ESTIMATING AN EQUILIBRIUM MODEL OF LIMIT ORDER MARKETS

THOMAS G. RUCHTI – [PRELIMINARY DRAFT]

ABSTRACT. I estimate an equilibrium model using limit order data, finding parameters that describe information and liquidity preferences for trading. As a case study, I estimate the model for Google stock surrounding an unexpected good-news earnings announcement in the 3rd quarter of 2009. I find a substantial decrease in asymmetric information prior to the earnings announcement. I also simulate counterfactual dealer markets and find empirical evidence that limit order markets perform more efficiently than do their dealer market counterparts.

1. INTRODUCTION

In this paper, I study equilibrium trading in limit order markets. Limit order markets are widely used in financial exchanges around the world, including New York Stock Exchange, NASDAQ, Stockholm Stock Exchange, and Paris Bourse, among others. A limit order market allows for direct interaction between traders, without a market-making dealer. Traders choose between market orders–which execute against existing orders in the book–and limit orders–which enter the book with a limit number of shares and limit price and await execution with a market order. Despite widespread use, empirical studies of limit order markets have been hampered by data availability and complexity of existing models. This paper pushes forward the empirical analysis of these markets by specifying and estimating a structural econometric model of equilibrium trading in a limit order market. As a case study, I estimate this model on the limit order book for Google stock surrounding a surprisingly good earnings announcement in the 3rd quarter of 2009.

Date: November 12, 2012.
Key words and phrases. limit order book, market microstructure, equilibrium, oligopoly, dynamic programming, indirect inference.
Division of Humanities and Social Sciences, Caltech, Pasadena, CA 91125. E-mail: truchti-at-caltech.edu.
This paper is a part of my dissertation at Caltech. I am especially grateful to Jakša Cvitanić, Ben Gillen, John Ledyard, and Matt Shum for encouragement and guidance. For helpful discussion and advice, I also thank Shmuel Baruch, Jonathan Brogaard, Claude Courbois, Federico Echenique, Peter Foley, Gustav Haitz, Andrea Bui Kanady, Andrei Kirilenko, Arseniy Kukanov, Sangmok Lee, Emiliano Pagnotta, and Jean-Laurent Rosenthal. For data access I appreciate the assistance of the LOBSTER project and NASDAQ OMX.

I develop an econometric framework based on work of Bernhardt and Hughson (1997) and Biais, Martimort, and Rochet (2003), in which the authors study a model of imperfect competition among liquidity suppliers under adverse selection. My estimation methodology incorporates an insight by Baruch (2008), a paper subsumed by Back and Baruch (2012), that liquidity suppliers' choice of supply schedules can be recursively characterized as a dynamic programming problem. This insight is key to my econometric procedure, allowing me to apply tools and methodologies developed for estimating dynamic optimization models to the explicitly non-dynamic (static) model of equilibrium in limit order markets. To my knowledge, this is the first paper performing structural econometric analysis in such a setting.

I use this econometric framework to address two important questions. First, is information transmitted evenly and quickly across the market? I present a measure of information asymmetry, the Informedness Ratio. A higher ratio means that insiders have more information than the market, with a lower ratio meaning asymmetry is less. The informedness ratio is high a day before the examined earnings announcement, settling at a lower level the day of the announcement, and continuing to fall the day after. The fall in the informedness ratio the day of the announcement indicates information asymmetry and uncertainty fall before the announcement is made. Second, what are the welfare implications of the switch from dealer to limit order markets, pervasive in real-world financial markets? I simulate counterfactual dealer markets using real data, and show that limit order markets perform as well as, if not much better than, one would expect dealer markets to perform.

I address an important question in the microstructure literature involving pervasive change in modern exchanges to a limit order-based system. Historically in most markets, trade was handled by a single monopolist dealer. If a trader wanted to buy or sell a certain number of shares, a dealer would quote terms, and the trade could be executed. Now most exchanges use a limit order book. Traders have the option of placing market orders, which execute against existing shares in the book, or placing a limit order, with a limit price and quantity, which enters the book and waits for a market order to execute against it. I use counterfactual examples to assess the change in market performance moving from two dealers to the imperfect competition found in limit order markets. Arguments could be made that limit order markets reduce restrictions on offer schedules, improving terms for trade, increasing welfare.[1] However, a dealer could bring a level of expertise to an asset, allowing a dealer to facilitate the market in times that decentralized traders would be unwilling to make a market for an asset. Using my model, I simulate counterfactual dealer markets from the data. Not surprisingly, more liquidity suppliers leads to higher welfare. As a robustness check, I find the model fits better with five liquidity suppliers than it does with two or ten–when traditional markets typically had one, but no more than two, dealers per asset. This means that the market behaves effectively as if it has five dealers. More dealers unambiguously means higher welfare, hence I conclude that limit order markets are more efficient than a counterfactual dealer market would be.

I touch on several different literatures related to this work in section 2. Section 3 explains the limit order book as a dynamic program and the risk-averse active trader setting used here. Section 4 discusses methods of simulation and estimation. I discuss my results on earnings announcements in section 5. In section 6, I perform robustness checks of the model specification and conduct an analysis of the move from dealer markets to limit order markets. Section 7 concludes.

2. Literature Review

I take an equilibrium model of imperfect competition in a common value setting from Bernhardt and Hughson (1997), which is further analyzed for equilibrium in Biais, Martimort, and Rochet (2003)[2] and apply it in an empirical setting. Their analysis involves strategic risk-neutral liquidity suppliers competing in schedules for the business of a risk-averse agent who is privately informed about the value of the asset and about hedging needs in a common value setting. Biais et al. extend previous analyses concerning what creates the bid-ask spread, patterns in trading volumes, and oligopolistic incentives.[3]

[1]See Biais, Foucault, and Salanié (1998).

[2]Madhavan (1992) uses the same risk-averse trader setting, but assumes competitive liquidity suppliers. For a recent model on imperfect competition in a contrasting *private* value setting, see Vives (2011).

[3]See Kyle (1985) for the seminal batch-trading model, Glosten (1989) for the monopolist dealer problem, Glosten (1994) for an early treatment of limit order markets and the limiting case as the number of traders

Baruch (2008) makes a key analytical insight, by applying dynamic programming methods to analyze the same model. From an econometric point of view, this insight proves useful in making these models amenable to estimation. While dynamic programming methods are inherently valuable to a variety of problems in structural analysis, these methods are often employed in explicitly dynamic models, where agents make intertemporal decisions. Using Baruch's intuition, I can solve for equilibrium price schedules using dynamic programming methods in price. Dynamic programming has become a crucial component of the structural econometrics tool box, beginning with Miller (1984), Wolpin (1984), Pakes (1986), and Rust (1987), and is used to approach a variety of problems. While my model is not explicitly dynamic, I use similar intuition in finding the equilibrium in a limit order market. By simulating supply curves from the data, I try to match parameters in the true model to actual data. This process is the principal of Indirect Inference, as in Gouriéroux, Monfort, and Renault (1993) and Smith (1993).[4]

While theory is concerned with microstructure and its implications on liquidity provision and information transmission, the related empirical literature focuses on reduced-form analysis of optimal order placement and dynamics between limit orders and market orders.[5] Biais, Hillion, and Spatt (1995)[6] study the early Paris Bourse computerized limit order exchange. They find that thin books elicit more limit orders whereas market depth results in market orders, or immediate trades.[7] Dufour and Engle (2000) use a model to assess the role of waiting time between transactions in the process of price formation. They find that as the speed of trades increases, price adjustment also increases, hence an increased presence of

grows, and Bernhardt and Hughson (1997) for the duopoly case. For related and important models in which nature chooses whether the trader is an informed trader or a liquidity-motivated trader, see Copeland and Galai (1983) and Glosten and Milgrom (1985). Some work has sought to reconcile the two-sided limit order problem, as in Roşu (2009). For a survey of the theoretical literature, see O'Hara (1998), or more recently, Vives (2010).

[4]Goettler and Gordon (2011) is a recent application of Indirect Inference in structural estimation.

[5]For a survey of empirical work, see Hasbrouck (2007).

[6]Ranaldo (2004) studies similar questions, but using ordered probit to empirically investigate order submission strategies. The author finds patient traders become more aggressive when their own side of the book is thicker, when the spread is wider, and when volatility is momentarily high. For an experimental study of liquidity in an electronic limit order market, see Bloomfield, O'Hara, and Saar (2005).

[7]This finding is consistent with modern theoretical and numerical simulation results. See Foucault, Kadan, and Kandel (2005), Goettler, Parlour, and Rajan (2005), and Pagnotta (2010).

informed traders. Ahn, Bae, and Chan (2001) study the relationship between market depth and transitory volatility. They find evidence that limit order traders enter the market, placing orders where liquidity is needed. This finding is in support of the notion that a limit order market is in equilibrium at any one point in time, with liquidity suppliers waiting to fill gaps in the market.

The literature of structural empirical work on limit order markets is smaller. Sandås (2001) present economic restrictions on price schedules offered in a competitive setting. He finds that there is insufficient depth in limit order books relative to theoretical predictions. Hollifield, Miller, and Sandås (2004)[8] examine optimal order placement and find order submission is a montone function for a trader's valuation of the asset. However, their model assumes traders trade only one unit of the asset, and the authors reject their private value trading restrictions for the order placements of traders with moderate private values. Hollifield, Miller, Sandås, and Slive (2006) study the gains from trade in a limit order market, comparing efficiency in a perfectly liquid market and a market with a monopolist to the actual gains from trade. Kelley and Tetlock (2012) estimate a model of strategic trader behavior that incorporates endogenously informed traders and discretionary liquidity traders. They show that these discretionary traders make up most trading volume, but that from 2001 to 2010, informed trading increasingly contributes to volume and stock price discovery. Their analysis exploits variation in trading and volatility correlated with time of day and public news arrival under a linear pricing equilibrium. I build on the structural estimation literature by estimating equilibrium trading in a limit order model with endogenous liquidity provision.

3. Model

In limit order markets, traders choose between placing market orders and limit orders. Market orders execute against orders in the book at the best price posted, what is called *walking the book*, or *taking liquidity*. Limit orders specify a limit price and quantity, where unexecuted portions of a limit order enter the book, what is called *supplying liquidity*. The limit order market allows traders to interact directly. Posting liquidity to the book and

[8]For an earlier investigation in optimal order strategies, see Harris and Hasbrouck (1996).

taking it have a timing component however. A market order executes with an order that came before it. The active trader making the market order may have information newer to the market than information liquidity suppliers had when offering shares. If this information asymmetry fully characterized trade, there would be no market (see Milgrom and Stokey (1982), Grossman and Stiglitz (1980)). Instead, active investors have incentives to trade beyond inside information. An active investor may wish to hedge a position, reducing exposure to an asset. This balance between opposing motivations is a key tension in information-based models of limit order markets.

There are n risk-neutral uninformed liquidity suppliers and a single informed active trader trading a single asset. Liquidity suppliers submit limit orders and the active trader observes all bids and offers and submits a marketable order. The asset is then liquidated at $v = \alpha + \epsilon$, where α, distributed normally, is the signal the informed active trader receives, and ϵ is noise. $I \sim N(\mu_I, \sigma_I^2)$ is the informed active trader's inventory of the asset. Latent supply, S_0 is used for estimation and is described in more detail later. Table 1 provides model paramters and definitions.

TABLE 1. Model Parameters

n	number of liquidity suppliers
σ_α	standard deviation of informed trader's signal
σ_ϵ	standard deviation of noise of informed trader's signal
μ_I	mean of the distribution of the informed trader's inventory of the asset
σ_I	standard deviation of the distribution of the informed trader's inventory
σ_{S_0}	standard deviation of latent supply

Without loss of generality, I focus my discussion on the offer side of the book. Liquidity suppliers could be viewed as institutional traders, for example, Goldman Sachs, etc. Active traders could be employees at the traded company who have access to information relevant to the company's stock performance that is unavailable to lqiuidity suppliers. However, active traders do not only have information motivations for trading, but also non-information hedging needs based on their current holding of the stock. Liquidity suppliers do not know whether a market order posted by an active trader arises from information or liquidity motivations. This information asymmetry between the active trader and liquidity suppliers is a source of adverse selection in the model. Generally, hedging-motivated trades would be profitable for the liquidity suppliers, but information-based trades would not. Thus, liquidity suppliers face adverse selection in deciding how many shares to supply at any one price. Higher information asymmetry will make liquidity suppliers wary, and they will respond by posting fewer shares. Less information asymmetry means more liquidity will be available and more hedging needs will be met.

3.1. **Notation.** Liquidity suppliers play a static game. To accommodate the data, I take prices as discrete. p_{ask} is the minimum price at which liquidity is offered and p_{max} is the maximum, hence prices are $p_{ask} \equiv p_1, \ldots, p_M \equiv p_{max}$, and are taken as given. A strategy for the i^{th} liquidity supplier is $S^i : \{p_1, \ldots, p_M\} \to \mathcal{R}$ where $S^i_r \equiv S^i(p_r)$ represents the total number of shares offered by i through price p_r. The limit order book changes rapidly within a short period of time. At the lowest price, p_1, we assume there is some additional, *latent supply*, S_0, which represents impatient trades, hidden orders, or spillover from the buy side of the market. While this does not change implications of the model, it is important for econometric implementation, as it generates variation in the limit order book across trading episodes.[9] Because the set of prices is discrete, the function $s^i_r \equiv S^i_r - S^i_{r-1}$ is well defined. I also define $s^{-i}_r \equiv \sum_{j \neq i} s^j_r$, $S_r \equiv \sum_{i=1,\ldots,n} S^i_r + S_0$, and $s_r \equiv \sum_{i=1,\ldots,n} s^i_r$.

