Limit Order Markets under Asymmetric Information

Ayan Bhattacharya and Gideon Saar*

Version: June 2021

Abstract

We develop a model of dynamic limit order markets under asymmetric information that can be simplified enough to be solved analytically. We use the trader arrival and information environment of the traditional sequential trade models but swap the dealer-based trading core of these models for a dynamic limit order market. We find that informed traders tend to "make" liquidity in illiquid markets and "take" liquidity from more liquid markets. The arrival of marketable and limit orders as well as the passage of time may convey information, resulting in repricing of orders in the book and generating the frequent cancellations and resubmissions that have become a staple of modern markets.

*Ayan Bhattacharya (Ayan.Bhattacharya@baruch.cuny.edu) is at the Zicklin School of Business, Baruch College, The City University of New York, and Booth School of Business, University of Chicago. Gideon Saar (gs25@cornell.edu) is at the Johnson Graduate School of Management, SC Johnson College of Business, Cornell University. We thank Shmuel Baruch, Sabrina Buti (AFA discussant), Thierry Foucault, Ciamac Moallemi, Maureen O'Hara, Roberto Pascual, Ioanid Rosu, and Bart Yueshen (INSEAD Finance Symposium discussant) for helpful comments, as well as seminar participants at Cornell University, the Market Microstructure Online Seminars Asia-Pacific, the Microstructure Exchange, the University of Alberta, the University of Chicago, the 2021 American Finance Association meeting, the 2021 North American Summer Meeting of the Econometric Society, and the 2021 INSEAD Finance Symposium.

1 Introduction

The literature on market structure and trading underwent a paradigm shift with the advent of adverse selection models. Two frameworks in particular proved insightful and versatile: the Glosten-Milgrom-Easley-O'Hara sequential trade models (Glosten and Milgrom (1985), Easley and O'Hara (1987, 1992)) and the Kyle (1985) strategic trader framework. The simple structure of these models made them readily amenable to extensions that incorporated various features of the economic environment, and this led to a rich vein of literature that shed new light on the role of adverse selection in dealer markets.

While dealer markets were the default mode for organizing trade in decades gone by, most trading centers today are structured as limit order markets. Despite this fact, our insights into the trading process continue to come primarily from the older dealer models because the limit order market has proved surprisingly hard to model theoretically. Important early progress in understanding this distinct market structure without asymmetric information was made in the dynamic models of Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2005) and Rosu (2009). Incorporating asymmetric information into these dynamic models, however, has proved particularly challenging, and the papers that have studied this market environment (Goettler, Parlour, and Rajan (2009), Ricco, Rindi, and Seppi (2018), Rosu (2020)) have resorted to numerical solutions to provide insights.

In this paper, we develop a dynamic limit order market model under asymmetric information that can be simplified enough to lend itself to a constructive solution. In particular, we adopt the trader arrival and asymmetric information environment of the Glosten-Milgrom-Easley-O'Hara framework (henceforth, GMEO or traditional sequential trade models) but swap the dealer-based trading core of these models for a dynamic limit order market in the spirit of Rosu (2009). Our focus is on studying how giving each arriving informed trader the ability to use a limit or a marketable order affects market outcomes and the manner in which information is impounded into prices.[1] This approach of embedding a dynamic limit order market within the fairly standard adverse selection

[1]We use the term "marketable" order to denote any order that results in immediate execution upon arrival in the market. This includes traditional market orders as well as limit buy (sell) orders priced at or above (below) the ask (bid) price in the market.

setup provides both novel insights into the manner in which asymmetric information impacts modern markets and a framework that we believe is amenable to being extended in many directions.

Two contributions form the basis of our approach. First, we demonstrate how the prices of limit orders in equilibrium can be expressed in terms of the wait time for execution of these orders. Second, we develop a recursive procedure that enables us to find the wait times of all limit orders using the arrival rates that uninformed traders expect for the various order types. The presence of informed traders in the market causes these arrival rates to depend on the uninformed traders' beliefs about the asset value. As a result, uninformed traders revise the prices of limit orders that rest in the book often as they learn new information. Therefore, our equilibrium structure generates the frequent cancellations and resubmissions that have become a staple of electronic limit order markets over the past two decades.

The model yields a number of novel empirical implications. We show that informed traders tend to submit marketable orders when the spread is narrow and limit orders when it is wide. A marketable order will have a larger permanent price impact when there is more depth on the same side of the book, while a limit order will have a larger permanent price impact when there is less depth on the same side of the book. In other words, informed traders tend to "make" liquidity in illiquid markets and "take" liquidity from more liquid ones, which implies that the likelihood of finding private information in the limit order book is higher when markets are illiquid.

With informed traders choosing between marketable and limit orders, their choices change the rate at which each order type arrives in the market. More frequent arrival of a certain order type could indicate that informed traders are using it, while less frequent arrival of the same order type, by the same token, could indicate that informed traders are not using it. Beyond the frequency of orders, even the passage of time without the arrival of any new order can cause beliefs to change and orders in the book to be repriced to reflect those changed beliefs.

Our model generates several empirical implications that are new to the literature, and therefore can be used to test the economic mechanism that is at the core of the model. Beyond the novel implications, however, we hope that providing a simplified framework that can handle asymmetric information constructively would attract other researchers to extend our model in a myriad of ways, making accessible new insights into limit order markets that have become ubiquitous in our trading landscape.

The rest of this paper is organized as follows. Section 2 reviews related theoretical papers, and Section 3 sets up the model. Section 4 solves for the equilibrium at every instant of time to find the wait times of limit orders and their prices. Section 5 finishes our characterization of the equilibrium by showing how beliefs are updated over time, and discusses how market participants learn from the order flow. Section 6 provides our conclusions.

2 Literature Review

Our paper seeks to embed a limit order market structure in an information-based sequential trade model environment. Traditional sequential trade models (e.g., Glosten and Milgrom (1985), Easley and O'Hara (1987, 1992)) assume that informed and uninformed traders arrive to the market according to an exogenous process, the value of a risky asset is typically either high or low, and informed traders are endowed with private information about this value. Uninformed traders have their demand exogenously determined, and informed traders buy or sell depending on whether the asset's value is high or low, with each trader typically submitting an order for one unit of the asset. The information environment precludes identification of the informed traders, and market participants therefore learn information by observing the order flow as trading evolves. Our paper borrows all these features to replicate the information environment of the GMEO framework.

The distinctive feature of our setup involves the specification of a different market structure—a dynamic limit order market rather than a dealer market—and therefore new order types and choices that informed traders have in such a market structure. The literature on price formation in limit order markets under asymmetric information originated in the static equilibrium models of Glosten (1994), Rock (1996), and Seppi (1997).[2] Glosten (1994) shows that the price paid by a marginal buy (sell) order is the "upper (lower) tail" conditional expectation, and suggests that an open limit order market design is "robust"

[2]For a survey of the literature on limit order markets see Parlour and Seppi (2008).

in the sense that it mimics competition between anonymous dealer markets.[3]

Most dynamic equilibrium models of limit order markets, while providing valuable insights about trader interactions in these markets, do not incorporate asymmetric information. Parlour (1998) models a one-tick market with sequential arrival of long-lived limit orders, demonstrating how liquidity dynamics in limit order markets can generate serial correlation in the order flow. Foucault (1999) models the risk of being picked off in a limit order market when limit orders survive for just one period, generating implications for the equilibrium mix of limit and market orders. Foucault, Kadan, and Kandel (2005) and Rosu (2009) study how the interplay of waiting costs and exogenous trading needs impacts spreads, market resiliency, order-flow autocorrelation, and other temporal properties of trading. While Foucault, Kadan, and Kandel (2005) do not allow cancellation of orders, Rosu (2009) demonstrates how adding this feature to a continuous-time model simplifies the analysis. We believe that Rosu's specification of a limit order market is particularly attractive, and base our modeling of this market structure on his framework.[4]

As far as we know, Goettler, Parlour, and Rajan (2009), Ricco, Rindi, and Seppi (2018), and Rosu (2020) are the only three papers that model fully dynamic limit order markets with informed traders making complex decisions.[5] These three papers, however, require numerical solutions, while our main goal is to present a framework that lends itself to a constructive, analytic solution. Goettler, Parlour, and Rajan (2009) was the first paper to develop such a model, focusing on the information acquisition of the informed traders and the impact of volatility on trading strategies. Ricco, Rindi, and Seppi (2018) present a dynamic limit order market with history-dependent Bayesian learning. They simplify

[3]Back and Baruch (2007) employ a continuous-time model to show that other market designs (for example, uniform price auctions) possess this robustness property as well.

[4]In an early contribution, Cohen, Maier, Schwartz, and Whitcomb (1981) focus on how a bid-ask spread arises endogenously from order choices of traders in a limit order market. Goettler, Parlour, and Rajan (2005) solve numerically for the equilibrium in a multi-period model in which traders with private valuations choose between limit and market orders, focusing on the welfare of traders. Buti and Rindi (2013) study a dynamic limit order market without asymmetric information in which traders can place iceberg (partially hidden) orders, while Buti, Rindi, and Werner (2017) study the market environment when a limit order book co-exists with a dark pool.

[5]Some papers on limit order markets introduce asymmetric information but not in a fully dynamic framework (e.g., Kumar and Seppi (1994), Chakravarty and Holden (1995), Kaniel and Liu (2006), Pagnotta (2010), Baruch, Panayides, and Venkataraman (2017), and Brolley and Malinova (2020)). Such papers often specify a hybrid model with both market makers and limit orders, which simplifies the pricing rule in that limit orders are priced without the need to consider their execution likelihood or time to execution.

the strategies of traders (for example, by not allowing cancellation and repricing), but incorporate market opening and closing effects, showing that price-discovery dynamics can have a pronounced non-Markovian flavor. In both models, however, traders cannot reprice limit orders at will, which means that asymmetric information does not result in frequent revision of limit order prices in the book as it does in our model.

In Rosu (2020), the asset's fundamental value evolves according to a diffusion process and traders entering the market can observe the instantaneous value of the asset by paying a cost. The informational advantage of a trader who chooses to pay decays over time as the value of the asset evolves and subsequent traders buy "fresher" information. In such a framework, Rosu shows that increasing the fraction of informed traders improves liquidity. These three papers utilize distinct modeling assumptions, address somewhat different questions, and require numerical methods to solve the models. Our distinct setup (Glosten-Milgrom-Easley-O'Hara meets Rosu (2009)) and our emphasis on an analytic solution enable us to contribute novel insights to this literature.

3 A Model of Limit Order Trading

Figure 1 describes our objective: combining the information environment of the traditional sequential trade models with the limit order market structure of Rosu (2009). Some assumptions we make are meant to replicate the GMEO environment, while other assumptions are made to implement Rosu's market structure. Specifically, our uninformed and informed traders arrive one at a time under a predefined exogenous process to trade one unit of an asset like in the traditional sequential trade models. To replicate the essence of the GMEO framework, we assume that the information environment precludes direct identification of the informed traders. Our key departure from this framework is in replacing its core market structure: a dealer who sets bid and ask prices as expected values of the asset conditional on the information in arriving market orders. In our limit order market, the rules of the limit order book and the incentives of traders combine to create an array of orders at various prices in the book. Marketable orders execute against these resting limit orders according to the price priority rule enforced by the limit order market's execution engine.

We start by briefly describing the limit order market structure. We then specify the assumptions made to replicate the GMEO environment, and conclude by discussing the notion of equilibrium.

3.1 Model Details

The Limit Order Market Structure

Following Rosu (2009), the market is organized as a continuous-time electronic limit order book: traders may submit limit orders (that enter the limit order book) or marketable orders (that execute against resting limit orders upon arrival in the market and therefore never enter the limit order book). These two types of orders express two levels of trading urgency in that marketable orders are limit orders that are priced to yield immediate execution: buy orders priced at or above the ask price and sell orders priced at or below the bid price. Marketable orders include traditional market orders that are not price contingent as a subset, although many electronic limit order books do not accept traditional market orders and require all orders to be price contingent. We refer to immediately executable orders throughout the paper as "marketable" orders.

A trade takes place only when an arriving marketable order executes against a limit order that is resting in the book. The limit order book trading mechanism (or execution engine) follows the usual price-priority rule. In other words, a buy (sell) limit order at a higher (lower) price is executed before a buy (sell) limit order at a lower (higher) price. This price priority is hard-coded into the execution engine and cannot be altered by traders. We use the term ask (bid) price to denote the lowest-priced (highest-priced) limit sell (buy) order in the market.

Once a trader arrives in the market and submits a limit order, the trader continuously monitors the market and is able to reprice the order at will until either the submitted limit order executes or the trader chooses to cancel the order and leave the market. Pre-trade and post-trade transparency in the market imply that resting limit orders are observable to all market participants and information about trades (and the prices at which they execute) is publicly available. Still, order submission is anonymous: the identity of traders who submit orders is unknown to other traders. The price grid for posting limit orders is continuous (the tick size in the limit order book is zero).

