ARTICLE IN PRESS
JID: EOR [m5G;April 10, 2018;5:55]
European Journal of Operational Research 000 (2018) 1–11

Contents lists available at ScienceDirect
European Journal of Operational Research
journal homepage: www.elsevier.com/locate/ejor

Interfaces with Other Disciplines

Investment in high-frequency trading technology: A real options approach☆

Laura Delaney

Department of Economics, City, University of London, Northampton Square, London, EC1V 0HB, UK

A R T I C L E I N F O

Article history:
Received 27 July 2016
Accepted 17 March 2018
Available online xxx

Keywords:
Finance
High frequency trading
Fragmented markets
Real options

A B S T R A C T

This paper derives an optimal timing strategy for a regular slow trader considering investing in a high-frequency trading (HFT) technology. The market is fragmented, and slow traders compete with fast traders for trade execution. Given this optimal timing rule, I then characterise the equilibrium level of fast trading in the market as well as the welfare-maximising socially optimal level. I show that there is always a unique cost of investment such that the equilibrium level of fast trading and the socially optimal level coincide. Finally I discuss potential policy responses to addressing equilibrium and social optimality misalignment in HFT.

1. Introduction

Over the last decade, the state of financial markets has changed considerably. In the first instance, markets have become highly fragmented. There are now more than 50 trading venues for U.S. equities – 13 registered exchanges and 44 so called Alternative Trading Systems (see Biais, Foucault, & Moinas, 2015 and O'Hara & Ye 2011). Hence, traders must search across many markets for quotes and doing so can be costly as it may delay full execution of their orders.

In response to the increase in market fragmentation, so called high frequency trading (HFT) technologies have been developed to reduce the associated costs borne by traders. HFT is a type of algorithmic trading that uses sophisticated computer algorithms to implement vast amounts of trades in extremely small time intervals. For example, traders can buy colocation rights (the placement of their computers next to the exchange's servers) which gives them fast access to the exchange's data feed, they can invest in smart routers which can instantaneously compare quotes across all trading venues and then allocate their orders accordingly, or they can invest in high-speed connections to the exchanges via fiber optic cables or microwave signals. Proprietary trading desks, hedge funds, and so called pure-play HFT outlets are investing large sums of money into such technologies in an effort to outpace the competition. Indeed, according to Hoffman (2014), recent estimates suggest that HFTs are now responsible for more than 50% of trading in U.S. equities.

In a recent paper, O'Hara (2015) details the many ways in which market microstructure has changed over the past decade and calls for a new approach to research in this area which "reflects the new realities of the high frequency world". Nevertheless, there has been a growth in the finance literature on HFT in recent years (see a survey by Foucault 2012). Much of the literature is empirical and, on the whole, the consensus has been that HFT improves liquidity through lower bid-ask spreads (Hendershott, Jones, and Menkveld 2011 and Hasbrouck & Saar, 2013); is highly profitable (Menkveld, 2013 and Baron, Brogaard, Hagströmer, & Kirilenko, 2016); and facilitates price discovery (Hendershott & Riordan, 2013 and Brogaard, Hendershott, & Riordan, 2014). The theoretical literature has also been growing in recent years. For example, Hoffman (2014) presents a stylised model of HFT in a limit order market where agents differ in their trading speed; Pagnotta and Philippon (2015) propose a model in which trading venues invest in speed and compete for traders who choose where and how much to trade; and Biais et al. (2015) develop a model of equilibrium investment in a HFT technology in a Glosten and Milgrom (1985) type framework.

While the finance literature strengthens our understanding of the nature and implication of HFT in many dimensions, to the best of my knowledge, the decision to be fast is always taken to be exogenous. However, the decision to be fast or slow is a real investment like any other. In particular, there is uncertainty associated with the payoff generated from investing in the HFT technology and the investment involves an upfront investment cost which is sunk. Moreover, the slow trader can adopt the technology at any

☆ Thanks to Jacco Thijssen, Saqib Jafarey, Carol Alexander, two anonymous referees, and seminar participants at the Young Finance Scholars Conference 2016 and the Bachelier Finance Society Conference 2016 for their much appreciated comments and suggestions.

E-mail address: laura.delaney.1@city.ac.uk

https://doi.org/10.1016/j.ejor.2018.03.025
0377-2217/

Please cite this article as: L. Delaney, Investment in high-frequency trading technology: A real options approach, European Journal of Operational Research (2018), https://doi.org/10.1016/j.ejor.2018.03.025

future point in time with no terminal date. This adds a dynamic aspect to the investment decision which is not accounted for in other models. Hence, it makes sense to endogenise the investment decision and determine how the optimality of investment timing has implications for HFT in the marketplace. To this end, the problem has a place in the operations research literature.

In this paper I use a real options approach to determine analytically the optimal time for financial market traders to invest in a HFT technology such that the market is fragmented and slow traders compete with HFTs for trade execution. Given this optimal timing strategy, I then compare the equilibrium level of fast trading in the market with the welfare-maximising socially optimal level. While optimal investment timing has been well-developed in the operations research literature through its application to many different types of problems (see, for example, Banerjee, Gucbilmez, & Pawlina, 2014; Battauz, Donno, & Sbuelz, 2015; Delaney & Thijssen, 2015; Munoz, Contreras, Caamano, & Correia, 2011), it is the first application of the approach in a HFT environment.