In this common value environment, profitability of a limit order is dependent on probability of execution and the expected value of selling the asset conditional on execution. If marketable, the active trader's bid walks up the book, picking off offers until it reaches its limit price or is filled. Private information observed by the active trader leads to adverse selection. Profitability of a liquidity supplier's offer at a price p_r is a function of that price, liquidity offered up to that price by all liquidity suppliers, liquidity offered at that price by the agent, and liquidity offered at that price by all other liquidity suppliers.

3.2. **Dynamic Programming Characterization of the Limit Order Book.** Because a single trader's order walks up the book, I can treat the limit order book, which consists

[9] For a detailed discussion of latent supply, see the equilibrium solution section of the appendix, 8.3.

in equilibrium, of the liquidity suppliers optimally chosen supply schedules, as the solution to a dynamic programming problem in price. This is a key insight of Baruch (2008). While profitability of shares offered at a lower price affects profitability of shares offered at a higher price, profitability of shares offered at lower prices are unaffected by shares offered at higher prices, hence at the price p_r, I define the value function,

$$V_L(p_r, S_{r-1}) = \max_{s^i_m,\ m=r,\dots,M} \sum_{m=r}^{M} u_L(p_m, S_{m-1}, s^i_m, s^{-i}_m) \tag{1}$$

such that

$$S_{m+1} = S_m + s^i_m + s^{-i}_m \tag{2}$$

Here $u_L(p_r, S_{r-1}, s^i_r, s^{-i}_r)$ represents profitability to the liquidity supplier.[10] The state variable is S_{r-1}, the total volume supplied at prices lower than the current price, p_r. The corresponding Bellman equation is,

$$V_L(p_r, q_r) = \max_{s^i_r}[u_L(p_r, S_{r-1}, s^i_r, s^{-i}_r) + V_L(p_{r+1}, S_{r-1} + s^i_r + s^{-i}_r)] \tag{3}$$

This recursively characterizes the optimal supply schedule $\{s^{i*}_1, s^{i*}_2, \dots, s^{i*}_M\}$. At maximal price p_M, the Bellman equation simplifies to

$$V_L(p_M, S_{M-1}) = \max_{s^i_M} u_L(p_M, S_{M-1}, s^i_M, s^{-i}_M) \tag{4}$$

For all values of S_{M-1}, I can solve for s^{i*}_M at the maximal price. Plugging in this strategy to the Bellman equation at p_{M-1}, the strategy s^{i*}_{M-1} satisfies,

(5)

$$V_L(p_{M-1}, S_{M-2}) = \max_{s^i_{M-1}}[u_L(p_{M-1}, S_{M-2}, s^i_{M-1}, s^{-i}_{M-1}) + V_L(p_M, S_{M-2} + s^i_{M-1} + s^{-i}_{M-1} + s^{i*}_M + s^{-i*}_M)]$$

Given the finiteness of prices, this dynamic programming problem can be solved for $s^{i*} \equiv \{s^{i*}_1, s^{i*}_2, \dots, , s^{i*}_M\}$ by backward induction, starting at the highest price p_M.

[10]The form of profitability is described later in section 3.4 and appendix 8.1.

3.3. **Risk Averse Active Trader.** In this paper I use a setting with a risk averse, *informed* trader facing risk neutral, *uninformed* liquidity suppliers, found in many papers, namely Glosten (1989) for the monopolist dealer setting, Madhavan (1992) for a competitive setting, Bernhardt and Hughson (1997) for duopoly, and Biais, Martimort, and Rochet (2003) for the further oligopoly setting. In this setting, active traders placing market orders bring new information to the market and are informed. Liquidity suppliers, however, have orders already in the book when a market order is placed, and are uninformed. As in Copeland and Galai (1983), active traders coming to the market–after limit orders are already posted–bring information to the market that liquidity suppliers do not yet have.

The active trader sees a signal, α, where $v = \alpha + \epsilon$ and ϵ is normally distributed with mean 0 and standard deviation σ_ϵ. In addition, the active trader has inventory of the asset I and knows this. The active trader uses this information to decide an optimal order, q, maximizing the expectation of the utility of wealth, where

$$u_A(q|I,\alpha) \equiv -e^{-\gamma W}, \tag{6}$$

$$W \equiv \underbrace{(q+I)v}_{\text{value of shares owned}} - \underbrace{\sum_{r=0}^{M} \max\{0, \min\{q, S_r\} - S_{r-1}\} p_r}_{\text{payment to suppliers}}, \text{[11]} \tag{7}$$

and γ is the coefficient of risk aversion. The active trader does not observe v directly, but receives a noisy signal α. A large α means that the underlying asset's value is likely high, and the active trader's order will walk the book picking off stale asks. Conversely, if I is large and positive, the active trader wants to sell, and when I is large and negative, the active trader wants to buy. The active trader is incentivized by α and I, so a high α (news) may mean buying even when I (exposure) is also high. The uninformed risk neutral liquidity suppliers know the distribution of $\alpha \sim N(\mu_\alpha, \sigma_\alpha)$ and $I \sim N(\mu_I, \sigma_I)$, but a liquidity supplier can only infer a posterior distribution on (α, I) from trades.

[11]For ease of notation, and following the definition of latent supply, $p_0 = p_1$.

For illustration purposes, I first consider the case where prices are continuous. A supplier posts an offer to sell at a price and an informed active trader's buy order walks the book, picking off these shares. Following rational expectations, the seller must exhibit no regret in selling shares to the profile of traders that bought them. In Figure 1, a supply schedule generates a purchase choice of q^* shares. The optimal choice by the active trader is the intersection of supply and incentives to buy shares, $\frac{\alpha - \text{price}}{\gamma\sigma_\epsilon^2} - I$. Given that $\gamma\sigma_\epsilon^2$ is common knowledge, the supplier knows the slope of this curve, and can back out the intercept, revealing the profile of active traders (α, I). The liquidity supplier considers the profile of traders who would buy at a given price and quantity and responds with the optimal number of shares.[12]

Figure 1. Inference and Continuous Solution

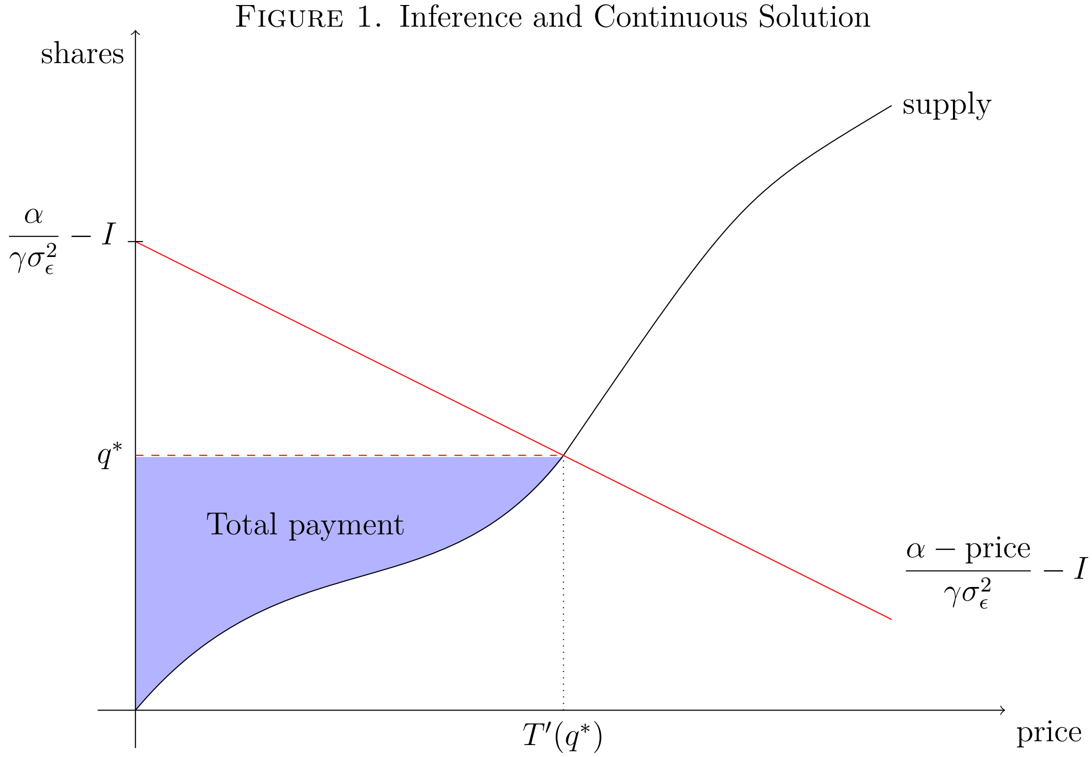

To put it more formally, liquidity suppliers compete in transfer schedules $T(\cdot) : \mathcal{R}^+ \to \mathcal{R}^+$ giving the payment for any quantity, q, demanded by an active trader. As a result, the

[12]This reasoning is similar to strategies conditional on winning the prize in auctions (see Milgrom and Weber (1982)), and vote pivotality in juries (see Feddersen and Pesendorfer (1998)). The liquidity supplier must offer enough shares to take advantage of the hedging needs of traders reaching that far in the book while balancing information asymmetry faced at that point as well.

active trader chooses q maximizing $E[-e^{-\gamma W}]$, from (6). This is equivalent to maximizing $E[W|\alpha, I] + \frac{\gamma}{2}V[W|\alpha, I]$. Variance of holdings is $(q+I)^2\sigma_\epsilon^2$. From transfer schedules this is equal to, $(q+I)\alpha - T(q) - \frac{\gamma}{2}(q+I)^2\sigma_\epsilon^2$. This yields an interior solution for supply schedules where $q^* = \frac{\alpha - T'(q^*)}{\gamma\sigma_\epsilon^2} - I$. Here the supply curve in Figure 1 is $T'(q)$, and $T'(q^*)$ is simply the price at which the last portion of q^* is supplied.

In typical data, supply schedules are not continuous functions of price, but rather step functions. This introduces complications to the model, as I show here.[13] Solving for equilibrium q, the best response quantity demanded by an active trader when utility is $u(q|\alpha, I) \equiv -e^{-\gamma W}$, is again equivalent to the active trader choosing a q maximizing $E[W|\alpha, I] + \frac{\gamma}{2}V[W|\alpha, I]$. Wealth is given by equation 7, and the trader maximizes,

$$\underbrace{(q+I)\alpha}_{\text{expected value of shares}} \underbrace{-\sum_{r=0}^{M}\max\{0, \min\{q, S_r\} - S_{r-1}\}p_r}_{\text{payment to suppliers}} - \underbrace{\frac{\gamma}{2}(q+I)^2\sigma_\epsilon^2}_{\text{volatility penalty}}. \tag{8}$$

In equation 8, v in wealth as in equation 7 is replaced by α, because σ_ϵ is expectation 0. However, σ_ϵ shows up in the risk aversion variance of wealth penalty. Because of the discrete nature of these supply curves, either the active trader reaches a first order condition, at some point between S_r and S_{r+1} for some r, or not, at S_{r+1} for some r. In the first case, $\alpha - p_m - \frac{\gamma}{2}\sigma_\epsilon^2(2q + 2I) = 0$, where p_r is s.t. $S_r < q < S_{r+1}$. In the second, the first order condition is not reached, and $q = S_{r+1}$ s.t. $\alpha - p_r - \frac{\gamma}{2}\sigma_\epsilon^2(2q + 2I) > 0$ and $\alpha - p_{r+1} - \frac{\gamma}{2}\sigma_\epsilon^2(2q + 2I) < 0$, hence the active trader has exhausted all gains from trade at price p_r, but would lose money trading at p_{r+1} for the next available liquidity in the book. The two cases are as follows,

Case 1: interior solution, where p_r is s.t. $S_{r-1} < q < S_r$

$$q = \frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I \tag{9}$$

[13]A treatment of price discreteness for bidding in multiunit auctions is found in Kastl (2006), McAdams (2008), and Hortaçsu and Kastl (2012).

Figure 2. Inference and Interior Solution

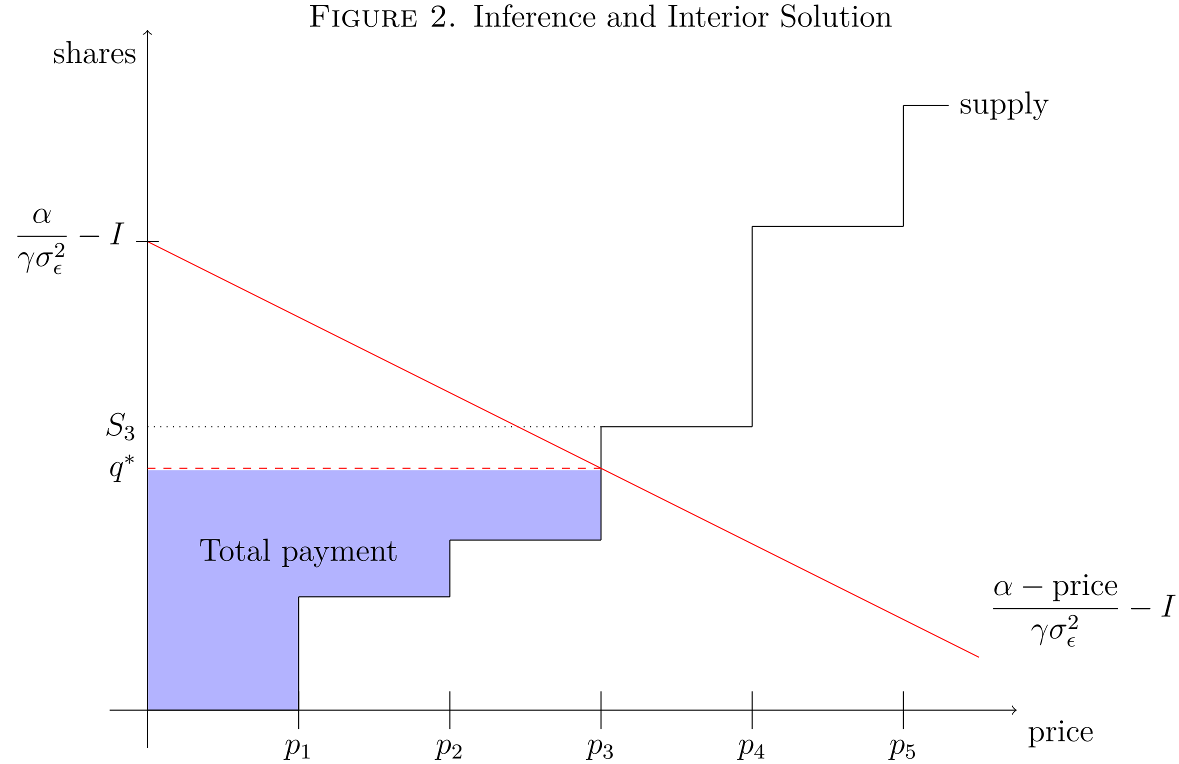

As can be seen, given a supply schedule and an active trader places an order, liquidity suppliers can infer a statistic, $\frac{\alpha}{\gamma\sigma_\epsilon^2} - I$, for α and I which contains all the information the liquidity suppliers can infer about α and I.