Constructing the Information Environment

A single risky asset is traded in the limit order market. As typical in traditional sequential trade models, the liquidation value of the asset, $\tilde{v}$, is a Bernoulli random variable that takes either a high value, $v+\sigma$, or a low value, $v-\sigma$. While Nature chooses the realization of the random variable prior to the beginning of trading, the realized value is not commonly known to all market participants. The parameter of the Bernoulli distribution, $q_0 \equiv \mathrm{P}_0\left[\tilde{v}=v+\sigma\right]$, is public information, and thus the expected value of the asset at the start of trading $(t=0)$ is

$$v_0 = q_0(v+\sigma) + (1-q_0)\,(v-\sigma). \tag{1}$$

There are five types of traders in our market: impatient uninformed buyers (denoted by bi), impatient uninformed sellers (si), patient uninformed buyers (bp), patient uninformed sellers (sp), and informed traders (I). The four types of uninformed traders and the informed traders arrive in the market according to independent Poisson processes with constant rates (or intensities) λ_{bi}, λ_{si}, λ_{bp}, λ_{sp} and λ_I, respectively. Like in the traditional sequential trade models, these exogenous arrival rates imply that each trader behaves as if he or she will arrive in the market only once, and all orders are for one unit of the asset.

Impatient uninformed traders participate in the market to satisfy liquidity or hedging needs. We follow the convention in the traditional sequential trade models and assign them an order direction (buy or sell), and their impatience dictates that they use marketable orders. As such, these traders are identical to the uninformed traders in the GMEO framework. To ensure that, as in this framework, trading only takes place in the interval $[v-\sigma, v+\sigma]$, we give these impatient uninformed traders reservation prices: buyers are only willing to buy the asset at prices at or below $(v+\sigma)$, while sellers are only willing to sell the asset at prices at or above $(v-\sigma)$.[6]

To replace the dealer-market core of the GMEO framework with a limit order market,

[6] A reservation price is a limit on the price of a good or a service offered in an auction: on the demand side, it is the highest price that a buyer is willing to pay; on the supply side, it is the lowest price a seller is willing to accept for a good or a service. Auctions allow the setting of reservation prices because sellers (or buyers, if it is a buyers' auction) have a private estimate of the value of the object on auction and are unwilling to trade if that value is breached.

we substitute the dealer with a dynamic flow of liquidity providers called "patient" uninformed traders. They have an exogenous trading desire to buy or sell, but being patient means that they are not willing to pay for liquidity. Rather, they would like to be compensated for supplying liquidity with limit orders. Therefore, a patient buyer (seller), trades off the gain that can be made by setting a lower (higher) limit order price against the cost of waiting longer in the limit order book for execution.[7]

Formally, let $q_t \equiv P_t[\tilde{v} = v + \sigma]$ be the beliefs of the patient uninformed traders at time t, and $v_t = q_t(v+\sigma) + (1-q_t)(v-\sigma)$. Let r be the discount rate, and $\tilde{T}_{bp}$ ($\tilde{T}_{sp}$) be the expected time to execution of a limit buy (sell) order. The utility functions of the patient uninformed buyers and sellers, respectively, are

$$\begin{array}{l} EU_t^{bp} = \mathrm{E}\left[v_t - \pi_{t,bp} - r(\tilde{T}_{bp} - t) \mid q_t\right] \\ EU_t^{sp} = \mathrm{E}\left[\pi_{t,sp} - v_t - r(\tilde{T}_{sp} - t) \mid q_t\right] \end{array}, \tag{2}$$

where the limit order price $\pi_{t,bp}$ ($\pi_{t,sp}$) is the patient uninformed buyer's (seller's) choice variable.

The traditional sequential models assume that the dealer must break even on his trades. To make our setup similar to theirs, we assume that the patient traders require non-negative expected utility. This assumption gives an economic meaning to the notion of being patient even when traders have a desire to buy or sell because it implies that, if they had the choice between marketable and limit orders, patient traders would always choose limit orders to provide rather than demand liquidity.[8] Without loss of generality, we set $r = 1$ to simplify the exposition going forward.

In terms of information, a patient uninformed trader (*he*) initially knows only v_0. However, just like the dealer in the GMEO framework, he uses information from the market

[7] Patient traders who use limit orders to provide liquidity to impatient traders are also central to the limit order models of Foucault, Kadan, and Kandel (2005) and Rosu (2009).

[8] To see that this is the case, note that the requirement to at least break even on liquidity provision implies that bid and ask prices in the market would always bracket the expected value of the asset, as in the traditional sequential trade models. The break-even condition of the patient uninformed traders implies

$$\pi_{t,sp} \geq v_t + \mathrm{E}\left[r(\tilde{T}_{sp} - t) \mid q_t\right] > v_t - \mathrm{E}\left[r(\tilde{T}_{bp} - t) \mid q_t\right] \geq \pi_{t,bp},$$

and hence $ask_t > v_t > bid_t$. This means that the expected utility of a patient uninformed seller (buyer) from using a marketable order will be negative: $bid_t - v_t < 0$ ($v_t - ask_t < 0$).

environment to update his beliefs and learn about the true value of the asset. In principle, limit order markets allow traders to choose their order type. We create a simpler structure—with uninformed traders that are assigned to both supplying and demanding liquidity—because we want to focus on integrating informed traders who can choose between limit and market orders into our market in order to study how information is incorporated into prices in limit order markets.

The essence of an informed trader (*she*) in our model is that she knows the realized value of the random variable $\tilde{v}$. Informed traders do not have an exogenous reason to trade, but rather arrive in the market to exploit their informational advantage. To simplify the exposition in the paper, we assume that informed traders do not bear waiting costs. This is equivalent to saying that their horizon for trading is the same as the horizon of the random variable $\tilde{v}$. There is no conceptual difficulty in having the informed traders bear waiting costs, but it makes the model more complex without adding insights on how information is incorporated into the market. As such, we choose to have the simpler exposition in the paper, while providing the more general formulation in the Online Appendix.

Given that all trading takes place in the interval $[v-\sigma, v+\sigma]$ and each informed trader behaves as if she will arrive in the market only once, her decision on whether to buy or sell the asset is the same as in the traditional sequential trade models: buying when the true value is high and selling when it is low. Trading in a limit order market, however, she need to choose between a marketable and a limit order. Specifically, if the true value is $v+\sigma$ $(v-\sigma)$, her utility is $v+\sigma-\pi_{t,bI}$ $(\pi_{t,sI}-(v-\sigma))$ if she submits a limit order with price $\pi_{t,bI}$ $(\pi_{t,sI})$, and, if the other side of the book is not empty, $v+\sigma-ask_t$ $(bid_t-(v-\sigma))$ if she submits a marketable order. Without having to bear waiting costs, an informed trader prefers submitting a limit order, holding everything else equal. Why would an informed trader ever submit a marketable order? Because uninformed traders learn from the order flow. If choosing a limit order conveys "too much" information to the market, the price impact of a limit order could be larger than the cost of executing a trade with a marketable order.

A key feature of the GMEO framework is the assumption that the information environment precludes identification of the informed traders, and hence uninformed traders learn gradually from the order flow. We would like to replicate this information environment in our model, which requires us to discuss off-equilibrium beliefs. The reason this discussion is missing from the traditional sequential trade models is because they restrict informed traders to using only market orders, and the dealer always stands ready to accommodate them. In our limit order market, informed traders can supply liquidity, and marketable orders may not be executed if the other side of the book is empty.

For example, it could be the case that submitting a limit order would reveal to the market the informed trader's identity (and hence her information), thus preventing her from capitalizing on her private information. If the other side of the book is empty and therefore she cannot trade with a marketable order, the informed trader may choose to submit a limit order and reveal her identity because she can still make money by supplying liquidity to the impatient uninformed traders even if the value of the asset is known with certainty. In most instances, the expected profit from both capitalizing on her private information and providing liquidity is greater than the profit from liquidity provision alone. However, the aforementioned example requires us to specify off-equilibrium beliefs in order to maintain the GMEO environment in which informed traders are unidentified.

To simplify the analysis in the paper, we assume that revealing oneself as an informed trader entails a cost that is equal to or higher than the expected utility she can obtain from submitting a limit order in a market where everyone knows her private information. There is more than one story that makes this assumption reasonable. For example, revealing oneself as an informed trader invites scrutiny from regulators. Dealing with regulatory scrutiny is costly (e.g., the cost of lawyers and the opportunity cost of time invested in defending oneself) even if the private information itself is legal.[9]

This simple assumption ensures that informed traders do not reveal themselves and the essence of the GMEO information environment is maintained. Alternatively, one could achieve the same outcome by simply assuming that informed traders believe that

[9]In our model, an informed trader would reveal herself by choosing a limit order price that would not be chosen in equilibrium by an uninformed trader, which would lead to an immediate adjustment of the uninformed traders' beliefs and a large price impact. Officials at the Market Regulation division of the Financial Industry Regulatory Authority (FINRA) told us that they constantly scan market prices to detect such abnormal moves. They also maintain an anonymous hot-line, receiving many calls from market participants who wish to point their attention to certain price moves or order flow strategies. FINRA would then pursue an investigation, and is able to identify the trader behind each order. As such, the assumption that informed traders would choose not to submit a limit order at strange prices that would cause a price dislocation because they fear regulatory scrutiny is very realistic.

trading in the market stops if they are revealed as informed traders. In the Online Appendix, we demonstrate how the model can be solved without such an assumption, though at the cost of added complexity. Given that the added complexity does not result in additional insights and that our goal is to replicate the information environment in the traditional sequential trade models, we proceed in the paper making the simple assumption on the off-equilibrium beliefs that prevents informed traders from revealing themselves.

3.2 Equilibrium Concept

Preliminary Discussion

It is important to recognize that there are two distinct games that are nested within our model. The first one is essentially the game at each node when players update their information, but in our case it is every instant of time due to the use of Poisson processes for order arrivals. We denote by $\mathcal{G}_t$ this instantaneous game at time t. The second game is the long-horizon game $\mathcal{G} = (\mathcal{G}_t)_{t\in[0,+\infty)}$, which is simply the collection of all instantaneous games; it is how the instantaneous games are connected through time. This seems like a typical structure that would call for applying the Perfect Bayesian Equilibrium concept, but we follow Rosu (2009)'s approach to solving the instantaneous game, and this creates a subtle distinction. A Perfect Bayesian Equilibrium reduces to a Bayes Nash Equilibrium for the static game, which in turn reduces to a Nash equilibrium when one moves to the agent normal form of the incomplete information game. In our case, the equilibrium in the agent normal form game is a version of rationalizability that has been specifically formulated in Rosu (2009) for the instantaneous game between liquidity providers in limit order markets.

The equilibrium introduced in Rosu (2009) is not exactly Markov perfect: while player strategies are indeed Markov in Rosu's game, the equilibrium is not Markov because it does not yield a Nash equilibrium in subgames.[10] Instead, Rosu's equilibrium criteria reduces to the rationalizability solution concept in subgames (see, Bernheim (1984) and Pearce (1984)). Rosu's discussion about maintaining appropriate distances between or-

[10]It is easy to check that liquidity providers would keep undercutting each other in a cyclical fashion if they were to use the Nash best-response criterion.

ders to prevent undercutting by traders, which we formally model below, is a description of a rationalizable strategy that follows from the common knowledge of rationality in the game. Therefore, his limit order model implements a dynamic version of rationalizability.

In Rosu's paper, however, all traders have the same information about the asset. In order to introduce informed traders into the mix, we augment Rosu's equilibrium criteria with two additional requirements borrowed from the Perfect Bayesian Equilibrium to create a solution concept tailored for limit order markets under asymmetric information. First, traders choose their actions conditional on their beliefs at each node of play. This implies that the patient uninformed traders choose their limit order prices in each instantaneous game $\mathcal{G}_t$ conditional on their beliefs at time t. Second, trader beliefs over information sets are updated using Bayes' rule. This means that in the long-horizon game $\mathcal{G} = (\mathcal{G}_t)_{t\in[0,+\infty)}$, trader beliefs are updated from $\mathcal{G}_{t'}$ to $\mathcal{G}_{t''}$ ($t'' > t'$) using Bayes' rule. As such, we create a Perfect Bayesian extension of Rosu's equilibrium.

