha^{\star\star} \mid s')$. Given the above, we update the trader's expected payoff for the order of $z = 2$ submitted in state s and time t that partially executes as follows

$$\mathrm{U}_{t_{re}}(\alpha^\star \mid s) = \frac{n_{\alpha^\star,s}}{n_{\alpha^\star,s}+1}\mathrm{U}_t(\alpha^\star \mid s) + \frac{1}{n_{\alpha^\star,s}+1}\left[e^{-\rho(t''-t)}(\beta + \mathrm{v}_{t''} - p)x_z + e^{-\rho(t_{re}-t)}\mathrm{V}(s')\right]. \tag{3.11}$$

The continuation value of state s', denoted as $\mathrm{V}(s')$, allows the agent in the trading game to profitably execute the order of the unexecuted shares in the future. Suppose that the particular large trader with $\phi_\delta = 1$ (*i.e.*, one share available to trade) is randomly drawn to re-enter to market m at time w. As we have seen in Eq. (3.8) the value to a trader of being in state s' at time t_{re} can be decomposed into the payoff from full (partial) execution before re-entry and by the continuation payoff in case of re-entry to market m at time w. Even if the order is not fully executed at time w the agent will still choose the optimal action (*i.e.*, $\alpha^{\star\star} \in \underset{\alpha\in A(s'')}{\operatorname{argmax}}\,\mathrm{U}_w(\alpha \mid s'')$) in the new state $s'' = \left(mc_{m_L,w}, \delta, \alpha^{\star\star}, \phi_\delta\right)$. That is an agent's future optimal actions at different states, that result in the order being fully filled are used to determine the expected payoff from taking action $\alpha^\star$ in state s and time t.

Continuing with the example above for the trader of type $\hat{\delta} = (\hat{\mu}, \hat{\beta})$ who submitted the order of $\hat{z} = 1$ share at price $\hat{p} = p$ and time t'' by facing state $s''' = \left(mc_{m_L,t''}, \hat{\delta}, \emptyset, \hat{\phi}_\delta\right)$ we update his expected payoff as follows

[28] The value assigned to z is used for illustration purposes.

$$\mathrm{U}_{t''}(\hat{\alpha}^{\star} \mid s''') = \frac{n_{\hat{\alpha}^{\star},s'''}}{n_{\hat{\alpha}^{\star},s'''}+1}\mathrm{U}_{t''}(\hat{\alpha}^{\star} \mid s''') + \frac{1}{n_{\hat{\alpha}^{\star},s'''}+1}(\hat{\beta}+\mathrm{v}_{t''}-\hat{p})\hat{x}_z. \tag{3.12}$$

Here $\hat{\alpha}^{\star}$ is the optimal action take by the trader with type $\hat{\delta} = (\hat{\mu}, \hat{\beta})$ at time t while facing state s'''. For each pair $(\alpha,\ s)$ in the simulation process we assign an initial positive integer $n_0 = n_{a,s}$. As $n_{a,s}$ gets very large for some pairs, this integer reduces the convergence speed of the algorithm. To achieve faster convergence, after 3×10^8 order submissions we reset the counter to its initial value n_0. Moreover, by assigning optimistic initial payoffs to all agents, states that are only visited at the start of the optimization, once execution probabilities are updated, stop been part of the recurrent set $\mathscr{R}_{\delta}$.

Agents irrespective of type conditional on the market conditions of the LM have to estimate the fundamental value before taking an action. Recall that we denoted the expectation of the fundamental value of the asset at time t by $\hat{\mathrm{v}}_t = E(\mathrm{v}_t \mid mc_{m_L,t})$ (see discussion after Eq (3.6)). To be able to define a trader's action set $A(s)$ we set $\hat{\mathrm{v}}_t = \mathrm{v}_{t-\Delta_t} + \zeta_n(mc_{m_L,t})$ where $\zeta_n(mc_{m_L,t})$ is the extent by which an agent at time t revises the term $E(\mathrm{v}_t \mid mc_{m_L,t})$ given lagged value $\mathrm{v}_{t-\Delta_t}$. We define

$$\zeta_{n+1}(mc_{m_L}) = \frac{\hat{n}}{\hat{n}+1}\zeta_n(mc_{m_L}) + \frac{1}{\hat{n}+1}(\mathrm{v}_t - \mathrm{v}_{t-\Delta_t}) \tag{3.13}$$

where $\hat{n}$ counts the number of times market condition m is encountered during the simulation process.[29]

Assigning initial overly optimistic expected payoffs to the market participants each time a state is visited for the first time is very important for the optimization process. The above lead agents - for every new state state encountered in the simulation - to submit orders with very low probability of execution but with very high expected payoffs. Through the learning process described above traders learn that these orders are not executing and will move to orders that offer lower expected payoffs but have higher probability of executing. If the trader was pessimistic (*i.e.*, by submitting orders with higher execution probability but lower expected payoffs) he would not have been able to identify if more profitable orders will execute.

[29] As we will see in Subsection 3.3.3; since our focus lies on stationary equilibria, the time subscript on market conditions is omitted.

3.3.3 Existence and Convergence Criteria

To be able to compute the fixed point where an agent's expected payoffs actually match the realized ones we follow the approach proposed by Goettler, Parlour, and Rajan (2009); Bernales (2014); and Chiarella and Ladley (2016).[30] That is initially we run the optimization process as described in Subsection 3.3.2 for several billion order submissions. After this point we keep fixed the expected payoffs for each market participant with type δ (denoted as $\hat{\mathrm{U}}(\cdot)$) and we re-run the model for several billion more order submissions where the actual (realized) payoffs are calculated. Suppose under the 'mid-market' scenario a large trader with type $\delta = (\mu_L, \beta)$ enters to the market at time $t > 0$ with $z = 2$ shares to trade and his optimal action involves submitting an order to the DM at price $p = \mathrm{mp}_t$.[31] Assume further at time w_1 and time w_2 with $w_2 > w_1$ before the trader's random re-entry to the market the order submitted fully executes. Since the order resulted in a trade at time w_2 the trader leaves the market by obtaining a realized payoff. This realized payoff denoted by $\hat{\mathrm{V}}(s)$ is equal to $\hat{\mathrm{V}}(s) = e^{-\rho(w_1-t)}(\beta + \mathrm{v}_{w_1} - p)x_z + e^{-\rho(w_2-t)}(\beta + \mathrm{v}_{w_2} - p)x_z$. For each pair (α, s) encountered in the simulation we then compare the expected and the realized payoffs. Under both scenarios if the mean absolute difference between the expected and the realized payoffs weighted by the number of times the pair (α, s) is encountered in the simulations is less than 0.03 and the correlation between $\hat{\mathrm{U}}(\cdot)$ and $\hat{\mathrm{V}}(\cdot)$ is more than 0.99 we conclude that the model has converged and the fixed point is found. If the above criteria are not satisfied we repeat the optimization process and the tests until convergence is achieved. Our focus lies on symmetric and stationary equilibria where agents with the same type adopt the same strategies and these strategies do not depend on time; since an agent with type δ will choose the same action when facing the same state s more than once.

3.3.4 Numerical Parametrization

To support the values assigned to the parameters γ^{LT} and γ^{ST} (proportion of institutional and retail traders respectively) we follow approximately the empirical find-

[30] We do not prove uniqueness although - for each parametrization - we run several simulations using different random seeds to guarantee computational uniqueness (*i.e.*, the model converges to the same equilibrium).

[31] See footnote 28.

ings of Menkhoff and Schmeling (2010). Analysing the data set of the order book of an electronic FX market with complete order details and coded trader identities they show that information about the counterparty affects the future trading decisions of individuals. Using the total trading volume as a proxy for overall trade size they allocate each trader, depending on its size, into three groups (small, medium and large trader group). Given the total number of market participants, they observe that 25% ($\simeq$ 24.37%) are large and medium traders while the rest are small traders. Thus we set $\gamma^{LT} = 0.2$.

The binary variable $\mu = \{\mu_S, \mu_L\}$ is the agents' group indicator; we set $\mu = \mu_S$ if the agent is a small (retail) trader and $\mu = \mu_L$ if he is a large (institutional) one. Small traders can trade one share of the asset ($z = 1$) while large trader can trade at most three shares of the asset ($\bar{z} = 3$). Agents are part of the trading game by revising their unfilled orders until they are fully executed, that is $\phi_\delta = 0$. Assume a new large trader enters to the market at time t and state s and submits an order of z shares to the DM at price mp_t and direction x. The trader re-enters to the market at random time $t' > t$ while observing state s'. If his previous order is not executed ($\phi_\delta \neq 0$ or $\phi_\delta = z$), the trader can take further actions conditional on the new state s'. He can retain the order in the DM, cancel and re-submit a new order to the LM or to the DM or submit no order. Conversely if $\phi_\delta = 0$, the trader does not have any shares to trade and leaves the market forever.

All traders irrespective of type observe the fundamental value of the asset with a time lag Δ_t. We set $\Delta_t = 10$, that is traders observe the fundamental with a time lag of ten seconds. At any instant, agents of any type, can place their limit orders to the LM at any feasible price that lies in the interval $p \in [\hat{v}_t - k, \hat{v}_t + k] \cap \mathscr{P}_{m_L}$. We set the value of the parameter $k = 13$; that is, the price of a limit order at any t may be submitted in a range between thirteen ticks above and below the trader's expectation of the fundamental.[32] It is worth highlighting that conditional on the price mechanism applied to the DM under any scenario ('mid-market' and 'VWAP-market' scenarios) there are no orders in that venue away from the single price.

In the model, to ensure numerical tractability: (i) all agents care about the rel-

[32] The restriction imposed on the price placement of limit orders at any feasible price within k ticks away from the trader's expectation is based on two reasons (i) computational tractability; and (ii) order placement strategies adopted by agents. Limit orders submitted at ticks greater than $k = 13$ were rarely executed..

ative price of a trade, the difference between the price and the fundamental value v (see discussion after Eq. (3.6));and (ii) given all the possible states a trader with type δ might face, we consider only states that belong to the recurrent set such that $s \in R_\delta \subset \mathscr{S}_\delta$. All the above, allow us by restricting the state space of the trading game for each agent, to compute the fixed point. Hence, the state at any instant $t > 0$ is defined by the four-tuple

$$s = (\underbrace{\{\overset{\star}{\mathscr{L}}_{m_L,t}, \mathrm{v}_{t-\Delta_t}, q_{m_L,t}(\cdot)\}}_{= mc_{m_L,t}}, \delta, \alpha, \phi_\delta) \tag{3.14}$$

where $mc_{m_L,t}$ are the market conditions of the LM observed by the agent with type δ (see Eq. (3.5)). If the returning trader's order of z shares, submitted at time $t' < t$ and state s', has not been fully executed before his re-entry, α is the trader's previous action taken in state s' and ϕ_δ denotes the remaining shares the trader can trade. The variables in $\overset{\star}{\mathscr{L}}_{m_L,t}$ that characterize $\mathscr{L}_{m_L,t}$ at time t, are: (i) the bid price and the ask price; (ii) the depth at the bid quote and the ask quote (*i.e.*, $\ell^B_{m_L,t}$ and $\ell^A_{m_L,t}$); and (iii) the buy and the sell quantities in the $\mathscr{L}_{m_L,t}$ (*i.e.*, $\sum_{i=0}^N \ell^i_{m_L,t} > 0$ and $\sum_{i=0}^N \ell^i_{m_L,t} < 0$).

We set the support in ticks of the private value distribution $\mathscr{F}^\mu_\beta$ to be $\{-8, -4, 0, 4, 8\}$ with cumulative distribution $\{0.0, 0.2375, 0.7625, 1.0, 1.0\}$ for the small agents; and with cumulative distribution $\{0.35, 0.5, 0.5, 0.65, 1.0\}$ for large traders. The values assigned to the private value distribution $\mathscr{F}^\mu_\beta$ follow approximately the findings of Hollifield et al. (2006). By splitting the private value into five different intervals, they estimate the distributions of private values and the optimal order submission strategies for three sample stocks on the Vancouver exchange.[33] They observe that across the three stocks the mean probability of private value in the interval $(-2.5\%, +2.5\%)$ is 44%; in the intervals $(-5\%, -2.5\%]$ and $[2.5\%, +5\%)$ is 26%; and in the intervals $(-\infty, -5\%]$ and $[+5\%, +\infty)$ is 30%. Combining the private value distributions the number of shares each trader type can trade and the proportion of large and small traders that participate in the trading game we can see that for small agents (*i.e.*, $z = 1$) with a β of 0 it holds that $0.8 \times 0.525 \times z = 0.42$ while for large traders (*i.e.*, $\bar{z} = 3$) for the same β it holds that $0.2 \times 0.0 \times \bar{z} = 0.0$. Similarly for a large traders with a $\beta = |8|$ we have that $0.2 \times 0.7 \times \bar{z} = 0.42$ and for small traders $0.8 \times 0.0 \times z = 0.0$. Agents with extreme private values (*i.e.*, our large traders) under the risk off being

[33] The three sample stocks are Barkhor Resources (BHO), Eurus Resources (ERR), and War Eagle Mining Company (WEM).

picked off or to avoid high waiting costs are keen to trade quickly so to capture their exogenous motives to trade. Agents with zero private values (*i.e.*, our small traders) may only trade when the fundamental value of the asset is mispriced by posting more conservative limit orders (*i.e.*, no exogenous motivation to trade).

The fundamental value of the asset at any t, denoted as v_t, follows a Poisson process with rate λ_{v}. This parameter denotes the expected number of innovations of the fundamental value in one second. Similar to Goettler, Parlour, and Rajan (2009) low volatility case, we set $\lambda_{\mathrm{v}} = 0.125$, that is the fundamental value increases or decreases by σ ticks every 8 seconds on average. We set $\sigma = 1$, so when an innovation occurs the fundamental value increases or decreases with equal probability by one tick.

New agents irrespective of type arrive randomly to the market following a Poisson process with intensity λ_N. To support the value assigned to λ_N, we use the empirical findings of Conrad, Wahal, and Xiang (2015). By sampling all common stocks in the US with a stock price and market cap greater than \$1 and \$100 million respectively, Conrad, Wahal, and Xiang (2015) examine the relation between high frequency quotations and the behaviour of stock prices for the period 2009 − 2011. By splitting firms into different size quantiles and by averaging over the entire time series, they observe that the average number of trades for stocks in the largest size quantile is almost one trade per second.[34] Given the above and conditional that each trade involves two traders we set $\lambda_N = 2$. That is, a trader enters to the market on average every 0.5 seconds.

In the model, an agent irrespective of type δ and scenario with $\phi_\delta \neq 0$ (more shares to trade) re-enters the market and revises his previous order according to a Poisson process with intensity λ_R. We set $\lambda_R = 0.1$, that is traders re-enter to the market and revise their orders on average every 10 seconds. This is approximately in line with the empirical findings of Hautsch and Huang (2011), where the median cancellation time of one asset (WCRX) among the 200 biggest stocks listed on NASDAQ is 10 seconds.[35]

In order to capture the risk of an agent submitting a limit order that might not result to a trade and thus to be able to find the fixed point we introduce the discount

[34] Stocks in the two largest size quantiles accounted for over 80% of the market cap.
[35] In terms of market capitalization on October 1, 2010.

rate denoted by ρ. In this model we set the discount rate ρ equal to 0.03.[36] The discount rate is common under any scenario and for both the LM and the DM and captures the agents preference to execute sooner than later. Moreover we set the smoothing factor λ equal to 0.3. This parameter is used to set the matching price of the DM under the 'VWAP-market' scenario (see Eq. (3.3)).[37]

3.4 Results

As described in Subsection 3.3.3 to ensure that the fixed point is computationally unique we run simulations using different random seeds. In this Section we introduce the model's results after the convergence criteria have met for 9×10^8 order submissions.

3.4.1 Market Quality

In this section we analyse how the introduction of the single price DM under the 'mid-market' and the 'VWAP-market' scenarios affects the quality of the LM. To do this we introduce different quality measures such as the bid-ask spread, the number of limit orders (total and effectively traded) at the bid and at the buy side of the books as a whole (the buy side is symmetric to the sell side). Table (3.1) depicts an analysis of the aforementioned measures. As can be seen from the table across the two scenarios there is an increase in liquidity supply in the LM under the 'VWAP-market' scenario. For example the number of limit orders at the bid and at the buy side of the limit order book related to the LM increase from 3.771 to 3.800 and from 17.227 to 20.057 ticks respectively. Moreover this increase in liquidity competition results in the reduction of the spread (*i.e.*, from 1.525 to 1.454 ticks) and the number of effectively traded limit orders (*i.e.*, limit orders submitted and executing without any modification) at the bid price and at the buy side of the book as a whole. In contrast to the 'mid-market' scenario limit order submitters under the 'VWAP-market' scenario are keen to cancel their unexecuted limit orders more frequently. To justify

[36] We simulated the trading game by setting $\rho = 0.05$ and $\rho = 0.01$ and the results we derived were qualitatively similar to our benchmark case although with lower values of ρ agents spent more time in the market increasing the model's running time.

