Modelling the Dynamics of Intermediated Exchange

Robert Engle

Robert Ferstenberg

Acknowledgments

- Abel Noser “Ancerno”
- WRDS: CRSP and TAQ
- NYU Volatility Institute
- Stern Computing

References

- Frazzini, A., Israel R., Moskowitz T. J. (2012). Trading Cost of Asset Pricing Anomalies. *Working Paper.*
- Angel, J. J., Spatt, C., & Harris, L. (2015). Equity Trading in the 21st Century: An Update. *Quarterly Journal of Finance, 5*(1).
- Engle, R., Ferstenberg, R., & Russell, J. (2012). *Measuring and Modeling Execution Cost and Risk* (Vol. 38(2)). Journal of Portfolio Management.
- Frazzini, A., Israel, R., & Moskowitz, T. J. (2012). *Trading Costs of Pricing Anomalies.* working papaer.
- Grinold, R., & Kahn, R. (1999). *Active Portfolio Management* (2 ed.). McGraw-Hill.
- Harris, L. (2003). *Trading & Exchanges.* New York: Oxford University Press, Inc.
- Hasbrouck, J. (2007). *Empirical Market Microstructure.* New York: Oxford University Ppress, Inc.
- ITG Inc. (2010). *ITG Global Trading Cost Review 2010Q1.* ITG Inc.
- Keim, D. B., & Madhavan, A. (1995). Anatomy of the trading process Empirical evidence on the behavior of institutional traders. *Journal of Financial Economics, 37*, 371-398.
- Keim, D. B., & Madhavan, A. (1996). The Upstairs Market for Large-Block Transactions: Analysis and Measurement of Price Effects. *The Review of Financial Studies, 9*(Spring), 1-36.
- Keim, D. B., & Madhavan, A. (1998). The Cost of Institutional Equity Trades. *Financial Analysts Journal, 54*, 50-69.
- Khandani, A. E., & Lo, A. W. (2008). What happened to the quants in August 2007? Evidence from factors and transactions data. *NBER Working Paper No. 14465*.
- Perold, A. (1988). The implementation shortfall: Paper versus reality. *Journal of Portfolio Management, 14*(3), 4-9.

Overview

- We model the dynamics of intermediated exchange for equities using the Ancerno dataset
- As an example application of our model, we decompose the components of realized cost for institutional trades and find a 40% reduction in forecasting error relative to existing methods
  - Resulting from a combination of extra degrees of freedom in the models and additional explanatory variables
- We additionally compare the forecasts from our model to the Amihud illiquidity measure

THE MECHANICS OF INTERMEDIATED EXCHANGE

Introduction

The Mechanics Of Intermediated Exchange

- We begin by classifying market participants into two categories:
  - Beneficial owners: these are investors who access public markets to either acquire positions that they will hold for one or more days or to liquidate positions that they have held for one or more days
    - Examples include institutions, hedge funds, and retail
  - Intermediaries: these were predominantly specialists/market makers pre-decimalization and are now predominantly HFT post-decimalization
    - Their primary objective is to provide intra-day immediacy between beneficial buyers and sellers and to collect a toll for facilitating trade
    - They carry relatively small overnight positions as a fraction of total volume
- The distinction between beneficial owners and intermediaries cannot be exact
  - Several counter-examples in each category
    - E.G.: hedge funds and retail can and do initiate and liquidate within the day
  - It is important only as a basis for subsequent assumptions that we will make about the Ancerno dataset
- For a given symbol/date, we assume that the shares sold by beneficial owners in aggregate must be close to the shares acquired by beneficial owners in aggregate since the intermediaries take relatively small overnight positions

Introduction

The Mechanics Of Intermediated Exchange

- The diagram on the right shows the sign and relative magnitude of aggregate Implementation Shortfall (IS) across all beneficial buyers and sellers for a given date/symbol under different return scenarios
  – $IS = b\left(\frac{\tilde{p}_x}{p_o} - 1\right)$
  – Where:
    - b = order side, +1 for buy, -1 for sell.
    - $\tilde{p}_x$= average execution price
    - p_o = arrival price
- The average execution price aggregated across all beneficial buy orders for a symbol/date will generally be greater than, or at least equal to, the average execution price aggregated across all beneficial sell orders
  – Dependent on degree of intermediary participation
  – Some individual buy orders can execute at prices lower than some sell orders and vice-versa, but not in aggregate
- Similarly for the aggregate market VWAP: aggregate beneficial buy orders will execute at prices at or above the VWAP and aggregate beneficial sell orders at or below
  – Some individual orders can execute at better than VWAP but not in aggregate

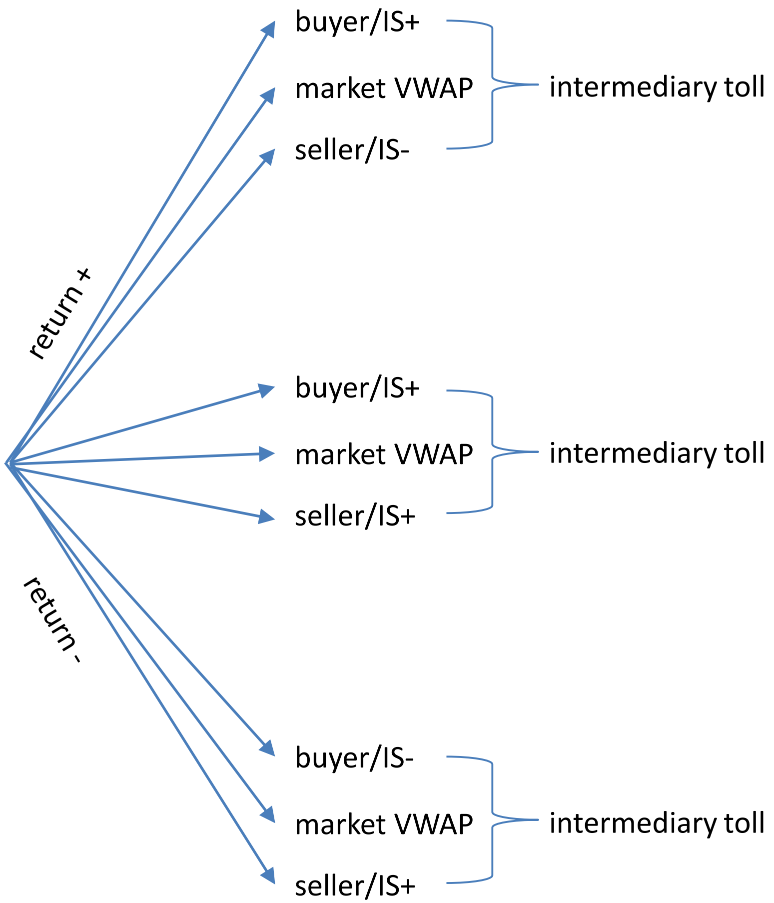

Introduction

The Mechanics Of Intermediated Exchange

- The various combinations of order side and market return scenarios can be categorized into four states as shown in the table below
  - Categorized by order side = sign return
- We use the terms "liquidity demander" and "liquidity supplier" as categorical variables
  - Not necessarily "informed" vs. "noise" traders as is typically described in the literature
  - Demanders chase price, crossing the spread and paying a premium
  - Suppliers are chased, moving back from the quote and receiving the premium
  - We will see that overall market return is a major determinant of the outcome