Formal Definition

To economize on notation, we use Λ and u to refer to the arrival rates and preferences of all trader types, respectively. Denote the tuples of limit order sellers and limit order buyers waiting in the book at some arbitrary time t by x_t and y_t, respectively, with $(p_i)_{i\in x_t}$and $(p_k)_{k\in y_t}$ as the prices of these limit sell and buy orders. Then, the instantaneous game at time t can be represented as

$$\mathcal{G}_t = \left(x_t, y_t, (p_i)_{i\in x_t}, (p_k)_{k\in y_t}, q_t, \Lambda, u\right). \tag{3}$$

The *players* in $\mathcal{G}_t$ are the patient uninformed traders and informed traders who are present in the market at t, represented by the variables x_t and y_t. Thus, traders who have not yet arrived by t or who have already exited the market are not players in $\mathcal{G}_t$. In terms of *actions* available to the players in $\mathcal{G}_t$, a patient uninformed trader chooses either to submit a limit order (and needs to decide on his limit order price) or leave the market without trading. An informed trader chooses between submitting a limit order, submitting a marketable order, or leaving the market without trading. The infinite horizon game $\mathcal{G} = (\mathcal{G}_t)_{t\in[0,+\infty)}$ inherits its structure from $\mathcal{G}_t$, and we implicitly assume the asset pays off at the end of game $\mathcal{G}$.[11] Using this structure, we define the equilibrium as follows.

Definition 1. (*Equilibrium in games $\mathcal{G}_t$ and $\mathcal{G}$*). An equilibrium in the instantaneous game $\mathcal{G}_t$ is:

(i) a set of limit order prices $(p_i)_{i \in x_t}$ and $(p_k)_{k \in y_t}$, and

(ii) the choices of patient uninformed traders, conditional on their beliefs q_t, between submitting limit orders and leaving the market without trading, and

(iii) the choices of informed traders between submitting limit orders, submitting marketable orders, and leaving the market without trading,

such that, given common knowledge of rationality among the traders, they have no incentives to reprice their limit orders or change their choices.

An equilibrium in the long-horizon game $\mathcal{G}$ is a sequence of equilibria in the instantaneous games $(\mathcal{G}_t)_{t \in [0,+\infty)}$ such that for any $t^{''} > t^{'}$,

1. $q_{t''} \in \mathcal{G}_{t''}$ is obtained from $q_{t'} \in \mathcal{G}_{t'}$ using a Bayesian update process given the orders submitted by informed and uninformed traders that arrive according to Λ in the interval $[t^{'}, t^{''})$, and
2. Traders update their choices, canceling and resubmitting limit orders or leaving the market altogether, in accordance with these updated beliefs to effect the change from $\mathcal{G}_{t'}$ to $\mathcal{G}_{t''}$.

Common knowledge of rationality in the equilibrium definition means that all market participants know that no trader will choose an action that is dominated by another action (in terms of utility) in the trader's action space, and they all know that they know this, and so on. Rosu (2009) provides an elegant discussion of how common knowledge of rationality would cause traders to reprice their orders in $\mathcal{G}_t$ if prices were to deviate from the equilibrium prices. The basic idea is that if a trader were to choose a limit order price that is not an equilibrium price, it would trigger a repricing response from another trader (because doing so would increase that trader's utility), which in turn would trigger a repricing response from another trader, and so on. Given common knowledge of rationality, traders should be able to reason through this chain; thus, equilibrium prices must satisfy the properties specified in Definition 1.

[11] We could equivalently assume that the asset pays off after an exponentially-distributed interval of time. Either way, our goal is to focus on a stationary equilibrium.

The repricing procedure that Rosu sketches (and we describe and solve in detail using wait times) is a form of iterated elimination that follows from the notion of rationalizability. It is important to emphasize that in the instantaneous game $\mathcal{G}_t$, this iterative repricing process is essentially a thought experiment that allows us to pin down equilibrium prices. No repricing actually occurs because limit orders in the book are priced such that no trader has an incentive to reprice his or her order.

The Perfect Bayesian Equilibrium element we borrow for the definition of equilibrium in the long-horizon game $\mathcal{G}$ embeds the equilibrium in each instant of time ($\mathcal{G}_t$) in a Bayesian rational framework that extends over time. Informed and uninformed traders submit their orders when they arrive in the market, and uninformed traders use Bayes' rule to update their beliefs on the value of the asset, which can lead to cancellations and resubmissions of resting limit orders in the book to reflect the updated beliefs. In other words, the process of moving from $\mathcal{G}_{t'}$ to $\mathcal{G}_{t''}$ may involve traders actually repricing their limit orders (if equilibrium prices change) or even canceling their orders and leaving the market altogether. This equilibrium definition mirrors the Bayesian framework of the sequential trade models, except that the dealer-based core of these models is swapped for a limit order market.[12]

4 Equilibrium in the Instantaneous Game: Wait Times and Prices

In this section, we construct the equilibrium in the instantaneous game $\mathcal{G}_t$. Our first contribution is to show how trader strategies combine with the limit order market structure to create a relationship between wait times for the execution of limit orders and prices in the book. We then derive the arrival rates of orders from the perspective of the unin-

[12] In Rosu's framework, prices at time t are solely determined using information in $\mathcal{G}_t$. Nothing in our Perfect Bayesian extension changes this structure. The first requirement we add is that the limit order price, π_t, is chosen by a patient uninformed trader based on his beliefs at time t, q_t, such that no other trader would undercut him in $\mathcal{G}_t$ (the rationalizability argument). The second requirement we add is that traders' beliefs over information sets (q_t) are updated using Bayes' rule as we move from $\mathcal{G}_{t'}$ to $\mathcal{G}_{t''}$ ($t'' > t'$). Note that these two requirements do not create a situation in which traders consider $\mathcal{G}_{t''}$ when choosing limit order prices in $\mathcal{G}_{t'}$. In other words, traders need not consider the possibility that prices may change in the future when setting their limit order prices at each instantaneous game.

formed traders, which enables us to present our second contribution: a novel recursive procedure for calculating wait times of limit orders in the book. These steps complete the construction of the equilibrium in the instantaneous game. The Bayesian update process that characterizes the long-horizon game $\mathcal{G}$ is developed in Section 5.

4.1 Wait Times

Recall that x_t and y_t denote the tuples of limit sell and buy orders in the book at time t, respectively. A particular configuration of the book with n limit sell orders and l limit buy orders is represented by $(x_1, \ldots, x_n; y_1, \ldots, y_l)_t$, where x_1 is the limit sell order with the highest price and x_n is the limit sell order with the lowest price (or the ask price in the market), although we often use the shorthand notation $(x^n; y^l)_t$. We refer to a specific limit sell order within a configuration, x_m, by writing $(x_m|x^n; y^l)_t$, $m \leq n$; for example, $(x_4|x^4; y^2)_t$ refers to the ask price at time t in a book with four limit sell orders and two limit buy orders. The notation for the buy side of the book is symmetric (substituting y for x).

Definition 2. (*Wait time*). Wait time $w_t\left(x_m|x^n; y^l\right), m \leq n$, is a real-valued function that measures the expected value of the (random) time $\tilde{T}$ it takes for the order x_m to be executed, where the expectation is taken with respect to the information set of the uninformed traders at time t. The wait time for the buy side of the book $w_t\left(y_k|x^n; y^l\right), k \leq l$, is defined in an analogous fashion.

Essentially, we define a *wait time* to denote how long uninformed traders expect to wait before execution of a particular order (x_m) given a certain configuration of the book $(x^n; y^l)$.[13] While the starting configuration of the limit order book is noted in the definition of the wait time, we only know in advance the ending configuration on one side of the book. Specifically, if x_m is the particular order whose wait time we calculate, we know that the sell side of the book will contain x_{m-1} following the execution. However, our calculation of the wait time would need to take into account all possible states of the other

[13]We drop the word "expected" in front of *wait time* to simplify the writing but we always mean an expected wait time.

side of the book that can arise due to the random arrivals of limit buy and marketable sell orders during the the time it takes for our particular limit sell order to execute.

In the next section, we show how trader decisions give rise to a relationship between wait times and prices of limit order in the book. Development of the recursive formulations to calculate the wait times follows in Section 4.3.

4.2 Trader Strategies

The strategies of the impatient uninformed traders are simple: buyers buy if $ask_t \leq v + \sigma$ and sellers sell if $bid_t \geq v - \sigma$. The strategies of the patient uninformed traders are more complex, and require us to look at their utility functions and the definition of equilibrium in the instantaneous game $\mathcal{G}_t$. The expected utility function of the patient uninformed traders in equation (2) consist of two components: the first is the difference between the expected value of the asset and the limit order price, and the second is the expected wait time to execution of the limit order. The price priority rule of the limit order book implies that the wait time of a limit sell order at a higher price is at least as long as the wait time of a limit sell order at a lower price because the latter would need to execute before the former would get a chance at execution:

$$w_t\left(x_n|x^{n+m};y^l\right) \geq w_t\left(x_n|x^n;y^l\right),\ m \geq 1. \tag{4}$$

Say there is a resting limit sell order in the book. The arrival of a new limit sell order that is placed at a lower price pushes up the expected wait time and brings down the expected utility of the seller with the resting limit sell order. Of course, that seller could reprice his order by canceling and resubmitting at a lower price to reclaim priority. However, if the newly arrived seller places his order at a price that is sufficiently below that of the resting limit sell order, the seller with the resting limit sell order has no incentive to cancel and resubmit—the loss in utility from reducing the price would outweigh the gain in utility from the shorter wait time. Given common knowledge of rationality, traders should be able to reason through this sequence, and thus in equilibrium the distance between the two limit sell orders should equal the the gain in utility from obtaining execution priority.[14] This creates a link between wait times and prices, and allows us to pin down the gap between limit order prices in equilibrium—a state in which no patient seller has an incentive to reprice his limit order.

Definition 3. (*Price gap*). The price gap $g\left(x_n, x_{n+1}|x^{n+m}; y^l\right)$ denotes the non-negative difference in price between the n^{th} and $(n+1)^{th}$ limit sell orders when the book contains $n+m$ ($m \geq 1$) limit sell orders and l limit buy orders at time t.

Our intuitive discussion above is proved formally in the following proposition.

Proposition 1. *In equilibrium, the price gap between sell orders with consecutive priorities in a limit order book with configuration $\left(x^{n+m}; y^l\right)_t$ is given by*

$$g_t\left(x_n, x_{n+1}|x^{n+m}; y^l\right) = w_t\left(x_n|x^{n+m}; y^l\right) - w_t\left(x_{n+1}|x^{n+m}; y^l\right),\ m \geq 1. \tag{5}$$

The price gap between buy orders is defined in an analogous fashion.

All proofs can be found in the Appendix. The price gap (or simply, *gap*) in equation (5) essentially represents a thought experiment at an instant in time, which is why the limit order book configuration is the same in all terms of the equation.[15] This equation tells us what distance should be between two prices in the book to prevent an endless cycle of undercutting by the traders. Wait times (and price gaps) are indexed by t because, as we will show later in this section, they are a function of the uninformed traders' beliefs q_t. We postpone the discussion of how beliefs are updated over time to Section 5; here, we establish the equilibrium in the instantaneous game $\mathcal{G}_t$ given the beliefs q_t.

Of course, price gaps are not enough to pin down all equilibrium prices. To obtain the prices, we need to anchor the first price and have a way of determining when the price gaps stop, which we do by providing two boundary conditions. The lower boundary

[14]With a continuous price grid, multiple limit order traders would not find it optimal to wait at the same price but rather would post their orders at other price levels. These levels are set exactly to give traders on the same side of the book the same expected utility after taking into account the wait cost. Hence, time priority does not play a role in our model.

[15]The discount rate r does not appear explicitly in the right-hand-side of equation (5) because to simplify the exposition, we assume $r = 1$. All propositions and proofs in the paper can be rewritten with an arbitrary discount rate at the cost of introducing some complexity to the expressions, but none of the results or insights would change.

condition for the sell side of the book comes from the break-even condition of the patient uninformed liquidity providers: $\pi_{t,sp} \geq v_t + \mathrm{E}\left[(\tilde{T_{sp}} - t) \mid q_t\right]$. To rewrite the lower boundary condition in terms of our wait time notation, we introduce the notion of a *full book.*

Definition 4. (*Full book*). The sell (buy) side of the book is full if the next patient uninformed trader to arrive in the market would choose to leave without submitting a limit sell (buy) order because posting it at an equilibrium price satisfying equation (5) would result in violating his break-even condition. The number of orders in the sell (buy) side of the book when it is full is denoted by F_s (F_b) .

The lower boundary therefore is established to ensure that the expected gain from selling at a particular price is sufficient to cover the expected cost of waiting at that price. The upper boundary condition for the sell side of the book is essentially the limit order price chosen by a seller who arrives when the sell side of the book is empty. Choosing a higher price increases expected utility, but a higher price also means that more limit orders can undercut our seller, increasing his wait time and reducing his expected utility. The next proposition shows that the positive direct effect of a price increase on expected utility outweighs the negative indirect effect through the wait-time channel.