[37] We experiment with different values of λ (*i.e.*, $\lambda = 0.5$ and $\lambda = 0.6$) and the results are qualitatively similar.

mid-market			VWAP-market	
LM	DM		LM	DM
		Spread		
1.525	–		1.454	–
		N. of limit orders at the bid		
3.771	2.575		3.800	1.778
		N. of limit orders at the bid (effectively traded)		
1.013	1.759		0.915	1.526
		N. of limit orders at the buy side of the book		
17.227	2.575		20.057	1.778
		N. of limit orders at the buy side of the book (effectively traded)		
4.885	1.759		4.798	1.526
		Microstructure noise: $\lvert p - v_t \rvert$		
0.674	0.863		0.666	0.738
		Tracking Error		
0.830	0.839		0.789	0.648

Table 3.1: Table showing market quality and liquidity under both available scenarios. The buy side is symmetric to the sell side. The spread is the distance in ticks, between the best ask and the best bid. The number of effectively traded limit orders at the bid price and at the buy side of the book are the limit orders submitted and executed without any modification. The tracking error is the average absolute distance between the midprice and the fundamental value. The microstrucutre noise is defined to be the absolute mean difference between the execution price and the fundamental value of the asset. Due to the DMs matching mechanism the spread is equal to zero (single price markets). The number of limit orders at the bid price of the DMs (total and effectively traded) is equal to the number of limit orders at the buy side of the DMs (total and effectively traded). All the statistics are calculated by observing the market every 10 simulation minutes. Number of observations: 85249 (mid-market), 79264 (VWAP-market). The fixed point for each of the available scenarios is obtained independently.

the above suppose at a random time before the trader's re-entry to the market the fundamental value has moved in an adverse (favourable) direction. The LM's more competitive environment - under the 'VWAP-market' scenario - will lead that particular trader to cancel his unexecuted limit order more often and re-position himself, in terms of aggressiveness, by posting a more conservative (aggressive) limit order so as to stay protected (at competitive prices) (see statistics in Table (3.5)).

Comparing further the two scenarios we focus next next on the microstructure noise difference between the two single price DMs. The microstructure noise is defined to be the mean absolute distance between the execution price and the fundamental value of the asset. The smaller the distance the closer the execution price is to the fundamental value. The fact that both books related to the DM are LM-dependent (*i.e.*, at any instant the execution prices are derived from the LM) justify the above. The different matching criteria imposed on these two venues (VWAP is updated every time a trade occurs in the LM), the narrower spread (increased competition) and the smaller tracking error (reduction in informational inefficiency) of the LM under the 'VWAP-market' scenario move transactions closer to the fundamental value of the asset. Moreover as can be seen from the table, across both scenarios, the number of limit (resting) orders submitted and executed without any modification at the single price of both DMs (*i.e.*, effectively traded) increase in contrast to the LMs. The single price mechanism applied to both DMs increase the likelihood of finding stale limit orders at these venues.

In order to have a better understanding on the execution costs related to a market order we introduce the effective spread. Similar to Chiarella and Ladley (2016) the effective spread - used to capture microstructure features of market m - is defined as: $2(p - mp)x_z$ where p is the execution price, mp is the mid-price when the order is submitted and x_z is the order's submitted quantity to buy or to sell z units of the asset, with $x_z = 1$ ($x_z = -1$) if the transaction involves a market buy (sell) order.[38] The effective spreads of each trader type in the LM under both scenarios are illustrated in Table (3.2). As can be seen from the table under the 'mid-market' scenario the effective spreads of both groups of agents irrespective of private value are greater. That is under the 'mid-market' scenario the trading cost of submitting a market order that executes against an outstanding limit order on the other side of the book related to the LM is higher. The improvement of the LM's quality (as we have seen in Table (3.1)) and the increased competition that narrows the spread of the limit order book related to the market m_L under the 'VWAP-market' scenario reduces the cost associated to submitting a market order.

Under both scenarios the effective spreads of large traders with the most extreme private values ($\beta = |8|$) are greater than the rest of the market participants. Large traders with a β of $|8|$ are keen to trade quickly in either market, even when the fundamental is not mispriced to benefit from their exogenous values. In contrast the effective spreads of small traders with zero private valuation have the lowest values. These traders may only trade profitably when the fundamental value is mispriced (*i.e.*, place market orders when significant mispricing is present) since their main

[38] Given that the execution price of the DM under the 'mid-market' is the mid-price (midpoint) we will focus only on the effective spreads of traders when they submit a market order to the LM.

Scenario	Venue	Traders	Effective spread			
			Private Value $\|\beta\|$			
			0	4	8	Total
mid-market	LM	Small	1.251	1.306	–	1.285
		Large	–	1.359	1.446	1.427
VWAP-market	LM	Small	1.219	1.232	–	1.227
		Large	–	1.292	1.370	1.354

Table 3.2: Table showing the effective spreads in the LM by each trader type under the two available scenarios. The effective spread is defined to be $2(p-mp)x_z$ where p is the execution price, mp is the mid-price when the order is submitted and x_z is the order's submitted quantity to buy or to sell z units of the asset, with $x_z = 1$ ($x_z = -1$) if the transaction involves a market buy (sell) order. The effective spreads of small traders with a β of |8| and the effective spreads of large traders with β of 0 are not available due to the values assigned to the private value distribution $\mathscr{F}^{\mu}_{\beta}$ (for further details see Subsection 3.3.4). All the statistics are calculated by observing the market every 10 simulation minutes.

benefits are obtained from the trading activity. For example large traders with a β of |8| are executing buy (sell) markets orders in the LM under the 'mid-market' scenario 0.723 ticks above (below) the market's midprice while under the 'VWAP-market' scenario the same traders execute at 0.685 ticks. Small traders with a β of 0 on the other hand execute market buy (sell) orders 0.625 ticks and 0.609 ticks above (below) the market's midprice under the 'mid-market' and the 'VWAP-market' scenario respectively.

Scenario	Venue	Traders	Time to Execution [Time to Execution limit orders]
mid-market	LM	Small	22.112 [28.322]
		Large	2.187 [3.823]
	DM	Small	30.357 [37.205]
		Large	4.061 [7.944]
VWAP-market	LM	Small	24.962 [33.192]
		Large	1.986 [3.582]
	DM	Small	27.584 [34.599]
		Large	3.874 [7.222]

Table 3.3: Table showing average time interval between a trader entering the market and his order executing (*i.e.*, market and limit order). In brackets is the average time interval between a trader entering the market and his limit order executing. Statistics are collected from 9×10^8 order submissions.

Table (3.3) further supports our findings. As can be seen from the table irrespective of the scenario, small traders have longer waiting times in contrast to large traders.[39] The latter agents despite the larger number of shares they can trade are willing to obtain the gains from trade available in the market (*i.e.*, even when there is no mispricing) instead of waiting. Irrespective of the scenario, small and large traders due to the properties of the DMs (*i.e.*, non-displayed liquidity) and the single price mechanism imposed on these venues, have spent a longer time in these markets. In contrast to the 'mid-market' scenario small agents under 'VWAP-market' scenario, spend a longer time in the LM; the increased competition coupled with the decreased spread, the smaller microstructure noise and the tracking error of the LM drive these agents to cancel their outstanding limit orders more frequently and as a result to have longer waiting times. This environment seems to benefit large traders since their average time to execution in the LM is reduced even more. Moreover comparing the two scenarios all agents spend less time in the DM under the 'VWAP-market' scenario. As we will see in the next section, the different strategies adopted by the market participants will verify the above observation.

3.4.2 Trading Behaviour

In the previous Subsection we saw that in contrast to 'mid-market' scenario the presence of the VWAP DM improves the LM's market quality. Liquidity provision in the LM is increased; the quoted spread, the microstructure noise and the tracking error decrease the number of limit orders at the bid and at the buy side of the book related to the LM increase. Moreover, the effective spreads of all agent types are reduced. In this section we are going to analyse the trading behaviour of the agents under the two available scenarios (*e.g.*, order submission strategies) and the effects these strategies have in terms on benefits from trade and price of executed orders.

Order Placement Strategies

At any instant, agents with zero private values (*i.e.*, speculators) - given the state of the trading game they observe as defined in Eq. (3.14) - may only trade profitably if the fundamental value is mispriced. In contrast agents with the most extreme values (*i.e.*, $\beta = \pm 8$) are more likely to submit orders even when the asset is not mispriced since their main benefits arise from their exogenous private values.

[39] Recall from Subsection 3.3.4 that we assign a β of zero only to small traders and a β of $|8|$ only to large traders.

Table (3.4) displays the order placement strategies (*i.e.*, market orders submitted per trader type) adopted by the market participants under the 'mid-market' and the 'VWAP-market' scenarios. As can be seen from the table across the two scenarios small traders with zero private valuation prefer to trade by submitting limit orders to the LM (*i.e.*, percentage of market orders submitted close to 17%); large traders on the other hand, with $\beta = |8|$ demand liquidity by submitting the majority of the market orders. The above is driven by the different private value distributions assigned to these agents. Recall that in our model we only consider small traders with a β of 0 and $|4|$ and large traders with a β of $|4|$ and with the most extreme private values (*i.e.*, $\beta = |8|$).

Scenario	Venue	Traders	market orders submitted (%)		
			Private Value $\|\beta\|$		
			0	4	8
mid-market	LM	Small	18.60%	29.43%	–
		Large	–	11.16%	40.81%
	DM	Small	52.14%	27.03%	–
		Large	–	9.67%	11.16%
VWAP-market	LM	Small	17.08%	30.05%	–
		Large	–	11.02%	41.85%
	DM	Small	40.15%	35.58%	–
		Large	–	12.90%	11.37%

Table 3.4: Table showing the percentage of market orders submitted by each trader type under both scenarios. The percentages of small traders with a β of $|8|$ and the percentages of large traders with β of 0 are not available due to the values assigned to the private value distribution $\mathscr{F}^{\mu}_{\beta}$ (for further details see Subsection 3.3.4). Statistics are collected from 9×10^8 order submissions.

As shown further in Table (3.4) under both scenarios the demand for liquidity in the DM shifts across small and large traders. Small agents with zero private values (*i.e.*, small traders) considerably increase their demand for liquidity. For example the percentage of market orders submitted by agents with a β of zero almost triples (from 18.60% to 52.14%) under the 'mid-market' scenario and is more than double (from 17.08% to 40.15%) under the 'VWAP-market' scenario. In contrast the percentage of market orders submitted by large traders with a β of $|8|$ decreases from 40.81% to 11.16% under the 'mid-market' scenario and from 41.85% to 11.37% under the 'VWAP-market' scenario. The reasons for this shift in liquidity demand across large and small agents are (i) the price mechanism applied to both DMs that increases the likelihood for agents with zero private values of finding mispriced stale limit orders at the DM submitted by large traders; (ii) the competitive environment of both LMs that increase the probability - primarily of limit order submitters - of being picked off after unfavourable fundamental value movements coupled with the anonymity the DM offers (see Tables (3.1) and (3.5)); and (iii) the exogenous characteristics of large traders with the most extreme private values that drive these agents to trade quickly in the LM even if the fundamental value is not mispriced.

mid-market					VWAP-market			
LM		DM			LM		DM	
Small	Large	Small	Large		Small	Large	Small	Large
				Number of limit order cancellations per trader				
2.39	0.10	1.24	0.26		2.70	0.09	0.84	0.29
				Probability of being picked off				
35.29%	12.95%	26.60%	19.34%		37.43%	12.69%	22.86%	19.61%

Table 3.5: Table showing the number of limit orders cancellations per trader, divided by group. The probability of being picked off is calculated when $x(v_t - v_s) < 0$, where v_s is the fundamental value at submission time, v_t is the fundamental value at execution time and x_z is the limit order's quantity when submitted. Statistics are collected from 9×10^8 order submissions.

In contrast to the 'mid-market' scenario small and large traders with intermediate characteristics (*i.e.*, $\beta = |4|$) increase even more their demand for liquidity under the alternative scenario (*i.e.*, 'VWAP-market' scenario). As we will see later the different strategies adopted by these agents in terms of aggressiveness in the DM coupled with the LM's improved quality and the DM's price mechanism (see Subsection 3.4.1) under the 'VWAP-market' scenario will lead them to obtain better terms of trade.

Order Aggressiveness

Consider a small trader with type δ entering the market at time $t > 0$ and his optimal action in state s involves submitting an order to market m at price p with $x = -1$ (*i.e.*, sell order). Given that any order submitted in our model is marketable; if there is no order on the buy side of the book the order is entered into the book related to market m as a limit order. If at time $t' > t$ the order executes before the trader reentry then the trader obtains a realized payoff equal to $e^{-\rho(t'-t)}(p - \mathrm{v}_{t'} - \beta)$. Agents with zero private values should demand a sell price higher than v_t while agents with $\beta \neq 0$ should be willing to demand a lower sell price, even negative.

Table (3.6) displays the average price of sell limit orders submitted by each agent type under the 'mid-market' and the 'VWAP-market' scenarios. As shown in the table, under both scenarios small traders (*i.e.*, agents with $\beta = 0$ and $\beta = |4|$) - although a little more aggressive under the 'VWAP-market' scenario - submit more conservative limit orders to either market as compared to large traders. For example under the 'mid-market' scenario traders with the most extreme private values (*i.e.*, $\beta = |8|$) demand on average a sell price of 0.33 below the fundamental in the LM while speculators submit limit sell orders at prices 2.49 ticks above the fundamental in the same venue.

Scenario	Venue	Traders	Price of sell limit orders		
			Private Value $|\beta|$		
			0	4	8
mid-market	LM	Small	2.49	0.26	–
		Large	–	-0.01	-0.33
	DM	Small	0.62	-0.22	–
		Large	–	-0.39	-0.52
VWAP-market	LM	Small	2.40	0.19	–
		Large	–	-0.02	-0.27
	DM	Small	0.42	-0.15	–
		Large	–	-0.23	-0.29

Table 3.6: Table showing the average price of sell limit orders demanded by each trader type under both scenarios. The prices of small traders with a β of $|8|$ and the prices of large traders with β of 0 are not available due to the values assigned to the private value distribution $\mathscr{F}_{\beta}^{\mu}$ (for further details see Subsection 3.3.4). Statistics are collected from 9×10^{8} order submissions.

Agents irrespective of type conditional on the single price mechanism applied to each DM (LM - dependent) under the risk off being picked off in the LM are keen to move forward in terms of aggressiveness when they submit limit sell orders to the DM (see rows three, four, seven and eight in Table (3.6)). Most noticeably small traders with zero private values submit limit sell orders to the DM when the execution price is 0.62 ticks above the fundamental under the 'mid-market' scenario and 0.42 ticks above the fundamental under the 'VWAP-market' scenario. The above strategies - in terms of aggressiveness - adopted by all market participants though, given the unobservable to the market participants depth on either side of the book related to the DM; and the increased competition that these markets have at the single price at any instant (*i.e.*, see number of effectively traded limited orders at the bid price of the DMs in Table (3.1)) increase the time agents spend to the DM (see rows three and seven in Table (3.3)) and will lead them to obtain worse terms of trade when participating in that venue as compared to the LM.

Comparing now the DMs across the two scenarios. In contrast to the 'mid-market' scenario, agents with the nonzero private values (*i.e.*, small and large traders) adopt different strategies - in terms of aggressiveness - when they submit limit sell orders to the DM under the 'VWAP-market' scenario. For example while small (large) traders with a β of |4| under the 'mid-market' scenario they demand a sell price of 0.22 (0.39) ticks below the fundamental, under the 'VWAP-market' scenario the same type of agents submit limit orders to the DM when the price is 0.15 (0.23) ticks below the fundamental. In contrast to the 'mid-market' scenario, the improved quality of the LM under the 'VWAP-market' scenario; the different price mechanisms imposed to the DMs (VWAP is updated every time a trade occurs in the LM); and the lower quantities available at the single price of the book related to the DM under the 'VWAP-market' scenario will create a more favourable environment for the agents with nonzero private values in terms of execution time and benefits from trade when participating in that venue.

Benefits From Trade

Agents irrespective of type given the state of the trading game they observe, choose an action $\alpha = \left(m, x, p\right) \in A(s)$ so as to obtain the best possible terms of trade from the other active agents. Driven from their exogenous motives to trade and to avoid the risk of being picked off, agents with the most extreme private values (*i.e.*, large traders with a β of |8|) are keen to trade quickly at prices worse than the agents with the zero or with the intermediate private values. Table (3.7) depicts the average execution price of all sell orders (*i.e.*, the average distance between the execution price p and the fundamental value, denoted as v_t at time of execution t, for all executed market and limit orders). For a small trader for example with $\beta = 0$ the greater distance signifies an improvement in the terms of trade.