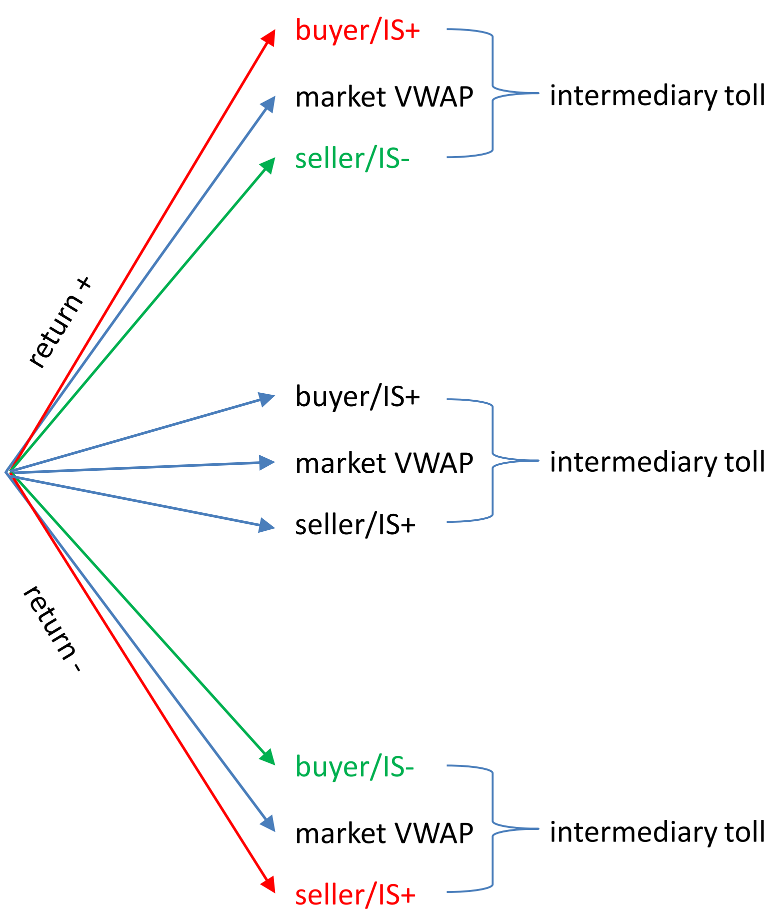

Order Side	VWAP Return	Sign of IS	Liquidity Category
Buy	+	+	Demander
Sell	+	-	Supplier
Buy	-	-	Supplier
Sell	-	+	Demander

Introduction

The Mechanics Of Intermediated Exchange

- In general, if one wants to examine the relation between a variable 'y' and a set of 'x' variables with the belief that the relation is different when 'y' is positive from when it is negative:

$$E(y|y>0,x) = x\beta^{+}$$
$$E(y|y<0,x) = x\beta^{-}$$

- It is easy to estimate such a relationship by least squares
- How would a forecast be made?

$$E(y) = E(y|x,y>0)P(y>0|x) + E(y|x,y<0)P(y<0|x)$$

- Proof:

$$\begin{aligned} E(y|x) &= \int_{-\infty}^{0} yf(y|x)dy + \int_{0}^{\infty} yf(y|x)dy \\ &= \int_{-\infty}^{0} yf(y|y<0,x)P(y<0|x)dy + \int_{0}^{\infty} yf(y|y>0,x)P(y>0|x)dy \\ &= E(y|y<0,x)P(y<0|x) + E(y|y>0,x)P(y>0|x) \end{aligned}$$

- If the two values of beta are different, then this is the efficient estimation method
  - It should be maximum likelihood under standard assumptions on the disturbances
- Forecasting however requires an additional logit estimate that should depend upon x
  - Endogeneity is not an issue

Introduction

Considerations For Model Estimation

- Without complete order data, our estimates must be made in terms of mean values
  - Not specific to any single symbol/date
- For example, the fraction of aggregate order volume that originates from beneficial owners (not total market volume) in a symbol/date that incurs positive IS should be at least 50%
  - We would expect the fraction across all symbol/dates in the Ancerno dataset to be greater than 50% because both sides can incur positive IS and because of missing order data from broker/dealers, etc.
    - It cannot be much less than 50% in the Ancerno dataset on average if measurements of aggregate cost across buyers and sellers is to be non-negative
  - We will see from the data that, when computed across all symbol/dates in each year, the fraction is very close to 50%, whether by frequency, share volume or notional volume
    - That measurement does not by itself prove that the fraction is close to 50% for each symbol/date.
    - Without complete order data from all beneficial owners, this cannot be proven definitively, but we believe that is strongly implied by the theory and our subsequent empirical results

Introduction

Considerations For Model Estimation

- To capture the aggregate effect of all orders that originate from beneficial owners, we aggregate orders in the Ancerno dataset by symbol/date/side
  - The shares are summed and the share weighted mean of execution prices are calculated by symbol/date/side
    - That yields at least one aggregated order per symbol/date (buy or sell) and at most two (buy and sell)

Introduction

Sample Bias: Censoring and Truncation

- We chose a combination of censoring and truncation:
  - Aggregate the raw orders by symbol/date/side
  - Censor orders with fraction of contemporaneous market volume > 1.0
    - Affects 1.2% of the sample
  - Truncate:
    - Buy orders with $\tilde{p}_x < p_v$ to p_v
    - Sell orders with $\tilde{p}_x > p_v$ to p_v
    - Affects about 48% of the aggregated sample
  - Censor the 1st and 100th percentiles of the resulting IS by year
- Clearly, mean IS will be shifted upward and centered around VWAP
  - Resulting values consistent with those reported in other sources