Case 2: boundary solution, $q = S_r$

$$\frac{\alpha - p_{r+1}}{\gamma\sigma_\epsilon^2} - I < q < \frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I \tag{10}$$

However, if $q^* = S_r$ for some r, then inference is not as clear for the liquidity suppliers. Now, $\frac{\alpha}{\gamma\sigma_\epsilon^2} - I$ can only be narrowed down to a range of profiles.

3.4. **Liquidity Suppliers and the Optimal Supply Problem.** I assume that offers at a given price are executed proportionally.[14] This means that at price p_r, the bid walks through the orders at a rate $\frac{s_r^i}{s_r^i + s_r^{-i}}$ for each supplier i until the order moves up the book to the next price. With these solutions I compute the utility of offering shares for the liquidity supplier at each price p_r, which is expected profitability given liquidity suppliers are risk-neutral.

[14]Because the model used here is not explicitly dynamic, assuming that some orders are executed after others would negate non-trivial symmetric equilibrium and likely would make the number of equilibria infinite.

FIGURE 3. Inference and Boundary Solution

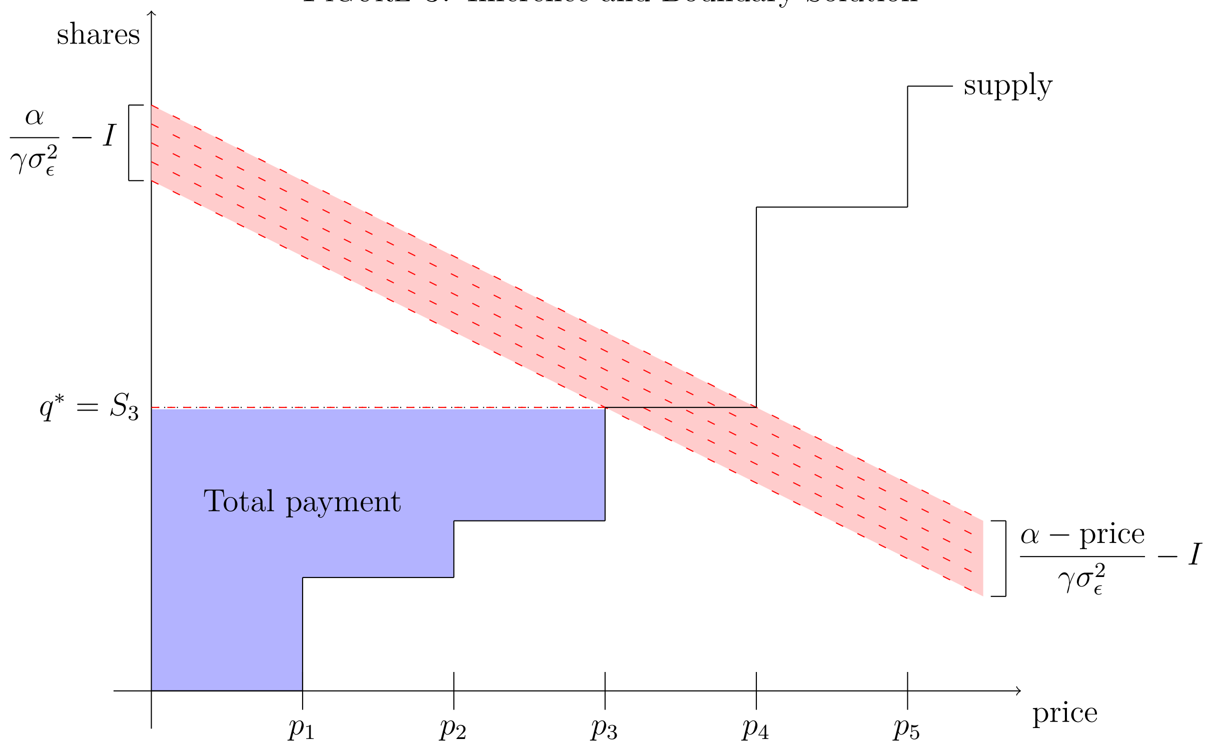

Either $q > S_{r-1} + s_r^i + s_r^{-i}$ or $S_{r-1} < q \leq S_{r-1} + s_r^i + s_r^{-i}$.

$$u_L(p_r, S_{r-1}, s_r^i, s_r^{-i}) = \tag{11}$$

$$E_{\alpha,\epsilon,I}\left[s_r^i \mathbf{1}_{\{S_{r-1}+s_r^i+s_r^{-i}\leq q\}}(p_r - v) + \frac{s_r^i}{s_r^i + s_r^{-i}}(q - S_{r-1})\mathbf{1}_{\{S_{r-1}<q<S_{r-1}+s_r^i+s_r^{-i}\}}(p_r - v)\right]$$

The expected profit is added to the continuation payoff for subsequent prices. For each number of shares offered, there is an expected profit for that price, and an expected payoff for subsequent supply. I consider a symmetric equilibrium for liquidity suppliers when the derivative of this expected profit with the derivative of continuation is 0.[15] I leave further derivation to the appendix, 8.1.

[15] While Back and Baruch (2012) show equilibrium theoretically for only a range of parameter values, I verify the equilibrium numerically.

4. Estimation

Here I describe my econometric method in which I use a nested fixed point algorithm, as in Rust (1987)[16]. Starting with parameter values of the structural model, I find equilibrium supply curves. Using a pseudo-model[17] (following the terminology in Gouriéroux, Monfort, and Renault (1993)), I fit the generated supply curves to those found in the data. Choosing new parameter values, I iterate until I achieve a best fit. The following sections describe the individual components of this analysis.

4.1. **Solving for Equilibrium.** Players compete in supply schedules at each price successively, as liquidity at lower prices affects the profitability of liquidity supplied later. In the optimal Markov strategy, liquidity suppliers' offers of shares at a given price are based soley on the amount of liquidity offered up to that point.[18] Suppliers consider a trade executing against their supply schedule. The market order picks off shares at prices at which liquidity is offered, until the order is met. A liquidity supplier infers what profile of information and liquidity incentives were faced given the trade reached that deep in the book, and decides how many shares to offer at that price. Suppliers determine their strategies starting at the limit of economically meaningful share depth and backward induct to find their optimal supply curves.

For interested readers, solving for equilibrium is described in detail in the appendix, 8.3.

4.2. **Indirect Inference.** For estimating model parameters, I use the simulated method of Indirect Inference, following seminal work in Gouriéroux, Monfort, and Renault (1993), and Smith (1993).[19] I estimate a pseudo-model for data and limit order supply curves simulated from the model. Realizing a best fit of the model means generating estimates as close together as possible. I find the set of parameters such that the pseudo-model's two sets of

[16]See Rust (1994) and Rust (1996) for further discussion.

[17]The pseudo-model is referred to as an *approximated model*, an *instrumental model*, and a *statistical model* by the literature. For clarity, I will use the term *pseudo-model* from here on.

[18]Given the nature of modeling competition in supply, if an optimal strategy exists for this problem, then a Markov optimal strategy exists.

[19]See Goettler and Gordon (2011) for another recent use of Indirect Inference for estimating equilibrium of a structural model.

estimations coincide.[20] Coefficients of the pseudo-model consist of a least squares regression of share quantities on four linear and nonlinear functions of price.

Let $\hat{\beta}$ denote parameters of the pseudo-model estimated from actual data. Analogously, let $\tilde{\beta}^{\lambda}(\theta)$ denote parameters estimated from data simulated from the limit order book under parameters θ (λ superscripts the specific simulation). The Indirect Inference estimator $\hat{\theta}^{\Lambda}$ [21] optimizes the following criterion,

$$\hat{\theta}^{\Lambda} \equiv \arg\min_{\theta \in \Theta} \left[\hat{\beta} - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda}(\theta)\right]' \Omega \left[\hat{\beta} - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda}(\theta)\right] \tag{12}$$

I explain details of the Indirect Inference method, demonstrate asymptotic normality of the estimator, and also prove the following proposition in the appendix, 8.4,

Proposition 2. Under assumptions 1–6, $\hat{\theta}^{\Lambda}$ is a consistent estimator of θ_0.

5. Empirical Results: A Snapshot of Google, October 14-16, 2009

In this section I estimate the model. I examine the limit order book for Google stock (trading on NASDAQ), using nine episodes bracketing the 2009 third quarter earnings announcement by Google. Each episode includes a series of observations of the limit order book, characterizing equilibrium in the limit order market at the time of the episode. Each set of observations are grouped together to form estimates of the model. I obtain a picture of how information asymmetry evolved before and after the announcement.

5.1. **Earnings Announcements, Asymmetric Information.** Earnings announcements convey meaningful information about a company. Markets price an asset consistently with market expectations, but the true profits and losses a company realizes over the quarter may shift these expectations, causing jumps in asset prices. Trading on this information could be very profitable and insiders trading on information could undermine the market. Leading up

[20] For all estimates, I try different starting values for the parameter search algorithm. In most cases, the outcomes for the model parameters are the same. When different starting values produced different parameter estimates, I chose the model parameters that led to the lowest criterion function values.

[21] Λ signifies the number of simulated supply curves analyzed by the pseudo-model. In this paper, I use 200 simulations for each criterion calculation.

to an earnings announcement, a Google employee, for example an upper level manager, who may be a significant Google shareholder, could have information about the announcement not yet known to the market. Uninformed institutional traders offer liquidity to this Google employee, keeping in mind that the employee may be trading based on liquidity preferences, e.g., reducing exposure to the asset, or inside information. If institutional traders anticipate heavy information trading, depth of the book will be shallow. However, if they anticipate liquidity trading, supply schedules will become steeper and more depth will fill the book.

There is an important literature studying returns and markets surrounding earnings announcements. Linnainmaa (2010) shows that using limit orders changes the inferences one can make about trading intentions. Examining several regularly identified investor trading patterns, he shows that most of these observed effects are due to the use of limit orders, and that much of the inferences about investors' trading abilities are due to limit orders' exposure to adverse selection risk. Christophe, Ferri, and Angel (2004) use NASDAQ data to examine short-selling prior to earnings announcements, and find a significant link between abnormal short-sales and post-announcement stock returns.[22] Similarly, Kaniel, Liu, Saar, and Titman (2012) consider large individual investor buys and sells on the New York Stock Exchange, and show corresponding abnormal returns following earnings announcements.[23] While these studies analyze price movements and trading volumes, as indirect evidence of information leakage, I use my structural model to directly quantify changes in information asymmetry surrounding an earnings announcement. As I show below, the implications of my results differ from implications one might draw from analyzing price movements and trading volume.

[22]Ball and Brown (1968) were the first to note the link between abnormal returns and unexpected earnings announcements. Foster, Olsen, and Shevlin (1984) replicated these findings regarding abnormal returns, uncovering similar market inefficiencies. Bernard and Thomas (1989) and Bernard and Thomas (1990) show support of the hypothesis of price response delay, rejecting that capital asset pricing systematically underestimates (overestimates) risk surrounding good (bad) news.

[23]Lee (1992) were early to find patterns regarding strategies of large and small investors. Bartov, Radhakrishnan, and Krinsky (2000), and Bhattacharya (2001) find that investor sophistication is negatively correlated with abnormal returns–investors with less sophistication underestimate the implications of a surprise earnings announcement. Battalio and Mendenhall (2005) show that investors making large trades respond to earnings forecast errors, while investors making small trades respond to a less-sophisticated signal. See Hirshleifer, Myers, Myers, and Teoh (2008) for contrasting analysis.

On Thursday, October 15, 2009, Google closed on NASDAQ at $530, well below the Wednesday close of $535.50. That afternoon, then CEO Eric Schmidt made the 3rd quarter earnings announcement, remarking,

> Google had a strong quarter–we saw 7% year-over-year revenue growth despite the tough economic conditions. While there is a lot of uncertainty about the pace of economic recovery, we believe the worst of the recession is behind us and now feel confident about investing heavily in our future.[24]

This was apparently a positive shock to the market; the stock reopened at $546.50 on Friday, Oct. 16, eventually closing at $550, while market returns were flat these days. Next, I examine this through the lens of the structural estimates from my equilibrium limit order model.

5.2. **Estimation Results.** I estimate the model on three days surrounding the announcement. The earnings announcement was made 4:30 p.m. EST, Oct. 15, a half-hour after the market closed, and estimates are taken from Oct. 14, 15, and 16, early, midday, and late during trading hours. I characterize information asymmetries and liquidity motivations for trading at each point in time.