Proposition 2. *If the sell side of the book is empty, a patient uninformed seller would place his order at the highest possible price, $p_t\left(x_1|x^1;y^l\right) = v + \sigma$. Moreover, the lowest-priced limit sell order when the sell side of the book is full must satisfy the condition*

$$p_t\left(x_{F_S}|x^{F_S};y^l\right) \geq v_t + w_t\left(x_{F_S}|x^{F_S};y^l\right). \tag{6}$$

If the buy side of the book is empty, a patient uninformed buyer would place his order at the lowest possible price, $p_t\left(y_1|x^n;y^1\right) = v - \sigma$. Moreover, the highest-priced limit buy order when the buy side of the book is full must satisfy the condition

$$p_t\left(y_{F_B}|x^n;y^{F_B}\right) \leq v_t - w_t\left(y_{F_B}|x^n;y^{F_B}\right). \tag{7}$$

The first patient uninformed seller who comes to an empty book would post the limit order at the highest possible price. If the second one arrives before the first one executes, he would post his limit order at the equilibrium price, $v + \sigma - g_t\left(x_1, x_2|x^2; y^l\right)$. As more patient sellers arrive, they post limit sell orders at lower and lower prices. The sell side of the book becomes full when, if one additional patient uninformed trader were to arrive and submit a limit order, the equilibrium price according to the gap formula would have violated the break-even condition. Any patient seller who arrives when the book is full would leave the market without submitting a limit order.

The condition in equation (6) could hold with a strict inequality when the book becomes full. A natural question in this case is why limit sell order $x_{F_{S+1}}$does not undercut limit sell order x_{F_S} by a small amount rather than leave the market. The answer is that any attempt to undercut it by less than $g_t\left(x_{F_S}, x_{F_S+1}|x^{F_S+1}; y^l\right)$ would not be an equilibrium. The other limit order traders would have an incentive to reprice their orders (the rationalizability argument) repeatedly. This process of repricing would stop and an equilibrium would prevail only if each order is priced such that the expected gain in utility by undercutting and obtaining execution priority is equal to the expected loss in utility from the lower price, which is described by the price gap function. An arriving patient trader can think through this chain of events, and his equilibrium strategy is simply to leave because he would not obtain higher expected utility in equilibrium by attempting to submit a limit order. Propositions 1 and 2 completely characterize the optimal strategies of patient uninformed traders.[16]

The trading strategy of the informed traders is driven by three considerations that reflect the information environment of the GMEO framework. First, each trader behaves as if they arrive to the market only once to trade a unit order, and therefore she would buy if the value of the asset is $v + \sigma$ and sell if the value of the asset is $v - \sigma$.[17] The second consideration is the assumption on off-equilibrium beliefs: revealing oneself as an informed trader entails a cost that is at least as high as the expected utility from submitting a limit order in a market in which the private information is publicly known. This means that the strategy of an informed trader must mimic the strategy of either a patient or

[16]For a limit order market without informed trading, Proposition 1 and 2 provide an alternative characterization of the equilibrium described in Rosu (2009). The main advantage of our approach is that it is constructive, and we can therefore calculate the prices that would prevail in the limit order book in equilibrium.

[17]As in traditional sequential trade models, this assumption eliminates any incentives to manipulate prices by trading in the opposite direction to her information.

an impatient uninformed trader in order not to be revealed. We make this assumption to replicate the essence of the GMEO framework in which informed traders hide in the order flow that comes from uninformed traders. The third consideration is that informed traders do not bear waiting costs. This means that, holding everything else equal, an informed trader would prefer to submit a limit order rather than a marketable order.

Therefore, if an informed trader who observes $v - \sigma$ arrives in the market and the sell side of the book is not full, she would submit a limit sell order with the same price that a patient uninformed trader would choose (satisfying the price gap equation) so that she is not revealed. If the same informed trader arrives in the market and the sell side of the book is full, she cannot submit a limit sell order without being revealed. She would then submit a marketable buy order if buy side of the book is not empty, and leave the market without trading if the buy side of the book is empty.[18] Thus, while informed traders in the GMEO framework submit marketable orders to hide behind the uninformed traders, here they choose to hide behind either patient or impatient uninformed traders by submitting either limit or marketable orders depending on what is more profitable for them given the state of the book.

Characterizing the strategy of the informed traders allows us to work out the equilibrium expected arrival rates of marketable and limit orders from the perspective of uninformed traders given various configurations of the book.

Proposition 3. *At time t,*

1. *If the buy side of the book is not full, uninformed traders expect the arrival rate of limit buy orders that enter the book to be $\lambda_{bp} + q_t \lambda_I$, and the arrival rate of marketable buy orders that execute limit sell orders to be λ_{bi}. If the buy side of the book is full, uninformed traders expect the arrival rate of limit buy orders that enter the book to be λ_{bp}, and the arrival rate of marketable buy orders that execute limit sell orders to be $\lambda_{bi} + q_t \lambda_I$.*
2. *If the sell side of the book is not full, uninformed traders expect the arrival rate of limit sell orders that enter the book to be $\lambda_{sp} + (1 - q_t)\lambda_I$, and the arrival rate of marketable sell orders that execute limit buy orders to be λ_{si}. If the sell side of the book is full, uninformed traders expect the arrival rate of limit sell orders that enter the book to be λ_{sp}, and the arrival rate of marketable sell orders that execute limit buy orders to be $\lambda_{si} + (1 - q_t)\lambda_I$.*

[18] The informed traders' strategy is part of the proof of Proposition 3 in the Appendix.

Proposition 3 formalizes the result that uninformed traders use their beliefs about the value of the asset together with an understanding of the trading strategy of the informed traders to infer the arrival rates of each type of order. Having derived the expected arrival rates, we can now demonstrate how these can be used to calculate the wait times of orders (and hence limit order prices) in order to complete the construction of the equilibrium in the instantaneous game $\mathcal{G}_t$.

4.3 A Recursive Formulation for Wait Times

One of the innovations in this paper is that we present a recursive formulation that can be used to compute the wait times for limit orders in the book. This recursive formulation works in limit order books with and without asymmetric information. In fact, it is simpler in markets without informed traders because the arrival rates of limit orders are independent of both time and the state of the other side of the book. For such markets, we show in the Online Appendix that our formulation for wait times in limit order markets reduces to a single recursive equation. The introduction of informed traders, however, means that our calculation of wait times needs to take into account every possible state of the other side of the book. This necessitates a multi-equation recursive system, which is somewhat more complex. At the end of the day, however, the exercise of solving for the wait times of limit orders on each side of the book is reduced to a system of $F + 1$ linear equations in $F + 1$ (where the number of equations is determined by how many orders can rest on the other side of the book) at each stage of the recursion, which means that the system can always be solved. Below we describe how the recursive process works, starting from the best price in the book and going towards the back of the book.

What is the wait time for execution of the lowest-priced limit sell order when the sell side of the book is full? The answer depends on the configuration of the other side of the book, and hence to answer this question we need to find the wait time for any possible such configuration. We discuss in great detail the first couple of steps in the derivation of these wait times to demonstrate the recursive logic that is at the core of our method. Once this logic is clear, the same process can be applied to solve for any wait time on both sides of the book. To economize on notation, we use $\sum \lambda = \lambda_{bi} + \lambda_{si} + \lambda_{bp} + \lambda_{sp} + \lambda_I$ to denote the sum of the arrival rates of all trader types. The expected time for the arrival of the next trader is $\frac{1}{\sum \lambda}$, and we can write the probability that this trader is, for example, an impatient uninformed buyer as $\frac{\lambda_{bi}}{\sum \lambda}$.

In the first step, we look for this wait time when the buy side of the book is also full, $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)$. It can be written as the sum of four terms depending on which order arrives next (marketable buy, marketable sell, limit buy, or limit sell), and we show each term in a separate line of equation (8) below. The first line of this equation tells us that, after an expected time interval of $\frac{1}{\Sigma\lambda}$, a marketable buy order arrives with probability $\frac{\lambda_{bi}+q_t\lambda_I}{\Sigma\lambda}$. The marketable order will come from either impatient uninformed traders (λ_{bi}) or from informed traders when the value of the asset is high ($q_t\lambda_I$), because the buy side of the book is full. If the marketable buy order arrives, our limit sell order executes and we are done. The second line of the equation tells us that, after an expected time interval of $\frac{1}{\Sigma\lambda}$, there is a probability of $\frac{\lambda_{si}+(1-q_t)\lambda_I}{\Sigma\lambda}$ that a marketable sell order will arrive that can come from either impatient uninformed traders (λ_{si}) or informed traders when the value of the asset is low ($(1-q_t)\,\lambda_I$) because the sell side of the book is full. In this case, a limit buy order is executed (the buy side of the book moves from being full to being not full with F_B-1 orders), and our limit sell order still needs to wait $w_t\left(x_{F_S}|x^{F_S};y^{F_B-1}\right)$ for execution.

The third line of the equation deals with the possibility that a patient uninformed buyer arrives with probability $\frac{\lambda_{bp}}{\Sigma\lambda}$. In this case, he leaves the market without submitting an order because the buy side of the book is full, and our limit sell order still has to wait $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)$ for execution. Lastly, the fourth line of the equation describes the possibility that a patient uninformed seller arrives with probability $\frac{\lambda_{sp}}{\Sigma\lambda}$. In this case, he leaves the market without submitting an order because the sell side of the book is full, and our limit sell order still needs to wait $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)$ for execution.

$$
\begin{aligned}
w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right) = & \left(\frac{\lambda_{bi}+q_t\lambda_I}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}\right) \\
& + \left(\frac{\lambda_{si}+(1-q_t)\lambda_I}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}+w_t\left(x_{F_S}|x^{F_S};y^{F_B-1}\right)\right) \\
& + \left(\frac{\lambda_{bp}}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}+w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)\right)
\end{aligned}
$$

$$+\left(\frac{\lambda_{sp}}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}+w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)\right). \tag{8}$$

These four terms exhaust all possibilities, and we can rearrange the terms to show that this is a simple linear equation for $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)$ in terms of $w_t\left(x_{F_S}|x^{F_S};y^{F_{B-1}}\right)$:

$$w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)=\frac{1}{\lambda_{bi}+\lambda_{si}+\lambda_I}+w_t\left(x_{F_S}|x^{F_S};y^{F_{B-1}}\right)\left(\frac{\lambda+(1-q_t)\,\lambda_I}{\lambda_{bi}+\lambda_{si}+\lambda_I}\right) \tag{9}$$

The next step, therefore, would be to calculate $w_t\left(x_{F_S}|x^{F_S};y^{F_B-1}\right)$, which represents the case that the configuration of the buy side of the book is not full but rather has room for one more limit order. Enumerating the four possibilities yields the following equation:

$$\begin{aligned}
w_t\left(x_{F_S}|x^{F_S};y^{F_B-1}\right)=&\left(\frac{\lambda_{bi}}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}\right)\\
&+\left(\frac{\lambda_{si}+(1-q_t)\lambda_I}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}+w_t\left(x_{F_S}|x^{F_S};y^{F_B-2}\right)\right)\\
&+\left(\frac{\lambda_{bp}+q_t\lambda_I}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}+w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)\right)\\
&+\left(\frac{\lambda_{sp}}{\Sigma\lambda}\right)\left(\frac{1}{\Sigma\lambda}+w_t\left(x_{F_S}|x^{F_S};y^{F_B-1}\right)\right)
\end{aligned} \tag{10}$$

This equation has the same structure as equation (8) except now a marketable buy order in the first line can only come from impatient uninformed buyers ($\frac{\lambda_{bi}}{\Sigma\lambda}$) because the buy side of the book is not full. A marketable sell order in the second line can still come from both impatient uninformed and informed traders (because the sell side of the book is full), executing a limit buy order and leaving F_B-2 orders on the buy side of the book. A limit buy order in the third line can come from both patient uninformed traders and informed traders (because the buy side of the book is not full) with probability $\frac{\lambda_{bp}+q_t\lambda_I}{\Sigma\lambda}$, entering the buy side of the book and making it full again. The fourth line of the equation describes the possibility that a patient uninformed seller arrives with probability $\frac{\lambda_{sp}}{\Sigma\lambda}$. In this case, he leaves the market without submitting an order, and our limit sell order still needs to wait $w_t\left(x_{F_S}|x^{F_S};y^{F_B-1}\right)$ for execution.