As shown in the table under the 'mid-market' scenario all market participants obtain better terms of trade in the LM in contrast to the DM. For example small (large) traders with a β of |4| execute sell orders in the LM 0.02 (0.10) ticks below the fundamental as compared to 0.40 (0.45) ticks below the fundamental when they participate in the DM. Under the 'VWAP-market' while small and large traders with

Scenario	Venue	Traders	Price of executed sell orders (benefits from trade)		
			Private Value $\lvert\beta\rvert$		
			0	4	8
mid-market	LM	Small	0.63	-0.02	–
		Large	–	-0.10	-0.42
	DM	Small	0.31	-0.40	–
		Large	–	-0.45	-0.66
VWAP-market	LM	Small	0.58	-0.02	–
		Large	–	-0.09	-0.39
	DM	Small	0.33	-0.24	–
		Large	–	-0.26	-0.35

Table 3.7: Table showing the average price of all sell market and limit orders executed by each trader type under both scenarios. The prices of executed sell orders for small traders with a β of $|8|$ and prices of executed sell orders for large traders with β of 0 are not available due to the values assigned to the private value distribution $\mathscr{F}_{\beta}^{\mu}$ (for further details see Subsection 3.3.4). Statistics are collected from 9×10^8 order submissions.

a $\beta \in \{0, \pm 4\}$ practice better terms of trade in the LM; large traders with the most extreme private values (*i.e.*, $\beta = |8|$) execute sell limit and market orders in the DM 0.35 ticks below the fundamental value as compared to 0.39 ticks below the fundamental when they participate in the LM. These traders' improvement in the terms of trade is incorporated in the relatively increased costs paid by the other agents - and primarily by the small traders with zero private valuation - that participate in the DM (compare rows seven and eight with rows five and six).

Apart from the small traders with zero private values that participate in the LM under the 'mid-market' scenario; across the two scenarios small and large agents with intermediate and the most extreme private values experience better terms of trade when they participate on either market m under the 'VWAP-market' scenario.[40] For example large traders that participate in the LM under the 'VWAP-market' scenario with a β of $|4|$ ($|8|$) execute 0.09 (0.39) ticks below the fundamental as compared to 0.10 (0.42) ticks below the fundamental when they participate in the LM under the 'mid-market' scenario. The reduced quoted spread, microstructure noise and tracking error of the LM under the 'VWAP-market' scenario (for further discussion about market quality see Subsection 3.4.1) seem to harm more the small agents

[40] The only agents that are not strictly better off in terms of benefits of trade are the small agents with a β of $|4|$ since in both scenarios they execute sell market and limit orders in the LM 0.02 ticks below the fundamental value.

with zero private value since we observe a decline in their terms of trade by 0.05 ticks.

3.4.3 Welfare Analysis

In this section using the analysis of gains from trade introduced by Hollifield et al. (2006) we examine the profitability of the traders under the 'mid-market' and the 'VWAP-market' scenarios.[41] We define a trader of type δ; profit for an order of z shares fully executing at time t' and market m as

$$GT_\delta = e^{-\rho(t'-t)} z(\beta + \mathrm{v}_{t'} - p)x. \tag{3.15}$$

where β is the trader's private valuation, $\mathrm{v}_{t'}$ is the fundamental value of the asset at the time of execution, p is the price of the order submitted and x is the submitted quantity.[42] Recall that the gains from an order of z shares executing are discounted back to the trader's first arrival to the market (in our example at time t) and at rate ρ. To be able to gain insight into the gains and losses for the trading process of the game we decompose the agent's profit as follows

$$GT_\delta = \beta + \beta(1 - e^{-\rho(t'-t)}) + e^{-\rho(t'-t)}\big(p - \mathrm{v}_t\big). \tag{3.16}$$

the intermediate term $\beta(1 - e^{-\rho(t'-t)})$ (waiting cost) captures the cost related to the trader's endogenous decision to delay the trade either due to lack of liquidity or unfavourable market conditions (*i.e.*, state observant decision). The last term $e^{-\rho(t'-t)}\big(p - \mathrm{v}_t\big)$ (money transfer) captures the trader's gains (losses) from the trading activity without including his private valuation. The 'money transfer' term is discounted back to the trader's first arrival to the market.

Table (3.8) displays the waiting costs and money transfer per trader type under the 'mid-market' and the 'VWAP-market' scenarios. As shown in the table, irrespective of the scenario, agents of any type trade faster in the LM in comparison to the DM. Most noticeably under both scenarios large traders with the most extreme private values (*i.e.*, $\beta = |8|$) trade much faster in the LM as can been seen by their lower

[41] Similar analysis is used by Bernales (2014) and Chiarella and Ladley (2016).

[42] In our model we do not pre-assume that a trader with a negative (positive) private valuation will sell (buy) the asset. All decisions are state-dependent and thus they depend on the prevailing market conditions of the LM at any instant.

Scenario	Venue	Traders	Waiting Cost			Money Transfer		
			Private Value $\|\beta\|$			Private Value $\|\beta\|$		
			0	4	8	0	4	8
mid-market	LM	Small	0.000	-0.301	–	0.313	-0.019	–
		Large	–	-0.353	-0.281	–	-0.097	-0.402
	DM	Small	0.000	-0.334	–	0.193	-0.371	–
		Large	–	-0.473	-0.573	–	-0.410	-0.621
VWAP-market	LM	Small	0.000	-0.267	–	0.274	-0.023	–
		Large	–	-0.327	-0.272	–	-0.090	-0.394
	DM	Small	0.000	-0.324	–	0.177	-0.231	–
		Large	–	-0.432	-0.554	–	-0.249	-0.346

Table 3.8: Table showing the average waiting costs and money transfer by each trader type under both scenarios. For full details on how these terms are defined see discussion after Eq. (3.16). The waiting costs and money transfer of small traders with a β of $|8|$ and of large traders with β of 0 are not available due to the values assigned to the private value distribution $\mathscr{F}^{\mu}_{\beta}$ (for further details see Subsection 3.3.4). Statistics are collected from 9×10^8 order submissions.

waiting costs (*e.g.*, under the 'mid-market' scenario from -0.573 to -0.281). Due to the number of shares large traders with intermediate private values hold, they trade slower in the LM (*i.e.*, higher waiting costs) as compared to the small traders with the same private values. For large traders with a β of $|8|$ this is less of a problem. These traders are keen to trade quickly by placing more aggressive orders to the LM since their benefits arise from their private values. Irrespective of the scenario, the different order placements strategies adopted by the large traders of any valuation when participating in the DM, completely shift the waiting costs for them (see Table (3.4)). For example under the 'VWAP-market' scenario while large traders with the most extreme values have on average a waiting cost of -0.554, large traders with a β of $|4|$ have a lower waiting cost equal to -0.432.

Comparing the two scenarios, in contrast to the 'mid-market' scenario all agents irrespective of type and venue choice trade faster under the 'VWAP-market' scenario. For example small trades with the intermediate private values (*i.e.*, $\beta = |4|$) under the 'mid-market' scenario, when participating in the LM have on average a waiting cost of -0.301 as compared to a waiting cost of -0.267 when the same type of agents participate in that venue under the 'VWAP-market' scenario. The reasons why all traders have lower waiting costs under the latter scenario are (i) the improved market quality of the LM; (ii) the different prices mechanishm applied to the DMs. Recall that VWAP is updated every time a trade occurs in the LM; and as a result of all the above (iii) the different strategies adopted by the market participants when they participate in either the DM or in the LM under the 'VWAP-market' scenario.

Irrespective of the scenario large traders with a β of $|8|$ incur higher money transfer than all the other trader types. For example as shown in Table (3.8), the money transfer value of large traders with the most extreme private values when participate in the LM under the 'mid-market' scenario is -0.402; in contrast the money transfer value of large traders with a β of $|4|$ is significantly smaller and equal to -0.097 ticks. The loss incurred to large traders with a $\beta = |8|$ from the trading activity is primarily captured by the small traders with zero private values. Recall that the latter agents under the risk of being picked off spend more time in the market and may only trade profitably when the fundamental value is mispriced. Apart from the large traders with a $\beta = |8|$ under the 'VWAP-market' scenario (recall from Table (3.7) that these traders experience an improvement in terms of trade when they participate in the DM); all other agents when participate in the DM, due to the properties of this venue (*i.e.*, no pre-trade transparency) face less favourable trading opportunities in terms of increased money transfer.

Under the 'VWAP-market' scenario: (i) large traders irrespective of private valuation; and (ii) small traders with intermediate private values (*i.e.*, $\beta = |4|$) when participating in the DM incur lower money transfer as compared to the 'mid-market' scenario. The improvement these type of agents experience - in the form of reduced money transfer - harms the rest of the market participants (*e.g.*, the small traders irrespective of valuation when participate in the LM and the small traders with zero private value when participate in the DM). As we will see next (analysis of Table (3.9)) the above observation coupled with the absence of waiting costs for the traders with zero private values; reduce these traders' profits under the 'VWAP-market' scenario.

Table (3.9) depicts the average profits per order by trader type under the two available scenarios. Irrespective of the scenario agents of any type make greater profits when participating in the LM in comparison to the DM. For example under the 'mid-market' scenario the average payoff of a small trader with $\beta = 0$, who participates in the LM is 0.313 ticks as compared to the 0.193 ticks when a trader with zero private value participates in the DM. The observable state of the LM, the low waiting costs coupled with the high money transfer of large traders with the most extreme private values benefit speculators (*i.e.*, small traders with zero private values). The latter traders make profits primarily from this money transfer. Conversely, the competitive environment of the LM and cognition limits that prevent agents from continuously monitoring the market; increase the picked off probabilities of these traders. The above coupled with the fact that speculators are particularly susceptible to this effect drive them to jump to the DM and execute orders at less favourable terms.

Scenario	Venue	Traders	Average Payoff			
			Private Value $\|\beta\|$			
			0	4	8	Total
mid-market	LM	Small	0.313	3.681	–	2.114
		Large	–	3.550	7.317	6.240
	DM	Small	0.193	3.296	–	0.969
		Large	–	3.117	6.806	4.958
VWAP-market	LM	Small	0.274	3.711	–	1.991
		Large	–	3.584	7.334	6.312
	DM	Small	0.177	3.445	–	1.441
		Large	–	3.320	7.100	4.963

Table 3.9: Table showing the average payoff by each trader type under both scenarios. For full details on how traders' payoffs are defined see text. The payoffs of small traders with a β of $|8|$ and of large traders with β of 0 are not available due to the values assigned to the private value distribution $\mathscr{F}_{\beta}^{\mu}$ (for further details see Subsection 3.3.4). Statistics are collected from 9×10^8 order submissions.

Across the two scenarios, the improvement of the LMs quality under the 'VWAP-market' scenario reduces the average payoff of speculators. Even though the trading cost of executing a market order in the LM is reduced for all market participants (see Table (3.2)); small traders with zero private values by primarily supplying liquidity in this venue are harmed the most. Moreover the DM's reduced microstructure noise - under the 'VWAP-market' scenario - that moves transactions closer to the fundamental value, decreases speculators profits as compared to the 'mid-market' scenario. Specifically, speculators' average payoff when they participate in the DM under the 'mid-market' scenario is 0.193 ticks as compared to the 0.177 when these traders participate in the DM under the 'VWAP-market' scenario.

Recall from Subsection 3.4.2 that under the 'VWAP-market' scenario large traders irrespective of valuation are strictly better off in terms of benefits from trade as compared to the 'mid-market' scenario. The above coupled with the improvement these traders experience in terms of reduced money transfer and waiting costs (see Table (3.8)) allows them to make greater profits when participating in any market m under the 'VWAP-market' scenario as compared to the 'mid-market' scenario.

3.5 Conclusion and Policy implications

Dark pools' increasing market share in global equity markets over the last decade coupled with the increased concerns of the exchange officials linked with 'dark' trading (*i.e.*, trading moving to off exchange platforms) are the driving forces of this study.[43] For the purposes of this study we build two scenarios where two markets - a 'lit' (LM) and a 'dark' (DM) - operate simultaneously. While under both scenarios we apply a single price mechanism to the dark pool, under the 'mid-market' scenario the 'dark' book matches orders at any instant at the mid-price (midpoint) of the 'open' limit order book. Under the 'VWAP-market' scenario the single price mechanism of the 'dark' book is the Volume-Weighted Average Price (VWAP). That is at any instant irrespective of the scenario dark pools match agents orders at a single price which is derived from the LM. We consider a dynamic continuous-time model of trade in a single financial asset were risk neutral agents submit orders in order to maximize their expected discounted payoff.

To be consistent with how real financial markets operate the trading game includes two groups of traders: large (institutional) traders who can trade at most three shares of the asset and small (retail) traders who trade one share of the asset. A factor that gives additional heterogeneity to the traders, is their intrinsic private value to trade the asset. Irrespective of the scenario - given the non-displayed properties of the DM - an agent's trading decisions depend on the prevailing market conditions of the LM. That is all actions in the trading game are state-dependent. Given that the model is analytically intractable (very large state space) we apply the numerical approach of Pakes and McGuire (2001) and Goettler, Parlour, and Rajan (2009) in order calculate the fixed point of the game.

Our findings have numerous policy implications as there is clear evidence that the presence of the VWAP DM compared to the midpoint DM improves the LM's market quality. Liquidity provision in the LM is increased; the quoted spread, the microstructure noise (*i.e.*, the absolute mean difference between the execution price and the fundamental value of the asset) and the tracking error (*i.e.*, the average absolute distance between the midprice and the fundamental value) decrease. The number of limit orders at the best quotes and in the limit order book of the LM as a

[43] For further discussion about dark pools and the concerns these markets raise to exchange officials see Section 3.1.

whole increase. As a result of the above, the effective spreads of all agent types are reduced. Moreover, as compared to the midpoint, the VWAP DM offers (i) a faster exchange to trade given that the average time interval between the trader's entry to the market and his order (market or limit) executing is reduced; and (ii) prices closer to the fundamental value of the asset given the DM reduced microstructure noise as compared to the midpoint DM.

In contrast to the 'mid-market' scenario, under the 'VWAP-market' scenario all agents with nonzero private values - irrespective of venue - experience an improvement in the form of reduced waiting costs. Moreover across the two scenarios, apart from the speculators (*i.e.*, small agents), all market participants obtain better terms of trade when participate in either venue under the 'VWAP-market' scenario. The above allow agents with nonzero private values irrespective of group to make greater profits when participate in any market under the 'VWAP-market' scenario as compared to the 'mid-market' scenario.

Chapter 4

A Nonparametric Analysis of Portfolio Choice

4.1 Introduction

In recent years there is an increasing number of papers that have developed and implemented - in a laboratory setting - *nonparametric* asset demand tests. These tests require the subject to choose a preferred option among financial assets subject to a standard budget constraint. Varian (1983, 1988) and Green and Srivastava (1986) derive necessary and sufficient conditions for a demand set to be consistent with the utility maximization model. Kübler, Selden, and Wei (2014) assume that state of nature probabilities are not fixed (*i.e.*, demand is continuous, differentiable function of prices, income and probabilities) and find restrictions on the contingent claim demand functions so to be consistent with state independent expected utility function. Choi et al. (2007) implement a portfolio choice experiment; given two commonly known states of the world, each participant was asked to purchase two Arrow-Debreu securities (*i.e.*, easy security offers one unit of account in one state and zero in the other) under different budget constraints. In their analysis, using Afriat's Theorem (Afriat, 1967) their objective is to check whether a subject's behaviour is consistent with the maximization hypothesis. By applying the experimental data collected by Choi et al. (2007); Polisson, Quah, and Renou (2017) develop a nonparametric approach (*i.e.*, lattice method) and test the consistency of the data with a wide range of models of choice under risk and under uncertainty (*i.e.*, expected utilty, rank dependent utility among others).[1]

In the majority of the aforementioned cases, the financial asset demand models make use of the classic Arrow-Debreu contingent claim setting; that is, under

[1] Other (functional form or nonparametric) asset demand tests have been derived from Dybvig (1983) and Echenique and Saito (2015).

a finite number of states, a subject possess preferences over state contingent consumption. Moreover markets are assumed to be complete (*i.e.*, same number of assets and states of the world). The focus of this study is to test this fundamental hypothesis used in the financial asset demand tests. To do that we design an experiment in which each subject is asked to purchase securities under four different scenarios in the asset space with contingent claims *fixed* across frames. Under the complete market case - for every budget line randomly selected in the contingent space - we create four different budget lines (*i.e.*, one for each scenario) in the asset space where subjects are asked to make choices. Subjects are facing *risk* where the probabilities of each state occurring are objectively known. There are two states of the world and each scenario is associated with two securities each of them promises different payoffs in the two available states (scenario I: Arrow-Debreu securities, scenarios II, III, IV: ordinary securities).

In our analysis, using Afriat's theorem (Afriat, 1967) we derive conditions that are necessary and sufficient for the dataset so as to be consistent with utility maximization within and across the available scenarios. Afriat's Theorem states that given any finite data set, if the Generalized Axiom of Revealed Preference (GARP) holds then there exists a utility function that rationalizes the data that is continuous, increasing and concave.[2] The above result proves to be very useful because in implementing our test we adopt the *nonparametric* revealed preference approach of Afriat (1967); Diewert (1973) and Varian (1982)); that is, the model's utility function that represents subjects' preferences does not reply on any functional form assumptions.