Introduction

Sample Bias: Estimating the Dynamics of Intermediated Exchange

- Our goal is to both estimate the dynamics of intermediated exchange and to mitigate the bias in sample population imbalance of supply and demand orders
  - We will perform separate regressions for the liquidity demand and liquidity supply cross sections of the dataset and subsequently test if imbalance bias is mitigated
- By estimating the model separately on the individual supply and demand cross sections, we are implicitly including endogenous variables on the RHS, raising a couple of potential issues:
  - The distributions of the dependent variables will be truncated
    - Subsequent out-sample forecasts suggest that empirically, this is not a problem
  - The fit statistics will be erroneously inflated and are useful for comparison purposes only
    - We rely on tests of in-sample forecast error to compare model alternatives
- This structure additionally alters the interpretation of what is actually being modeled
  - The LHS essentially becomes absolute return between average execution price and open price, making this a model of the marginal contribution of incremental volume to absolute excess return
  - Another interpretation might be that this is a model of volatility conditional on historical volatility and incremental volume

Model Estimation

Outline

- Estimation of the dynamics of intermediated exchange will be presented in three phases:
  1. Dataset features and summary statistics
  2. Evidence from summary statistics
  3. Evidence from estimation

Model Estimation

Study Time Period and Symbol Universe

- Time period: 1999 through 2013 Q3
  - Chosen to include the pre-decimalization period, the quant meltdown of 2007, the financial crisis and after
- The study symbol universe was limited to US listed common stock in US domiciled companies found in the CRSP dataset:
  - No ETFs, ADRs, REITs, Closed end funds, etc.
  - Some Ancerno data is lost due to symbol mis-match with CRSP and TAQ

Model Estimation: Dataset

Explanatory Market Data Variables

- The explanatory market data variables were extracted from WRDS/CRSP and WRDS/TAQ and are defined as follows:
  - Open and close price
    - Denoted by p_o, p_c
    - Extracted from CRSP
  - VWAP of all trades after the open and through the close
    - Denoted by p_v
    - Extracted from TAQ
      - The process is described in the appendix
  - Market volume
    - Denoted by V
    - Extracted from CRSP
  - Volatility: the lagged 21 day sample zero mean standard deviation of open/vwap return scaled to the duration of a full US trading day (6.5 hours), not annualized
    - Denoted by σ
  - AD$: the lagged 21 day sample mean notional dollar volume scaled to a full US trading day
    - This variable also serves as a proxy for capitalization
    - Denoted by $\bar{V}$
  - Spread: the lagged 21 day sample mean of closing bid/ask spread (or last quote) divided by the open price
- We will frequently use the ratio of Fill Notional, denoted by q, to AD$ as an explanatory variable
  - The product of the Ancerno fill quantity and the CRSP open price
  - Denoted by $\frac{q}{\bar{V}}$

Model Estimation

Summary Statistics

Year	Sample Count	Unique Symbol Count	Open Price	AD$	Spread	Volatility	Fill Shares	Fill Notional	Fill ADV
1999	609,502	4,919	$ 35	$ 38,968,518	1.6%	0.021	51,711	$ 2,346,806	14%
2000	652,744	4,929	$ 35	$ 56,567,502	1.4%	0.029	67,076	$ 3,108,018	13%
2001	629,376	4,339	$ 27	$ 44,517,599	1.1%	0.026	89,559	$ 2,819,669	13%
2002	826,015	4,127	$ 25	$ 34,158,041	0.9%	0.022	100,055	$ 2,627,508	13%
2003	785,901	4,045	$ 25	$ 29,793,219	0.5%	0.017	77,430	$ 2,040,741	12%
2004	898,640	4,200	$ 29	$ 31,417,519	0.3%	0.015	98,983	$ 3,031,619	13%
2005	981,188	4,141	$ 32	$ 36,735,634	0.2%	0.014	83,050	$ 2,706,548	12%
2006	869,041	3,823	$ 34	$ 51,532,704	0.2%	0.013	54,272	$ 1,831,301	7%
2007	903,072	3,943	$ 37	$ 71,813,352	0.2%	0.014	47,616	$ 1,724,355	6%
2008	821,714	3,618	$ 31	$ 88,818,739	0.3%	0.026	45,445	$ 1,378,001	5%
2009	793,031	3,553	$ 26	$ 68,444,455	0.3%	0.024	39,530	$ 889,875	4%
2010	773,875	3,317	$ 32	$ 78,677,863	0.1%	0.015	28,971	$ 802,110	4%
2011	925,026	3,558	$ 34	$ 62,556,438	0.1%	0.016	63,990	$ 2,054,547	6%
2012	1,039,528	3,442	$ 34	$ 58,352,789	0.1%	0.015	69,138	$ 2,344,418	7%
2013	804,315	3,491	$ 40	$ 58,947,878	0.1%	0.013	63,534	$ 2,423,595	7%

Model Estimation

Summary Statistics

- The mean IS after aggregation, truncation and censoring is similar in progression and magnitude to other studies
  - E.G.: ITG website