Our wait time $w_t\left(x_{F_S}|x^{F_S};y^{F_B-1}\right)$ can therefore be expressed as a function of $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)$, which we already had in equation (9), and a new wait time, $w_t\left(x_{F_S}|x^{F_S};y^{F_B-2}\right)$. We could then create an equation for $w_t\left(x_{F_S}|x^{F_S};y^{F_B-2}\right)$, similar in structure to equation (10), that would yield the wait time as a function of the one we had in the previous equation and a new term, $w_t\left(x_{F_S}|x^{F_S};y^{F_B-3}\right)$. This process is repeated for any number of orders $F_B - i$ on the buy side of the book, $0 < i < F_B$, and the resulting equation can be simplified to obtain

$$w_t\left(x_{F_S}|x^{F_S};y^{F_B-i}\right) = \frac{1}{\lambda_{si}+\lambda_{bp}+\lambda_{bi}+\lambda_I} + w_t\left(x_{F_S}|x^{F_S};y^{F_B-i-1}\right)\left(\frac{\lambda_{si}+(1-q_t)\lambda_I}{\lambda_{si}+\lambda_{bp}+\lambda_{bi}+\lambda_I}\right) + w_t\left(x_{F_S}|x^{F_S};y^{F_B-i+1}\right)\left(\frac{\lambda_{bp}+q_t\lambda_I}{\lambda_{si}+\lambda_{bp}+\lambda_{bi}+\lambda_I}\right). \tag{11}$$

The last equation we need is therefore for $i = 0$, which means the buy side of the book is empty. This equation is simpler because an arriving marketable sell order does not lead to a change in the (empty) buy side of the book, and it can be rearranged to give

$$w_t\left(x_{F_S}|x^{F_S};y^{0}\right) = \frac{1}{\lambda_{bi}+\lambda_{bp}+q_t\lambda_I} + w_t\left(x_{F_S}|x^{F_S};y^{1}\right)\left(\frac{\lambda_{bp}+q_t\lambda_I}{\lambda_{bi}+\lambda_{bp}+q_t\lambda_I}\right) \tag{12}$$

The set of equations in (9), (11), and (12) is a system of $F_B + 1$ linear equations in $F_B + 1$ unknowns, and therefore it has a unique solution for the values of $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right), \ldots,$ $w_t\left(x_{F_S}|x^{F_S};y^{0}\right)$. Given any possible state of the other side of the book, we have now found the wait time for the execution of the ask price in the market when the sell side of the book is full.

We construct the wait time for all other orders on the sell side of the book in a similar fashion, starting with the wait time for the limit sell order in the slot just above the ask price, $w_t\left(x_{F_S-1}|x^{F_S-1};y^{F_B}\right)$. The construction is very similar to that of equation (8) with two differences. The first difference is that the sell side of the book is not full, which means informed traders who observe a low value would use a limit order rather than a marketable order. This means that marketable sell orders arrive with probability $\frac{\lambda_{si}}{\Sigma\lambda}$ and limit sell orders arrive with probability $\frac{\lambda_{sp}+(1-q_t)\lambda_I}{\Sigma\lambda}$. The second difference is that if a limit sell order arrives, it becomes the new ask in the market, and hence we first need to wait $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)$ for it to execute for our order to become again the the ask price. Since we have already found $w_t\left(x_{F_S}|x^{F_S};y^{F_B}\right)$ by solving the set of equations in (9), (11), and (12), we get a similar recursive structure as before, which simplifies to

$$\begin{aligned} w_t\left(x_{F_S-1}|x^{F_S-1};y^{F_B}\right) =& \frac{1}{\lambda_{bi}+\lambda_{si}+q_t\lambda_I} + w_t\left(x_{F_S-1}|x^{F_S};y^{F_B}\right)\frac{\lambda_{sp}+(1-q_t)\lambda_I}{\lambda_{bi}+\lambda_{si}+q_t\lambda_I} \\ &+ w_t\left(x_{F_S-1}|x^{F_S-1};y^{F_B-1}\right)\frac{\lambda_{si}}{\lambda_{bi}+\lambda_{si}+q_t\lambda_I} \end{aligned} \tag{13}$$

Solving a system of F_B+1 linear equations in F_B+1 unknowns provides us with a unique solution for wait times $w_t\left(x_{F_S-1}|x^{F_S-1};y^{F_B}\right), \ldots, w_t\left(x_{F_S-1}|x^{F_S-1};y^{0}\right)$.

The same structure can be used to solve for all wait times in every level of the book on both sides. The only three variables we need in order to implement this method are q_t, F_B, and F_S. The first one, the beliefs of the uninformed traders, is known at time t. The other two variables are found from the break-even conditions in Proposition 2 using the guess-and-verify method or simply by starting with $F_B = 1$, and $F_S = 1$ and continuing to increment on each side until the break-even conditions are violated.

The recursive systems of equations seem a bit complex, but they are straightforward to implement and give the exact equilibrium wait times and limit order prices $(p_i)_{i\in x_t}$ and $(p_k)_{k\in y_t}$ that we require for the equilibrium in Definition 1. The trading strategies of the patient uninformed and informed traders that we use to derive the equilibrium price gaps and arrival rates take care of the equilibrium requirements that traders have no incentives to reprice their limit orders or change their choices. This completes the construction of the equilibrium in the instantaneous game $\mathcal{G}_t$.

Several important questions are still left unanswered by the construction of equilibrium in the instantaneous game. For example, how do traders learn about the true value? What happens to prices in the book as orders arrive and bring new information? Can new information cause limit order traders to leave the market without trading? What information is known to traders when trade execution occurs? To answer these questions, we proceed in the next section to discuss the equilibrium in the long-horizon game $\mathcal{G} = (\mathcal{G}_t)_{t\in[0,+\infty)}$.

5 Equilibrium in the Long-Horizon Game: Learning from Order Flow[19]

The definition of equilibrium in the long-horizon game $\mathcal{G}$ borrows from the Perfect Bayesian Equilibrium concept the requirement that beliefs are updated using Bayes rule. It also specifies that traders update their choices, canceling and resubmitting their limit orders or leaving the market altogether, in accordance with these updated beliefs to effect the change from $\mathcal{G}_{t'}$ to $\mathcal{G}_{t''}$, $t'' > t'$. In Section 5.1, we show how uninformed traders' beliefs are updated when orders arrive, and present empirical implications on the manner in which the state of the book affects the permanent price impact of orders. Section 5.2 discusses what the change in beliefs means for market interactions, and Section 5.3 takes a closer look at how time itself may impact the repricing of limit orders, providing additional empirical implications on this process.

5.1 Bayesian Updating of Uninformed Traders' Beliefs

Proposition 3 implies that uninformed traders know in equilibrium which order type an arriving informed trader would use by observing the configuration of each side of the book (full versus not full). What uninformed traders do not know, however, is whether the informed trader would want to buy or sell, which is why they employ Bayesian learning to update their beliefs about the asset's value. The order type and direction of an arriving order, in conjunction with the state of the book, are used to determine whether beliefs are updated and how. Depending on the parameters of the economy, the passage of time without the arrival of an order may also be used by the uninformed traders to update their beliefs. In this section we first address how traders learn from the arrival of orders.

As a concrete example of the Bayesian updating process when an order arrives, consider the case in which the buy side of the book is full and uninformed traders hold the prior q_t when a marketable buy order (MB) arrives at time t. To see how the arrival of the order would convey information to the limit order market, we use Bayes' Rule to calculate

[19]We thank Bart Yueshen and Ciamac Moallemi for insightful discussions that resulted in substantial revisions to this section.

the uninformed traders' posterior beliefs:

$$\mathrm{P}_t\left[\tilde{v}=v+\sigma|MB\right]=\frac{\mathrm{P}_t\left[MB|\tilde{v}=v+\sigma\right]\mathrm{P}_t\left[\tilde{v}=v+\sigma\right]}{\mathrm{P}_t\left[MB|\tilde{v}=v+\sigma\right]\mathrm{P}_t\left[\tilde{v}=v+\sigma\right]+\mathrm{P}_t\left[MB|\tilde{v}=v-\sigma\right]\mathrm{P}_t\left[\tilde{v}=v-\sigma\right]}. \tag{14}$$

Both impatient uninformed traders and informed traders submit marketable buy orders when the buy side of the book is full and $\tilde{v}=v+\sigma$, while only impatient uninformed traders submit marketable buy orders when $\tilde{v}=v-\sigma$. Hence, $\mathrm{P}_t\left[MB|\tilde{v}=v+\sigma\right]=\frac{\lambda_{bi}+\lambda_I}{\Sigma\lambda}$ and $\mathrm{P}_t\left[MB|\tilde{v}=v-\sigma\right]=\frac{\lambda_{bi}}{\Sigma\lambda}$. Substituting these values into equation (14) above, we get the posterior beliefs

$$\mathrm{P}_t\left[\tilde{v}=v+\sigma|MB\right]=\frac{q_t\left(\lambda_{bi}+\lambda_I\right)}{\lambda_{bi}+q_t\lambda_I}. \tag{15}$$

Similarly, the posterior beliefs conditional on the arrival of a marketable sell order are

$$\mathrm{P}_t\left[\tilde{v}=v+\sigma|MS\right]=\frac{q_t\lambda_{si}}{\lambda_{si}+(1-q_t)\lambda_I}. \tag{16}$$

The following proposition describes how limit order prices in the book change as beliefs are updated.

Proposition 4. *If the buy (sell) side of the book is full, arrival of a marketable buy (sell) order causes the posterior beliefs about the asset value to be go up (down) according to equation (15) (equation (16)), which leads to repricing of limit orders in the book such that gaps between resting buy (sell) orders widen and gaps between resting sell (buy) orders narrow. If the buy (sell) side of the book is not full, arrival of a marketable buy (sell) order does not cause a change in the beliefs of uninformed traders, but arrival of a limit buy (sell) order causes analogous effects on posterior beliefs (replacing λ_{bi} with λ_{bp} in equation (15) and λ_{si} with λ_{sp} in equation (16)) and price gaps.*

Not every order arrival results in beliefs updates. Orders that could possibly come from informed traders (e.g., limit orders when the same-side-book is not full) cause all limit order traders in the book to cancel and resubmit their orders at different prices. This repricing of limit orders in the book is driven by the revision in the uninformed traders' beliefs about the true value, which changes the inferred arrival rates of orders specified in Proposition 3. This leads to a revision in the wait times and the magnitude of equilibrium price gaps, resulting in limit orders being repriced.

Formally, the set of equilibrium limit order prices in the instantaneous game $\mathcal{G}_{t'}$ is determined by the uninformed traders' beliefs at t'. As beliefs are updated in the infinite horizon game $\mathcal{G}$ with the arrival of an order that could come from an informed trader, a new instantaneous game $\mathcal{G}_{t''}$ with $t'' > t'$ will have a different set of equilibrium prices. The transition between $\mathcal{G}_{t'}$ and $\mathcal{G}_{t''}$ implies that traders cancel and resubmit, or reprice, their limit orders in the book.[20]

Proposition 4 gives rise to empirical implications concerning order-flow information in limit order markets. One such implication is that the informativeness of marketable and limit orders depends on the size of the prevailing spread. When the book is full and therefore the bid-ask spread is narrow, informed traders use marketable orders. When the book is not full and the bid-ask spread is therefore wider, informed traders use limit orders. The size of the spread influences the desire of the patient uninformed traders to supply liquidity. The informed traders' order choice is driven by their need to hide in the order flow of the uninformed traders, and therefore the size of the spread affects the informativeness of incoming orders.

Empirical Implication 1. *Marketable orders have a larger (smaller) permanent price impact when the spread is narrow (wide). Limit orders have a larger (smaller) permanent price impact when the spread is wide (narrow).*

In general, testing the implications of any model involves making choices on the appropriate design for the empirical analysis, and we want to comment briefly on several relevant considerations. The first point is that Empirical Implication 1 is a time-series, not a cross-sectional, implication. Hence, by a "narrow" or "wide" spread we mean relative to the time-series properties of spreads. The second point is that, while we model a continuous price grid, empirical analysis needs to account for the minimum price increment. For example, stocks that are tick-constrained (i.e., trading with a one-tick spread) will more often exhibit the state of the book in which informed traders cannot undercut resting limit orders and therefore are more likely to submit marketable orders. As a re-

[20]Release of public information (e.g., an earnings announcement) will also result in repricing of limit orders in the book as prices in the new instantaneous game $\mathcal{G}_{t''}$ reflect changes in the fundamental attributes of the asset's true value (e.g., v and σ).

sult, there may be cross-sectional patterns in the use of marketable and limit orders by informed traders that depend on market structure frictions.

The third point is that the size of the spread does not have a one-to-one correspondence with whether each side of the book is full or not. The inference problem arises when one side of the book is full and the other side is not full. One could reduce the potential misclassification somewhat by introducing a sufficiently wide wedge between the two spread size categories, but misclassifications cannot be eliminated completely simply because the size of the spread only imperfectly distinguishes the four states of the book in our model (i.e., full on both sides, not full on both sides, or a mixture with one side full and the other side not full).

To try and reduce misclassifications, one could condition not just on the size of the spread but also on the frequency of orders in a given interval. The arrival rate of information-bearing orders is higher than that of non-information-bearing orders because it is the sum of uninformed and informed arrival rates. As such, if the narrow spread is not narrow enough to coax informed traders to submit marketable orders, for example, the frequency of marketable orders and their information content would be lower. This is summarized in the following empirical implication.

Empirical Implication 2. *When the spread is narrow (wide), a higher frequency of marketable (limit) orders should result in a larger permanent price impact than a lower frequency of marketable (limit) orders.*

To identify separately whether one or the other side of the book is full, one could look at depth relative to its time-series properties. Holding everything else equal, the buy or sell side of the book will have more shares (or limit orders) when it is full than when it is not full. Our model implies that a marketable order will convey more information when the same-side book has more depth and a limit order will convey more information when the same-side book has less depth.