Before turning to my structural estimates, I consider typical measures of information asymmetry; specifically, I look at the bid-ask spread and the amount of shares offered in the order book. Figure 4 shows an overall downward trend in bid-ask spreads in this three day period, with a spike in spread in the middle of the day following the announcement, consistent with Lee, Mucklow, and Ready (1993).[25] This hints at an overall reduction in information asymmetry during this period. Similarly, Figure 4 demonstrates the amount of shares in the first five ticks of the market are low early in the day, high late.

Next I turn to my structural estimates. Results and definitions of terms can be found in Table 2. The level of μ_I changes a great deal over this period of time, being lower (and negative, so greater in absolute magnitude, and more hedging preferences for trade) later

[24] `http://investor.google.com/pdf/2009Q3_earnings_google.pdf`

[25] Lee, Mucklow, and Ready (1993) show that liquidity providers are sensitive to changes in information asymmetry risk and use both spreads and depths to actively manage this risk.

Figure 4. Bid-Ask Spread and Depth

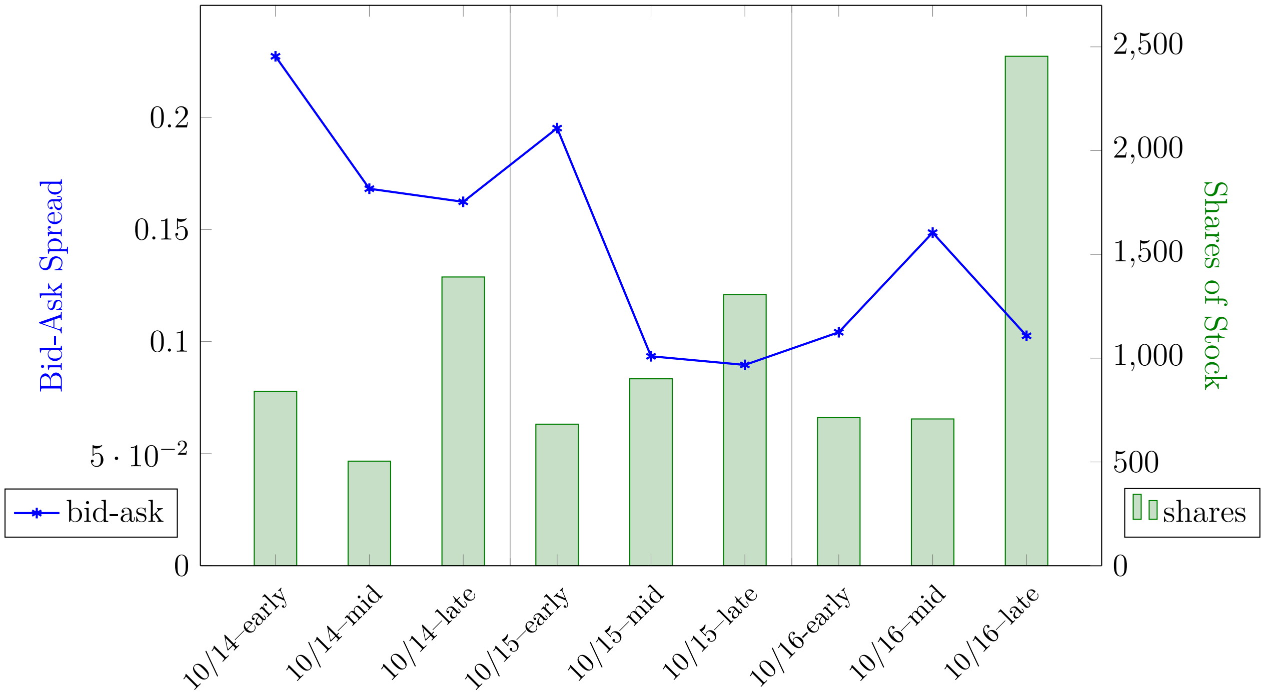

in the day than earlier, with a peak (and therefore less liquidity preferences for trade) in the middle of the day. However, σ_α changes rapidly, and beginning Oct.14, σ_ϵ is gradually increasing. As σ_α falls and σ_ϵ rises, the ratio of σ_α to σ_ϵ decreases, hence, asymmetric information drops significantly. Information asymmetry drops from Oct. 14 to Oct. 15, continuing to fall later on Oct. 15–as evidenced by the increased hedging incentive revealed in the market at that time in the day. We see this asymmetry fall again on Oct. 16, with a lower point in the middle of the trading day.

5.3. **Informedness Ratio.** To summarize the implications of parameter estimates on the extent of information asymmetry and adverse selection, I introduce the *informedness ratio*. The informedness ratio is defined as the ratio of the standard deviations of the asset value from the liquidity suppliers' perspective and the active trader's perspective. The active trader receives a signal about the asset's true value and the liquidity suppliers do not. Therefore,

Table 2. Earnings Announcements Results

This table shows fitted parameters of the true model for the sell side of the Google limit order book surrounding its 3rd quarter earnings announcement, made 4:30 p.m. EST, October 15, 2009. Parameters are estimated from the first five ticks of the book in a set of three samples, one on each side of a time of day, where samples are separated by 15 minutes. The dates and times chosen are early, midday and late on October 14, 15, and 16. Market open at 9 a.m., each early estimate is taken at 10:00 a.m., EST, unless otherwise specified, midday, 1:00 p.m., and late, 3:30 p.m., market closing at 4 p.m. σ_α is the standard deviation of the informed trader's signal about the asset's true value, and σ_ϵ is the standard deviation of the white noise around the informed trader's signal, both in \$. μ_I and σ_I are the mean and variance of the inventory of the informed trader, in shares of Google stock. σ_{S_0} is the standard deviation of a calibration variable representing hidden orders and impatient sellers, in shares of Google. Bootstrap standard errors in parentheses.

Date-Time	σ_α	σ_ϵ [a]	μ_I	σ_I	σ_{S_0}
10/14 early	0.7260	0.0415	-700.106	199.806	20.927
	(0.0076)	(0.0005)	(2.9226)	(6.8451)	(3.2822)
10/14 mid	0.3367	0.0306	-401.749	111.887	20.005
	(0.0066)	(0.0020)	(5.247)	(3.454)	(0.2524)
10/14 late	0.5720	0.0216	-1598.362	561.992	19.407
	(0.0013)	(0.00004)	(2.352)	(2.667)	(0.2983)
10/15 early	0.4375	0.0254	-947.812	148.192	12.030
	(0.0009)	(0.00009)	(0.5244)	(3.454)	(0.0154)
10/15 mid	0.4274	0.0339	-736.098	230.567	19.861
	(0.0026)	(0.0001)	(2.963)	(3.452)	(0.1252)
10/15 late	0.4571	0.0284	-2169.115	275.685	20.300
	(0.0058)	(0.0002)	(34.372)	(13.408)	(1.1364)
10/16 early[b]	0.6871	0.0481	-912.498	81.159	8.0851
	(0.0006)	(0.00003)	(0.1377)	(0.3464)	(0.0077)
10/16 mid	0.3855	0.0445	-704.534	253.961	20.851
	(0.0177)	(0.0011)	(9.602)	(17.688)	(0.3929)
10/16 late	0.2933	0.0204	-2376.286	725.681	20.496
	(0)	(0)	(0)	(0)	(0)
N[c]	51	–	–	–	–
S[d]	200	–	–	–	–

[a] The coefficient of risk aversion, γ never *stands alone*, meaning that it is estimated directly with σ_ϵ. I assume $\gamma = 1$ and show estimates for σ_ϵ. [b] This time was offset by a total of two (2) fifteen minute periods to 10:30 a.m.. The bid-ask spread was too wide to produce meaningful estimates, and I waited until a point in the day that the market was back in equilibrium. [c] For each set of estimates I take 3 samples. Each sample has a number of instances of the order book. This is the sample size. [d] I use simulations to fit data. I take a certain number of draws of latent supply and find the associated supply curves.

the ratio is simply the ratio of the standard deviation of the true value from its expectation, and the standard deviation of the white noise the active trader faces in the signal, which is a natural measure of how much less informed the liquidity suppliers are compared to the active trader. The true value of the asset is $v = \alpha + \epsilon$; assuming that the signal α is not correlated with white noise ϵ, the informedness ratio is defined as,

$$\frac{\sigma_v}{\sigma_\epsilon} = \frac{\sqrt{\sigma_\alpha^2 + \sigma_\epsilon^2}}{\sigma_\epsilon} \equiv \text{Informedness Ratio}^{26} \tag{13}$$

The informedness ratio has no absolute cardinal meaning. However, a higher value may be indicative of large asymmetry in information or a preponderance of information trading, while a lower value indicates more liquidity-motivated trading and hence, less information asymmetry.

The informedness ratio from my results is plotted in Figure 5. There is a pronounced downward trend in the informedness ratio over the three day period, implying a reduction in information asymmetry. The informedness ratio is rather high late on Oct. 14. It falls steeply by Oct. 15, settling around 17 that day. This shows that the market responded to asymmetry that is reduced the day of the earnings announcement. Whatever information traders were ready to trade on, prior to the earnings announcement, the market is no longer responding to it once the earnings announcement is made. However, this information asymmetry dissipates further, as evidenced by continued decline in the informedness ratio the day after the earnings announcement is made. It reaches 14 and then 8 after being double that the day before.

This plot indicates that the market behaves as if active traders were better informed (relative to uninformed liquidity suppliers) the day before the announcement, and that this informedness dissipated by the day of the announcement. The fall in the informedness ratio between Oct. 14 late and Oct. 15 early contrasts with the trend in the bid-ask spread, which peaked during Oct. 15 early, as shown in Figure 4. Overall, the downward trend

[26]The coefficient of risk aversion, γ never *stands alone*, meaning that it is estimated directly with σ_ϵ. I assume $\gamma = 1$ and show estimates for σ_ϵ.

FIGURE 5. Informedness Ratio

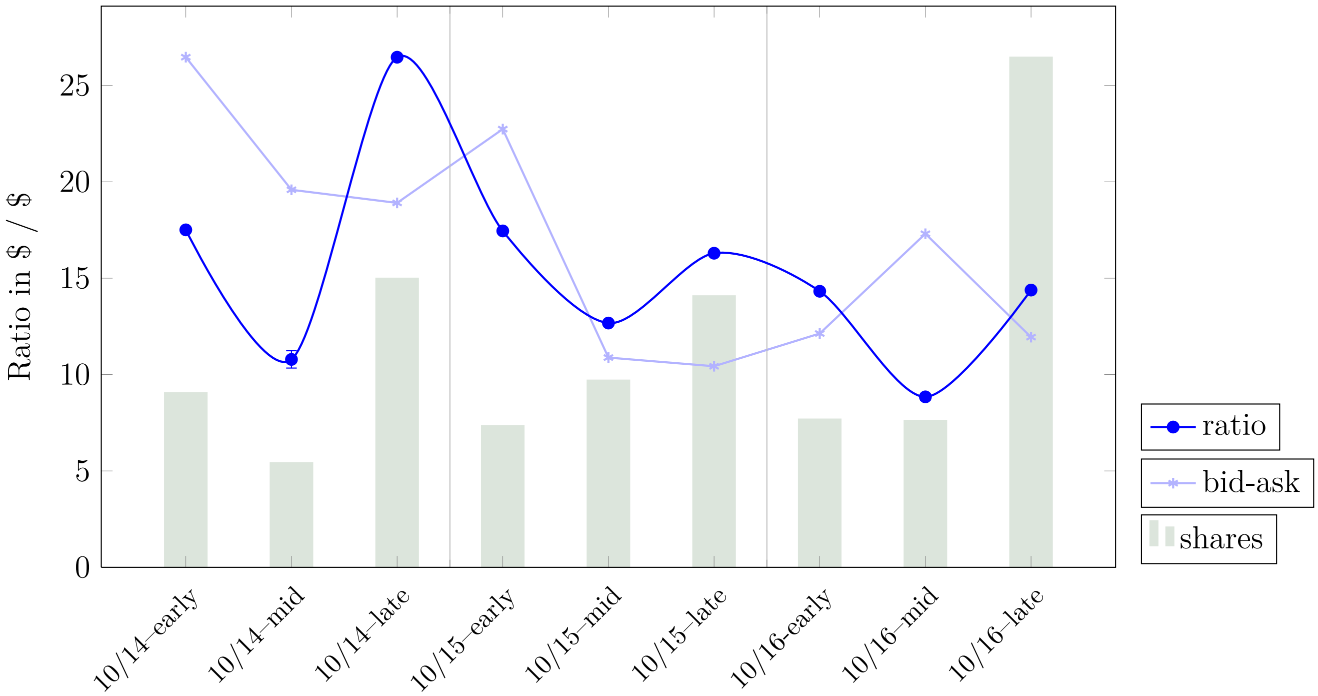

in the informedness ratio corroborates a similar trend in the bid-ask spread, however, the informedness ratio implies that the decline in information asymmetry began earlier.

6. Robustness Checks: Levels of Competition

In limit order data, individual-level order placement goes unobserved. However, the model specifies a number of traders. I estimate the model using $n = 5$, a choice resulting from a specification search over different values of n.

6.1. **Comparison of Fit.** I show that the market behaves as if there are more liquidity suppliers than two and fewer than ten. If the model were to fit well with $n = 5$ liquidity suppliers, and not as well with $n = 2$ or $n = 10$, this would be evidence the market behaves as if there are effectively five liquidity suppliers competing for the business of active traders. Model fit and estimates for $n = 2$, $n = 5$, and $n = 10$ are compared in Table 3.