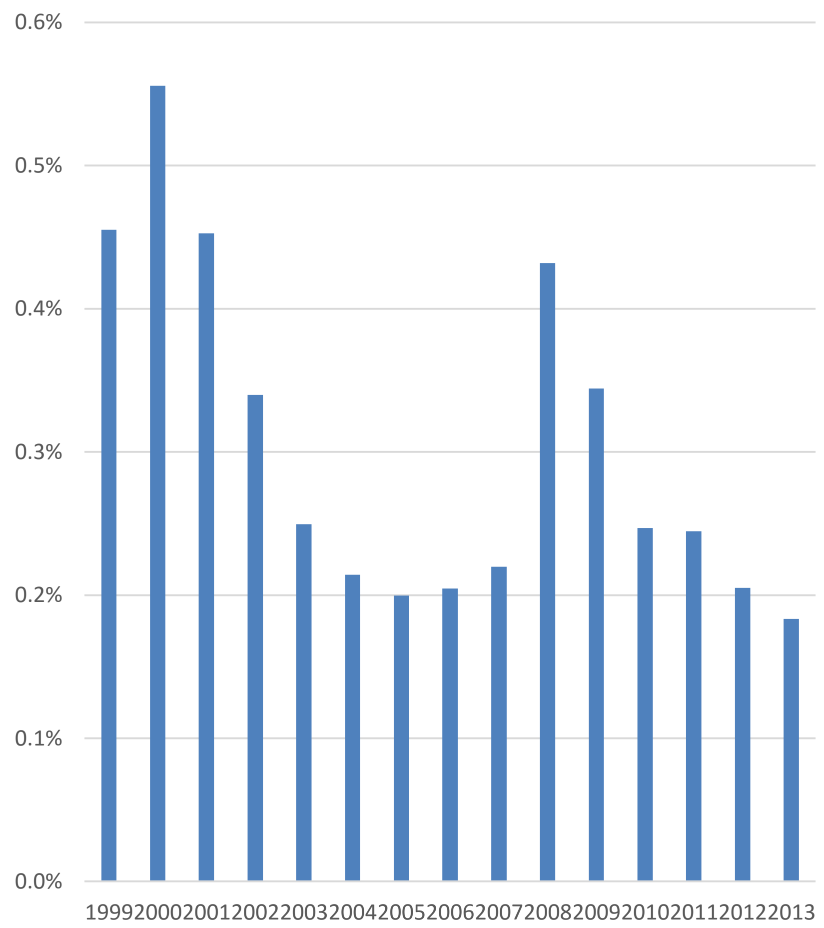

Model Estimation

Summary Statistics

- Some additional variable definitions are required for the slides that follow:
  - $\tilde{p}_x$ = execution price for a fill
  - $\tilde{r}_x = \frac{\tilde{p}_x}{p_o} - 1$: return of execution price relative to the open
    - Labelled OXRET in tables and graphs
    - IS = $b\tilde{r}_x$
  - $r_v = \frac{p_v}{p_o} - 1$: return of market vwap relative to the open
    - Labelled OVRET in tables and graphs

Model Estimation

Summary Statistics

- The mean values of $\tilde{r}_x$ for buy and sell orders by deciles of r_v
- We expect that $\tilde{r}_x$ for buy orders $\geq r_v \geq \tilde{r}_x$ for sell orders
  - $\tilde{r}_x$ for buy orders should be more positive/less negative
  - Returns for buy and sell orders should "bracket" r_v

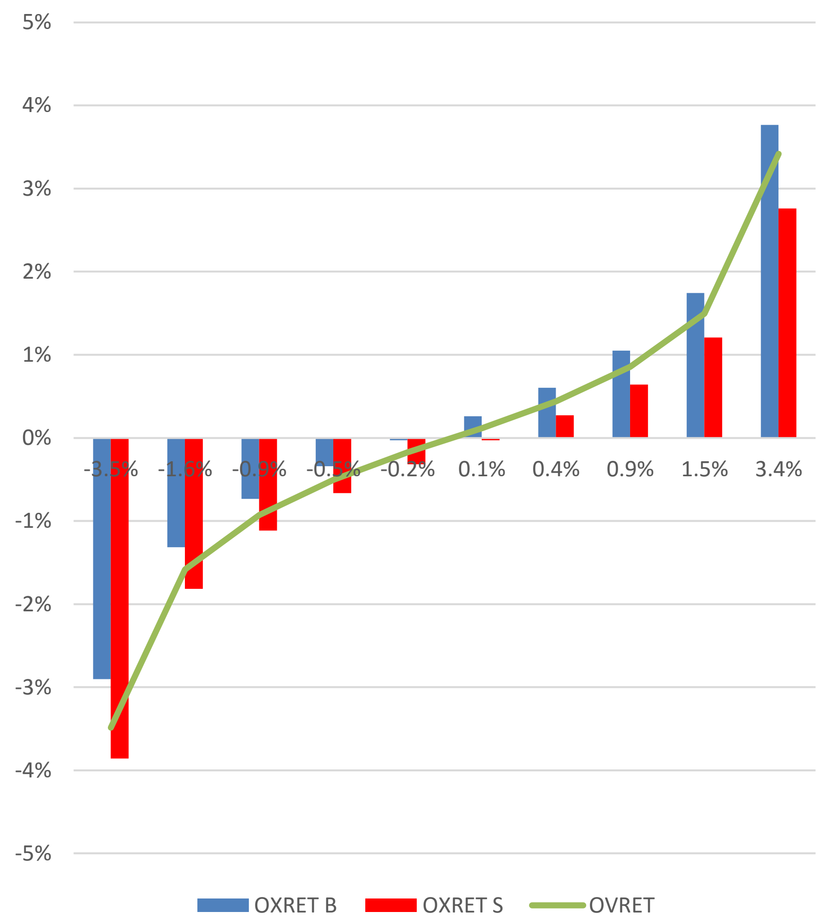

Model Estimation

Summary Statistics

- We next partition the data by the sign of realized IS
  - Orders with positive IS will be designated as liquidity demand and the complement as liquidity supply
    - A very small percentage are neither
- Summary statistics for these cross-sections will be presented in the slides that follow
- Variables for the liquidity demand cross-section of the data will have the suffix ".d", with ".s" for liquidity supply

Model Estimation

Summary Statistics

- Equal weighted mean IS by year by supply/demand cross section (EQ-IS.d, EQ-IS.s) compared with:
  - Mean volatility (σ):
    - Note how the trend through time corresponds closely with each cross section and with similar magnitude
  - Mean equal weighted IS (EQ-IS)
    - Note the relative magnitude of IS by supply/demand versus the overall mean
      - Illustrates what happens during liquidation under duress
- Illustrates how an institutional order that is worked over the course of the trading day will realize an average execution price that is highly correlated with market VWAP