Empirical Implication 3. *A marketable order will result in a larger permanent price impact when there is more depth than usual on the same side of the book than when there is less depth on that side. A limit order will result in a larger permanent price impact when there is less depth than usual on the same side of the book than when there is more depth on that side.*

We know of two papers that contain evidence consistent with these empirical implications. Putnins and Michayluk (2018) show increased limit order use by informed traders when spreads are wide, consistent with Empirical Implication 1. Kwan, Philip, and Shkilko (2021) empirically test the predictions of our model using machine learning techniques applied to limit order book data from the Australian Securities Exchange. They find that, consistent with Empirical Prediction 1, price discovery through marketable orders increases when the spread is narrow, while price discovery through limit orders increases when the spread is wide. They also find that, consistent with Empirical Prediction 2, marketable orders arriving frequently have larger price impacts compared to marketable orders arriving infrequently when the spread is narrow, and a similar result can be observed for limit orders when the spread is wide.

It is important to emphasize that our empirical implications involve the permanent price impact, not the immediate price impact, of orders. Empirical studies in market microstructure often define the permanent price impact of a sell order as the midquote prevailing at the time the order arrives minus the midquote after a certain interval of time that is sufficient for the reversal stemming from the temporary price impact to take effect. How would the permanent price impact manifest in our market? Say, for example, both sides of the book are not full. The arrival of a limit sell order that improves the best ask will result in an immediate price impact, lowering the midpoint between the bid and ask prices. This immediate price impact can be viewed as the sum of two components: the price revision that occurs in the prices of all resting limit orders to reflect the information content of the order (given that the book is not full), and a temporary effect that arises mechanically from the addition of a new limit order below resting limit orders.

If a marketable buy order arrives after the limit sell order and executes it, there would be no repricing of resting limit orders because the buy side of the book is not full. The midquote following the execution will rise mechanically because the previous limit sell order was executed, reversing its temporary price impact. However, the other resting limit sell orders would be positioned at a lower price than they were initially, reflecting the updated beliefs of uninformed traders brought about by the arrival of the limit sell order, which is how the permanent price impact is manifested.

Therefore, the process by which a permanent price impact emerges in a limit order market is more complex than that described in traditional sequential trade models of dealer markets. In the Glosten-Milgrom-Easley-O'Hara framework, all price impacts are permanent (i.e., there is no distinction between immediate and permanent price impacts). Incorporating informed traders in a limit order market gives rise to an immediate price impact that is comprised of both permanent and temporary elements, and thus comes closer to describing empirically-observed price patterns. This also implies that, unlike in traditional sequential trade models, transaction prices in our market do not converge to the true value because they are not expectations of the value of the risky asset. Rather, transaction prices would deviate from the true value even as beliefs converge due to waiting costs and the need to compensate liquidity suppliers.[21]

The empirical implications in this section focus on how the state of the book impacts the information content of arriving orders. The arrival of orders, in turn, brings about changes to the limit order book. The next section discusses how this process unfolds.

5.2 Changes in Beliefs and Market Interactions

In this section, we describe in more detail what the change in beliefs means for market interactions. In particular, we discuss the sequence of events that unfolds when an order arrives, the relationship between traders and orders, and what happens when the equilibrium number of orders goes down after an update in beliefs.

It is useful to think in terms of a sequence of steps that happen instantaneously when when beliefs change. The first step is that the new location of all limit orders is computed based on the information in the incoming order. This step essentially finds the equilibrium prices in the new instantaneous game. The second step is that traders use whatever protocol they agree on to decide who is in which slot in the book after repricing. In the third step, if the incoming order is a limit order, the new order is added to the book at a price distance computed using the new beliefs; if the incoming order is a marketable

[21]As uninformed traders update their beliefs, the book on one side will become empty at some point when the expected value approaches the true value. For example, if the expected true value v_t comes very close to $v - \sigma$, patient uninformed traders will stop supplying liquidity on the buy side of the book because there will be no limit prices that could satisfy the break-even condition. However, informed traders and patient uninformed sellers will continue to submit limit sell orders to the book, and trading will continue as impatient uninformed buyers submit their marketable orders. Thus, the arrival of limit sell and marketable buy orders will convey information about the true value and the convergence of the expected value of the asset to $v - \sigma$ will continue.

order, the limit order with the lowest price on the other side of the book is executed. This step is simply a manifestation of the limit order market rules. While the series of steps is straightforward, three aspects nonetheless deserve additional discussion: the sequencing, the relationship between traders and specific slots in the book that is alluded to in the second step, and the possibility that traders cancel their orders and leave the market.

The sequencing of the steps specifies that uninformed traders first update their beliefs upon the arrival of an order, and then the order is processed according to the market rules. This seems like a natural sequence because, if a limit order were to arrive and be placed in the book in a slot based on the old information, then immediately afterwards it would need to be canceled and resubmitted to reflect the new information. Given that these two things happen instantaneously, it seems reasonable to update beliefs first and then post the order. If a marketable order arrives when the same-side-book is full, this sequence of events means that traders in the book reprice their orders and therefore the bid and ask prices change, and only then the marketable order executes at a price that reflects its information content. This is similar in nature to the outcome in traditional sequential trade models where execution prices are "regret-free" in the sense that they reflect the information content of the arriving order.

The second issue that merits discussion is the relationship between specific traders and slots in the book. The equilibrium definition of the instantaneous game $\mathcal{G}_t$ is done in terms of a set of limit order prices $(p_i)_{i \in x_t}$ and $(p_k)_{k \in y_t}$, but it does not require an identification of a particular limit order price with a specific trader. Our uninformed patient traders are indifferent between the slots because their expected utility at each slot is the same in equilibrium.[22] By indifferent we mean that patient uninformed traders could choose any protocol they like to work out the priorities (or who is in which slot) following a cancellation and resubmission—we do not need to take an explicit stand on the nature of the protocol. For example, one such protocol could be that each trader on the sell side of the book sticks to his original slot after repricing, and a new limit order is added below the ask price to create the new ask. However, our model can accommodate any protocol on which the patient uninformed traders agree. Informed traders in the book must follow the protocol chosen by the patient uninformed traders in order not to be revealed as informed. Hence, all traders with limit orders in the book—both uninformed and informed—follow

[22]This is also a key feature of the model in Rosu (2009).

the same protocol. The exact nature of the protocol, or said differently, the identity of the trader associated with each slot, is immaterial to the determination of equilibrium prices in the book.

The third issue that requires clarification is what happens if an incoming order changes beliefs in such a way that a full book in the new equilibrium contains fewer orders than the number of orders resting in the book. For example, say the sell side of the book is full and consists of 20 orders, while the buy side of the book is not full. The arrival of a limit buy order will cause an upward revision in the expected value of the asset, and it could happen that the updated break-even condition and price gaps in the new equilibrium allow for only 19 orders on the sell side of the book. In such a case, the limit sell order with the lowest price would be canceled and the trader would leave the market. Why would the trader agree to leave the market without trading? A patient uninformed trader would leave because remaining in the book would imply negative expected utility and he can have zero expected utility by leaving the market.[23] An informed trader would need to follow exactly the same protocol as the patient uninformed traders if she wishes not to be revealed, and hence would also leave the market without trading.[24]

5.3 Information and the Passage of Time

Beyond the arrival of orders, the passage of time itself without the arrival of an order may cause uninformed traders to update their beliefs. It is useful to discuss time in the context of a slightly more general model in which the intensity of informed trading when buying

[23]Which trader has to leave is part of the same protocol that patient uninformed traders agree on, and it can simply be that the trader who happens to be in the slot with the lowest limit sell order price that breaches the break-even condition is the one leaving the market.

[24]She would not cancel the limit order and simultaneously submit a marketable order because this would also reveal her to the market (marketable orders otherwise arrive according to a Poisson process). As we discuss in the Online Appendix, revelation of the true value essentially causes one side of the book to become empty in the relevant price range ($[v - \sigma, v + \sigma]$). For example, a revelation that the value is $v - \sigma$ means that the break-even condition for the patient uninformed buyers would be violated for all prices above $v - \sigma$, hence they would cancel all their limit buy orders. The sequence of steps whereby prices adjust to information before execution takes place means that the informed trader's marketable order would therefore not execute. While the informed trader could in principle have placed a limit sell order at a price that reflects symmetric information about the true value to gain from providing liquidity to impatient uninformed buyers, our assumption about off-equilibrium beliefs implies that this is a worse outcome for her than simply leaving the market without trading.

can differ from their intensity when selling. Such a situation may arise in equity markets, for example, because of short-sale constraints (as in Diamond and Verrecchia (1987)) or because holding a short position in the asset exposes the informed traders to the risk of being short-squeezed. We denote the arrival rate of informed traders when they observe a high asset value λ_{IH} and their arrival rate when observing a low asset value λ_{IL}.

Let N_{MB}, N_{MS}, N_{LB} and N_{LS} denote the number of marketable buy, marketable sell, limit buy, and limit sell orders, respectively, that arrive in an interval $\delta > 0$. When the asset value is $v + \sigma$, informed traders will buy using either marketable or limit orders, and the assumption that trader types arrive according to independent Poisson processes means that the probability that no order arrives during the interval is therefore $\mathrm{P}_t \left[N_{MB} = N_{MS} = N_{LB} = N_{LS} = 0 | \tilde{v} = v + \sigma\right] = e^{-\delta\left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IH}\right)}$. Similarly, informed traders would sell using either marketable or limit orders when $v - \sigma$, and the probability that no order arrives during the interval is $\mathrm{P}_t \left[N_{MB} = N_{MS} = N_{LB} = N_{LS} = 0 | \tilde{v} = v - \sigma\right]$ $= e^{-\delta\left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IL}\right)}$. Therefore, the change in beliefs over the interval δ is

$$q(t+\delta) - q(t) = \frac{q(t)(1-q(t))\left[e^{-\delta\left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IH}\right)} - e^{-\delta\left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IL}\right)}\right]}{q(t)e^{-\delta\left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IH}\right)} + (1-q(t))e^{-\delta\left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IL}\right)}}$$

Given that the Poisson process is right-continuous and that there are no arrivals during the interval, we can take the right limit as δ goes to zero to obtain

$$\begin{aligned} q'(t) =& \lim_{\delta \to 0} \frac{q(t+\delta) - q(t)}{\delta} \\ =& -q(t)(1-q(t))\left[\left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IH}\right) - \left(\lambda_{bi}+\lambda_{bp}+\lambda_{si}+\lambda_{si}+\lambda_{IL}\right)\right] \\ =& -q(t)(1-q(t))\left(\lambda_{IH}-\lambda_{IL}\right). \end{aligned} \tag{17}$$

Thus, beliefs are updated continuously to reflect the absence of arrival. Say, for example, that some informed traders are not willing to risk a short squeeze and hence $\lambda_{IH} > \lambda_{IL}$. Given that uninformed traders expect more orders to arrive when informed traders buy than when they sell, the absence of an arrival means a lower likelihood that the value of the asset is $v + \sigma$ and hence beliefs drift downward. The discussion above in conjunction with Proposition 4 gives rise to the following result.

Proposition 5. *When $\lambda_{IH} > \lambda_{IL}$ ($\lambda_{IH} < \lambda_{IL}$), the passage of time without the arrival of an order implies that q_t is updated downward (upward), which leads to repricing of limit orders in the book such that gaps between resting sell (buy) orders widen and gaps between resting buy (sell) orders narrow. If $\lambda_{IH} = \lambda_{IL}$, the passage of time without the arrival of an order does not change q_t.*

The repricing of limit orders follows the first two steps from Section 5.2: the new location of all limit orders is computed based on the updated beliefs, and then traders reprice their orders following whatever protocol they agree on to decide who is in which slot in the book after repricing.

To simplify the exposition in the paper, we assumed in the previous sections that the informed traders' arrival rate is the same when they buy and when they sell. However, the equilibrium in the instantaneous game and the recursive procedure we introduce to compute the expected wait times can also be implemented when informed traders have different buying and selling arrival rates.[25] It is clear from looking at equation (17) that when $\lambda_{IH} = \lambda_{IL} = \lambda_I$, the drift disappears and beliefs remain the same without an order arrival. The intuition for this is simple: the total arrival rate of orders is the same irrespective of whether the asset value is $v + \sigma$ or $v - \sigma$, and hence the lack of arrival cannot tell us anything about the value of the asset. Only when the total arrival rate of orders differs depending on whether the asset value is high or low is the lack of arrival informative.

Hypotheses about the differential arrival rates of informed traders can be tested empirically by looking at repricing (cancellations and resubmissions) of limit orders in the book when no new orders are arriving. Since only some order types are informative in our model (depending on the state of the book), the drift in beliefs only occurs when informative orders do not arrive, while beliefs do not change with the arrival of orders that are known to come only from uninformed traders. This enable us to provide more refined hypotheses from combining Proposition 4 and Proposition 5.

[25]The derivation of the expected wait time in the more general case with $\lambda_{IH} \neq \lambda_{IL}$ is provided in the Online Appendix.