Table 3. Robustness Checks

This table shows comparative fitted parameters of the true model for market structures with two, five, and ten liquidity suppliers for the sell side of the Google limit order book surrounding its 3rd quarter earnings announcement, made 4:30 p.m. EST, October 15, 2009. The market for Google opens at 9 a.m., EST, and unless otherwise noted, early refers 10:00 a.m., mid refers to 1:00 p.m., and late refers to estimates being made for data at 3:30 p.m., market closing at 4:00 p.m. σ_α is the standard deviation of the informed trader's signal about the asset's true value, and σ_ϵ is the standard deviation of the white noise around the informed trader's signal, both in \$. μ_I and σ_I are the mean and variance of the inventory of the informed trader, in shares of Google stock, where greater magnitude means more liquidity preferences. σ_{S_0} is the standard deviation of a calibration variable representing hidden orders and impatient sellers, in shares of Google. Fit shows the mean-squared error of the estimates, and a lower number is a better fit. The model fits better with number of liquidity suppliers being five rather than two, hence markets behave as if they have five dealers, as opposed to two. This is evidence that modern limit order markets are more efficient than hypothetical dealer markets–characterized by one, or maybe two dealers–would be. Standard errors in parentheses.

	n	fit	σ_α	σ_ϵ	μ_I	σ_I	σ_{S_0}
10/14 late	2	52,244.2	0.4253	0.0391	-2633.938	814.957	20.607
	5	7651.4	0.5720	0.0216	-1598.362	561.992	19.407
	10	10,863.5	0.2794	0.0227	-527.943	176.344	17.314
10/15 early	2	214.0	0.3754	0.0579	-1454.298	415.548	20.388
	5	6.0	0.4375	0.0254	-947.812	148.192	12.030
	10	55.4	0.6434	0.0698	-285.119	190.829	22.548
10/15 mid	2	6296.4.8	0.3608	0.0436	-1853.380	542.495	20.210
	5	90.0	0.4274	0.0339	-736.098	230.567	19.861
	10	367.5	0.3772	0.0327	-366.799	105.985	20.116
10/15 late	2	95,737.8	0.2632	0.0209	-2193.646	265.143	20.225
	5	83.9	0.4571	0.0284	-2169.115	275.685	20.300
	10	297.5	0.1925	0.0433	-692.216	294.226	20.538
10/16 early[a]	2	9631.1	0.4286	0.0424	-2361.003	683.236	20.436
	5	687.7	0.6871	0.0481	-912.431	80.932	8.004
	10	581.4	0.2484	0.0214	-463.436	137.343	19.778

[a] This time was offset by a total of two (2) fifteen minute periods to 10:30 a.m.. The bid-ask spread was too wide to produce meaningful estimates, and I waited until a point in the day that the market was back in equilibrium.

Table 3 contains fitted θs for five dates and times examined earlier. I also include θs for $n = 2$ and $n = 10$ to contrast specification fits. The model fits are better in four of the five dates and times. Market data is fitted by the model consistently better for $n = 5$ over $n = 2$ with only one instance of a fit worse than $n = 10$.

6.2. **Comparing Limit Order Markets and Dealer Markets.** The model I estimate fits better with $n = 5$ than it does with $n = 2$ or $n = 10$. Looking at $n = 2$, corresponds to estimating a dealer market. I investigate the change in welfare moving from dealer to limit order markets. Historically, NASDAQ, New York Stock Exchange, and others, depended on dealers or market makers to facilitate the market for an asset. More recently, modern markets have employed a direct trading scheme in the form of a limit order market. In dealer markets, traders never directly interact. Instead, a trader would make known to a dealer that she was interested in selling shares of a stock. The dealer would post terms for trade, and the trader would sell as much as she wished according to those terms. The dealer would then sell those shares on the other side of the market. This way a trader could always trade and never had to wait for terms to be met as long as the dealer was there to take the other side of any transaction. Limit order markets are centralized, but trade need not go through a single intermediary.

While the move from dealer markets to limit order markets is pervasive, it is not clear they are better in terms of welfare. It may be the case that while limit order markets bypass the middleman, that liquidity is more available under a dealer market scheme. However, it may be that the ability of any trader to play the role of dealer allows for greater liquidity provision.

The difficulty of characterizing the limit order book theoretically makes this welfare question very difficult to analyze. Pagnotta (2010) presents a computational model that nests dealer and limit order markets, allowing for a comparison of utilities. Finding that welfare is not noticeably better by adding a dealer, he deduces that the limit order market does better than its historical counterpart.

The model estimated here also nests dealer and limit order markets, and is amenable to a welfare comparison of the two. Because active traders and liquidity suppliers trade with each other, the model is essentially a dealer market. If liquidity suppliers offer a large number of shares early, then the informed traders they face at higher prices in the book will likely have a strong signal about the price of the asset, decreasing expected profits. In a monopolistic setting, a liquidity supplier bears the entirety of this loss, and so will offer liquidity at higher prices, leading to a large bid-ask spread. However, adding liquidity suppliers increases incentives to offer liquidity at low prices. With a large number of suppliers, there is close to efficient liquidity provision and a small bid-ask spread. I showed earlier that the model fit is better for $n = 5$ than it is for $n = 2$. I demonstrate here that increasing the number of liquidity suppliers in the market raises welfare. Because a large number of dealers is unambiguously good, and the model fits better for $n = 5$, data confirms Pagnotta's results; limit order markets are better for welfare than dealer markets.

6.3. **Counterfactual Utilities.** I calculate profitability for liquidity suppliers and certainty-equivalent utility for the active trader. These values will be useful in considering counterfactual utilities, comparing social welfare in dealer and limit order markets. Interested readers can find relevant calculations in the appendix, 8.2.

I generate counterfactual markets, varying the number of liquidity suppliers $n = 1, \ldots, 50$, ranging from a monopoly to relative competition. Using these simulated markets, I compute expected profits to liquidity suppliers and the certainty-equivalent utility increase for active traders. Table 4 presents counterfactual estimates for five dates and times examined in the earnings announcements section. For each example, profits to liquidity suppliers go down as the number of competitors increases. Conversely, certainty-equivalent utility goes up for active traders as the number of liquidity suppliers competing for their business increases. Comparing profits and utilities in absolute number is misleading because of the nature of the active trader's exponential utility. However, profits to liquidity suppliers are a transfer of utility from the active trader, so relative changes in the sum of utilities matter. Moving from a more monopolistic setting to a more competitive one is better for the market for this asset. Hence, I show with data that limit order markets are more efficient than their dealer market predecessors.

Comparing certainty-equivalent utility gain to the active trader at different points in time, the active trader sees a much lower gain Oct. 14 late than Oct. 16 early, despite volatility being similar, 0.5724 and 0.6888, respectively. At both points in time, active traders are exposed to a large amount of risk in the volatility of the asset. The informedness ratio is high Oct 14 late, 26.463, compared to 14.320 on Oct 16 early, so liquidity suppliers are offering a less favorable supply schedule. The increased hedging active traders can do on Oct 16 early is due to reduced asymmetric information in the market. Similarly, on Oct. 15 late, there is an even greater gain to liquidity suppliers and active traders. Despite volatility being relatively high, 0.4601, the informedness ratio was 16.295, and there was a large amount of depth in the book. Many traders were posting liquidity, leading to steep supply curves, allowing for a large amount of hedging needs to be met.

Following robustness checks, I determined that the model fit data better with $n = 5$ than specifications with fewer or more liquidity suppliers.[27] This number represents the oligopolistic incentives liquidity suppliers face in the limit order market. The limit order market does not operate as a duopoly, nor does it operate under perfect competition. In addition, counterfactual utilities show that competition among liquidity suppliers is unambiguously good for welfare. The model specified with $n = 2$ is essentially a dealer market. Moving from this hypothetical dealer market to $n = 5$, the robust specification for estimating the data, increases counterfactual utilities. Hence, data confirms Pagnotta's result that moving from dealer markets to limit order markets increases welfare.

7. Conclusion

I develop a method of estimating an equilibrium model of the limit order book using market data.

[27]While there may be settings in which a value of $n > 5$ would be appropriate, it is clear that n is at least larger than $n = 2$, evidence limit order market structure is more efficient than dealer markets.

Table 4. Counterfactual Profits and Utilities

This table shows counterfactual profits to liquidity suppliers and utility to informed active traders across varied market structures. I study here the sell side of the limit order book for Google stock on dates October 14, 15, and 16, surrounding the 3rd quarter earnings announcement by Google in 2009. The market for Google opens at 9 a.m., EST, and unless otherwise noted, early refers 10:00 a.m., mid refers to 1:00 p.m., and late refers to estimates being made for data at 3:30 p.m., market closing at 4:00 p.m. Liquidity suppliers are considered the uninformed dealers while the informed active trader takes liquidity from the limit order book. The number of suppliers is varied from 1 (monopoly) to 50 (relative competition). Total profits are listed for liquidity suppliers and decrease with competition, whereas certainty-equivalent utility as defined for the informed active trader increases with competing liquidity suppliers. Because profits to liquidity suppliers are zero-sum, this is evidence that more competition, and so more liquidity suppliers, in these markets leads to higher levels of efficiency. Bootstrap standard errors are in parentheses.

		Number of Suppliers					
		1	2	5	10	20	50
10/14 late	suppliers	1.525	0.8029	0.3480	0.1887	0.1185	0.0691
	active trader	70.043	80.275	79.311	82.383	84.483	85.287
10/15 early	suppliers	27.034	11.739	4.172	2.357	1.260	0.6383
	active trader	61.843	121.451	141.205	145.073	147.294	148.882
10/15 mid	suppliers	29.129	11.999	3.291	1.556	0.8809	0.4941
	active trader	163.773	198.182	205.612	210.075	211.387	214.921
10/15 late	suppliers	134.179	49.909	25.937	20.708	16.926	14.156
	active trader	1385.844	1527.734	1541.288	1574.801	1596.956	1610.083
10/16 early[a]	suppliers	62.334	24.417	7.341	4.339	2.437	1.245
	active trader	634.427	693.385	707.728	707.093	718.011	723.144

[a] This time was offset by a total of two (2) fifteen minute periods to 10:30 a.m.. The bid-ask spread was too wide to produce meaningful estimates, and I waited until a point in the day that the market was back in equilibrium.

As an application, I examine the market for Google around the time of an earnings announcement. The informedness ratio, a model-consistent measure of information asymmetry, begins decreasing even before the earnings announcement. This is consistent with the notion that some information related to the announcement was already present in the market before the announcement itself. A second application examines the welfare implications of dealer vs. limit order markets.

My method lends itself to a number of further applications. Initial public offerings involve important market structure changes over time. In addition, mergers would be an interesting phenomenon to investigate. These I leave for future research.

References

Ahn, H., K. Bae, and K. Chan (2001): "Limit orders, depth, and volatility: evidence from the Stock Exchange of Hong Kong," *The Journal of Finance*, 56(2), 767–788.

Back, K., and S. Baruch (2012): "Strategic Liquidity Provision in Limit Order Markets," *forthcoming, Econometrica.*

Ball, R., and P. Brown (1968): "An empirical evaluation of accounting income numbers," *Journal of accounting research*, pp. 159–178.

Bartov, E., S. Radhakrishnan, and I. Krinsky (2000): "Investor sophistication and patterns in stock returns after earnings announcements," *Accounting Review*, pp. 43–63.

Baruch, S. (2008): "The Limit Order Book and the Break-Even Conditions, Revisited," *University of Utah Working Paper.*

Battalio, R., and R. Mendenhall (2005): "Earnings expectations, investor trade size, and anomalous returns around earnings announcements," *Journal of Financial Economics*, 77(2), 289–319.

Bernard, V., and J. Thomas (1989): "Post-earnings-announcement drift: delayed price response or risk premium?," *Journal of Accounting research*, 27, 1–36.

——— (1990): "Evidence that stock prices do not fully reflect the implications of current earnings for future earnings," *Journal of Accounting and Economics*, 13(4), 305–340.

Bernhardt, D., and E. Hughson (1997): "Splitting orders," *Review of Financial Studies*, 10(1), 69–101.

Bhattacharya, N. (2001): "Investors' trade size and trading responses around earnings announcements: An empirical investigation," *Accounting Review*, pp. 221–244.

Biais, B., T. Foucault, and F. Salanié (1998): "Floors, dealer markets and limit order markets," *Journal of Financial Markets*, 1(3), 253–284.

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

Biais, B., D. Martimort, and J. Rochet (2003): "Competing mechanisms in a common value environment," *Econometrica*, 68(4), 799–837.

Bloomfield, R., M. O'Hara, and G. Saar (2005): "The Şmake or takeŤ decision in an electronic market: Evidence on the evolution of liquidity," *Journal of Financial Economics*, 75(1), 165–199.

Christophe, S., M. Ferri, and J. Angel (2004): "Short-Selling Prior to Earnings Announcements," *The Journal of Finance*, 59(4), 1845–1876.

Copeland, T., and D. Galai (1983): “Information effects on the bid-ask spread,” *Journal of Finance*, pp. 1457–1469.

Dufour, A., and R. Engle (2000): “Time and the price impact of a trade,” *The Journal of Finance*, 55(6), 2467–2498.

Feddersen, T., and W. Pesendorfer (1998): “Convicting the innocent: The inferiority of unanimous jury verdicts under strategic voting,” *American Political Science Review*, pp. 23–35.

Foster, G., C. Olsen, and T. Shevlin (1984): “Earnings releases, anomalies, and the behavior of security returns,” *Accounting Review*, pp. 574–603.

Foucault, T., O. Kadan, and E. Kandel (2005): “Limit order book as a market for liquidity,” *Review of Financial Studies*, 18(4), 1171–1217.

Glosten, L. (1989): “Insider trading, liquidity, and the role of the monopolist specialist,” *Journal of Business*, pp. 211–235.

——— (1994): “Is the electronic open limit order book inevitable?,” *Journal of Finance*, pp. 1127–1161.

Glosten, L., and P. Milgrom (1985): “Bid, ask and transaction prices in a specialist market with heterogeneously informed traders,” *Journal of Financial Economics*, 14(1), 71–100.

Goettler, R., and B. Gordon (2011): “Does AMD spur Intel to innovate more?,” *Journal of Political Economy*, 119(6), 1141–1200.