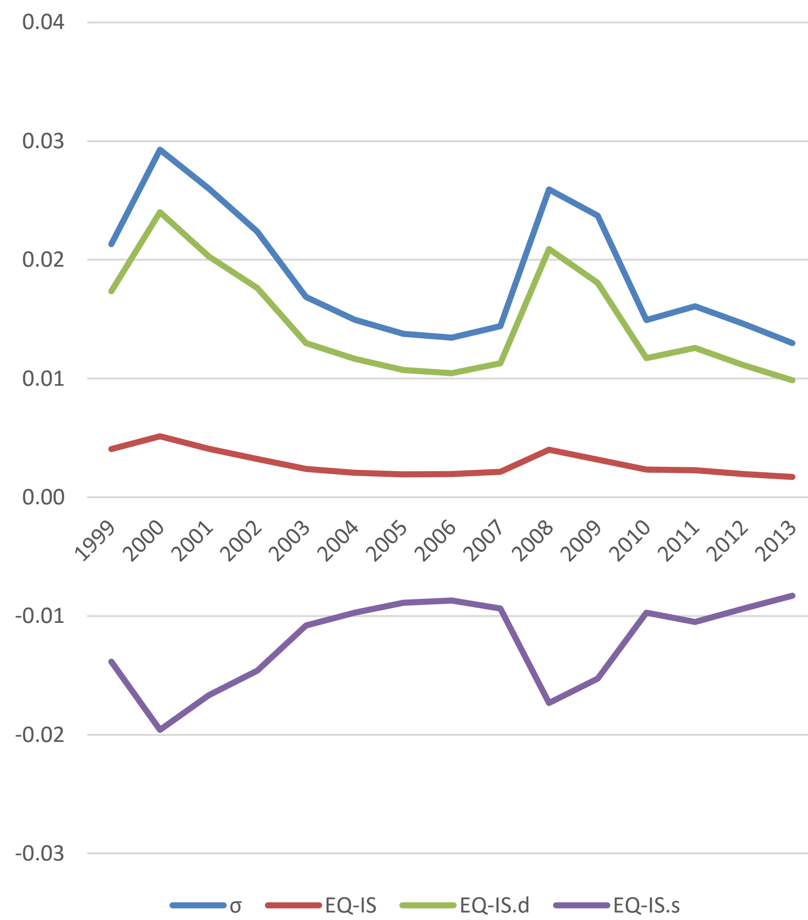

Model Estimation

Summary Statistics

- The ratio of the mean of -IS.s to the mean of IS.d
- The ratio is less than 1.0 across all years which is consistent with the claim that supplier IS will be lower in magnitude (closer to zero) than demander IS
  - Subsequently tested via estimation
- The spread between EQ-IS.d and EQ-IS.s has narrowed since decimalization

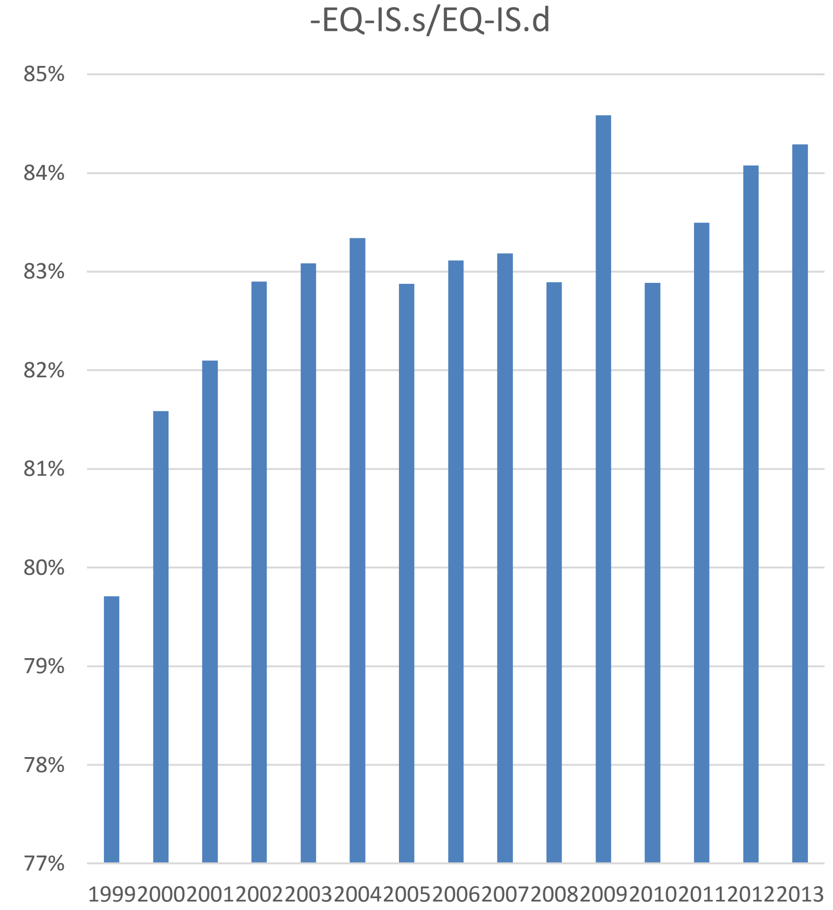

Model Estimation

Summary Statistics

- The relative population of demand orders by year
  - Counted by frequency, shares, and notional value
  - NB: the numerator does not include orders that incurred zero IS but the denominator does
- As expected, the relative fraction is slightly above 50% for these cross-sections

Model Estimation

Outlier Censoring

- Outliers are first removed from the entire CRSP market data set by year as follows:
  - The 1-%tiles are calculated of:
    - The ratio of daily volume to lag daily volume
    - The (log) return between the close price and open price
  - The top 1-%tile of the ratio of daily volume to lag daily volume is removed
  - The bottom (most negative) and top (most positive) 1-%tiles of the open/close return are removed
- The remaining CRSP data is merged with the Ancerno data

Model Estimation

Regression Equation

- $\frac{\tilde{r}_x}{\sigma} \sim 0 + \delta b + \gamma b \frac{v}{\sigma} + \beta \frac{r_m}{\sigma} + \epsilon$
- This model is estimated for all symbols and days in a year and separately for the demand and supply cross sections
- The intercept is constrained to be zero
  - Improves in/out sample forecast error
  - Serendipitously eliminates the need for economic interpretation of the intercept