There are a number of novel results generated by the model, all of which arise from the optimal timing policy derived, in particular, the inclusion of a value of waiting to invest into the slow trader's value function. The results are as follows. (i) It is optimal to wait longer to invest if the level of high frequency trading in the market increases, and early adoption is optimal if the slow trader's probability of finding a liquid venue decreases. (ii) It is also optimal to wait longer if the uncertainty of the profit process increases and, if the probability of finding a liquid venue is low, if the discount rate increases and/or the shortfall in the expected rate of return from holding the option to invest decreases. However, if the probability of finding a liquid venue is high, an increase in the discount rate and a decrease in the shortfall make early adoption optimal. These comparative static results with respect to the discount rate and the shortfall are novel from a real options perspective, which I discuss in a later section of the paper. (iii) There is always a unique equilibrium level of fast trading in the market. (iv) The equilibrium level of fast trading decreases in the probability of finding a liquid venue. (v) There can be either under-investment or over-investment in equilibrium relative to the trading industry welfare-maximising socially optimal level. Over-investment arises when the cost of investing in the technology is low, and under-investment arises when the cost is high. (vi) There is a unique cost of investment such that the socially optimal level of fast trading and the equilibrium level coincide. (vii) Increases in the discount rate and/or uncertainty over the profit process alleviate the extent of over-investment in equilibrium and exacerbate the extent of under-investment. However, over-investment is alleviated and under-investment exacerbated by decreases in the shortfall in the expected rate of return from holding the option to invest.

I discuss the efficacy of using Pigouvian taxes or subsidies to align the equilibrium level of HFT with the socially optimal level. In the case of under-investment, subsidising slow traders' investment cost by an amount equal to the difference between the marginal effect of the level of HFT on utilitarian welfare at the socially optimal level and the equilibrium level appears to be an appropriate policy response to aligning these levels, but in the case of over-investment, taxing HFTs by an amount equal to this difference may not be the most effective response because it will not alter the prevailing level of HFT activity in the market. Instead, when there is over-investment, subsidising HFTs to exit the market by refraining from more fast trading, by an amount equal to the size of this difference, may be a more effective response in that case.

Finally I present and discuss some empirical implications which are generated by the model.

The remainder of this paper is organised as follows. The set-up of the model is described in the next section. In Section 3 the solution to the optimal stopping problem is given, as well as a brief discussion on the comparative statics. Section 4 characterises and compares the market equilibrium and welfare-maximising socially optimal levels of fast trading, as well a providing some discussion on policy implications. Section 5 provides some empirical implications and Section 6 concludes. All proofs are placed in the appendix.

2. The model

Consider a risk-neutral market trader contemplating investment in a HFT technology. Time is continuous, the horizon is infinite, and indexed by $t \in [0, \infty)$. The trader discounts the future at the risk-free rate $r > 0$. Investing in the technology incurs a sunk cost $I > 0$.

The objective is to determine the optimal time to invest in the technology so that the trader's discounted expected payoff from investing is maximised. To solve for this optimal stopping problem, we must determine the expected present value of trading profits for the trader as if he were (i) fast and (ii) slow. These value functions depend on the trading environment which I describe in following subsection, and then I derive the value functions in accordance with this environment.

2.1. The trading environment

The market is fragmented and, as such, liquidity conditions vary across venues so that traders must search for quotes which can lead to delayed execution of orders. At every instant a fraction of the trading venues are "liquid". In the context of this model, a trading venue is liquid if the trader's order, when sent to the exchange venue, is fully executed.

At any instant, the trader can only send an order to one trading venue and his choice of venue is random. HFTs have extremely fast connection speeds to the market and can observe all venues instantaneously. Thus, if the trader is fast, he always finds a liquid one immediately upon order submission. He is indifferent between liquid venues. However, if the trader is slow, he must search for liquid trading venues and finding one can take time. Thus, at each instant he executes his trade with probability λ. Otherwise, with complementary probability, he must continue to search for a liquid venue.[1]

Obtaining empirical estimates of this probability λ is possible using the Fragulator software supplied by http://www.fragmentation.fidessa.com. This software provides comprehensive data and statistics on market fragmentation in the U.S., Europe, Japan, China, and Australia. One such statistic is the Fidessa Fragmentation Index (FFI) which provides a measure of how different stocks are fragmenting across primary markets and alternative venues. The FFI can range from 1 to V_N, where V_N is the number of venues trading a given stock. An FFI = 1 indicates trading residing on one venue, whereas an FFI = V_N indicates trading for the stock is spread evenly across all venues. Therefore, FFI/V_N would provide an appropriate empirical estimate of λ.

2.2. Valuations

The trading activities of a slow trader yields a stream of profits X^S in perpetuity, and the activities of a fast trader yields a profit stream X^F, such that both processes depend on a stochastic process $(X_t)_{t \geq 0}$ which is a geometric Brownian motion of the form:

$$dX = \mu X dt + \sigma X dW, \tag{1}$$

[1] This notion of a fragmented market environment has the flavour of that in Biais et al. (2015).