Given that the power of the experiment depends on the number of observations each subject is facing; it is natural to expect that hardly any behaviour at the level of individual subject, will be *exactly* consistent with the maximization hypothesis.[3] Given the aforementioned, we allow in our analysis for departures from rationality using Afriat's *critical cost efficiency index* or CCEI (see Afriat, 1972; Afriat, 1973, and; Varian, 1990). This index, denoted by e, takes values from 0 to 1, with $e = 1$ if a subject passes the test exactly. Moreover we determine the power of the test/model using the approach proposed by Bronars (1987). By introducing artificial subjects

[2] For further analysis see Section 4.2.

[3] In our analysis, we find that only the choice behaviour of one subject is consistent with the model of choice. (see Section 4.5).

that make uniformly random choices in the asset space, we estimate the probability of the dataset generated by these subjects being inconsistent - at a given efficiency index - with the model of choice. All of the above allows us to evaluate the difference between the pass rate and the model's precision (*i.e.*, the probability of a random dataset passing the test) at a given efficiency index by introducing the *Selten's index of predictive success* (Selten, 1991).

Our main findings include: (i) At a cost efficiency index 0.90 almost 90% of the subjects are consistent with the utility maximization model under scenario I (Arrow-Debreu securities). The more complex payoff structure of the ordinary securities in scenarios II, III and IV, reduces this percentage. As expected the less cost inefficient we allow subjects to be, the less subjects are consistent with the model of choice; (ii) the probability of a random dataset to be consistent with the utility maximization model across all scenarios at an efficiency index of 0.85 is effectively zero. The above leads to the conclusion that the power of the model is close to perfect. The CCEI scores - and as a result the power of the experiment depends on the number of observations each subject is facing; (iii) the test's predictive success across and within scenarios is within the positive range. That is, the test has explanatory power against the alternative hypothesis of random choice; (iv) given the CCEI scores at any given efficiency index within and across scenarios; some subjects are consistent within and across scenarios while others are consistent within but not across. (heterogeneous behaviour).

This study is organised as follows. Section 4.2 emphasises on the *nonparametric* approach while Section 4.3 on the equivalence between the financial asset setting and the contingent claim setting. Section 4.4 describes the experimental design and procedures. Section 4.5 depicts the results of our analysis, and Section 4.6 concludes.

4.2 Revealed Preference

For our purposes in order to be able to understand what the model of choice says about the data, we use as empirical strategy the revealed preference methodology (*e.g.*, Samuelson, 1938, 1948; Afriat, 1967, and; Varian, 1982). We follow the tradition of Afriat (1967) and we suppose that we have access to a finite dataset, such that:

Let $\Omega = \{(p^t, x^t)\}_{t=1}^{T}$, be a finite set of observations. At every observation $t = 1, 2, \ldots, T$, we interpret $x^t = \left(x_1^t, x_2^t, \ldots, x_\ell^t\right) \in \mathbb{R}_+^\ell$ as the observed bundle (the demand bundle) of ℓ goods chosen by an individual consumer at prices $p^t = \left(p_1^t, p_2^t, \ldots, p_\ell^t\right) \in \mathbb{R}_{++}^\ell$. We define the individual's budget set, given income $w > 0$, as $\mathscr{B}^t = \left\{x \in \mathbb{R}_+^\ell : p \cdot x \leq w\right\}$. Let $U : \mathbb{R}_+^\ell \to \mathbb{R}$ denote a local nonsatiated utility function. We say that the function U *rationalizes* the finite set $\Omega = \{(p^t, x^t)\}_{t=1}^{T}$, if $\forall t$ $(t \in \{1, 2, \ldots, T\})$ it holds $U(x^t) \geq U(x)$, $\forall x \in \mathscr{B}^t = \left\{x \in \mathbb{R}_+^\ell : p^t \cdot x \leq p^t \cdot x^t\right\}$. That is, the bundle x^t is the one that maximizes the consumer's utility function U (*i.e.*, $w = p^t \cdot x^t$), compared to any other bundles that belong in $\mathscr{B}^t$.

Given the above setting, we can ask what are the necessary and sufficient conditions for the data set $\Omega = \{(p^t, x^t)\}_{t=1}^{T}$ to be consistent with an individual consumer that maximizes the utility function U. The existence of the local nonsatiated utility function U that *rationalzies* the data depends upon what information can be *revealed* from Ω. Given a set $\mathscr{O}$, that contains the observed bundles, such that $\mathscr{O} = \{x^t\}_{t=1}^{T}$ and the price vector $p^t = \left(p_1^t, p_2^t, \ldots, p_\ell^t\right) \in \mathbb{R}_{++}^\ell$, at every $t = 1, 2, \ldots, T$ we define the following:

Definition 4.1. *For the observed bundles x^t, $x^s \in \mathscr{O}$ we say that x^t is directly revealed preferred to the bundle x^s (denoted as $x^t \succeq^* x^s$) if $p^t \cdot x^s \leqslant p^t \cdot x^t$.*

In words, x^t is *directly revealed preferred* to the bundle x^s if the bundle x^s was affordable when the bundle x^t was chosen. Replacing $p^t \cdot x^s \leqslant p^t \cdot x^t$ with strict inequality we derive $p^t \cdot x^s < p^t \cdot x^t$. That is, the bundle x^s was not only affordable but also was cheaper when x^t was chosen. More formally:

Definition 4.2. *For the observed bundles x^t, $x^s \in \mathscr{O}$ we say that x^t is directly revealed strictly preferred to the bundle x^s (denoted as $x^t \succ^* x^s$) if $p^t \cdot x^s < p^t \cdot x^t$.*

Definition 4.3. *For the observed bundles x^t, $x^s \in \mathscr{O}$ we say that x^t is revealed preferred to the bundle x^s (denoted as $x^t \succeq x^s$) if there is some sequence $i, j, \ldots, m$ such that $x^t \succeq^* x^i, x^i \succeq^* x^j, \ldots, x^l \succeq^* x^m$, and $x^m \succeq^* x^s$.*

Definition 4.4. *For the observed bundles x^t, $x^s \in \mathscr{O}$ we say that x^t is revealed strictly preferred to the bundle x^s (denoted as $x^t \succ x^s$) if there exist x^i and x^j such that $x^t \succeq^* x^i, x^i \succ^* x^j, x^j \succeq^* x^s$.*

The relation $\succ$ is the *transitive closure* of the relation $\succeq^*$. It implies, based on the above definitions, that the bundle x^t can be revealed preferred to the bundle x^s either directly or through a chain. When a given data set is *rationalizable* by a local nonsatiated utility function, its revealed preference must satisfy the Generalized Axiom of Revealed Preference (GARP). GARP is a no-cycling condition that states if a bundle x^t is *revealed preferred* to a bundle x^s (*i.e.*, $x^t \succeq x^s$) it cannot be the case that x^s is *directly revealed strictly preferred* to the bundle x^t (*i.e.*, $x^s \not\succ^* x^t$). More formally:

Definition 4.5. *A data set $\Omega = \{(p^t, x^t)\}_{t=1}^T$ obeys the Generalized Axiom of Revealed Preference (GARP) if given a chain such that $x^t \succeq^* x^i, x^i \succeq^* x^j, \ldots, x^l \succeq^* x^m, x^m \succeq^* x^s \Longrightarrow x^s \not\succ^* x^t$.*

Hence, it is intuitive that any finite data set will be *rationalizable* by a local nonsatiated utility function if the revealed preference relations satisfy the GARP condition. We now state the problem as follows

Definition 4.6. *A data set $\Omega = \{(p^t, x^t)\}_{t=1}^T$ is rationalizable if there exists a local nonsatiated utility function $U : \mathbb{R}_+^\ell \to \mathbb{R}$, such that at every observation $t = 1, 2, \ldots, T$ it holds*

$$U(x^t) \geq U(x) \ \forall x \in \mathscr{B}^t = \left\{x \in \mathbb{R}_+^\ell : p^t \cdot x \leq p^t \cdot x^t\right\}$$

In words, the data set is rationalizable if there exists a local nonsatiated utility function such that the bundle chosen at every observation t, provides a utility weakly greater than any other affordable bundle. The Theorem of Afriat (1967) states that if a data set obeys GARP then there exists a continuous, strictly increasing and concave utility function that *rationalizes* the given data set.[4]

Theorem 4.1. *Let the data set $\Omega = \{(p^t, x^t)\}_{t=1}^T$. The following statements about Ω are equivalent:*

1. *The data set Ω is rationalizable by a nonsatiated utility function $U : \mathbb{R}_+^\ell \to \mathbb{R}$.*
2. *The data set Ω obeys the Generalized Axiom of Revealed preference (GARP).*

[4] The term GARP is equivalent to Afriat's *cyclical consistency* condition (*e.g.*, Varian, 1982, see Fact 1).

3. *Associated to each observation* $t = 1, 2, \ldots, T$ *there exist numbers* $u^t \in \mathbb{R}$ *and* $\lambda^t \in \mathbb{R}_{++}$, *such that*

$$u^{t'} \leqslant u^t + \lambda^t p^t \cdot (x^{t'} - x^t),\ \forall t, t' = 1, 2, \ldots, T$$

4. *There exists a continuous, strictly increasing and concave utility function* $U : \mathbb{R}^{\ell}_{+} \to \mathbb{R}$ *that rationalizes the data set* Ω.

Proof. See Afriat (1967), Diewert (1973), Varian (1982) and Fostel, Scarf, and Todd (2004). □

In words, Afriat's Theorem states that given any finite data set, if GARP holds then there exists a utility function that rationalizes the data that is continuous, increasing and concave. Moreover, it provides two alternative tests for utility maximization. Both of these empirical tests include checking if GARP holds either directly or by solving a system of linear inequalities (*i.e.*, Afriat inequalities).[5]

The purpose of this paper is to study decision making under risk in the financial asset setting with contingent claims fixed across frames. Under certain assumptions, we are able to transform our observations into the contingent consumption setting and focus on the application of revealed preference in order to test the model of choice under risk and under uncertainty; given $s = 1, 2, \ldots, S$ states of the world, we let the vectors $\boldsymbol{x}^t = (x^t_1, x^t_2, \ldots, x^t_S) \in \mathbb{R}^S_+$ and $\boldsymbol{p}^t = \left(p_1, p_2, \ldots, p_S\right) \in \mathbb{R}^S_{++}$ to be the individual's contingent consumption and state prices, at observation t.[6] Given the above, we can ask what conditions are necessary and sufficient for a given data set to be rationalizable by a local nonsatiated utility function. To be more specific, the utility function $U : \mathbb{R}^S_+ \to \mathbb{R}$ rationalizes the data set $\Omega = \{(\boldsymbol{p}^t, \boldsymbol{x}^t)\}^T_{t=1}$ if for every observation $t = 1, 2, \ldots, T$ it holds that

$$U(\boldsymbol{x}^t) \geq U(\boldsymbol{x})\ \forall \boldsymbol{x} \in \mathscr{B}^t = \left\{\boldsymbol{x} \in \mathbb{R}^s_+ : \boldsymbol{p}^t \cdot \boldsymbol{x} \leq \boldsymbol{p}^t \cdot \boldsymbol{x}^t\right\} \tag{4.1}$$

That is, $\boldsymbol{x}^t$ maximizes the the individual's utility function U compared to any other bundle of contingent consumption (*i.e.*, $\boldsymbol{x}$) that belong in $\mathscr{B}^t$.

[5] The latter test require for each observation $t = 1, 2, \ldots, T$ to find numbers (u^t, λ^t) that solve the inequality introduced in statement 3 of Theorem 4.1.

[6] See for example Varian (1982, 1983, 1988), Green and Srivastava (1986), Diewert (2012), Kübler, Selden, and Wei (2014),Echenique and Saito (2015) and Polisson, Quah, and Renou (2017).

4.3 From the Financial Asset Setting to the Contingent Claim Setting

Lets consider an economy with $k \in \{1,2,\ldots,K\}$ financial assets and and $s \in \{1,2,\ldots,S\}$ states of the world. An Arrow-Debreu (AD) security $k \in K$ (see Arrow, 1964, and; Debreu, 1987) is the one that promises one unit of consumption in state $s \in S$ and nothing in any other state. Denote the payoff an AD security k offers in state s by $\tilde{\xi}_{sk}$. It holds that $\tilde{\xi}_{sk} = 1$ for $s = k$ and $\tilde{\xi}_{sk} = 0$ for $s \neq k$, *i.e.*, the index of an AD security is identical to the state in which this security pays off (state claim). Hence, given S states, denote the payoff matrix of the K AD securities by $\tilde{\boldsymbol{\Xi}} \in \mathbb{R}^{S\times K}$, such that

$$\tilde{\boldsymbol{\Xi}}_{S\times K} = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix} \tag{4.2}$$

Throughout this paper we focus on a complete system of financial markets with S states and $K = S$ financial assets. Under this assumption, a complete system of AD securities offers the same trading opportunities to individuals as a contingent claims system (Arrow, 1964).

An ordinary (complex) security, is a type of asset that provides a more complex structure than an AD security. An ordinary security is one that potentially pays off in more than just a single state. In a similar fashion as before, we denote as ξ_{sk} the payoff an ordinary security k offers in state s. Thus, given K and S ordinary securities and states respectively, the payoff matrix $\boldsymbol{\Xi} \in \mathbb{R}^{S\times K}$ takes the following form

$$\boldsymbol{\Xi}_{S\times K} = \begin{pmatrix} \xi_{11} & \xi_{12} & \cdots & \xi_{1K} \\ \xi_{21} & \xi_{22} & \cdots & \xi_{2K} \\ \vdots & \vdots & \ddots & \vdots \\ \xi_{S1} & \xi_{S2} & \cdots & \xi_{SK} \end{pmatrix} \tag{4.3}$$

Denote by z_k and $\boldsymbol{z} = (z_1, z_2, \ldots, z_K) \in \mathbb{R}^K_+$ the quantity of the k asset held by an individual and the portfolio of ordinary securities respectively. Denote the price of asset k by p_k, with $\boldsymbol{p} = (p_1, p_2, \dots, p_K) \in \mathbb{R}^K_{++}$ denoting the corresponding price vector. Given the matrix Ξ, the term $\sum_{k=1}^{K} \xi_{sk} z_k$ determines the payoff that portfolio $\boldsymbol{z}$ offers in state $s \in \{1, 2, \dots, S\}$. Denote the individual's contingent claims in state s by x_s such that $\boldsymbol{x} \in \mathbb{R}^S_+$, $\forall s \in S$. Assume that the individual's preferences over state contingent consumption are represented by the utility function $U(\boldsymbol{x})$. Assume further that the probability of state s is commonly known to be π_s, with $\pi_s \in \Delta(S)$.[7] Holding income $w > 0$, the decision maker's asset optimization problem is given by

$$\begin{aligned}
&\max_{\boldsymbol{x} \in \mathbb{R}^S_+, \boldsymbol{z} \in \mathbb{R}^K_+} \mathscr{U}(\boldsymbol{x}; \boldsymbol{\pi}) = \sum_{s=1}^{S} \pi_s u(x_s) \\
&\qquad \text{subject to} \\
&\qquad\qquad \boldsymbol{p} \cdot \boldsymbol{z} = w \text{ and } \boldsymbol{x} = \Xi \boldsymbol{z} \\
&\qquad\qquad \Rightarrow \sum_{k=1}^{K} p_k z_k = w \text{ and } x_s = \sum_{k=1}^{K} \xi_{sk} z_k, \text{ for all } s = 1, 2, \dots S
\end{aligned} \tag{4.4}$$

Given a system of AD and a system of ordinary securities, a portfolio's payoff in any given state $s = 1, 2, \dots, S$, is determined by the term $\sum_{k=1}^{K} \tilde{\xi}_{sk} z_k$ and $\sum_{k=1}^{K} \xi_{sk} z_k$ respectively. With AD securities, this payoff depends only on the quantity $(z_{k=s})$ of the respective AD security that pays off in state s (see Matrix (4.2)). With ordinary securities, this payoff depends on the combination of all ordinary securities held by an agent. This comes from the fact that all securities held and which form the portfolio may promise payoffs in more than a single state (see Matrix (4.3)). Hence, in both sets the return on the portfolio decides the amount of purchasing power transferred to the agent.