Model Estimation

Parameter Estimates

year	δ_d		γ_d		β_d		R_d^2	δ_s		γ_s		β_s		R_s^2
	coef	t-stat	coef	t-stat	coef	t-stat		coef	t-stat	coef	t-stat	coef	t-stat	
1999	0.88	569	0.002	41	0.09	42	0.53	-0.66	-465	-0.001	-31	0.06	32	0.49
2000	0.88	607	0.001	28	0.11	54	0.53	-0.68	-503	-0.001	-23	0.08	41	0.51
2001	0.83	568	0.001	43	0.12	62	0.52	-0.67	-482	-0.001	-32	0.09	48	0.50
2002	0.82	660	0.002	43	0.15	117	0.56	-0.68	-571	-0.001	-32	0.11	88	0.54
2003	0.79	653	0.001	44	0.20	126	0.57	-0.66	-575	-0.001	-23	0.14	98	0.55
2004	0.81	709	0.001	44	0.22	121	0.56	-0.66	-593	-0.002	-48	0.17	98	0.55
2005	0.79	725	0.002	59	0.24	139	0.57	-0.66	-630	-0.001	-35	0.18	107	0.55
2006	0.79	707	0.001	40	0.25	130	0.57	-0.66	-600	-0.001	-28	0.19	103	0.55
2007	0.81	714	0.002	44	0.19	124	0.56	-0.68	-603	-0.001	-28	0.14	95	0.54
2008	0.82	681	0.004	45	0.20	155	0.58	-0.69	-568	-0.002	-24	0.16	123	0.56
2009	0.75	679	0.006	50	0.27	164	0.60	-0.64	-574	-0.002	-23	0.22	132	0.58
2010	0.77	688	0.001	19	0.29	181	0.61	-0.64	-561	0.000	-7	0.21	134	0.58
2011	0.76	720	0.002	41	0.34	206	0.61	-0.64	-603	-0.001	-26	0.25	157	0.59
2012	0.76	771	0.001	31	0.34	201	0.59	-0.64	-654	-0.001	-18	0.26	165	0.58
2013	0.76	666	0.002	39	0.33	165	0.59	-0.65	-568	-0.001	-25	0.26	135	0.57

Model Estimation

Parameter Estimates

- The t-statistics for all variables are plotted on the right
- Note the dominance of volatility as a fixed cost
  - Followed by market return then order size
    - Similar findings in Frazzini, Israel and Moskowitz [2012]
- The coefficients of the fixed and variable cost terms are plotted on the next slide
  - Variable cost would be more stable and significant as $\sqrt{\frac{q}{\overline{V}}}$