Empirical Implication 4. *When the spread is narrow, the passage of time without the arrival of marketable orders should result in lowering (raising) the prices of limit orders in the book if the intensity of informed trader buying is higher (lower) than their intensity of selling. When the spread is wide, the same effects should be associated with the passage of time without the arrival of limit orders.*

As in Empirical Implication 3, greater depth on a particular side of the book can be used to indicate a full book, resulting in the following empirical implication.

Empirical Implication 5. *When there is more depth than usual on both sides of the book, the passage of time without the arrival of marketable orders should result in lowering (raising) the prices of limit orders in the book if the intensity of informed trader buying is higher (lower) than their intensity of selling. When there is less depth than usual on both sides of the book, the same effects should be associated with the passage of time without the arrival of limit orders. If there is more depth than usual on one side of the book and less depth than usual on the other side of the book, the same effects should be associated with the passage of time without the arrival of marketable orders from the side of the book with more depth and limit orders from the side of the book with less depth.*

The repricing of limit orders in the book as informative orders arrive according to Proposition 4 and the repricing caused by the passage of time without the arrival of an order imply a great amount of cancellations and resubmissions of limit orders as traders reposition their orders to reflect their new beliefs. As such, a limit order market under asymmetric information differs from traditional sequential trade models, where just the best bid and ask prices (or the dealer's quote) change, and also from a limit order market under symmetric information, where limit orders are never repriced but rather rest in the book until they execute. In a model with continuous prices and no cost of repricing, it is clear why repricing would occur continuously. Price discreteness would imply that repricing occurs only when beliefs move by more than a threshold dictated by the minimum price increment. Similarly, costs associated with monitoring, submission, and cancellation of limit orders would also inhibit the frequency of repricing that occurs both with and without the arrival of a new order. Still, the process of updating beliefs in limit order markets under asymmetric information should result in very active markets.

Empirical Implication 6. *Information asymmetry in a limit order market drives frequent cancellations and resubmission of limit orders throughout the book in response to the information conveyed in the order flow.*

This empirical implication may appear to trivially follow from the structure of the model: traders dynamically learn about the asset value and hence reposition their limit orders. Indeed, it is simply an attribute of the equilibrium. We highlight it because many market observers have noted the high frequency with which limit orders are revised in today's limit order markets, and the correspondingly short duration of these orders (see, for example, Hasbrouck and Saar (2009), Hasbrouck and Saar (2013) and the literature on high-frequency trading). Some have contended that frequent cancellations and resubmissions are an abnormality that must be due to nefarious activity on the part of high-frequency traders, and could be detrimental to market integrity. We want to make the point that frequent cancellations and resubmissions are a natural attribute of limit order markets under asymmetric information: it is the optimal response on the part of uninformed limit order traders to the presence of informed trading in the order flow.[26]

6 Conclusions

On the whole, some things remain the same but others change dramatically in a sequential trade model when you replace the dealer with a dynamic limit order market. Order flow still brings information to the market, and the expected value of the asset adjusts toward the true value in a manner that is similar to the adjustment in traditional sequential trade models. Learning about private information from the order flow is an essential property of both market structures simply because uninformed traders rationally learn from their environment. Uninformed traders incur execution costs when submitting marketable orders in the limit order market exactly as they do in a dealer market. For the no-trade theorem to be violated, some traders need to have non-informational reasons to

[26]Back and Baruch (2013) provide another explanation by modeling a game between liquidity suppliers in a limit order market. They focus on a symmetric mixed-strategy equilibrium, and interpret the mixed strategies as a manifestation of the actions of liquidity providers who manage their exposure to undercutting by rapidly canceling their quotes and replacing them with new randomly chosen ones.

trade immediately and these traders pay for the liquidity they demand. Beyond these two results, however, the trading environment looks very different.

First, informed traders use both limit and marketable orders in a manner that depends on the state of the book. If taking liquidity is cheap, they take liquidity; if it is expensive, they supply liquidity. Thus, market participants need to use the state of the book, order direction, and the passage of time in their inference about the asset's value. Second, transaction prices are not simply conditional expectations of the asset's value given the order flow as in the traditional sequential trade models. Rather, they reflect both the updated beliefs of uninformed traders and the compensation that needs to be paid to patient traders to induce them to provide liquidity.

Third, we see that many limit orders, not just those at the top of the book, get revised (canceled and resubmitted) as traders learn information. They need not be revised with every order, but they are revised with every order that could potentially come from an informed trader, and when long stretches of time pass without the arrival of informative orders. These constant revisions are now a ubiquitous feature of electronic limit order markets in which submission and cancellation of orders is (almost) costless, and we provide one rationale for such behavior.

A reasonable question is why we have not observed such frequent cancellations and resubmissions until the last couple of decades. The likely answer is that frictions in the trading environment used to inhibit the rational price revisions of limit orders in response to order flow. Throughout most of the history of U.S. equity markets, for example, the large tick size (an eighth of a dollar) prevented effective repricing of limit orders, and it was rather costly to monitor, cancel, and resubmit limit orders. As frictions were reduced by decimalization and the rapid developments in trading technology, we began observing more such price revisions of limit orders.[27] If many of these price revisions stem from the rational response of uninformed traders to adverse selection, as they do in our model, then enacting policies that discourage them may not be beneficial.

Our goal—to nest a limit order market in a sequential trade model—committed us to focusing on a rather simple setting to provide transparency into the inner workings of

[27] Remaining frictions may be causing limit orders closer to the best prices in the book to be revised more frequently than limit orders deeper in the book because there is less urgency in updating the prices of limit orders that are unlikely to execute soon.

the limit order market. At the end of the day, some implications of the model may seem rather stark. In particular, our results include price revisions of all orders in the book, and perfect identification of the order type that informed traders use in each state of the book. In actual markets, minimum price increments and uncertainty over the parameters of the stock or trader arrival (and hence over what constitutes a “full” book) would work to produce more subtle effects. Still, it is important to identify the implications without these additional frictions to appreciate their importance for the market environment.

Our emphasis on simplicity in exposition also means that the model does not include some features that have been investigated elsewhere. For example, we do not investigate the horizon of private information (Kaniel and Liu 2006), information acquisition (Goettler, Parlour, and Rajan 2009, Rosu 2020), market opening and closing effects (Ricco, Rindi, and Seppi 2018), or information decay costs (Rosu 2020). The benefit of the approach we take is in creating a robust, intuitive framework for limit order markets that could be extended along many dimensions.

Regulatory interventions coupled with technological progress have been transforming markets worldwide, and in particular have propelled the rise of the electronic limit order book market structure. Just as the adverse selection models of dealer-intermediated trading spurred an enormous research effort into the microstructure of dealer markets, we hope that our simple model of limit order markets can help the effort to gain a better understanding of today’s most important market structure.

Appendix

Proof. (Proposition 1). The expected utility of a patient uninformed seller in equation (2) has two components: (i) an expected profit component from price, $E\left[\pi_{t,sp} - v_t\right]$, and (ii) an expected wait time component, $E\left[\left(\tilde{T}_{sp} - t\right)\right]$. If the limit sell order in the n^{th} lowest priority execution slot lowers its price to gain priority over the limit sell order in the $(n+1)^{th}$ slot, the gain in expected utility from the wait time component is $w_t\left(x_n|x^{n+m};y^l\right) - w_t\left(x_{n+1}|x^{n+m};y^l\right)$. At the same time, the loss in expected utility from the price component is equal to the price reduction. Thus, when

$$g_t\left(x_n, x_{n+1}|x^{n+m};y^l\right) \geq w_t\left(x_n|x^{n+m};y^l\right) - w_t\left(x_{n+1}|x^{n+m};y^l\right),$$

the seller in the n^{th} priority slot has no incentive to cancel and resubmit his limit order to claim priority. However, if the inequality above is strict, the seller in the $(n+1)^{th}$ priority slot can improve his expected utility by raising the limit order price. In equilibrium, neither seller should want to cancel and resubmit, and we thus get equation (5). □

Proof. (Proposition 2). We prove the proposition for the sell side book. The arguments for the buy side book are symmetric.

From equation (5), we have

$$p_t\left(x_1|x^2;y^l\right) - p_t\left(x_2|x^2;y^l\right) = w_t\left(x_1|x^2;y^l\right) - w_t\left(x_2|x^2;y^l\right).$$

Since $w_t\left(x_1|x^2;y^l\right) \geq w_t\left(x_1|x^1;y^l\right)$ from inequality (4), we can re-write the equality above as

$$p_t\left(x_1|x^2;y^l\right) - p_t\left(x_2|x^2;y^l\right) \geq w_t\left(x_1|x^1;y^l\right) - w_t\left(x_2|x^2;y^l\right).$$

In equilibrium, a patient uninformed seller at a particular execution slot should have no incentive to reprice his order when a new patient seller places an order in the limit book, which means $p_t\left(x_1|x^1;y^l\right) = p_t\left(x_1|x^2;y^l\right)$. Thus, we can re-write the inequality above as

$$p_t\left(x_1|x^1;y^l\right) - p_t\left(x_2|x^2;y^l\right) \geq w_t\left(x_1|x^1;y^l\right) - w_t\left(x_2|x^2;y^l\right).$$

Using an inductive argument, this can be generalized to

$$p_t\left(x_1|x^1;y^l\right) - p_t\left(x_n|x^n;y^l\right) \geq w_t\left(x_1|x^1;y^l\right) - w_t\left(x_n|x^n;y^l\right).$$

The above inequality shows that if a patient uninformed seller in an empty book lowers his price, the gain in expected utility from having a shorter wait time cannot be greater than the loss in expected utility from having a lower execution price. Hence, in equilibrium a patient uninformed sellers prefers a higher limit order price to a lower limit order price, and since $v + \sigma$ is the highest possible price in the empty book, he chooses this price.

Condition (6) follows from the fact that patient sellers require non-negative expected utility. From equation (2),

$$\mathrm{E}\left[p_t\left(x_{F_S}|x^{F_S};y^l\right) - v_t - w_t\left(x_{F_S}|x^{F_S};y^l\right)\right] \geq 0,$$

which implies $p_t\left(x_{F_S}|x^{F_S};y^l\right) \geq v_t + w_t\left(x_{F_S}|x^{F_S};y^l\right)$. □

Proof. (Proposition 3). The informed traders' utility function in (**??**) together with the relation $ask_t > v_t > bid_t$ derived in footnote 8 imply that informed traders always prefer limit orders over marketable orders because they do not bear waiting costs. The assumption we make on the off-equilibrium beliefs of the informed traders means that they do not want to get identified. This implies that, first, informed traders may use limit orders only in scenarios when patient uninformed traders are expected to use limit orders, and, second, they would only post limit order prices that could be chosen by the patient uninformed papers. Since patient uninformed traders require non-negative utility, they do not submit limit orders when the same-side-book is full. Consequently, an informed trader would also not choose a limit order when the same-side-book is full, and by (**??**) would then choose a marketable order. Thus, informed traders are expected to choose marketable orders on a particular side of the book only when that side is full, and limit orders otherwise.

At time t, the probability that the asset value is $v + \sigma$ from the perspective of uninformed traders is q_t. Therefore, from their perspective the probability that the informed traders use buy orders is q_t, and the probability they use sell orders is $(1 - q_t)$. Which means that if the buy side of the book is not full, uninformed traders expect the arrival rate of limit buy orders to be $\lambda_{bp} + q_t\lambda_I$, and the arrival rate of marketable buy orders to be λ_{bi}. If the buy side of the book is full, uninformed traders expect the arrival rate of limit buy orders to be λ_{bp}, and the arrival rate of marketable buy orders to be $\lambda_{bi} + q_t\lambda_I$. The arguments for the sell side are symmetric. □

Proof. (Proposition 4). Equations (15) and (16) show how the arrival of orders lead to revisions in the uninformed traders' beliefs, q_t. Therefore, we need to show how price gaps are affected by changes in q_t. Gaps measure the difference between wait times at consecutive priority slots, so we have

$$g_t\left(x_{F_s-i-1}, x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right) = w_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j}\right) - w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right). \quad (18)$$

Thus, the rest of the proof focuses on showing that gaps shrink as q_t increase,

$$\frac{\partial\left(w_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j}\right) - w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right)\right)}{\partial q_t} < 0. \quad (19)$$

Following the arguments in Section 4.3, we may express the wait time for x_{F_s-i} when there are $F_s - i$ sell orders and $F_b - j$ buy orders as

$$\begin{aligned} w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right) =& \frac{1}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I} + w_t(x_{F_s-i+1}|x^{F_s-i+1};y^{F_b-j})\frac{\lambda_{sp}+(1-q_t)\lambda_I}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I} \\ &+ w_t(x_{F_s-i}|x^{F_s-i};y^{F_b-j-1})\frac{\lambda_{si}}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I} \\ &+ w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j+1}\right)\frac{\lambda_{bp}+q_t\lambda_I}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I}. \end{aligned} \quad (20)$$

Similarly, the wait time for x_{F_s-i-1} if there were $F_s - i$ sell orders and $F_b - j$ buy orders would be

$$w_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j}\right) = \frac{1}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I} + w_t(x_{F_s-i}|x^{F_s-i};y^{F_b-j})\frac{\lambda_{sp}+(1-q_t)\lambda_I}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I}$$

$$
\begin{aligned}
&+ w_t(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j-1})\frac{\lambda_{si}}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I} \\
&+ w_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j+1}\right)\frac{\lambda_{bp}+q_t\lambda_I}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I}.
\end{aligned} \tag{21}
$$

Equations (20) and (21) are recursive expressions. Thus, the value for $w_t(x_{F_s-i+1}|x^{F_s-i+1};y^{F_b-j})$ in equation (20), and $w_t(x_{F_s-i}|x^{F_s-i};y^{F_b-j})$ in equation (21), are plugged in from the previous round of the recursion, and are effectively constant in these calculations.