A necessary and sufficient condition that ensures an equivalence between the two aforementioned sets arises from the complete market assumption. Our assumption of completeness of markets, guarantees that the matrix of payoffs Ξ has full rank, that is rank$(\Xi) = S$. The full rank condition tell us that across S states, the payoff matrix Ξ of the K ordinary securities has S linearly independent columns. Denote the contingent claim vector of AD prices by $\tilde{\boldsymbol{p}} = (p_1, p_2, \dots, p_s)$. Then, by the no-arbitrage theorem on asset prices, there exists a vector $\tilde{\boldsymbol{p}} \gg 0$ that correctly prices

[7] Here $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots, \pi_S) \in \Delta(S)$, where $\Delta(S) = \left\{ \boldsymbol{\pi} \in \mathbb{R}^S : \boldsymbol{\pi} \gg 0 \text{ and } \sum_{s=1}^{S} \pi_s = 1 \right\}$.

all available K ordinary securities. Given the above, it holds that

$$
\begin{aligned}
&\tilde{\boldsymbol{p}}\cdot\boldsymbol{x}=\boldsymbol{p}\cdot\boldsymbol{z}\\
&\Rightarrow\sum_{s=1}^{S}\tilde{p}_s\left(\sum_{k=1}^{K}\xi_{sk}z_k\right)=\sum_{k=1}^{K}p_k z_k\\
&\Rightarrow\tilde{p}_1(\xi_{11}z_1+\cdots+\xi_{1K}z_K)+\cdots+\tilde{p}_S(\xi_{S1}z_1+\cdots+\xi_{SK}z_K)=\\
&\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad=p_1z_1+\cdots+p_Kz_K\\
&\Rightarrow\begin{cases}\sum_{s=1}^{S}\tilde{p}_s\xi_{s1}=p_1\\ \sum_{s=1}^{S}\tilde{p}_s\xi_{s2}=p_2\\ \qquad\vdots\\ \sum_{s=1}^{S}\tilde{p}_s\xi_{sK}=p_K\end{cases}\\
&\Rightarrow\begin{pmatrix}\xi_{11}&\xi_{21}&\cdots&\xi_{S1}\\ \xi_{12}&\xi_{22}&\cdots&\xi_{S2}\\ \vdots&\vdots&\ddots&\vdots\\ \xi_{1K}&\xi_{2K}&\cdots&\xi_{SK}\end{pmatrix}\begin{pmatrix}\tilde{p}_1\\ \tilde{p}_2\\ \vdots\\ \tilde{p}_K\end{pmatrix}=\begin{pmatrix}p_1\\ p_2\\ \vdots\\ p_K\end{pmatrix}
\end{aligned}
\tag{4.5}
$$

$$
\overset{(4.3)}{\Rightarrow}\qquad \boldsymbol{p}=\boldsymbol{\Xi}^{\mathsf{T}}\tilde{\boldsymbol{p}}\Leftrightarrow\boldsymbol{p}^{\mathsf{T}}=(\boldsymbol{\Xi}^{\mathsf{T}}\tilde{\boldsymbol{p}})^{\mathsf{T}}\Leftrightarrow\boldsymbol{p}^{\mathsf{T}}=\tilde{\boldsymbol{p}}^{\mathsf{T}}\boldsymbol{\Xi}
$$

or equivalently

$$
p_k=\sum_{s=1}^{S}\tilde{p}_s\xi_{sk}\quad\text{for all}\quad k=1,2,\ldots,K
$$

Thus now, instead of the asset optimization problem (see Eq.(4.4)) we focus only on demand tests for the contingent claim system. By setting $\boldsymbol{x}=\boldsymbol{\Xi}\boldsymbol{z}$ and $\boldsymbol{p}^{\mathsf{T}}=\tilde{\boldsymbol{p}}^{\mathsf{T}}\boldsymbol{\Xi}$ we turn our attention to the contingent consumption demand $\boldsymbol{x}\in\mathbb{R}^S_+$ that solves[8]

[8] Proposition A.1 in Appendix A shows the equivalence between the two sets.

$$\begin{aligned}
\max_{\boldsymbol{x}\in\mathbb{R}^S_+} \mathscr{U}(\boldsymbol{x};\boldsymbol{\pi}) &= \sum_{s=1}^{S} \pi_s u(x_s) \\
&= \sum_{s=1}^{S} \pi_s u\left(\sum_{k=1}^{K} \xi_{sk} z_k\right) \\
&= \pi_1 u(\xi_{11} z_1 + \ldots \xi_{1K} z_K) + \cdots + \pi_S u(\xi_{S1} z_1 + \ldots \xi_{SK} z_K)
\end{aligned} \tag{4.6}$$

subject to

$$\begin{aligned}
&\tilde{\boldsymbol{p}} \cdot \boldsymbol{x} = w \\
\Rightarrow &\sum_{s=1}^{S} \tilde{p}_s \left(\sum_{k=1}^{K} \xi_{sk} z_k\right) = w \\
\Rightarrow &\tilde{p}_1 (\xi_{11} z_1 + \ldots \xi_{1K} z_K) + \cdots + \tilde{p}_S (\xi_{S1} z_1 + \cdots + \xi_{SK} z_K) = w
\end{aligned}$$

4.4 Experimental Design and Procedures

4.4.1 Design

In this experiment each participant is asked to purchase securities - under four different scenarios in the asset space - under risk. There are two states of the world, and each scenario is associated with two securities each of them offer different payoffs in the two available states. For each of the four scenarios, there are 20 decision tasks; that is each subject is asked to complete a series of 80 decision tasks. Each task requires the subject to allocate tokens between two securities in the asset space. The complete market assumption and the fact that asset prices preclude arbitrage ensure equivalence between a contingent claim and a financial asset setting; allowing us to study decision making under risk in the latter with contingent claims fixed across frames.[9] Specifically, for every budget line randomly selected in the contingent space (CS), we are able to create for each of the four scenarios, an equivalent environment in the asset space (AS) where subjects are asked to make choices under risk. To be precise, in scenario I each subject is asked to divide the budget between two AD securities. In scenarios II, III or IV subjects are asked to divide the budget between two ordinary securities (OS). Table 4.1 illustrates the structure of the payoff

[9] See Section 4.3 for further information.

matrices for each scenario available. The order the scenarios are presented to every participant is random.

scenario I AD	scenario II OS	scenario III OS	scenario IV OS
$\tilde{\Xi}=\begin{pmatrix}1 & 0\\0 & 1\end{pmatrix}$	$\Xi_1=\begin{pmatrix}4 & 0\\0 & 1\end{pmatrix}$	$\Xi_2=\begin{pmatrix}4 & 1\\0 & 1\end{pmatrix}$	$\Xi_3=\begin{pmatrix}4 & 1\\0 & 2\end{pmatrix}$

Table 4.1: Payoff matrices associated with each available scenario. Scenario I considers Arrow-Debreu securities (AD) while scenarios II-IV consider ordinary securities (OS); *e.g.*, in scenario IV security A pays off 4 tokens in state 1 and 0 in state 2 and security B pays off 1 token in state 1 and 2 tokens in state 2.

To provide intuition, we present a simple example of how the budget lines in the AS for all the scenarios, are generated. Suppose the computer has drawn for one participant in the CS, the x-intercept to be 50.00 tokens and the y-intercept to be 20.00 tokens. Denote the x- and y-intercepts by X^{CS} and Y^{CS} respectively such that $X^{CS}=50.00$ and $Y^{CS}=20.00$. Being consistent with the notation used in Section 4.3, and by setting the individual's wealth to 100 (*i.e.*, $w = 100$), we infer the relative prices of the two Arrow-Debreu securities such that $\tilde{p}_1=2$ and $\tilde{p}_2=5$. Using for each security and scenario all the payoff relevant information (see Table 4.1) and the fact that $\boldsymbol{p}^{\mathsf{T}}=\tilde{\boldsymbol{p}}^{\mathsf{T}}\Xi$ (see concluding equation in (4.5), where $\boldsymbol{p}=(p_1,p_2)\gg 0$) we calculate the corresponding prices and intercepts (*i.e.*, X^{AS} and Y^{AS}) in the AS. Figure 4.1, using the example above, represents graphically the budget line generated in the CS (Figure 4.1a), and the budget lines in the AS under the 4 available scenarios (Figure 4.1b).

To ensure that the budget lines cross frequently in the CS, we use as our starting point, the same procedure of budget line selection as in Choi et al. (2007).[10] That is, given two states of the world and two associated AD securities we program z-Tree (Zurich Toolbox for Ready-made Economic Experiments) to generate for each subject 20 independent budget lines in the CS, by picking randomly the x- and y-intercepts.[11] Both axes are scaled from 0 to 100 tokens. Additionally, both intercepts are drawn randomly uniformly between 10 and 100 tokens, with a restriction

[10] Given sufficient relative price variations, we can test consistency of subjects' behavior under the model of choice.

[11] See Fischbacher (2007) for a description of the z-Tree software.

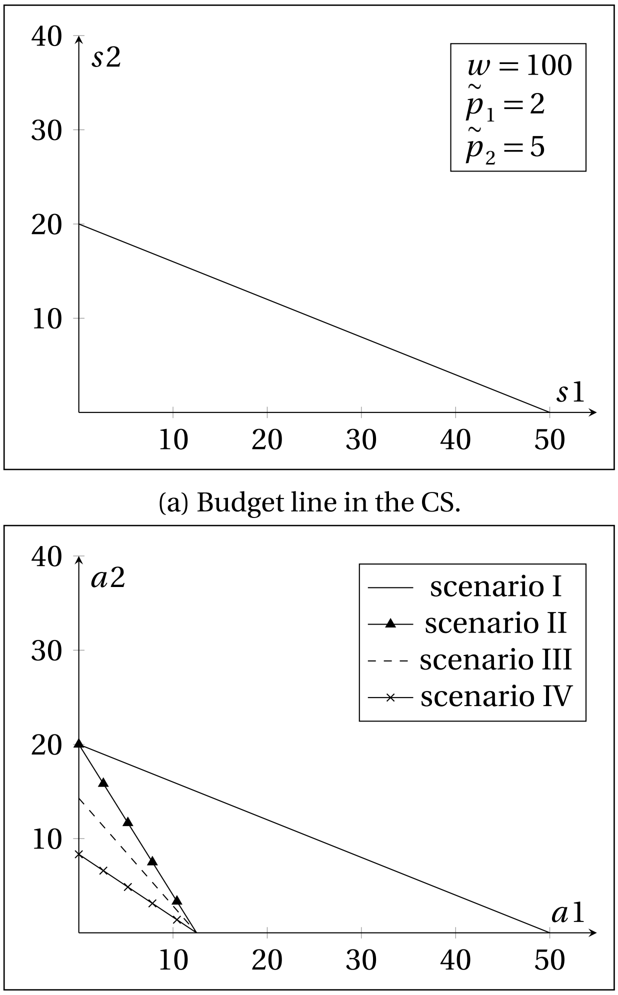

(a) Budget line in the CS.

(b) Budget lines in the AS.

Figure 4.1: Budget line drawn in the contingent space with two AD securities and two states, with X^{CS} = 50.00 and Y^{CS} = 20.00 (Fig. 4.1a). Budget lines generated in the asset space that offer an equivalent environment to the participants as in the contingent space, under the four available scenarios (Fig. 4.1b).

imposed that at least one of the intercepts is at or above the 50 token level. Denote the demand for the AD security that promises a payoff of one unit in state s by x_s. By setting the individual's wealth equal to 100 (*i.e.*, $w = 100$), given a random draw in the CS, we then can infer the relative prices ($\tilde{\boldsymbol{p}} = (\tilde{p}_1, \tilde{p}_2) \gg 0$) of the AD securities.[12]

In this experiment, the graphical interface used follows the lines of Loomes (1991) and Loomes and Pogrebna (2014). That is, subjects see a graphical representation

[12] Here, $\tilde{\boldsymbol{p}}_{s=1(2)}$ is the price of the AD security that pays off in state 1(2); and the budget constraint is $\tilde{p}_1 x_1 + \tilde{p}_2 x_2 = w$.

of two centred lines where they are asked to choose any affordable portfolio of two securities under risk. Given a decision task and any available scenario, the two lines represent the x- and y- intercepts in the AS; that is, the two lines map the closed intervals $[0, X^{AS}]$ and $[0, Y^{AS}]$ respectively. Figure 4.2, represents graphically the two lines used to map the budget line from the AS under scenario IV, for the example described above.

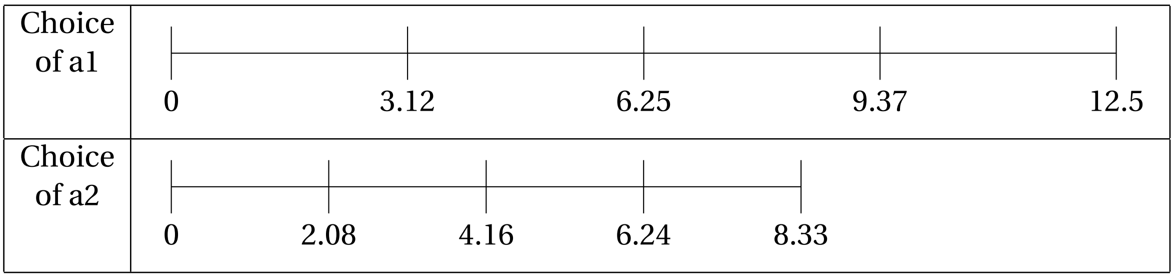

Figure 4.2: Display used to represent the budget line generated in the asset space under scenario IV (see Figure 4.1b). Both lines represent the minimum and maximum number of tokens a subject can allocate between the two securities in the given space.

4.4.2 Procedures

This experiment was conducted in the Leicester Experimental Economics (LExEcon) Laboratory in the Department of Economics at the University of Leicester. The 67 subjects that participated in the experiment, were recruited from undergraduate classes at University of Leicester. Each experimental session - apart from one - consisted of 10 subjects, each of them completed the experiment at an individual computer terminal.[13] Instructions were read aloud by an experimenter and two practise tasks were provided to the subjects. In order to move to the actual experiment, subjects had to answer correctly a question related to how the payoff in tokens, is calculated in a hypothetical case.[14] Each experimental session lasted approximately one hour and eight minutes. Payoffs were calculated in terms of tokens and then converted into money at an exchange rate of £0.5 per token. Subjects were paid an average of £18 (including a £2 participation fee), privately at the end of the session.

[13] Out of the seven sessions conducted, only the second session included 7 subjects.

[14] See a copy of the instructions and of the question (Figure B.1) in Appendix B.

In each of the 80 decision tasks, a subject had to choose how to allocate tokens between a RED asset and a BLUE asset in the AS. Feasible allocations for both assets were represented by two centred lines (see Figure 4.2); with the top (bottom) line corresponding to the RED (BLUE) asset. The granularity of these lines was set to 0.01 tokens. On each of these lines, a cursor (dot) was placed to represent the allocation to the corresponding asset (see Figure C.1 in Appendix C for an illustration). To choose an allocation, subjects were asked to click and drag either cursor to their desired allocation. The movement of one cursor automatically adjusted the position of the other to satisfy the budget constraint.[15] Subjects were able to adjust the position of both cursors within the feasible sets as many time as they wish, with no time restrictions. At the start of each decision task, the initial allocation of tokens, and the corresponding cursors, was random and independent across tasks.

In this experiment, subjects faced risk since each task included two possible events (event 1 and event 2). In each decision task, the tokens a subject allocated in each asset had different returns depending on the scenario and the two events. With probabilities assigned to each event being objectively known; at the end of each task, the computer randomly selected one of the two events with equal probability ($\pi = {}^1/_2$), independently across tasks. Applying the random lottery mechanism, subjects were not informed of the event that was drawn by the computer at the end of each task. At the end of the experiment, the computer randomly selected with equal probability one decision task for each subject. Conditional on the allocation of tokens, the scenario and the event randomly drawn by the computer, the subject was paid the amount he earned in that specific task.

4.5 Results

We first test for utility maximization on the given experimental data and we present the exact pass rates in Table 4.2. As can be seen from the table across the 80 decision tasks - each subject had to face - only 1 out of the 67 subjects passes GARP on every scenario (*i.e.*, 'all scenarios'). That is, only the choice behaviour of one subject is consistent with the utility maximization model. Moreover, subjects under scenarios I, II and II performed better than under scenario IV. More precisely, of the 67

[15] In any decision task, subjects could choose any combination of $(x_1, x_2) \geq 0$ in the asset space, that satisfied the constraint $p_1 x_1 + p_2 x_2 = 1$.

subjects 31 pass GARP under scenario I, 26 pass under scenarios II and III, 20 pass under scenario IV (see columns one, two, three and four of Table 4.2).

	scenario I	scenario II	scenario III	scenario IV	all scenarios
GARP	31/67	26/67	26/67	20/67	1/67

Table 4.2: Exact Pass Rates

Due to the fact that GARP offers an 'exact' test, *i.e.*, either the dataset satisfies GARP or it does not, we allow in our analysis for departures from rationality. To measure the degree of GARP violations we use Afriat's *critical cost efficiency index* or CCEI (see Afriat, 1972; Afriat, 1973, and; Varian, 1990).[16] To be more specific the departure involves adjusting the budget constraint $\mathscr{B}^t$ such that $\forall\, t$ ($t \in \{1,2,\ldots,T\}$) the budget is now defined to be

$$\mathscr{B}^t(e) = \left\{x \in \mathbb{R}^s_+ : p^t \cdot x \le e p^t \cdot x^t\right\} \tag{4.7}$$

for some $e \in [0,1]$.[17] That is, x^t is directly revealed preferred to x at efficient level e if $p^t \cdot x \le e p^t \cdot x^t$. Assume that a particular subject has chosen under the contingent consumption setting the bundle $x^1 = (17.5, 31.25)$ at prices $p^1 = (1.25, 2.5)$ and the bundle $x^2 = (35, 12.5)$ at prices $p^2 = (2.5, 1)$. The above situation is displayed in Figure 4.3. We can observe that the bundle x^1 is revealed preferred to the bundle x^2 (*i.e.*, $x^1 \succ x^2$) at prices p^1 since $p^1 \cdot x^2 < p^1 \cdot x^1$.[18] Moreover, x^2 is revealed preferred to x^1 (*i.e.*, $x^2 \succ x^1$) at prices p^2 since $p^2 \cdot x^1 < p^2 \cdot x^2$. That is the given data violates GARP and cannot be rationalizable by a nonsatiated utility function. Shifting (shrinking) the budget lines as displayed in Figure 4.3 by some factor (*i.e.*, $e = {}^{30}/_{40} = 0.75$) it holds that $x^1 \succeq^*_e x^2$ since $p^1 \cdot x^2 \le e p^1 \cdot x^1$ and $x^2 \not\succeq^*_e x^1$ since $p^2 \cdot x^1 \not< e p^2 \cdot x^2$. That is when the CCEI is equal to 0.75 the data satisfies GARP and will be rationalizable by a local nonsatiated utility function.