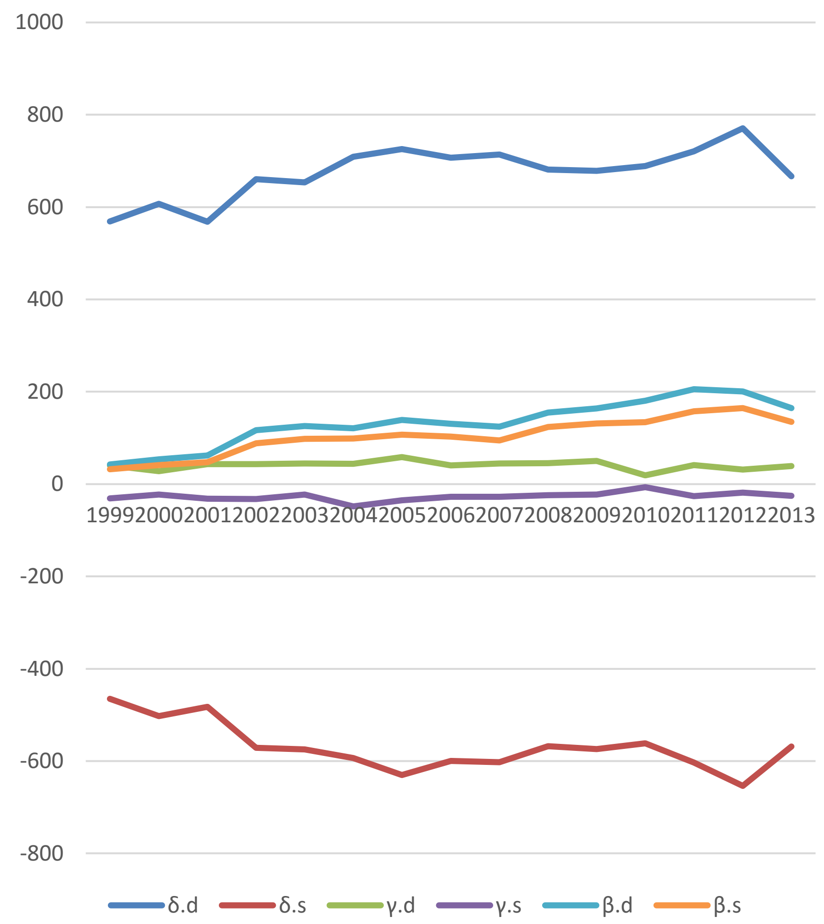

Model Estimation

Parameter Estimates

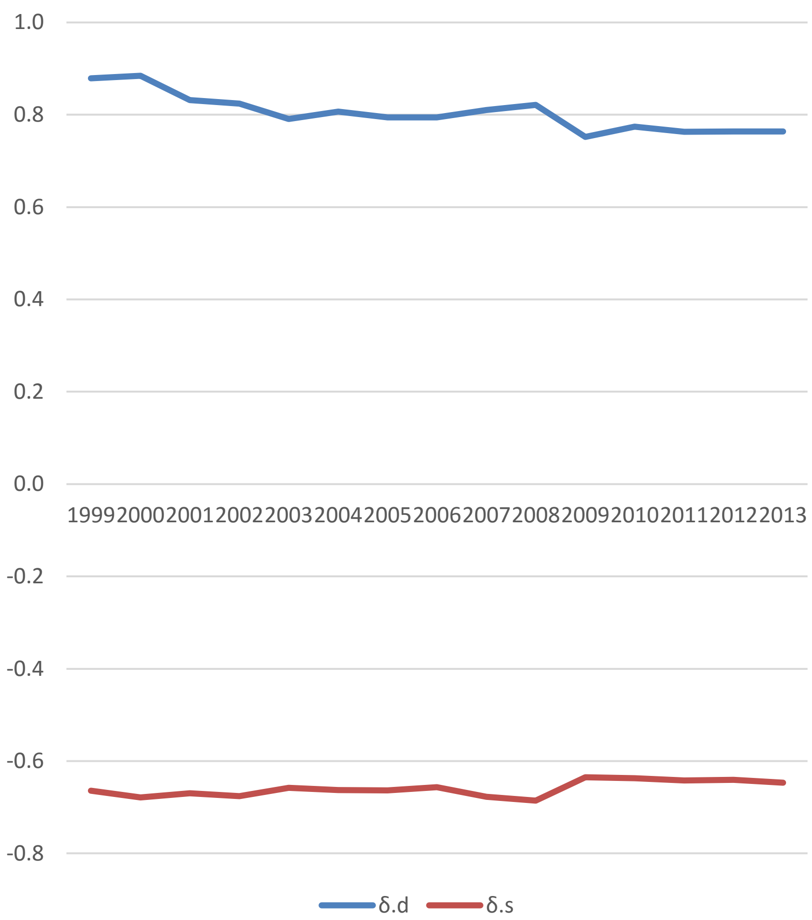

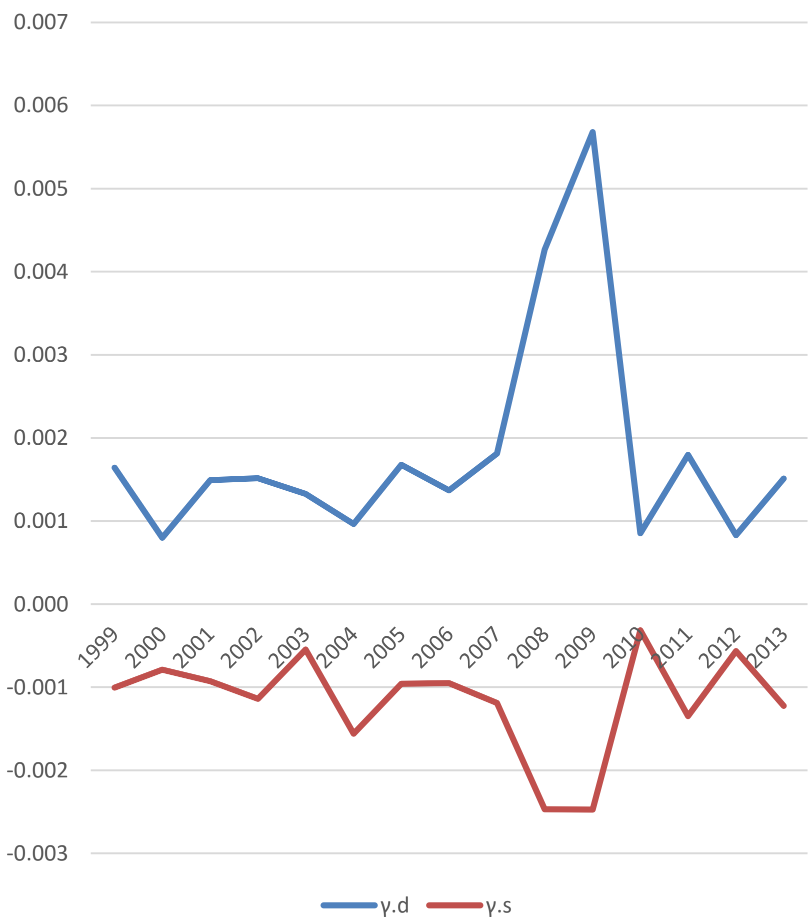

Model Estimation

Test of Model Functional Form

- As a test of how well the functional form of the model fits the data we perform forecasts where:
  - The supply/demand function is chosen by the sign of realized IS
  - Parameters from each year are applied to in-sample data
  - Parameters from the lagged year are applied to data from the next year
    - “quasi” out of sample
- The plots of mean forecast and realization by year are shown on the next slide
  - The “quasi” out of sample forecasts illustrate the relative stability of the parameter estimates and how the bias in the sample population of supply/demand orders is mitigated
  - Note also how the magnitude of forecast mean supply IS is smaller than demand

Model Estimation

Mean Forecast by Year by Supply/Demand Cross-section

In-Sample Using Sign IS

Out-Sample Using Sign IS

POST-TRADE ANALYSIS

Post-Trade Analysis

Overview

- As an example application of our model we perform in-sample forecasts of IS to decompose the contributing components of realized cost:
  - Fixed
  - Variable
  - Market
  - Supply/Demand
- Before proceeding we first need to describe how forecasts can be made with the supply/demand functions

Post-Trade Analysis

Making Forecasts With the Supply/Demand Functions

- We can use the realized sign of IS to choose either the supply or demand function or we can express the forecast of IS as an expectation of two possible outcomes:
  - $E[\tilde{r}_f] = \omega_d \tilde{r}_{f.d} + \omega_s \tilde{r}_{f.s}$
- A natural application for logit estimation:
  - $order\_liquidity\_state \sim \zeta_1 \sqrt{\frac{q}{V}} + \zeta_2 b \frac{r_v}{\sigma}$
  - Where:
    - $order\_liquidity\_state = \begin{Bmatrix} 0 \; when \; order \; is \; liquidity \; supplier \\ 1 \; when \; order \; is \; liquidiy \; demander \end{Bmatrix}$
    - $\frac{q}{V}$ = order quantity over contemporaneous market volume
    - $\frac{r_v}{\sigma}$= realized symbol open to vwap return standardized by lagged volatility
  - $\rho_d = \left(1 + e^{-\left(\zeta_1 \sqrt{\frac{q}{V}} + \zeta_2 b \frac{r_v}{\sigma}\right)}\right)^{-1}$
  - $\omega_d = \rho_d$
  - $\omega_s = (1 - \rho_d)$
- Essentially adding more degrees of freedom to the model while allowing contemporaneous realized return and volume to be used as explanatory variables

Post-Trade Analysis

Mean Forecast by Year

- The model generally under-forecasts the magnitude of IS because of the “dampening” effect of the logit forecast
  - Neither 0 or 1
- However the forecast RMSE is lower than if the realized sign of IS is used to choose either the supply or demand function

Post-Trade Analysis

Decomposition of Mean Forecast by Year by Supply/Demand

- Fixed cost dominates
  - The aggregate execution price across all participants will be close to market VWAP
  - The benchmark price is the open
  - Hence, the mean absolute open to aggregate execution price will look like volatility as we measure it

Post-Trade Analysis

Is the Complexity of the Extra Degrees of Freedom Justified?

- Clearly, it is not necessary to estimate separate supply/demand models to produce these decompositions
  - The data could be pooled into one set of coefficients
  - The plots on the right show the relative RMSE for the pooled model and the partitioned model
  - The extra degrees of freedom, and the additional information from using realized contemporaneous explanatory variables resulted in a relative decline in RMSE between forecast and realized of 40%

Post-Trade Analysis

Is the Complexity of the Extra Degrees of Freedom Justified?