Goettler, R., C. Parlour, and U. Rajan (2005): “Equilibrium in a dynamic limit order market,” *The Journal of Finance*, 60(5), 2149–2192.

Gouriéroux, C., A. Monfort, and E. Renault (1993): “Indirect Inference,” *Journal of Applied Econometrics*, 8(S1), 85–118.

Grossman, S., and J. Stiglitz (1980): “On the impossibility of informationally efficient markets,” *The American Economic Review*, 70(3), 393–408.

Harris, L., and J. Hasbrouck (1996): “Market vs. limit orders: the SuperDOT evidence on order submission strategy,” *Journal of Financial and Quantitative Analysis*, 31(2).

Hasbrouck, J. (2007): *Empirical market microstructure: The institutions, economics and econometrics of securities trading*. Oxford University Press, USA.

Hirshleifer, D., J. Myers, L. Myers, and S. Teoh (2008): “Do individual investors cause post-earnings announcement drift? Direct evidence from personal trades,” *The Accounting Review*, 83(6), 1521–1550.

Hollifield, B., R. Miller, and P. Sandås (2004): “Empirical analysis of limit order markets,” *Review of Economic Studies*, 71(4), 1027–1063.

Hollifield, B., R. Miller, P. Sandås, and J. Slive (2006): “Estimating the Gains from Trade in Limit-Order Markets,” *The Journal of Finance*, 61(6), 2753–2804.

Hortaçsu, A., and J. Kastl (2012): "Valuing dealers' informational advantage: A study of Canadian treasury auctions," *forthcoming, Econometrica.*

Kaniel, R., S. Liu, G. Saar, and S. Titman (2012): "Individual investor trading and return patterns around earnings announcements," *The Journal of Finance*, 67(2), 639–680.

Kastl, J. (2006): "Discrete Bids and Empirical Inference in Divisible Good Auctions," *Review of Economic Studies.*

Kelley, E., and P. Tetlock (2012): "How (Un)Informed Is Trading?," *Working Paper.*

Kyle, A. (1985): "Continuous auctions and insider trading," *Econometrica*, pp. 1315–1335.

Lee, C. (1992): "Earnings news and small traders: An intraday analysis," *Journal of Accounting and Economics*, 15(2-3), 265–302.

Lee, C., B. Mucklow, and M. Ready (1993): "Spreads, depths, and the impact of earnings information: An intraday analysis," *Review of Financial Studies*, 6(2), 345–374.

Linnainmaa, J. (2010): "Do limit orders alter inferences about investor performance and behavior?," *The Journal of Finance*, 65(4), 1473–1506.

Madhavan, A. (1992): "Trading mechanisms in securities markets," *Journal of Finance*, pp. 607–641.

McAdams, D. (2008): "Partial identification and testable restrictions in multi-unit auctions," *Journal of Econometrics*, 146(1), 74–85.

Milgrom, P., and N. Stokey (1982): "Information, trade and common knowledge," *Journal of Economic Theory*, 26(1), 17–27.

Milgrom, P., and R. Weber (1982): "A theory of auctions and competitive bidding," *Econometrica: Journal of the Econometric Society*, pp. 1089–1122.

Miller, R. (1984): "Job matching and occupational choice," *The Journal of Political Economy*, pp. 1086–1120.

O'Hara, M. (1998): *Market Microstructure Theory.* Blackwell.

Pagnotta, E. (2010): "Information and liquidity trading at optimal frequencies," *NYU Stern Working Paper.*

Pakes, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica*, 54(4), 755–784.

Ranaldo, A. (2004): "Order aggressiveness in limit order book markets," *Journal of Financial Markets*, 7(1), 53–74.

Roşu, I. (2009): "A dynamic model of the limit order book," *Review of Financial Studies*, 22(11), 4601–4641.

Rust, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica*, pp. 999–1033.

——— (1994): "Structural estimation of Markov decision processes," *Handbook of Econometrics*, 4, 3081–3143.

——— (1996): "Numerical dynamic programming in economics," *Handbook of Computational Economics*, 1, 619–729.

Sandås, P. (2001): "Adverse selection and competitive market making: Empirical evidence from a limit order market," *Review of Financial Studies*, 14(3), 705–734.

Smith, A. (1993): "Estimating nonlinear time-series models using simulated vector autoregressions," *Journal of Applied Econometrics*, 8(S1), S63–S84.

Vives, X. (2010): *Information and learning in markets: the impact of market microstructure*. Princeton University Press.

——— (2011): "Strategic supply function competition with private information," *Econometrica*, 79(6), 1919–1966.

Wolpin, K. (1984): "An estimable dynamic stochastic model of fertility and child mortality," *The Journal of Political Economy*, pp. 852–874.

8. Appendix

8.1. Liquidity Suppliers and the Optimal Supply Problem. Beginning with the utility of offering shares,

$$u_L(p_r, S_{r-1}, s_r^i, s_r^{-i}) = E_{\alpha,\epsilon,I}\left[s_r^i \mathbf{1}_{\{S_{r-1}+s_r^i+s_r^{-i}\le q\}}(p_r - v) + \frac{s_r^i}{s_r^i + s_r^{-i}}(q - S_{r-1})\mathbf{1}_{\{S_{r-1}<q<S_{r-1}+s_r^i+s_r^{-i}\}}(p_r - v)\right] \tag{14}$$

If $S_{r-1}+s_r^i+s_r^{-i} \le q$, then following the boundary conditions, $\frac{\alpha - p_{r+1}}{\gamma\sigma_\epsilon^2} - I \le S_{r-1}+s_r^i+s_r^{-i} \le \frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I$, or $S_{r-1}+s_r^i+s_r^{-i} < \frac{\alpha - p_{r+1}}{\gamma\sigma_\epsilon^2} - I$. The condition $S_{r-1} < q < S_{r-1}+s_r^i+s_r^{-i}$, results in an interior solution, so $q = \frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I$, so I have that $S_{r-1} < \frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I < S_{r-1} + s_r^i + s_r^{-i}$. Because ϵ is not correlated with I or α, I replace v with α. I can see that (14) is equal to,

$$s_r^i E\mathbf{1}_{\left\{S_{r-1}+s_r^i+s_r^{-i}\le\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}-I\right\}}(p_r-\alpha)+ \frac{s_r^i}{s_r^i + s_r^{-i}}E\mathbf{1}_{\left\{S_{r-1}<\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}-I<S_{r-1}+s_r^i+s_r^{-i}\right\}}(p_r-\alpha)\left(\frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I - S_{r-1}\right) \tag{15}$$

If I let $\phi_\alpha(x) = \frac{1}{\sqrt{2\pi}\sigma_\alpha}e^{-\frac{1}{2}\left(\frac{x-\mu_\alpha}{\sigma_\alpha}\right)^2}$ and $\phi_I(x) = \frac{1}{\sqrt{2\pi}\sigma_I}e^{-\frac{1}{2}\left(\frac{x-\mu_I}{\sigma_I}\right)^2}$ be the *pdf*s of α and I, then (15) is,

$$s_r^i \int_{-\infty}^{\infty}\int_{\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r}^{\infty}(p_r-\alpha)\phi_I(I)\phi_\alpha(\alpha)d\alpha dI+ \frac{s_r^i}{s_r^i+s_r^{-i}}\int_{-\infty}^{\infty}\int_{\gamma\sigma_\epsilon^2(S_{r-1}+I)+p_r}^{\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r}\left[\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2} - I - S_{r-1}\right](p_r-\alpha)\phi_I(I)\phi_\alpha(\alpha)d\alpha dI \tag{16}$$

And to find the optimal liquidity provision at each price p_r, a liquidity supplier chooses s_r^i such that (16) is maximized. I find a first order condition, and the derivative of (16) with respect to s_r^i is,

$$\begin{aligned}
(17)\quad & \int_{-\infty}^{\infty}\int_{\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r}^{\infty}(p_r-\alpha)\phi_I(I)\phi_\alpha(\alpha)d\alpha dI \\
&+\gamma\sigma_\epsilon^2 s_r^i\int_{-\infty}^{\infty}\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)\phi_\alpha(\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r)\phi_I(I)dI \\
&+\frac{s_r^{-i}}{(s_r^i+s_r^{-i})^2}\int_{-\infty}^{\infty}\int_{\gamma\sigma_\epsilon^2(S_{r-1}+I)+p_r}^{\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r}\left[\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}-I-S_{r-1}\right](p_r-\alpha)\phi_I(I)\phi_\alpha(\alpha)d\alpha dI \\
&-\gamma\sigma_\epsilon^2\frac{s_r^i}{s_r^i+s_r^{-i}}\int_{-\infty}^{\infty}\left[s_r^i+s_r^{-i}\right]\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)\phi_\alpha(\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r)\phi_I(I)dI
\end{aligned}$$

Which reduces to,

$$\begin{aligned}
(18)\quad & \int_{-\infty}^{\infty}\int_{\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r}^{\infty}(p_r-\alpha)\phi_I(I)\phi_\alpha(\alpha)d\alpha dI \\
&+\frac{s_r^{-i}}{(s_r^i+s_r^{-i})^2}\int_{-\infty}^{\infty}\phi_I(I)\int_{\gamma\sigma_\epsilon^2(S_{r-1}+I)+p_r}^{\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r}\left[\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}-I-S_{r-1}\right](p_r-\alpha)\phi_\alpha(\alpha)d\alpha dI
\end{aligned}$$

and that makes my second derivative

$$\begin{aligned}
(19)\quad & \gamma\sigma_\epsilon^2\int_{-\infty}^{\infty}\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)\phi_\alpha(\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r)\phi_I(I)dI \\
&+\frac{2s_r^{-i}}{(s_r^i+s_r^{-i})^2}\int_{-\infty}^{\infty}\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)\phi_I(I)\phi_\alpha(\gamma\sigma_\epsilon^2(S_{r-1}+s_r^i+s_r^{-i}+I)+p_r)dI
\end{aligned}$$

8.2. **Counterfactual Utilities.** I look at certainty-equivalent utility for the active trader, summed up over the different ranges of q, the number of shares the active trader chooses to buy. From these utilities I subtract utility from inventory holdings, as a baseline.

Different types of traders are categorized into the same cases as they are in other calculations, but I compute the optimal choice over orders here. I consider a supply schedule $S=\{S_1,\dots,S_M\}$, and $P=\{p_1,\dots,p_M\}$ given.

Case 1: p_r is s.t. $S_{r-1}<q<S_r$

$$q=\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}-I \tag{20}$$

Case 2: $q=S_r$

$$\frac{\alpha-p_{r+1}}{\gamma\sigma_\epsilon^2}-I<q<\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}-I \tag{21}$$

I also consider that the active trader is trying to maximize $-e^{-\gamma W}$ where W is wealth, from equation (7). This is equivalent to maximizing $E[W|\alpha,I]+\frac{\gamma}{2}V[W|\alpha,I]=(q+I)\alpha-\sum_{r=1}^{M}\max\{0,\min\{q,S_r\}-S_{r-1}\}p_r-\frac{\gamma}{2}(q+I)^2\sigma_\epsilon^2$. To find the average certainty-equivalent utility change for active traders, I must integrate over the spaces of α and I, taking into account latent supply S_0. Presented are the equations for M and all other $r = 1, \dots, M-1$ simultaneously. Calculating the counterfactual utility involves taking,

$$(q+I)\alpha - \sum_{r=1}^{M} \max\{0, \min\{q, S_r\} - S_{r-1}\} p_r - \frac{\gamma}{2}(q+I)^2 \sigma_\epsilon^2 \tag{22}$$

and summing over the the range of q.

Expected utility for those traders choosing to take liquidity at price p_M is,

$$\begin{aligned} E_{\alpha,I} \Bigg\{ & \mathbf{1}_{\{S_M \le \frac{\alpha - p_M}{\gamma\sigma_\epsilon^2} - I\}} \left[\alpha(S_M + I) - \frac{\gamma\sigma_\epsilon^2}{2}(S_M+I)^2 - \sum_{k=1}^{M} s_k p_k - p_1 S_0 \right] \\ & + \mathbf{1}_{\{S_{M-1} < \frac{\alpha - p_M}{\gamma\sigma_\epsilon^2} - I < S_M\}} \\ & \times \left[\alpha \left(\frac{\alpha - p_M}{\gamma\sigma_\epsilon^2} \right) - \frac{\gamma\sigma_\epsilon^2}{2} \left(\frac{\alpha - p_M}{\gamma\sigma_\epsilon^2} \right)^2 - \sum_{k=1}^{M-1} s_k p_k - p_1 S_0 - p_M \left(\frac{\alpha - p_M}{\gamma\sigma_\epsilon^2} - I - S_{M-1} \right) \right] \Bigg\} \end{aligned} \tag{23}$$

and for traders taking more liquidity than at price p_r,

$$\begin{aligned} E_{\alpha,I} \Bigg\{ & \mathbf{1}_{\{\frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I \le S_r \le \frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I\}} \left[\alpha(S_r + I) - \frac{\gamma\sigma_\epsilon^2}{2}(S_r+I)^2 - \sum_{k=1}^{r} s_k p_k - p_1 S_0 \right] \\ & + \mathbf{1}_{\{S_{r-1} < \frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I < S_r\}} \\ & \times \left[\alpha \left(\frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} \right) - \frac{\gamma\sigma_\epsilon^2}{2} \left(\frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} \right)^2 - \sum_{k=1}^{r-1} s_k p_k - p_1 S_0 - p_r \left(\frac{\alpha - p_r}{\gamma\sigma_\epsilon^2} - I - S_{r-1} \right) \right] \Bigg\} \end{aligned} \tag{24}$$