[16] Some seminar papers that use the CCEI, the VEI or other indices to measure a model's fit are Choi et al. (2007); Echenique, Lee, and Shum (2011); Choi et al. (2014); Dean and Martin (2016) and Polisson, Quah, and Renou (2017).

[17] The value of e determines by how much the budget should be shrunk so to satisfy GARP. When $e = 1$ the data set is fully rationalizable by a local nonsatiated utility while when $e = 0$ GARP is trivially satisfied.

[18] Recall that the transitive closure relation $\succ$ implies that a bundle can be directly revealed preferred to another bundle either directly or through a chain.

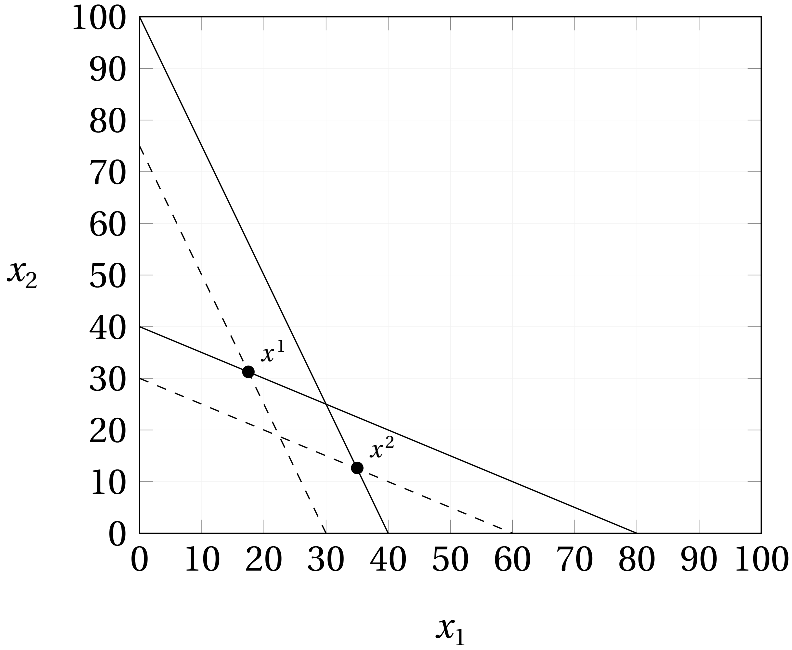

Figure 4.3: Calculating the CCEI

4.5.1 Distributions of Efficiency Indices and Bronars Power

Given the number of decision tasks each subject had to face and that the results presented in Table 4.2 are sharp, we next focus on the different efficiency levels at which subjects pass GARP. Figure 4.4 displays the empirical distribution (dashed black line) of the CCEI scores generated by the experimental data on every scenario. As can be seen from the figure at an efficiency index of 0.9 approximately 40% of the subjects are consistent with the utility maximization model. Moreover by allowing a 20% efficiency loss (*i.e.*, subject are now 10% more cost inefficient as compared to the 0.9 efficiency index) we can observe that this percentage increases to 65%.

To be able to justify the aforementioned (*i.e.*, the power of our experiment is determined by the number of decision tasks each subject had to face) and to understand whether the observed distribution of the CCEI scores imply relative success or failure of the model to explain the experimental data we follow the approach proposed by Bronars (1987). The Bronars (1987) approach introduces artificial subjects that make uniformly random choices from a collection of budget sets. Being consistent with the notation used in Section 4.3, we denote the quantities of assets by z_k and their prices by p_k with $\boldsymbol{z} \in \mathbb{R}_{+}^{K=2}$ and $\boldsymbol{p} \in \mathbb{R}_{++}^{K=2}$ denoting the corresponding vectors. Using the number of observations a subject had to face in the experiment

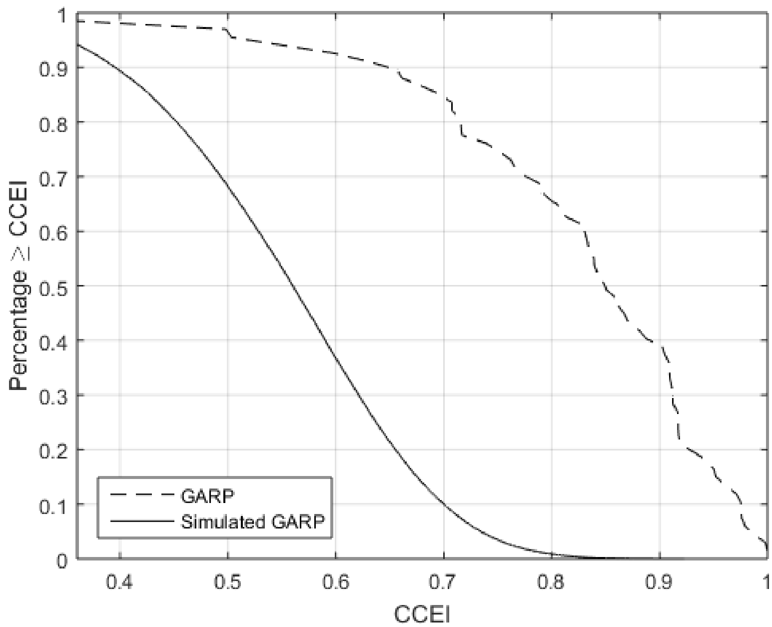

Figure 4.4: CCEI Distributions on every scenario ('all scenarios')

and the number of securities we first generate uniformly asset holdings. Given a set of prices - randomly generated in the experiment - under any available scenario in the AS; the asset holdings and the actual subjects' choices for each subsequent task; we then generate uniformly random portfolios of the $K = 2$ available securities. Finally by setting $\boldsymbol{x} = \boldsymbol{\Xi} \boldsymbol{z}$ and $\boldsymbol{p}^{\top} = \tilde{\boldsymbol{p}}^{\top} \boldsymbol{\Xi}$ with $\boldsymbol{x} \in \mathbb{R}_{+}^{S=2}$ and $\tilde{\boldsymbol{p}} \in \mathbb{R}_{++}^{S=2}$ we were able to transform these uniformly random choices so as to focus on the contingent consumption setting. That is from a collection of 20 budget lines randomly generated in the CS for each actual subject in the experiment; each artificial subject made 80 choices in the AS each of them transformed to a claim contingent. With the process described above we generate a random sample of 100,500 artificial subjects.[19]

Figures 4.4 and 4.5 compare the distributions of the CCEI scores generated by the 100,500 artificial subjects (solid black lines) and the CCEI scores generated by the actual 67 subjects (dashed black lines) under scenarios I-IV and on every scenario. Choosing the 0.95 efficiency level we observe that the percentage of the ran-

[19] We compared these results to a sample of 33500 random subjects. By applying five one-sided Kolmogorov-Smirnov tests to the CCEI distributions of the two random generated samples we observed at a 5% significant level the estimated area does not change. More precisely the test statistic for the K-S test under scenarios I-IV and all scenarios was 0.0066, 0.0064, 7.2637e-04, 0.0046 and 0.0026.

dom subjects that have CCEI scores at or above this threshold ranges from 4.4% under scenario II to 19.4% under scenario III (Figure 4.5). The percentage of the actual subjects with a CCEI score 0.95 or higher increases significantly to 64.1% under scenario II (smallest) and to 74.6% under scenario I (highest). Moreover, while more than 83% of the actual subjects on average - under scenarios I-IV - have a CCEI score of 0.90 or higher this percentage falls to approximately 25% on average for the random subjects (Figures 4.5a - 4.5d).

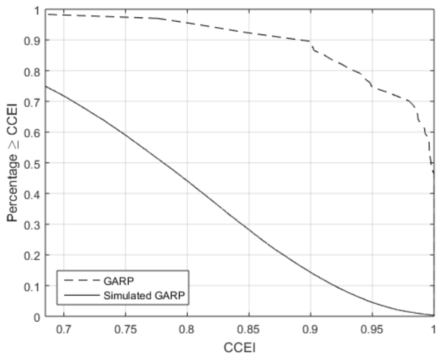

(a) scenario I

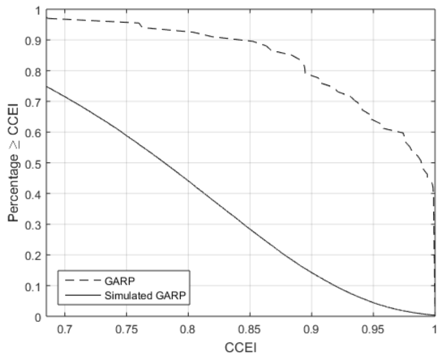

(b) scenario II

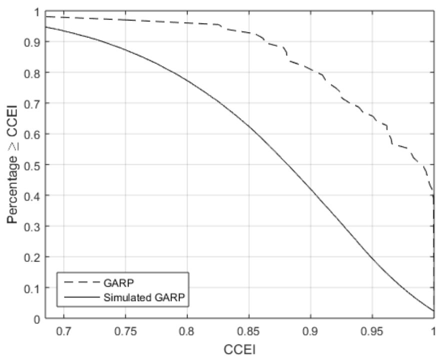

(c) scenario I

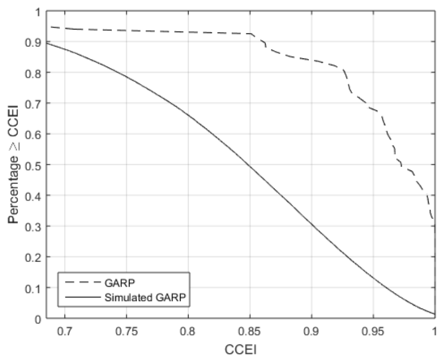

(d) scenario I

Figure 4.5: CCEI Distributions by scenario

Comparing now the CCEI distributions depicted in Figures 4.4 and 4.5 of both the actual and the artificial subjects we can observe for example that at any $e = \{0.90, 0.95, 1.00\}$ the actual subjects' CCEI scores are smaller under the 'all scenarios' case as compared to any of the scenarios I, II, III or IV. The above strongly suggests that the CCEI scores - and as a consequence the power of the experiment - depend on the number of observations each subject is facing.[20]. Conditional on the

[20] Recall from Section that the experimental design is balanced so each subject has to complete

aforementioned efficiency levels the likelihood of an artificial subject passing GARP on every scenario is very close to zero (see Figure 4.4). This difference between the CCEI distributions of the actual and the artificial subjects signifies that the test's consistency is not a result of a random behaviour and subjects clearly understood the procedures.

The set of feasible allocations a subject can choose under a particular scenario in the AS is given by a budget line (see for example Figure 4.2 in Subsection 4.4.1); that is the subject is strained to exhaust a fixed amount of purchasing power (*i.e.*, is choosing a combination $(z_1, z_2) \geq 0$ such that $p_1 z_1 + p_2 z_2 = m$). Suppose at some $t \in \{1, 2, \ldots T\}$ the subject is observed to have chosen under scenario i a portfolio $\boldsymbol{z}^t \in \mathbb{R}^K_+$ such that $\boldsymbol{z}^t \in \vartheta \mathscr{B}^t_i$, where $\vartheta \mathscr{B}^t_i$ is the upper boundary of $\mathscr{B}^t_i = \{\boldsymbol{x} \in \mathbb{R}^S_+ \mid \exists \boldsymbol{z} \in \mathbb{R}^K_+ : \boldsymbol{p} \cdot \boldsymbol{z} \leq \boldsymbol{p} \cdot \boldsymbol{z}^t \text{ and } \boldsymbol{x} = \boldsymbol{\Xi} \boldsymbol{z}^t\}$, where $i \in \{1, 2, 3, 4\}$. Denote an element that belongs to the set $\mathscr{B}^t_i$ under scenario i by $\boldsymbol{x}$. It holds that $\boldsymbol{x} \in \vartheta \mathscr{B}^t_i$ if $\nexists\, \boldsymbol{y} \in \mathscr{B}^t_i$ such that $\boldsymbol{y} > \boldsymbol{x}$.[21] Given a particular randomly selected budget line in the CS, the payoff matrices associated to each of the four scenarios generate budget sets in the AS, such that $\mathscr{B}_1 \subset \mathscr{B}_2 \subset \mathscr{B}_3 \subset \mathscr{B}_4$. Thus in the simulated experiment, artificial subjects should attain at any efficiency level higher CCEI scores under scenarios III and IV as compared to the scenarios I and II. Figure 4.5 supports the above.

The results suggest that subjects are consistent within scenarios but not across. To justify the above we conduct the Wilcoxon sign ranked test (Wilcoxon, 1945). That is we test the null hypothesis that the median difference between the subjects' CCEI scores within (within and across) scenarios is equal to zero. The p-values are depicted in Table 4.3.

Scenarios	II	III	IV
I	0.052	0.1176	**0.002**
	I	III	IV
II	–	0.932	0.960
	I	II	IV
III	–	–	0.503

Scenarios	all
I	**0.000**
	all
II	**0.000**
	all
III	**0.000**
	all
IV	**0.000**

Table 4.3: Wilcoxon sign ranked test within (within and across) scenarios (p-values)

20 decision tasks under scenario I, II, III and IV

[21] Given any two vectors $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^K_+$: $x > y$ if $x \neq y$ and $x_i \geq y_i\ \forall i$.

As shown in the table, at the 5% significance level we reject the null hypothesis that the median difference between the subjects' CCEI scores within and across scenarios is equal to zero. Moreover looking within the scenarios we can observe that the null hypothesis is rejected at the 5% significance level for scenarios I and IV; given the payoff matrices associated with each scenario, under scenario IV subjects had to choose a portfolio that included ordinary securities with more complex payoffs' structure than the AD ones (*i.e.*, scenario I).

4.5.2 Predictive success

To be able to test the predictive success of the model for each and across all scenarios we aggregate the pass rates and the power of the test into a single axiomatic measure proposed by Selten (1991). Selten's *index of predictive success* is described as $p = r - a \in [-1, 1]$ where p is the measure of predictive success, $r \in [0, 1]$ is the 'hit rate' (relative frequency of correct predictions) and a is the area which is one minus the the model's power (the percentage of the feasible space explained by the model). Table 4.4 depicts the hit rates under scenarios I to IV and across all scenarios disaggregated by the efficiency levels 0.90, 0.95 and 1.00.

		Hit rate		
		0.90	0.95	1.00
GARP	scenario I	60/67 (89.5%)	50/67 (74.6%)	31/67 (46.2%)
	scenario II	53/67 (79.1%)	43/67 (64.1%)	26/67 (38.8%)
	scenario III	55/67 (82.0%)	45/67 (67.1%)	26/67 (38.8%)
	scenario IV	57/67 (85.0%)	46/67 (68.6%)	20/67 (29.8%)
	all scenarios	27/67 (40.2%)	12/67 (17.9%)	1/67 (1.4%)

Table 4.4: Hit Rates by scenario and Efficiency Level

As expected moving to higher efficiency levels the hit rate is reduced in all cells. For example, under scenario IV 57/67 subjects obey GARP at the 0.90 efficiency level, 46/67 at the 0.95 efficiency level and 20/67 if we do not allow the decision maker to make mistakes. Scenario I presents the highest hit rates under all displayed efficiency levels of 89.5%, 74.6% and 46.2% respectively (*i.e.*, under scenario I, if we allow subjects to be 10% cost inefficient 89.5% of the individual choices can be explained by the utility maximization model). This is due to the simplified portfolio structure that the corresponding scenario offers to subjects (*i.e.*, AD securities).

Given any efficiency level the hit rate percentages are reduced significantly across scenarios ('all scenarios').

		Area (1 - Power)		
		0.90	0.95	1.00
GARP	scenario I	14.3%	4.5%	0.3%
	scenario II	14.2%	4.4%	0.3%
	scenario III	41.9%	19.4%	2.3%
	scenario IV	30.6%	13.0%	1.3%
	all scenarios	0.0%	0.0%	0.0%

Table 4.5: Area by scenario and Efficiency Level

Table 4.5 disaggregates the second component of the Selten's index, the area. The percentages in the cells close to zero indicate that the test's statistical power with respect to the alternative hypothesis of uniform random choice is close to one (*i.e.*, a very small number of the artificial subjects out of 100,500 pass the test).[22] For example, under scenario III (I) the statistical power of the test at 0.90, 0.95 and 1 efficiency level is 58.1 (85.7), 80.6 (95.5) and 97.7 (99.7) respectively. The notable high percentages of scenarios III and IV - when we allow subjects to be 10% and 5% cost inefficient - are justified by the payoff matrices associated with these two scenarios and as a consequence by the choice set available under the AS space. The values displayed in table 4.5 and in Figures 4.4 and 4.5 coincide.