Differentiating $w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right)$ with respect to q_t, we get

$$
\begin{aligned}
&\frac{\partial w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right)}{\partial q_t} = -w_t(x_{F_s-i+1}|x^{F_s-i+1};y^{F_b-j})\frac{(\lambda_I+\lambda_{bp}+\lambda_{bi}+2\lambda_{si})\lambda_I}{(\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I)^2} \\
&- w_t(x_{F_s-i}|x^{F_s-i};y^{F_b-j-1})\frac{\lambda_{sp}\lambda_I}{(\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I)^2} \\
&- w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j+1}\right)\frac{(\lambda_{bi}+\lambda_{si})\lambda_I}{(\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I)^2} - \frac{\lambda_I}{(\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I)^2} \\
&+ \frac{\partial w_t(x_{F_s-i}|x^{F_s-i};y^{F_b-j-1})}{\partial q_t}\frac{\lambda_{si}}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I} \\
&+ \frac{\partial w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j+1}\right)}{\partial q_t}\frac{\lambda_{bp}+q_t\lambda_I}{\lambda_{bp}+\lambda_{bi}+\lambda_{si}+q_t\lambda_I}.
\end{aligned} \tag{22}
$$

Wait times are positive quantities, so we may rewrite equation (22) as

$$
\begin{aligned}
\frac{\partial w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right)}{\partial q_t} =& - K^i_{i,j} + \frac{\partial w_t(x_{F_s-i}|x^{F_s-i};y^{F_b-j-1})}{\partial q_t}C^i_{i,j+1} \\
&+ \frac{\partial w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j+1}\right)}{\partial q_t}C^i_{i,j-1},
\end{aligned} \tag{23}
$$

where $K^i_{i,j}$, $C^i_{i,j+1}$ and $C^i_{i,j-1}$ are positive constants, and $C^i_{i,j+1}, C^i_{i,j-1} < 1$.

Similarly,

$$
\begin{aligned}
\frac{\partial w_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j}\right)}{\partial q_t} =& - K^i_{i+1,j} + \frac{\partial w_t(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j-1})}{\partial q_t}C^i_{i+1,j+1} \\
&+ \frac{\partial w_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j+1}\right)}{\partial q_t}C^i_{i+1,j-1},
\end{aligned} \tag{24}
$$

where $K^i_{i+1,j}$, $C^i_{i+1,j+1}$ and $C^i_{i+1,j-1}$ are positive constants, and $C^i_{i+1,j+1}, C^i_{i+1,j-1} < 1$.

When the book has $F_s - i$ sell orders, x_{F_s-i-1} has to wait for x_{F_s-i} to execute and exit the market before a chance at execution. Therefore,

$$w_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j}\right) > w_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right) \quad \text{for all } \ i,j, \tag{25}$$

which means that $-K^i_{i,j} > -K^i_{i+1,j}$. On the other hand, $C^i_{i,j+1} = C^i_{i+1,j+1}$ and $C^i_{i,j-1} = C^i_{i+1,j-1}$.

Subtracting equation 23 from equation 24 and using the definition of price gaps, we obtain

$$\begin{aligned}\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j}\right)\right)}{\partial q_t} &= -\left(K^i_{i+1,j}-K^i_{i,j}\right)\\ &\quad +C^i_{i,j+1}\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j-1}\right)\right)}{\partial q_t}\\ &\quad +C^i_{i,j-1}\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j+1}\right)\right)}{\partial q_t}.\end{aligned} \tag{26}$$

Similar relations can be derived for $\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j-1}\right)\right)}{\partial q_t}$ and $\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j+1}\right)\right)}{\partial q_t}$, and then for $\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j-2}\right)\right)}{\partial q_t}$ and $\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j+2}\right)\right)}{\partial q_t}$, and then for $\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j-3}\right)\right)}{\partial q_t}$ and $\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j+3}\right)\right)}{\partial q_t}$ and so on, and substituted into equation (26). If such substitution is undertaken repeatedly, we are left with an expression of the form

$$\begin{aligned}\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b-j}\right)\right)}{\partial q_t} &= -\sum_{l=1}^{N} K_l\\ &\quad +\left(\prod_{m=1}^{N} C_m\right)\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{0}\right)\right)}{\partial q_t}\\ &\quad +\left(\prod_{n=1}^{N} C_n\right)\frac{\partial\left(g_t\left(x_{F_s-i-1}|x^{F_s-i};y^{F_b}\right)\right)}{\partial q_t},\end{aligned} \tag{27}$$

where the K_l are of the form $\left(K^i_{i+1,j}-K^i_{i,j}\right)$, and C_m and C_n are of the form $C^i_{i,j+1}$ and

$C^i_{i,j-1}$ in equation (26). Since each C_m and C_n are less than 1, the product of such terms approaches zero as N becomes large. Therefore, lines 2 and 3 in the above equation approach zero, and $\frac{\partial\left(g_t\left(x_{F_s-i}|x^{F_s-i};y^{F_b-j}\right)\right)}{\partial q_t}$ is negative, as we set out to prove. $\square$

References

Back, K., and S. Baruch, 2007. Working Orders in Limit Order Markets and Floor Exchanges. Journal of Finance, 62, 1589-1621.

Baruch, S., and L. Glosten, 2013. Fleeting Orders. Working paper, University of Utah.

Baruch, S., M. Panayides, and K. Venkataraman, 2017. Informed Trading and Price Discovery before Corporate Events. Journal of Financial Economics 125, 561-588.

Bernheim, B. D., 1984. Rationalizable Strategic Behavior. Econometrica 52, 1007-1028.

Brolley, M., and K. Malinova, 2020. Informed Liquidity Provision in a Limit Order Market. Forthcoming in the Journal of Financial Markets.

Buti, S., and B. Rindi, 2013. Undisclosed Orders and Optimal Submission Strategies in a Limit Order Market. Journal of Financial Economics 109, 797-812.

Buti, S., B. Rindi, and I. Werner, 2017. Dark Pool Trading Strategies, Market Quality, and Welfare. Journal of Financial Economics 124, 244-265.

Chakravarty, S., and C. W. Holden, 1995. An Integrated Model Of Market And Limit Orders. Journal of Financial Intermediation 4, 213-241.

Cohen, K. J., S. F. Maier, R. A. Schwartz, and D. K. Whitcomb, 1981. Transaction Costs, Order Placement Strategy, and Existence of the Bid-Ask Spread. Journal of Political Economy 89, 287-305.

Diamond, D. W., and R. E. Verrecchia, 1987. Constraints on short-selling and asset price adjustment to private information. Journal of Financial Economics 18, 227-311.

Easley, D., and M. O'Hara, 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69-90.

Easley, D., and M. O'Hara, 1992. Time and the Process of Security Price Adjustment. Journal of Finance 47, 577-605.

Foucault, T., 1999. Order Flow Composition and Trading Costs in a Dynamic Limit Order Market. Journal of Financial Markets 2, 99-134.

Foucault T., O. Kadan, and E. Kandel, 2005. Limit Order Book as a Market for Liquidity. Review of Financial Studies 18, 1171-1217.

Glosten, L. R., 1994. Is the Electronic Open Limit Order Book Inevitable? Journal of Finance 49, 1127-1161.

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 13, 71-100.

Goettler, R., C. Parlour, and U. Rajan, 2005. Equilibrium in a Dynamic Limit Order Market. Journal of Finance 60, 2149-2192.

Goettler, R., C. Parlour, and U. Rajan, 2009. Informed Traders and Limit Order Markets. Journal of Financial Economics 93, 67-87.

Hasbrouck, J., and G. Saar, 2009. Technology and Liquidity Provision: The Blurring of Traditional Definitions. Journal of Financial Markets 12, 143-172.

Hasbrouck, J., and G. Saar, 2013. Low-Latency Trading. Journal of Financial Markets 16, 646-679.

Kaniel, R., and H. Liu, 2006. So What Orders Do Informed Traders Use? Journal of Business 79(4), 1867-1913.

Kumar, P., and D. Seppi, 1994. Limit and Market Orders with Optimizing Traders. Working paper, Carnegie Mellon University.

Kwan, A., R. Philip, and A. Shkilko, 2021. The Conduits of Price Discovery: A Machine Learning Approach. Working paper, University of Sydney.

Kyle, A., 1985. Continuous Auctions and Insider Trading. Econometrica 53, 1315–1335.

Pagnotta, E., 2010. Information and Liquidity Trading at Optimal Frequencies. Working paper, Imperial College Business School.

Parlour, C., 1998. Price Dynamics in Limit Order Markets. Review of Financial Studies 11, 789-816.

Parlour C. A., and D. Seppi, (2008). Limit Order Markets: A Survey, Chapter 3 in Handbook of Financial Intermediation & Banking, edited by A. W. A. Boot and A. V. Thakor, 63-96.

Pearce, D. G., 1984. Rationalizable Strategic Behavior and the Problem of Perfection. Econometrica 53, 1029-1050.

Putnins, T., and D. Michayluk, 2018. When do Informed Traders Provide Liquidity? Working paper, University of Technology Sydney.

Ricco R., B. Rindi, and D. Seppi, 2018. Information, Liquidity, and Dynamic Limit Order Markets. Working paper, Carnegie Mellon University.

Rock, K., 1996. The Specialist's Order Book and Price Anomalies. Working paper, Harvard University.

Rosu, I., 2009. A Dynamic Model of the Limit Order Book. Review of Financial Studies 22, 4601-4641.

Rosu, I., 2020. Liquidity and Information in Limit Order Markets. Forthcoming in the Journal of Financial and Quantitative Analysis.

Seppi, D., 1997. Liquidity Provision with Limit Orders and a Strategic Specialist. Review of Financial Studies 10, 103–150.

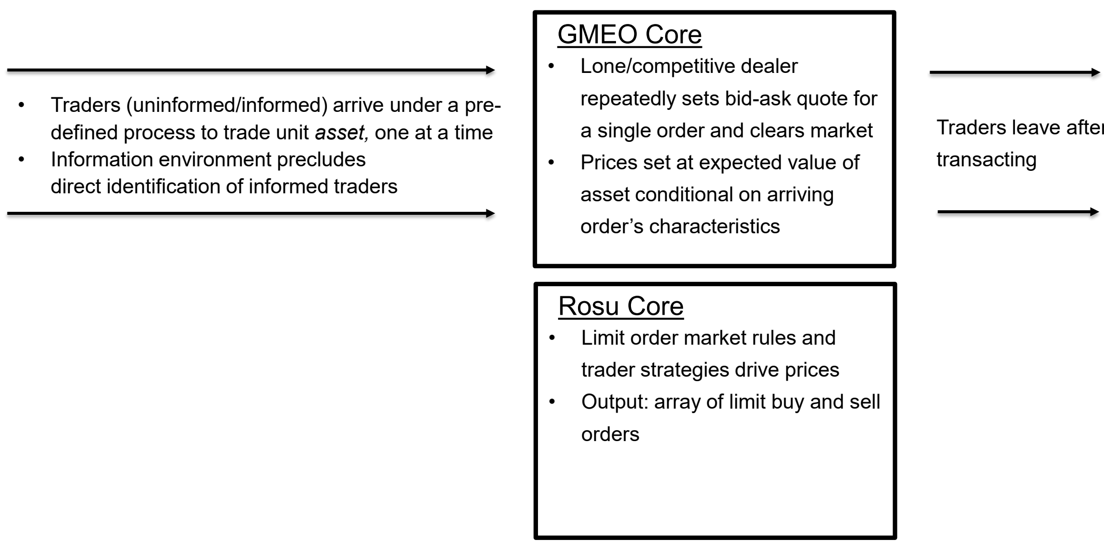

Figure 1: **Modeling objective: GMEO meet Rosu (2009).** This figure illustrates our modeling objective: using the trader arrival and information environment of the Glosten-Milgrom-Easley-O'Hara framework, but replacing its dealer-market trading core with the dynamic equilibrium limit order market introduced by Rosu (2009).