I sum up the above for p_M and $p_1, \dots, p_{M-1}$, respectively, then I add utility from latent supply,

$$E_{\alpha,I} \left\{ \mathbf{1}_{\{0 \le \frac{\alpha - p_1}{\gamma\sigma_\epsilon^2} - I \le S_0\}} \left[\alpha \left(\frac{\alpha - p_1}{\gamma\sigma_\epsilon^2} \right) - \frac{\gamma\sigma_\epsilon^2}{2} \left(\frac{\alpha - p_1}{\gamma\sigma_\epsilon^2} \right)^2 - p_1 \left(\frac{\alpha - p_1}{\gamma\sigma_\epsilon^2} - I \right) \right] \right\} \tag{25}$$

while subtracting utility of original endowment,

$$E_{\alpha,I} \left\{ \mathbf{1}_{\{0 \le \frac{\alpha - p_1}{\gamma\sigma_\epsilon^2} - I\}} \left[\alpha I - \frac{\gamma\sigma_\epsilon^2}{2} I^2 \right] \right\} \tag{26}$$

Whereas disregarding utility from latent supply means ignoring (25) and (26) and subtracting

$$E_{\alpha,I} \left\{ \mathbf{1}_{\{S_0 \le \frac{\alpha - p_1}{\gamma\sigma_\epsilon^2} - I\}} \left[\alpha(S_0 + I) - \frac{\gamma\sigma_\epsilon^2}{2}(S_0+I)^2 - p_1 S_0 \right] \right\}. \tag{27}$$

These give us respectively, in integral form,

$$\int_{\gamma\sigma_\epsilon^2(S_M+I)+p_M}^{\infty}\int_{-\infty}^{\infty}\left[\alpha(S_M+I)-\frac{\gamma\sigma_\epsilon^2}{2}(S_M+I)^2-\sum_{k=1}^{M}s_kp_k-p_1S_0\right]\phi_\alpha(\alpha)\phi_I(I)dId\alpha$$
$$+\int_{\gamma\sigma_\epsilon^2(S_{M-1}+I)+p_M}^{\gamma\sigma_\epsilon^2(S_M+I)+p_M}\int_{-\infty}^{\infty}\left[\alpha\left(\frac{\alpha-p_M}{\gamma\sigma_\epsilon^2}\right)-\frac{\gamma\sigma_\epsilon^2}{2}\left(\frac{\alpha-p_M}{\gamma\sigma_\epsilon^2}\right)^2\right.$$
$$\left.-\sum_{k=1}^{M-1}s_kp_k-p_1S_0-p_M\left(\frac{\alpha-p_M}{\gamma\sigma_\epsilon^2}-I-S_{M-1}\right)\right]\phi_\alpha(\alpha)\phi_I(I)dId\alpha \tag{28}$$

$$\int_{\gamma\sigma_\epsilon^2(S_r+I)+p_r}^{\gamma\sigma_\epsilon^2(S_r+I)+p_{r+1}}\int_{-\infty}^{\infty}\left[\alpha(S_r+I)-\frac{\gamma\sigma_\epsilon^2}{2}(S_r+I)^2-\sum_{k=1}^{r}s_kp_k-p_1S_0\right]\phi_\alpha(\alpha)\phi_I(I)dId\alpha$$
$$+\int_{\gamma\sigma_\epsilon^2(S_{M-1}+I)+p_r}^{\gamma\sigma_\epsilon^2(S_r+I)+p_r}\int_{-\infty}^{\infty}\left[\alpha\left(\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}\right)-\frac{\gamma\sigma_\epsilon^2}{2}\left(\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}\right)^2\right.$$
$$\left.-\sum_{k=1}^{M-1}s_kp_k-p_1S_0-p_r\left(\frac{\alpha-p_r}{\gamma\sigma_\epsilon^2}-I-S_{r-1}\right)\right]\phi_\alpha(\alpha)\phi_I(I)dId\alpha \tag{29}$$

$$\int_{\gamma\sigma_\epsilon^2I+p_1}^{\gamma\sigma_\epsilon^2(S_0+I)+p_1}\int_{-\infty}^{\infty}\left[\alpha\left(\frac{\alpha-p_1}{\gamma\sigma_\epsilon^2}\right)-\frac{\gamma\sigma_\epsilon^2}{2}\left(\frac{\alpha-p_1}{\gamma\sigma_\epsilon^2}\right)^2-p_1\left(\frac{\alpha-p_1}{\gamma\sigma_\epsilon^2}-I\right)\right]\phi_\alpha(\alpha)\phi_I(I)dId\alpha \tag{30}$$

$$\int_{\gamma\sigma_\epsilon^2I+p_1}^{\infty}\int_{-\infty}^{\infty}\left[\alpha I-\frac{\gamma\sigma_\epsilon^2}{2}I^2\right]\phi_\alpha(\alpha)\phi_I(I)dId\alpha \tag{31}$$

$$\int_{\gamma\sigma_\epsilon^2(S_0+I)+p_1}^{\infty}\int_{-\infty}^{\infty}\left[\alpha(S_0+I)-\frac{\gamma\sigma_\epsilon^2}{2}(S_0+I)^2-p_1S_0\right]\phi_\alpha(\alpha)\phi_I(I)dId\alpha. \tag{32}$$

8.3. **Solving for Equilibrium, detail.** For an n player symmetric equilibrium, I start at the highest price, p_M and work backward. Given existing supply S_{M-1}, the optimal supply solves the F.O.C. given in (18) where $r = M$, and $s_M^{-i} = (n-1)s_M^i$. While it is possible that a F.O.C. will not be reached, my computational methods avoid this. If the algorithm reaches the boundary at 0, naturally the liquidity supplier offers 0 shares. If, however, the boundary is at the computed supply matrix maximum, then the matrix is expanded and I restart the entire process. This ensures that liquidity is determined endogenously in the model. I call optimal supply for individual i at M, s_M^{i*}. Finding this value for all possible S_{M-1} and generating the value function at that price and quantity, $V(p_M, S_{M-1}) \equiv u(p_M, S_{M-1}, s_M^{i*}, (n-1)s_M^{i*})$, I can then solve for s_{M-1}^{i*} using (18) added to the derivative of $V(p_{M-1}, S_{M-2} + n\cdot s_{M-1}^i)$ for all S_{M-2}, generating $V(p_{M-1}, S_{M-2}) \equiv u(p_{M-1}, S_{M-2}, s_{M-1}^{i*}, (n-1)s_{M-1}^{i*}) + V(p_M, S_{M-2} + n\cdot s_{M-1}^{i*}(q), M)$. Working backward from

$M - 2$ to 1, I find a symmetric equilibrium discrete price supply function for all i. The function is in the form of a supply matrix.

Constructing a supply curve from a supply matrix involves interpolation. I begin with latent supply S_0 and compute optimal response in symmetric equilibrium at p_1. Moving forward, taking into account supply S_1, the optimal response at p_2 is computed. This process continues successively to p_M. My interpolation tool is a piecewise cubic hermite interpolating polynomial.

From the lowest ask price, I assume a level of latent supply, S_0, present in the book before the best ask, just below the best ask. This is to accommodate a variety of things. First of all, the model is made for one side of the market. While a story about institutional and active traders is fitting, it is restrictive. In response, I might see some spillover from the buy side of the market with active traders posting quickly executable liquidity. Second, there may be some impatient liquidity suppliers who do not play as part of an equilibrium that looks like what I consider here. Third, there are hidden orders present in the market at many points in time. While it is difficult or impossible to recreate these hidden orders as econometrician, I can allow for some liquidity to be hidden in the form of latent supply. I calculate the optimal supply schedule at each latent supply and interpolating according to calculated optimal values walking up the book. This brings us to the architecture of the supply matrix. Because of interpolation methods, I must have a bound on total aggregate supply. In addition, I need a delimeter over this space to give us a finite number of operations. In my calculations, I define limiting supply endogenously in my procedure so as to always reach a F.O.C. I call this $\bar{S}$, representing the maximum liquidity that can be offered by a single trader at each of the M prices, $n \cdot \bar{S}$ for all. The procedure is illustrated in figure 6.

Selection of lowest offer price and highest offer price is not arbitrary. The lowest price is determined by the bid-ask spread. The highest offer price is a point at which offering shares is no longer profitably meaningful, as discussed above. One might consider a market setting in which offering liquidity at any high price, no matter how high, would be profitable. However, those outcomes are improbable and economically irrelevant.

8.4. **Indirect Inference Method, Proofs, and Asymptotic Normality.** Here I outline in more detail my implementation of Indirect Inference in finding parameter estimates.

Over the course of the trading week, and even within the trading day, the composition of the market for a single asset changes. An individual instance of an order book tells us about a particular point in time. However, executions, limit orders, and cancellations can all change the shape of the order book. The same parameters could describe books with very different shapes, depending on the prices traders choose to offer liquidity. Hence, several different instances of the order book around the same point in time look different, but tell the same economic story. Additionally, placing or removing a single order does

FIGURE 6. Supply Matrix Procedure

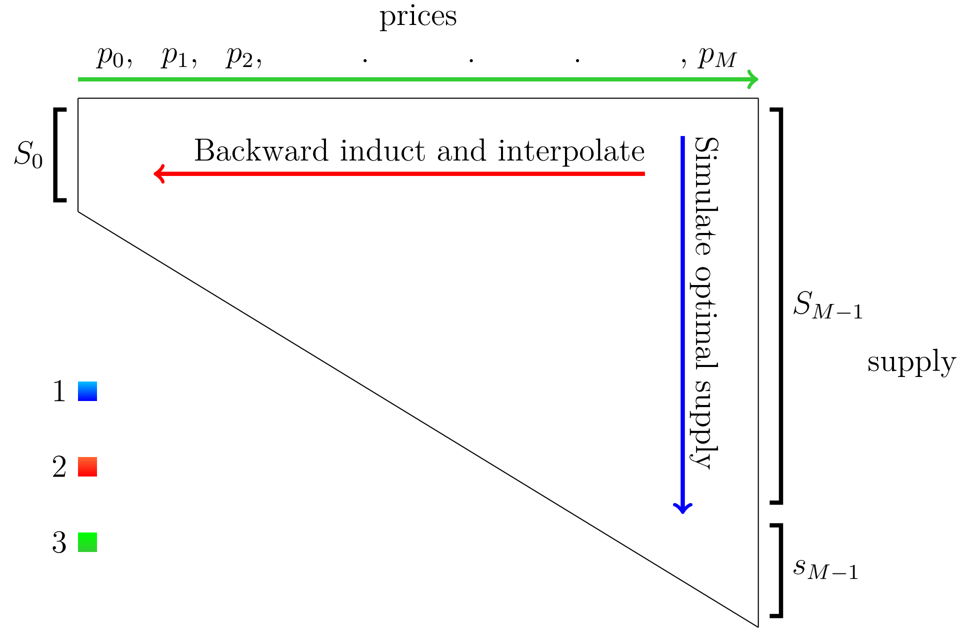

not immediately dissipate market conditions. I consider the order book from time T_i to T_j. I want to find $\theta = \{\sigma_\alpha, \sigma_\epsilon, \mu_I, \sigma_I, \sigma_{S_0}\}$–refer to Table 1 for parameter definitions–which describes equilibrium supply schedules over that period of time.

I look at a period of time (T_i, T_j), where $\triangle \equiv T_j - T_i$. I take K order book samples $t_1, t_2, \ldots, t_k, \ldots, t_K$, each of size L, where $t_k \equiv T_i + \frac{\triangle}{K+1}k$. Keeping in mind that total depth is M, a single supply curve is $(S_k^\ell, P_k^\ell) \equiv ([S_{k,1}^\ell, S_{k,2}^\ell, \ldots, S_{k,M}^\ell], [p_{k,1}^\ell, p_{k,2}^\ell, \ldots, p_{k,M}^\ell])$ for each sample, $k = 1, \ldots, K$ and all instances of the order book within each sample, $\ell = 1, \ldots, L$.

I introduce a pseudo-model for supply,

$$S_{k,m}^\ell = g(p_{k,m}^\ell) + \nu_{k,m}^\ell,\ \nu_{k,m}^\ell \sim IID(0, \sigma_\nu^2) \tag{33}$$

and I assume

$$S_{k,m}^\ell = \beta_1 + \beta_2 p_{k,m}^\ell + \beta_3 p_{k,m}^{\ell\,2} + \beta_4 log(p_{k,m}^\ell) + \nu_{k,m}^\ell + \eta_{k,m}^\ell,\ \eta_{k,m}^\ell \sim IID(0, \sigma_\eta^2) \tag{34}$$

so, given

A 1. $E(\eta_{k,m}^\ell | \nu_{k,m}^\ell, p_{k,m}^\ell) = 0$

$$S_{k,m}^\ell \simeq \beta_1 + \beta_2 p_{k,m}^\ell + \beta_3 p_{k,m}^{\ell\,2} + \beta_4 log(p_{k,m}^\ell) + \eta_{k,m}^\ell \tag{35}$$

or in matrix notation

$$\mathcal{S}_k^L = X_k^L \beta_k + \eta_k^L,\ \eta_k^L \sim IID(0, \sigma_\eta^2 I) \tag{36}$$

where $\mathcal{S}_k^L \equiv [S_{k,1}^1, S_{k,2}^1, \ldots, S_{k,M}^1, S_{k,1}^2, \ldots, S_{k,M}^L]'$, $X_k^L \equiv [\mathbf{1} \;\; P_k^L \;\; P_k^{L^2} \;\; log(P_k^L)]$, $P_k^L \equiv [p_{k,1}^1, p_{k,2}^1, \ldots, p_{k,M}^1, p_{k,1}^2, \ldots, p_{k,M}^L]'$, and $\eta_k^L \equiv [\eta_{k,1}^1, \eta_{k,2}^1, \ldots, \eta_{k,M}^1, \eta_{k,1}^2, \ldots, \eta_{k,M}^L]'$. My criterion is that of reducing the sum of squared residuals,

$$\max_{\beta \in B} C_{k,L}(\mathcal{S}_k^L, P_k^L, \beta) \equiv \max_{\beta \in B} - \sum_{\ell=1}^{L} \sum_{m=1}^{M} (s_{k,m}^\ell - x_{k,m}^\ell \beta)^2. \tag{37}$$

I assume this criterion tends asymptotically to a limit,

A 2. $\lim_{L \to \infty} C_{k,L}(\mathcal{S}_k^L, P_k, \beta) = C_{k,\infty}(Q_0, P_{k,\infty}, \theta_0, \beta)$

I assume that the limit criterion is continuous in β and has a unique maximum,

A 3. $\beta_{k,\infty} \equiv \arg\max_{\beta \in B} C_{k,\infty}(Q_0, P_{k,\infty}, \theta_0, \beta)$

I use the ordinary least squares estimator, $\hat{\beta}(\mathcal{S}_k^L, P_k^L) = [X_k^{L\prime} X_k^L]^{-1} X_k^{L\prime} \mathcal{S}_k^L$. Because prices are bounded, and each element of X_k^L is a continuous function of price, $\exists$ vector c s.t. $E\, x_{k,m}^\ell x_{k,m}^{\ell\;\prime} = c$, $\forall \ell$, m. Here $x_{k,m}^\ell \equiv [1 \;\; p_{k,m}^\ell \;\; p_{k,m}^{\ell\,2} \;\; log(p_{k,m}^\ell)]$. By the law of large numbers, $\frac{1}{L}(X_k^{L\prime} X_k^L) \to c$ and so, $\operatorname*{plim}_{L \to \infty} \frac{1}{L}(X_k^{L\prime} X_k^{L\prime}) = S_{X_k^{L\prime} X_k^L}$ exists.