In order to see if the variation of the hit rate across the available scenarios is statistically significant we build the 95% confidence intervals with respect to the the relative size of subjects who pass the test under scenarios I to IV and across scenarios ('all scenarios') given $e = \{0.90, 0.95, 1.00\}$. Recall that hit rates, as displayed in Table 4.4, signify the relative frequency of correct predictions. This frequency follows the binomial distribution with estimates the proportion of successful events. The results are depicted in Figure 4.6.

Under any efficiency level, what emerges immediately is that there is no significant difference between scenarios II, III and IV. Scenario I that corresponds to AD securities, offers higher hit rates at any efficiency level. By relaxing the assumption that decision makers can make mistakes; and focus on the 0.90 efficiency level, scenario IV presents the second highest hit rate. Finally, there are significant differ-

[22] Theoretically the area should be a strictly positive number. However the cells across scenarios ('all scenarios') indicate an estimated area of zero, that is zero artificial subjects pass the test.

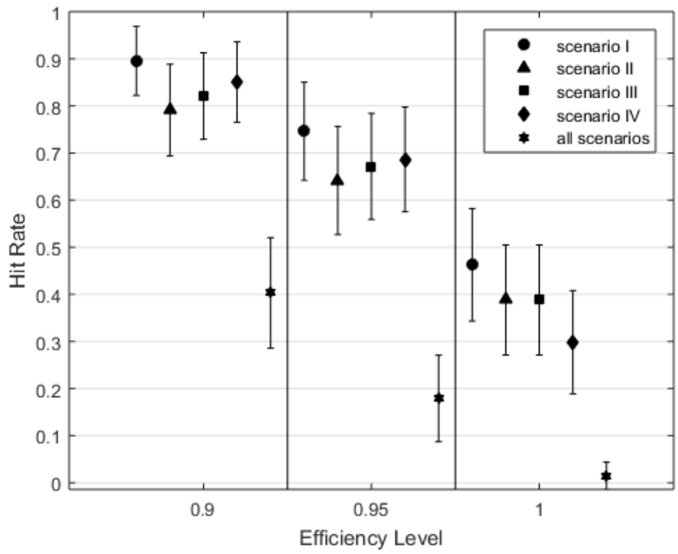

Figure 4.6: 95% Confidence Intervals on hit rates by scenario and efficiency index $e = \{0.90, 0.95, 1.00\}$

ences if we focus on the hit rates across scenarios ('all scenarios'). Even if we relax the test by 10%, 'all scenarios' presents a hit rate significantly lower than scenarios I, II, III and IV.

Selten index	0.90	0.95	1.00
scenario I	0.752	0.701	0.459
scenario II	0.649	0.597	0.385
scenario III	0.401	0.477	0.365
scenario IV	0.544	0.556	0.285
all scenarios	0.402	0.179	0.014

Table 4.6: Predictive success by scenario and Efficiency Level

The Selten index for the threshold levels 0.90, 0.95 and 1.00 is presented in Table 4.6. What we can observe immediately is that at any $e = \{0.90, 0.95, 1.00\}$ - under scenarios I and II - the Selten axiomatic measure is almost determined by the hit rates (see Table 4.4). All cells have indices within the positive range, indicating the power of the test against the alternative hypothesis of random choice. Naturally, the less cost inefficient we allow subjects to be (*i.e.*, efficiency levels 0.95 and 1.00) the indices under all cases are smaller but still superior to the alternative hypothesis. It is worth highlighting that across scenarios ('all scenarios') the index of predictive success is completely characterized by the hit rate under any e; even though these indices are low ranked.

4.6 Conclusion

We develop and implement in an experimental setting a financial asset demand test using revealed preference conditions. The complete market assumption and the fact that asset prices exclude arbitrage allow us to study decision making under risk in the financial asset setting. Subjects purchase securities under four different scenarios given the two states of the world with contingent claims *fixed* across frames. Subjects are facing risk, *i.e.*, where the probabilities assigned to each state are objectively known. Each scenario is associated with a set of securities each of them offer different payoffs in the two available states. Under scenario I subjects are choosing between two Arrow-Debreu securities, under scenarios II, III ,IV they are choosing between two ordinary securities. In our analysis, by applying Afriat's Theorem (Afriat, 1967) we derive conditions that are necessary and sufficient for the dataset so as to be consistent with utility maximization within and across the available scenarios.

At a cost efficiency index 0.90 almost 40% of the subjects are consistent with the utility maximization model, across all scenarios. Naturally, this percentage increases significantly if our focus is within the scenarios. For example under scenario I the percentage of subjects who obey GARP is close to 90%. The test's predictive success across and within scenarios is well within the positive range. That is, the test has explanatory power against the alternative hypothesis of random choice. Moreover, the dataset reveals some heterogeneity among subjects, *i.e.*, some subjects are consistent within and across scenarios while others are consistent within but not across. This heterogeneity among subjects might arise from the different payoff matrices associated with each scenario or from the fact that subjects under different scenarios choose different rules.

Chapter 5

Conclusion

This thesis studied the behaviour of different types of individuals (*i.e.*, traders, experimental subjects) in financial markets under different environments. In Chapters 2 and 3 we modelled the behaviour of different groups of traders in dynamic (limit) order book based financial markets. In a competitive setting where two limit order books operated simultaneously (*i.e.*, an open market and a 'dark' market - Multi-Markets scenario) chapter 2 studied the effects that the 'dark' market had on the 'lit' market's quality, trading behaviour and welfare of market participants. Analysis of the model had shown that - in contrast to the Single-Market scenario (*i.e.*, one single 'lit' market) the 'lit' market's quality is reduced (*i.e.*, reduced liquidity, quoted spread increased, the effective spreads of all market participants increased) under the Multi-Markets scenario. The different trading strategies adopted by the agents led them to make higher profits when they participated in the 'lit' market under the Multi-Markets scenario. An area in which chapter 2 can be expanded would be endogenous information acquisition where agents can choose to buy information by paying a cost. Another area in which this chapter can be advanced is by changing the characteristics of the asset (*i.e.*, volatility of changes in the fundamental value) so as to observe if agents will employ different trading strategies.

Chapter 3 addressed the policy concerns related to 'dark' trading. This chapter studied whether or not 'dark' trading increased the traders' profits - by adopting different trading strategies - and the market quality of the incumbent exchange (*i.e.*, 'lit' market). An innovative setup was used containing different groups of traders (*i.e.*, small and large agents) and two different dark pools in terms of price mechanism applied. Analysis of the model had shown that the 'VWAP-market' scenario (*i.e.*, VWAP dark pool) benefited the 'lit' market's quality while only a certain group of agents (speculators) made less profits in this setup as compared to the alternative scenario (*i.e.*, 'mid-market' scenario). An area in which chapter 3 can be expanded would be to allow large traders to split orders between markets; another area would be the endogenous determination of an agent's type given the prevailing market conditions.

On the other hand, Chapter 4 developed and implemented in an experimental setting a financial asset demand test using revealed preference conditions. In the experiment, subjects purchased securities under four different scenarios in the asset space with contingent claims fixed across frames. Subjects were facing risk since the probabilities assigned to each of the two states were objectively known. Under scenario I subjects were choosing between two Arrow-Debreu securities, under scenarios II, II,IV subjects were choosing between two ordinary securities. The chapter studied the consistency of the data with respect to the maximization hypothesis. Analysis of the test has shown that test's predictive success across and within scenarios was within the positive range. Moreover, the power of the test was close to perfect; and some heterogeneous behaviour had been observed for some subjects across and within scenarios. This heterogeneity among subjects might arise from the different payoff matrices associated with each scenario or from the fact that subjects under different scenarios choose different rules. To address the latter, will require to adopt a different approach (*i.e.*, parametric) in implementing the test. By imposing stronger assumptions over the functional forms and to be able to estimate parameters of interest (*i.e.*, CRRA, CARA).

Appendices

Appendix A

Proof of Proposition

Proposition A.1. *Consider K ordinary securities with an associated matrix of payoffs $\boldsymbol{\Xi}$, such that $\boldsymbol{\Xi} \in \mathbb{R}^{S\times K}$ where S is the number of states. Denote the contingent claim vector of AD prices and the ordinary securities' price vector by $\tilde{\boldsymbol{p}} = (\tilde{p}_1, \tilde{p}_2, \ldots, \tilde{p}_S)$ and $\boldsymbol{p} = (p_1, p_2, \ldots, p_S)$ respectively. Assume that $\boldsymbol{p}$ is related to $\tilde{\boldsymbol{p}}$ in a way such that $\exists\, \tilde{\boldsymbol{p}} \gg 0$ with $\boldsymbol{p}^{\mathsf{T}} = \tilde{\boldsymbol{p}}^{\mathsf{T}} \boldsymbol{\Xi}$ (i.e., prices exclude arbitrage). Given income $w > 0$, consider the sets*

$$CS = \{\boldsymbol{x} \in \mathbb{R}^S_+ \mid \tilde{\boldsymbol{p}} \cdot \boldsymbol{x} = w\}$$
$$AS = \{\boldsymbol{x} \in \mathbb{R}^S_+ \mid \exists \boldsymbol{z} \in \mathbb{R}^K_+ : \boldsymbol{p} \cdot \boldsymbol{z} = w \text{ and } \boldsymbol{x} = \boldsymbol{\Xi}\boldsymbol{z}\}$$

Then, if $\mathrm{rank}(\boldsymbol{\Xi}) = S$ (i.e., $\boldsymbol{\Xi}$ has full rank), the two sets are equal (i.e., $CS = AS$).

Proof. If $\boldsymbol{x} = \boldsymbol{\Xi}\boldsymbol{z}$, it holds that $\tilde{\boldsymbol{p}} \cdot \boldsymbol{x} = \tilde{\boldsymbol{p}}^{\mathsf{T}} \boldsymbol{x} = \tilde{\boldsymbol{p}}(\boldsymbol{\Xi}\boldsymbol{z}) = (\tilde{\boldsymbol{p}}\boldsymbol{\Xi})\boldsymbol{z} = \boldsymbol{p}^{\mathsf{T}}\boldsymbol{z} = \boldsymbol{p} \cdot \boldsymbol{z}$.

($CS \subseteq AS$): Given that the matrix $\boldsymbol{\Xi}$ has full rank, the payoffs of the K ordinary securities across states are linearly independent. That is, there exists a portfolio of ordinary securities that can uniquely represent any state contingent consumption. Denote by $\boldsymbol{z} \in \mathbb{R}^K_+$ the portfolio of the K ordinary securities such that $\boldsymbol{x} = \boldsymbol{\Xi}\boldsymbol{z}$, where $\boldsymbol{x} = (x_1, x_2, \ldots, x_s) \geq 0$ is the vector of contingent consumption. By the definition of CS, $\boldsymbol{x} \in CS$ and $\tilde{\boldsymbol{p}} \cdot \boldsymbol{x} = w \Leftrightarrow \boldsymbol{p} \cdot \boldsymbol{z} = w$ and hence $\boldsymbol{x} \in AS$.

($AS \subseteq CS$): By the definition of AS, given the vector of contingent consumption $\boldsymbol{x} \in AS$ there $\exists \boldsymbol{z} \in \mathbb{R}^K_+$ such that $\boldsymbol{p} \cdot \boldsymbol{z} = w$ and $\boldsymbol{x} = \boldsymbol{\Xi}\boldsymbol{z}$. That is $\boldsymbol{p} \cdot \boldsymbol{z} = w \Leftrightarrow \tilde{\boldsymbol{p}} \cdot \boldsymbol{x} = w$ where we can derive that $\boldsymbol{x} \in CS$. $\square$

Appendix B

Instructions to Subjects

Thank you for taking part in this experiment.

In this experiment you can earn money in addition to the Âč2 show up fee. The amount you will earn will depend partly on your decisions and partly on chance. Please read the instructions carefully because a considerable amount of money is involved. During the experiment, we will refer to tokens instead of pounds. Your earnings will be calculated in tokens and later paid to you in private, in pounds.

In this experiment, 2 tokens = £1.

If you have any questions at any point, please raise your hand and an experimenter will come to your desk.

By clicking NEXT, you will proceed to the instructions.

The experiment consists of 80 tasks. In each task there are two assets: asset RED and asset BLUE. You will need to choose how to allocate tokens between these two assets. In each task, assets RED and BLUE have some associated returns (how many tokens an asset will pay you for each token you allocated to that asset). Different tasks will offer different returns for each asset.

In addition, each task will include two possible events (event 1 and event 2). These events are drawn by the computer randomly with equal probability, independently across tasks. The tokens you have allocated to each asset will have different returns depending on the chosen event.

At the end of the experiment, the computer will draw randomly with equal probability one task out of 80. Your payment in this experiment is determined by your allocation of tokens, the returns and the event randomly drawn by the computer in that task.

You are provided detailed instructions of the experiment in the next screen.

First, the interface to make choices is introduced. The probabilities of each event appear in the box on the top-left. As you can see below, each event has a 50% probability. This probability will remain fixed throughout this experiment. The probability of either event being drawn is independent across tasks. The returns of the assets appear in the other box. In the example below, if event 1 is chosen by the computer, asset RED and BLUE will pay 1 and 0 tokens, respectively. If event 2 is chosen, asset RED and BLUE will pay 0 and 1 token, respectively.

The maximum quantities available of asset RED and BLUE appear in the box on the bottom-left. The maximum quantity available of each asset is the result of a random process. This quantity can vary between 2.5 and 100. In the example below, the maximum quantities of asset RED and BLUE are 100 and 50, respectively. In addition, the prices of asset RED and BLUE appear in the box on the bottom-right. These prices determine the allocations that are possible. As you can see below, the prices of asset RED and BLUE are 1 and 2, respectively. This implies that for each additional token of asset BLUE, 2 tokens need to be deducted from asset RED.

The interface below shows additional boxes where you will be allowed to choose an allocation. Every feasible allocation between 0 and the maximum number of tokens appears in the centre of the screen. Feasible allocations for asset RED and BLUE are represented by the top and bottom line, respectively. On each of these lines, you will have a corresponding dot, representing the allocation to that asset. You can choose an allocation by clicking and dragging either the RED or the BLUE dot to your desired allocation. You can click and drag either dot as many times as you need. Both dots are related to prices as explained in the previous screen. So, if you move one dot, the other will adjust automatically. The initial allocation of tokens, and the corresponding dots, is random; and independent across tasks.

You may practice the choice of an allocation of assets before you proceed to click NEXT.

Lastly, four additional boxes appear in the interface. On the right, a pie chart shows the percentage of expenditure allocated in asset RED and BLUE. Consequently, the pie chart is updated with every new allocation (every new position of the dots). The additional bottom-left box shows your current choice and the bottom-right box calculates the payoffs according to your current choice in the case that either event 1 or 2 is chosen. A summary of your final allocation is provided above the button OK, which you click to confirm your allocation and proceed with the next task.

Let's go through an example. The details of this example intend to help you understand how the interface works. This example does not intend to suggest how you should make your decisions. In this particular example, if event 1 is drawn, asset RED has a return of 1 and asset BLUE has a return of 0 tokens. If event 2 is drawn asset RED has a return of 0 and asset BLUE has a return of 1 token. One possible allocation can be A, where 27.82 tokens are allocated to asset RED and 20.26 tokens to asset BLUE. Another possible allocation can be B, where 63.16 tokens are allocated to asset RED and 5.14 tokens to asset BLUE.

Besides A and B, there are many feasible allocations among asset RED and asset BLUE that you can choose.

On the next screen, you will be given two practice tasks. In each practice task, you will see the two lines representing possible allocations. To choose an allocation, you can click either on the RED or BLUE dot and drag them along the lines towards your desired allocation of tokens; release the mouse button to choose the allocation. This process can be repeated as many times as you want until you are happy with your choice. To confirm your choice, click the OK button. You will then automatically proceed to the next task. The choices you make in the practice tasks will not be used for payment.

Click the START button to begin the practice tasks

[Subjects will take in part in two practice rounds]

Now you know how to use the interface. If you have any questions at this stage, please raise your hand and an experimenter will come to your desk.

The method to determine your payment in this experiment is as follows.

At the end of the experiment, the computer will randomly draw one task from the 80 tasks. All tasks have the same probability of being drawn. In that particular task, the computer will then randomly draw either event 1 or event 2 with equal probability. If event 1 is drawn, then your earnings equal to the return each asset offered for each token you allocated to that asset in the chosen event 1. The returns that assets RED and BLUE offered in the alternative event 2 will not be taken into account. If event 2 is drawn, then your earnings equal to the return each asset offered for each token you allocated to that asset in the chosen event 2. The returns that assets RED and BLUE offered in the alternative event 1 will not be taken into account.

At the end of the experiment, the selected task, assets' prices, the returns of each asset in both events, the probabilities assigned to each event, your own choice in that particular task, and the amount that will be paid are shown on the screen.

The tokens will then be exchanged for money 2 tokens = £1.

Summary

You will see 80 screens. Your task is to choose an allocation of tokens between two assets, asset RED and asset BLUE. You can do this by clicking the RED or the BLUE dot and dragging it towards your selected choice.

In each task, event 1 and event 2 have 50% probability assigned to them (this probability will remain fixed throughout this experiment). The choice of either event is independent across tasks.

At the end of the experiment, the computer will draw a task and an event randomly. You will be paid according to the number of tokens that you have allocated on each asset and the associated returns each asset offers in the chosen event. The returns that assets RED and BLUE offer in the alternative event will not be taken into account.

In this experiment, 2 tokens = £1.

Before you start the experiment, on the next screen you will be asked a question related to how the payoffs are calculated in a hypothetical case.

[Subjects will answer the questionnaire]

The section with the instructions of the experiment is over.

By clicking the START button, you consent to participate in this experiment. Even if you decide to take part in the experiment, you are free to withdraw at any time. Withdrawing from this particular experiment will not affect your relationship with the laboratory.

[Subjects will take part in the experiment]

B.1 Questionnaire

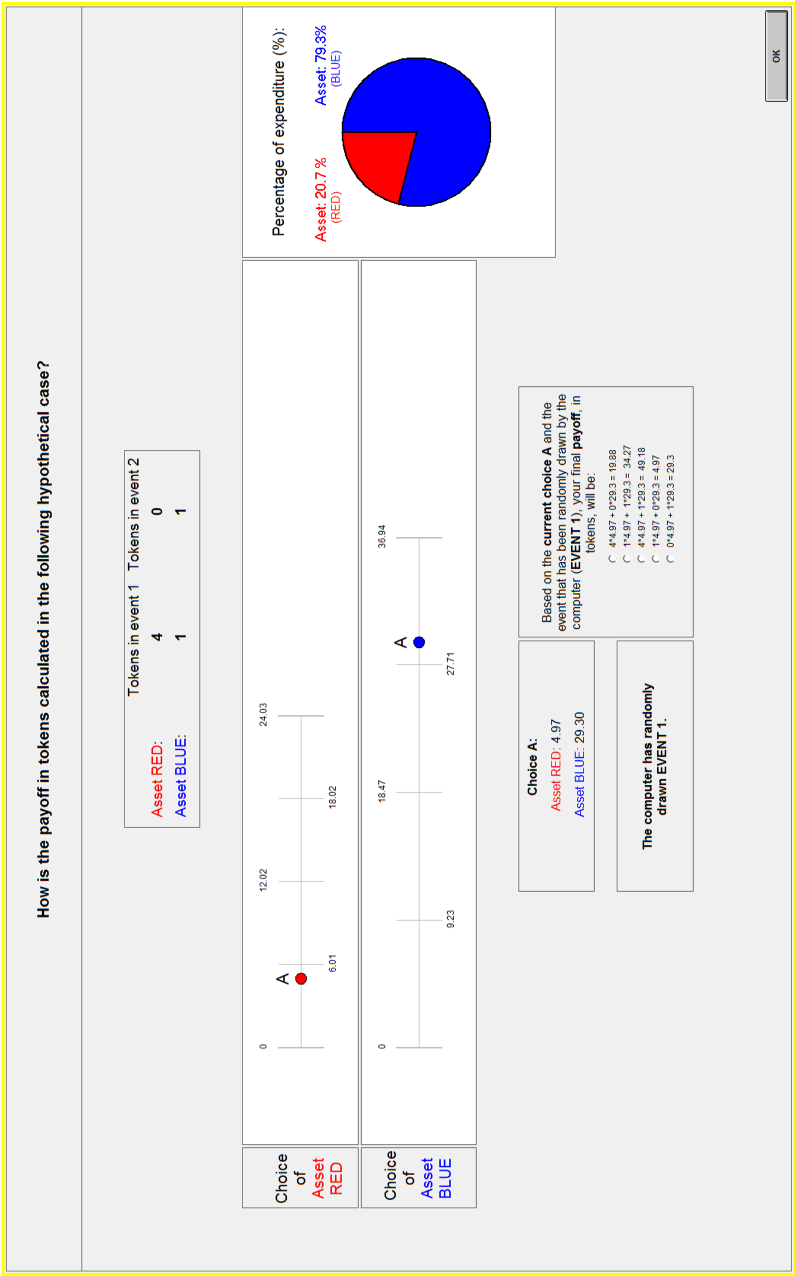

Figure B.1: Questionnaire

Appendix C

Example Screenshot

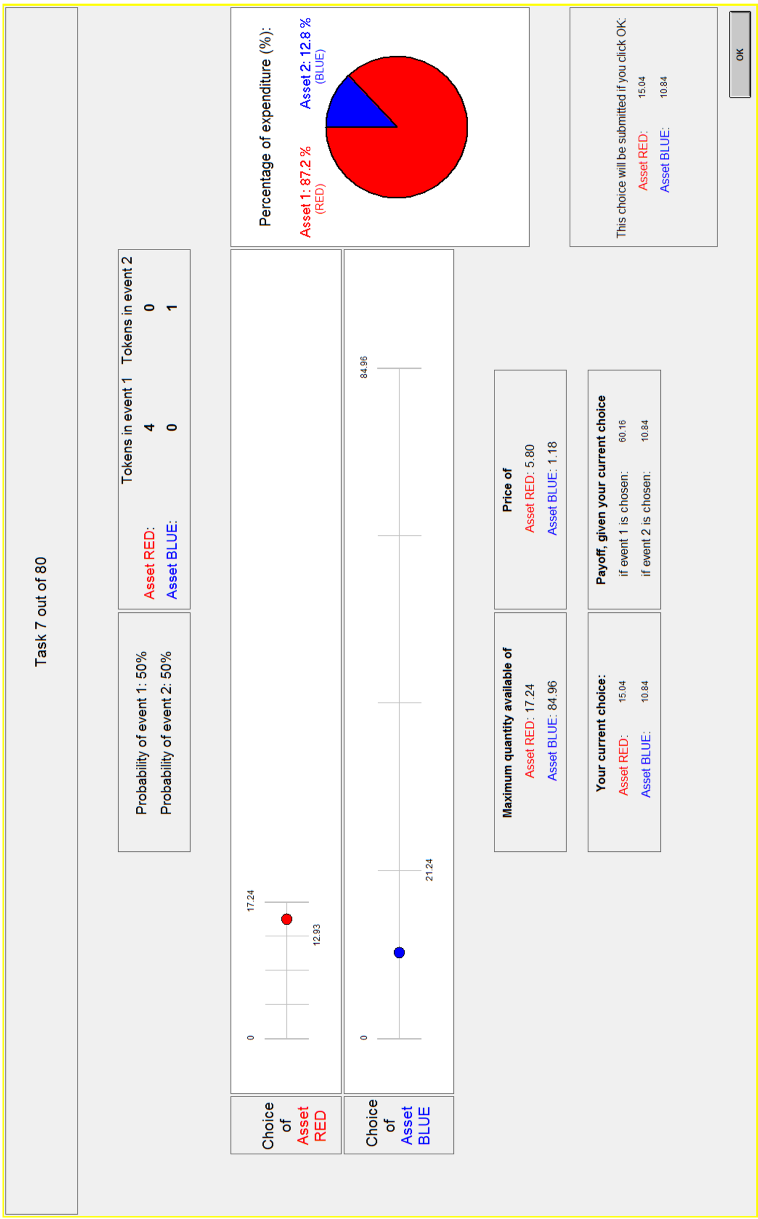

Figure C.1: Screenshot of an Individual Decision Problem

Bibliography

Afriat, S. N. (1972). “Efficiency estimation of production functions”. *International Economic Review*, pp. 568–598.

Afriat, S. N. (1967). “The construction of utility functions from expenditure data”. *International economic review* 8.1, pp. 67–77.

Afriat, S. N. (1973). “On a system of inequalities in demand analysis: an extension of the classical method”. *International economic review*, pp. 460–472.

Arrow, K. J. (1964). “The role of securities in the optimal allocation of risk-bearing”. *The Review of Economic Studies* 31.2, pp. 91–96.

Banks, E. (2014). *Dark Pools: Off-exchange Liquidity in an Era of High Frequency, Program, and Algorithmic Trading.* Springer.

Bernales, A. (2014). *Algorithmic and high frequency trading in dynamic limit order markets.* Working paper. University of Chile.

Biais, B., P. Hillion, and C. Spatt (1995). “An empirical analysis of the limit order book and the order flow in the Paris Bourse”. *the Journal of Finance* 50.5, pp. 1655–1689.

Boulatov, A. and T. J. George (2013). “Hidden and displayed liquidity in securities markets with informed liquidity providers”. *The Review of Financial Studies* 26.8, pp. 2096–2137.

Bronars, S. G. (1987). “The power of nonparametric tests of preference maximization”. *Econometrica: Journal of the Econometric Society*, pp. 693–698.

Brunnermeier, M. K. and L. H. Pedersen (2005). “Predatory trading”. *The Journal of Finance* 60.4, pp. 1825–1863.

Buti, S., B. Rindi, and I. M. Werner (2017). “Dark pool trading strategies, market quality and welfare”. *Journal of Financial Economics* 124.2, pp. 244–265.

Chiarella, C. and D. Ladley (2016). “Chasing trends at the micro-level: The effect of technical trading on order book dynamics”. *Journal of Banking & Finance* 72, S119–S131.

Choi, S. et al. (2007). “Consistency and Heterogeneity of Individual Behavior under Uncertainty”. *American Economic Review* 97.5, pp. 1921–1938.

Choi, S. et al. (2014). “Who is (more) rational?” *The American Economic Review* 104.6, pp. 1518–1550.

Chordia, T., R. Roll, and A. Subrahmanyam (2011). "Recent trends in trading activity and market quality". *Journal of Financial Economics* 101.2, pp. 243–263.

Conrad, J., S. Wahal, and J. Xiang (2015). "High-frequency quoting, trading, and the efficiency of prices". *Journal of Financial Economics* 116.2, pp. 271–291.

Dean, M. and D. Martin (2016). "Measuring rationality with the minimum cost of revealed preference violations". *Review of Economics and Statistics* 98.3, pp. 524–534.

Debreu, G. (1987). *Theory of value: An axiomatic analysis of economic equilibrium.* 17. Yale University Press.

Degryse, H., F. De Jong, and V. v. Kervel (2015). "The impact of dark trading and visible fragmentation on market quality". *Review of Finance* 19.4, pp. 1587–1622.

Diewert, W. E. (1973). "Afriat and revealed preference theory". *The Review of Economic Studies* 40.3, pp. 419–425.

Diewert, W. E. (2012). "Afriat's Theorem and some Extensions to Choice under Uncertainty". *The Economic Journal* 122.560, pp. 305–331.

Dybvig, P. H. (1983). "Recovering additive utility functions". *International Economic Review*, pp. 379–396.

Echenique, F., S. Lee, and M. Shum (2011). "The money pump as a measure of revealed preference violations". *Journal of Political Economy* 119.6, pp. 1201–1223.

Echenique, F. and K. Saito (2015). "Savage in the Market". *Econometrica* 83.4, pp. 1467–1495.

European Commission (2010). "Review of the markets in financial instruments directive (MiFID)". *Public Consultation, European Commission.*

Fischbacher, U. (2007). "z-Tree: Zurich toolbox for ready-made economic experiments". *Experimental economics* 10.2, pp. 171–178.

Foley, S. and T. J. Putniņš (2016). "Should we be afraid of the dark? Dark trading and market quality". *Journal of Financial Economics* 122.3, pp. 456–481.

Fostel, A., H. E. Scarf, and M. J. Todd (2004). "Two new proofs of AfriatâĂŹs theorem". *Economic Theory* 24.1, pp. 211–219.

Fudenberg, D. and J. Tirole (1991). *Game theory.* The MIT press.

Glosten, L. R. and P. R. Milgrom (1985). "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders". *Journal of financial economics* 14.1, pp. 71–100.

Goettler, R. L., C. A. Parlour, and U. Rajan (2005). “Equilibrium in a dynamic limit order market”. *The Journal of Finance* 60.5, pp. 2149–2192.

Goettler, R. L., C. A. Parlour, and U. Rajan (2009). “Informed traders and limit order markets”. *Journal of Financial Economics* 93.1, pp. 67–87.

Green, R. C. and S. Srivastava (1986). “Expected utility maximization and demand behavior”. *Journal of Economic Theory* 38.2, pp. 313–323.

Handa, P., R. Schwartz, and A. Tiwari (2003). “Quote setting and price formation in an order driven market”. *Journal of financial markets* 6.4, pp. 461–489.

Harris, L. (2013). *Maker-taker pricing effects on market quotations*. Working paper. USC Marshall School of Business.

Hasbrouck, J. (2002). “Stalking the âĂIJefficient priceâĂİ in market microstructure specifications: an overview”. *Journal of Financial Markets* 5.3, pp. 329–339.

Hautsch, N. and R. Huang (2011). *Limit order flow, market impact and optimal order sizes: evidence from NASDAQ TotalView-ITCH data*. Working paper. University of Vienna and Humboldt University of Berlin.

Hoffmann, P. (2014). “A dynamic limit order market with fast and slow traders”. *Journal of Financial Economics* 113.1, pp. 156–169.

Hollifield, B. et al. (2006). “Estimating the Gains from Trade in Limit-Order Markets”. *The Journal of Finance* 61.6, pp. 2753–2804.

Jain, P. K. (2005). “Financial market design and the equity premium: Electronic versus floor trading”. *The Journal of Finance* 60.6, pp. 2955–2985.

Kübler, F., L. Selden, and X. Wei (2014). “Asset demand based tests of expected utility maximisation”. *American Economic Review* 104, pp. 3459–3480.

Kwan, A., R. Masulis, and T. H. McInish (2015). “Trading rules, competition for order flow and market fragmentation”. *Journal of Financial Economics* 115.2, pp. 330–348.

Kyle, A. S. (1985). “Continuous auctions and insider trading”. *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.

Loomes, G. (1991). “Evidence of a new violation of the independence axiom”. *Journal of Risk and uncertainty* 4.1, pp. 91–108.

Loomes, G. and G. Pogrebna (2014). “Measuring individual risk attitudes when preferences are imprecise”. *The Economic Journal* 124.576, pp. 569–593.

Menkhoff, L. and M. Schmeling (2010). “Trader see, trader do: How do (small) FX traders react to large counterpartiesâĂŹ trades?” *Journal of International Money and Finance* 29.7, pp. 1283–1302.

Nimalendran, M. and S. Ray (2014). “Informational linkages between dark and lit trading venues”. *Journal of Financial Markets* 17, pp. 230–261.

Pakes, A. and P. McGuire (2001). “Stochastic algorithms, symmetric Markov perfect equilibrium, and the âĂŸcurseâĂŹof dimensionality”. *Econometrica* 69.5, pp. 1261–1281.

Parlour, C. A. (1998). “Price dynamics in limit order markets”. *The Review of Financial Studies* 11.4, pp. 789–816.

Polisson, M., J. K.-H. Quah, and L. Renou (2017). *Revealed preferences over risk and uncertainty*. Working paper. University of St Andrews.

Rieder, U. (1979). “Equilibrium plans for non-zero-sum Markov games”. *Game theory and related topics*, pp. 91–101.

Samuelson, P. A. (1938). “A note on the pure theory of consumer’s behaviour”. *Economica* 5.17, pp. 61–71.

Samuelson, P. A. (1948). “Consumption theory in terms of revealed preference”. *Economica* 15.60, pp. 243–253.

Securities and Exchange Commission (2005). “Regulation NMS”. *Federal register* 70, pp. 37496–37644.

Securities and Exchange Commission (2010). “Concept release on equity market structure; proposed rule”. *Federal register* 75, pp. 3593–3614.

Securities and Exchange Commission (2015). “Regulation of NMS Stock Alternative Trading systems; proposed rule”. *Federal register* 80, pp. 80998–81154.

Selten, R. (1991). “Properties of a measure of predictive success”. *Mathematical Social Sciences* 21.2, pp. 153–167.

Theissen, E. (2000). “Market structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets”. *Journal of Financial Markets* 3.4, pp. 333–363.

Varian, H. R. (1982). “The nonparametric approach to demand analysis”. *Econometrica: Journal of the Econometric Society*, pp. 945–973.

Varian, H. R. (1983). “Non-parametric tests of consumer behaviour”. *The review of economic studies* 50.1, pp. 99–110.

Varian, H. R. (1988). “Estimating risk aversion from Arrow-Debreu portfolio choice”. *Econometrica: Journal of the Econometric Society*, pp. 973–979.

Varian, H. R. (1990). “Goodness-of-fit in optimizing models”. *Journal of Econometrics* 46.1-2, pp. 125–140.

Wilcoxon, F. (1945). “Individual comparisons by ranking methods”. *Biometrics bulletin* 1.6, pp. 80–83.

Ye, M. (2010). *Non-execution and market share of crossing networks*. Working paper. University of Illinois.

Zhu, H. (2013). “Do dark pools harm price discovery?” *The Review of Financial Studies* 27.3, pp. 747–789.