From A 1, $E(\eta_{k,m}^\ell | x_{k,m}^\ell) = \mathbf{0}$, and therefore,

$$\operatorname*{plim}_{L \to \infty} \frac{1}{L} X_k^{L\prime} \eta_k^L = \operatorname*{plim}_{L \to \infty} \frac{1}{L} \sum_{\ell=1}^{L} \sum_{m=1}^{M} x_{k,m}^{\ell\prime} \eta_{k,m}^\ell = \mathbf{0} \tag{38}$$

This gives us that $\hat{\beta}_k = \hat{\beta}(\mathcal{S}_k^L, P_k^L)$ is a consistent estimator of $\beta_{k,\infty}$.

I introduce the *binding function* from Gouriéroux, Monfort, and Renault (1993),

$$b(Q, P_{k,\infty}, \theta) \equiv \arg\max_{\beta \in B} C_\infty(Q, P_{k,\infty}, \theta, \beta) \tag{39}$$

The binding function links the auxiliary parameters of the pseudo-model with the parameters of the true model. I have

$$\beta_{k,\infty} = b(Q_0, P_{k,\infty}, \theta_0) \tag{40}$$

The binding function is analytically intractable. I find it numerically, however. For a θ, $\bar{P}_k^L \equiv \frac{1}{L} \sum_{\ell=1}^{L} P_k^\ell$, $\bar{X}_k \equiv [\mathbf{1} \;\; \bar{P}_k^L \;\; \bar{P}_k^{L\,2} \;\; log(\bar{P}_k^L)]$ given, I consider Λ supply curves, $\tilde{S}^\lambda(\theta, \bar{P}_k^L)$, $\lambda = 1, \ldots, \Lambda$, generated from drawings of latent supply S_0. Latent supply represents noisy liquidity supply which may be some spillover from the other side of the order book, a few impatient suppliers, or hidden orders. S_0 is drawn from a truncated normal distribution with mean calibrated to the empirical supply curves, and variance $\sigma_{S_0}^2$. From the supply matrix determined by θ and each draw of S_0^λ, $\theta = 1, \ldots, \Lambda$, I interpolate to find the respective supply curve, $\tilde{\mathcal{S}}^\lambda(\theta, \bar{P}_k^L)$. For each curve I can find the respective estimates solving,

$$\max_{\beta \in B} C_{k,L}(\tilde{S}^\lambda(\theta, \bar{P}_k^L), \bar{P}_k^L, \beta) \tag{41}$$

for which I consider the OLS estimate

$$\tilde{\beta}^{\lambda}(\theta, \bar{P}_k^L) = [\bar{X}_k^{L\prime}\bar{X}_k^L]^{-1}\bar{X}_k^{L\prime}\tilde{\mathcal{S}}^{\lambda}(\theta, \bar{P}_k^L). \tag{42}$$

I assume that the vector $\bar{P}_k^L$ reaches asymptotically the population prices from sample k.

A 4. $\lim_{L\to\infty} \bar{P}_k^L = P_{k,\infty}$

And so, I define $\lim_{L\to\infty} \bar{X}_k^L = X_{k,\infty}$. This means that as L tends to infinity, this solution tends toward $\tilde{\beta}^{\lambda}(\theta, P_{k,\infty}) = [X'_{k,\infty}X_{k,\infty}]^{-1}X'_{k,\infty}\tilde{\mathcal{S}}^{\lambda}(\theta, P_{k,\infty})$, which solves $\max_{\beta\in B} C_{\infty}(Q, P_{k,\infty}, \theta, \beta)$. Therefore,

$$\lim_{L\to\infty} \tilde{\beta}^{\lambda}(\theta, \bar{P}_k^L) = b(Q, P_{k,\infty}, \theta) \tag{43}$$

and is therefore a consistent functional estimator of the binding function. I define the following,

$$b_K(\theta) \equiv [b(Q, P_{1,\infty}, \theta); b(Q, P_{2,\infty}, \theta), \dots, b(Q, P_{K,\infty}, \theta)] \tag{44}$$

$$\hat{\beta}_K^L \equiv [\hat{\beta}(\mathcal{S}_1^L, P_1^L); \hat{\beta}(\mathcal{S}_2^L, P_2^L); \dots; \hat{\beta}(\mathcal{S}_K^L, P_K^L)] \tag{45}$$

$$\tilde{\beta}_K^{\lambda,L}(\theta) \equiv [\tilde{\beta}^{\lambda}(\theta, \bar{P}_1^L); \tilde{\beta}^{\lambda}(\theta, \bar{P}_2^L); \dots; \tilde{\beta}^{\lambda}(\theta, \bar{P}_K^L)] \tag{46}$$

And make the assumptions

A 5. $\bar{P}_k^L \neq \bar{P}_j^L$, $\forall\ j \neq k$

A 6. $\bar{X}_k^{L\prime}\frac{\partial\tilde{S}^{\lambda}}{\partial\theta'}(\theta_0, \bar{P}_k^L)$ has full column rank, $\forall\ k = 1, \dots, K$.

Claim 1. *For number of samples* $K \geq \frac{|\theta|}{|\beta|}$, $\frac{\partial b_K}{\partial\theta'}(\theta_0)$ *is of full column rank.*

Proof. $\bar{X}_1^{L\prime}\frac{\partial\tilde{S}^{\lambda}}{\partial\theta'}(\theta_0, \bar{P}_1^L)$ has full column rank. It is therefore clear that $\tilde{\beta}^{\lambda}(\theta_0, \bar{P}_1^L) = [\bar{X}_1^{L\prime}\bar{X}_1^L]^{-1}\bar{X}_1^{L\prime}\tilde{\mathcal{S}}^{\lambda}(\theta_0, \bar{P}_1^L)$ has full column rank. Because $\lim_{L\to\infty}\tilde{\beta}^{\lambda}(\theta_0, \bar{P}_1^L) = b(Q, P_{1,\infty}, \theta_0)$, $\frac{\partial b}{\partial\theta'}(Q, P_{1,\infty}, \theta_0)$ is also of full column rank. The same is true for $\frac{\partial b}{\partial\theta'}(Q, P_{k,\infty}, \theta_0)$ for $k = 2, \dots, K$. Each matrix of derivatives is independent of all others because of assumption 5. I stack up binding functions on top of each other, and because $K \geq \frac{|\theta|}{|\beta|}$, I get a matrix of full column rank. This matrix is contained in $\frac{\partial b_K}{\partial\theta'}(\theta_0)$, so it is of full column rank.

□

Using both pseudo-model objects, (45) and (46), the Indirect Inference estimator, $\hat{\theta}_K^{\Lambda,L}$ optimizes the following criterion

$$\hat{\theta}_K^{\Lambda,L} \equiv \min_{\theta\in\Theta}\left[\hat{\beta}_K^L - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}_K^{\lambda,L}(\theta)\right]'\Omega\left[\hat{\beta}_K^L - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}_K^{\lambda,L}(\theta)\right] \tag{47}$$

Proposition 2. *Under assumptions 1–6, $\hat{\theta}_K^{\Lambda,L}$ is a consistent estimator of θ_0*

Proof. The F.O.C. are

$$\left[\frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\frac{\partial\tilde{\beta}_K^{\lambda,L\prime}}{\partial\theta}(\hat{\theta}_K^{\Lambda,L})\right]\Omega\left[\hat{\beta}_K^L - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda,L}(\hat{\theta}_K^{\Lambda,L})\right] = \mathbf{0} \tag{48}$$

$$\left[\frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\frac{\partial\tilde{\beta}_K^{\lambda,L\prime}}{\partial\theta}(\hat{\theta}_K^{\Lambda,L})\right]\Omega\left[\hat{\beta}_K^L - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda,L}(\theta_0) - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\frac{\partial\tilde{\beta}_K^{\lambda,L}}{\partial\theta'}(\theta_0)[\hat{\theta}_K^{\Lambda,L}-\theta_0]\right] \simeq \mathbf{0} \tag{49}$$

$$\sqrt{L}(\hat{\theta}_K^{\Lambda,L}-\theta_0) \simeq \left[\frac{\partial b_K'}{\partial\theta}(\theta_0)\Omega\frac{\partial b_K}{\partial\theta'}(\theta_0)\right]^{-1}\frac{\partial b_K'}{\partial\theta}(\theta_0)\Omega\sqrt{L}\left[\hat{\beta}_K^L - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda,L}(\theta_0)\right] \tag{50}$$

Because $\hat{\beta}_K^L$ and $\frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda,L}(\hat{\theta}_K^{\Lambda,L})$ are both consistent estimators of $\beta_{k,\infty}$, the estimator is consistent. □

I also show asymptotic normality of the estimator.

A 7. $\forall$ k and λ, $\displaystyle\operatorname*{plim}_{L\to\infty}\frac{\partial^2 C_{k,L}}{\partial\beta\,\partial\beta'}(\tilde{S}^\lambda(\theta,\bar{P}_k^L),\bar{P}_k^L,\beta_0) = \frac{\partial^2 C_{k,\infty}}{\partial\beta\,\partial\beta'}(Q_0,P_{k,\infty},\theta_0,\beta_0)$

Asymptotically, by assumption 7 and viewing $\mathcal{S}_k^L$ and P_k^L as simulations, so $\forall$ k,

$$\frac{\partial C_{k,L}}{\partial\beta}(\mathcal{S}_k^L,P_k^L,\hat{\beta}(\mathcal{S}_k^L,P_k^L)) = 0 \tag{51}$$

$$\sqrt{L}\frac{\partial C_{k,L}}{\partial\beta}(\mathcal{S}_k^L,P_k^L,\beta_{k,\infty}) + \frac{\partial^2 C_{k,L}}{\partial\beta\,\partial\beta'}(\mathcal{S}_k^L,P_k^L,\beta_{k,\infty})\sqrt{L}(\hat{\beta}(\mathcal{S}_k^L,P_k^L)-\beta_{k,\infty}) \simeq 0 \tag{52}$$

$$\sqrt{L}(\hat{\beta}(\mathcal{S}_k^L,P_k^L)-\beta_{k,\infty}) \simeq -\frac{\partial^2 C_{k,\infty}}{\partial\beta\,\partial\beta'}(Q_0,P_{k,\infty},\theta_0,\beta_{k,\infty})\sqrt{L}\frac{\partial C_{k,L}}{\partial\beta}(\mathcal{S}_k^L,P_k^L,\beta_{k,\infty}) \tag{53}$$

By a similar argument,

$$\sqrt{L}(\tilde{\beta}^\lambda(\theta,\bar{P}_k^L)-\beta_{k,\infty}) \simeq -\frac{\partial^2 C_{k,\infty}}{\partial\beta\,\partial\beta'}(Q_0,P_{k,\infty},\theta_0,\beta_{k,\infty})\sqrt{L}\frac{\partial C_{k,L}}{\partial\beta}(\tilde{S}^\lambda(\theta,\bar{P}_k^L),\bar{P}_k^L,\beta_{k,\infty}) \tag{54}$$

By 53 and 54, I have that,

$$\begin{aligned}\sqrt{L}\left[\hat{\beta}_k^L(\mathcal{S}_k^L,P_k^L) - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda,L}(\theta_0,\bar{P}_k^L)\right] &\simeq -\frac{\partial^2 C_{k,\infty}}{\partial\beta\,\partial\beta'}(Q_0,P_{k,\infty},\theta_0,\beta_{k,\infty})\\ &\times\left[\sqrt{L}\frac{\partial C_{k,L}}{\partial\beta}(\mathcal{S}_k^L,P_k^L,\beta_{k,\infty}) - \sqrt{L}\frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\frac{\partial C_{k,L}}{\partial\beta}(\tilde{S}^\lambda(\theta,\bar{P}_k^L),\bar{P}_k^L,\beta_{k,\infty})\right]\end{aligned} \tag{55}$$

Therefore, this difference is asymptotically normal with zero mean. Therefore the difference in stacks,

$$\sqrt{L}\left[\hat{\beta}_K^L - \frac{1}{\Lambda}\sum_{\lambda=1}^{\Lambda}\tilde{\beta}^{\lambda,L}(\theta_0)\right] \tag{56}$$

is also asymptotically normal with zero mean.