- This framework allows for the possibility of making out of sample forecasts where the sign of IS can be forecast on an expectation basis depending on the context of the trade
  - Investment or trading strategy, e.g.:
    - Growth vs. value, momentum vs. reversion
    - Index rebalances
  - Market context, e.g. liquidation under duress

COMPARISON TO THE AMIHUD ILLIQUIDITY MEASURE

The Amihud Illiquidity Measure

Overview

- The Amihud illiquidity measure (ILLIQ) is the ratio of symbol open-to-close return to market notional volume
  - Intended to measure the return per unit of volume traded
- We calculate the lagged 21 day average of ILLIQ for a symbol universe of 25 symbols with 5 symbols selected from each quintile (20%) of capitalization in the SPX and compare it to:
  - The mean value of realized IS found in Ancerno
    - We found an average of 440 aggregated orders for each symbol and capitalization quintile
  - The forecast value of IS for hypothetical trade sizes of:
    - $2MM notional: the approximate mean value in our aggregated orders
    - 10% ADV: similarly

The Amihud Illiquidity Measure

Summary Statistics

- Mean ILLIQ-21 by capitalization quintile by year in units of bp of return per $1MM volume
- Suggests that the frictional costs of trading have declined by two orders of magnitude since decimalization

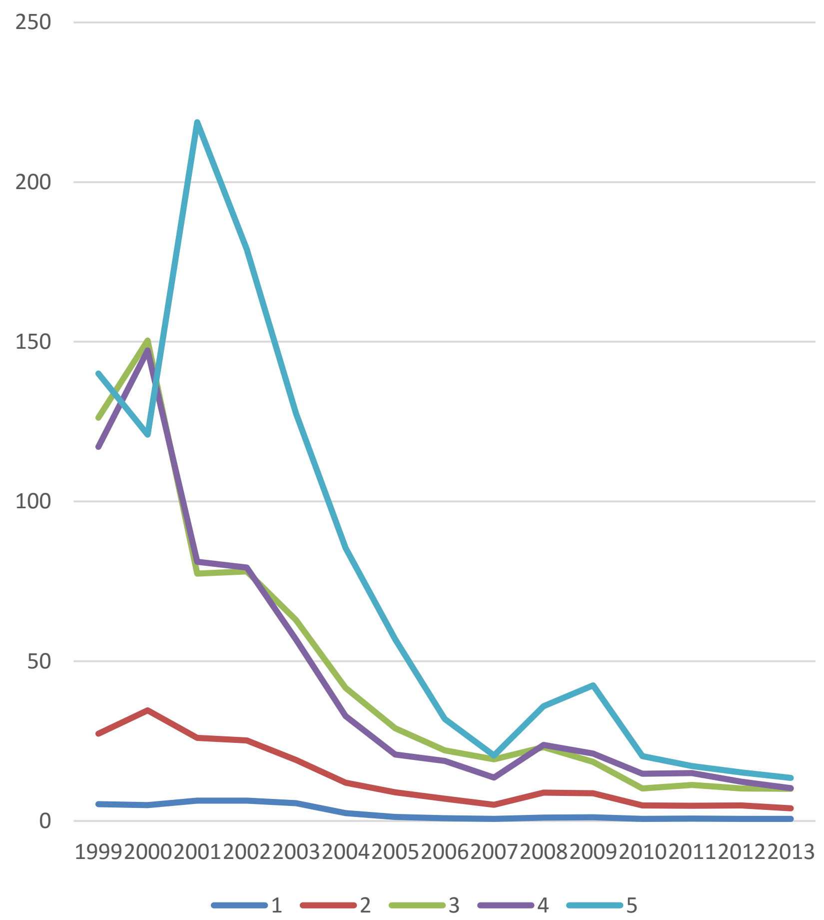

The Amihud Illiquidity Measure

Sample IS From Ancerno

- Mean sample IS by capitalization quintile by year in units of bp
  - Shows only a 1/3rd decline in frictional costs
- The next slide shows that at the same time, mean order size in Fill ADV units dropped by about 1/2

The Amihud Illiquidity Measure

Sample Fill ADV From Ancerno

- Mean Fill ADV by capitalization quintile by year
  - Relatively constant across capitalization quintiles

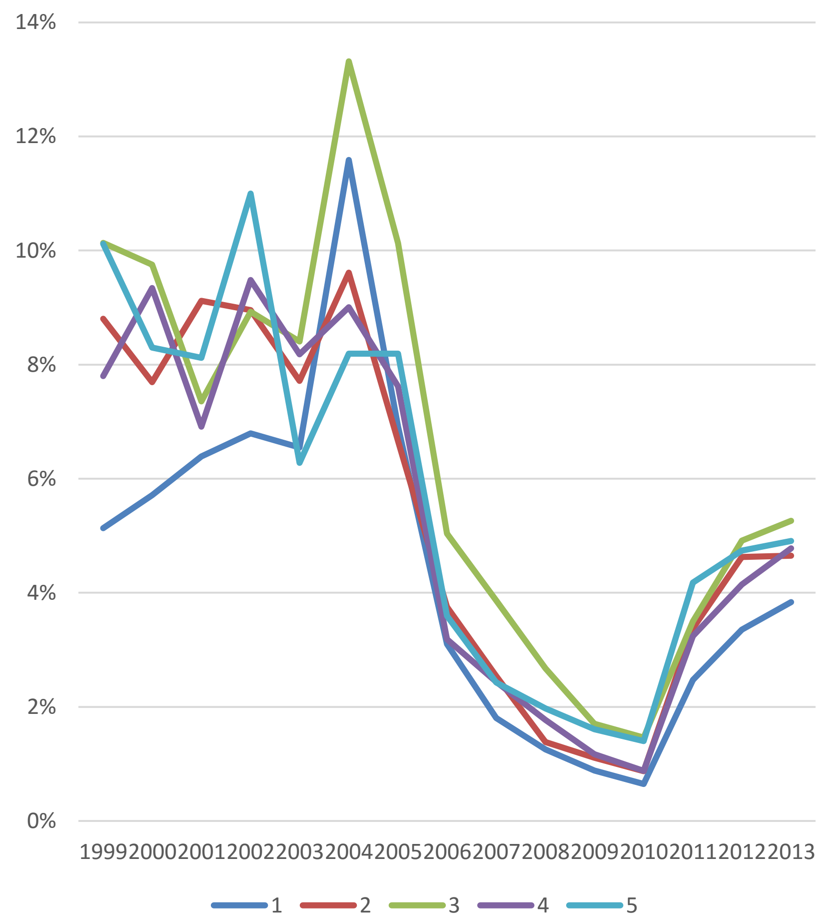

The Amihud Illiquidity Measure

Model Forecasts

- The next slide shows the forecasts from our model for hypothetical trades of size 10% ADV and $2MM
- They look very similar in magnitude and progression
  - To each other and to mean realized IS shown previously
  - Illustrates the dominance of fixed costs in realized IS for orders of the magnitude in our sample

The Amihud Illiquidity Measure

Forecasts For Fixed Trade Sizes Using Estimated Parameters

10% ADV

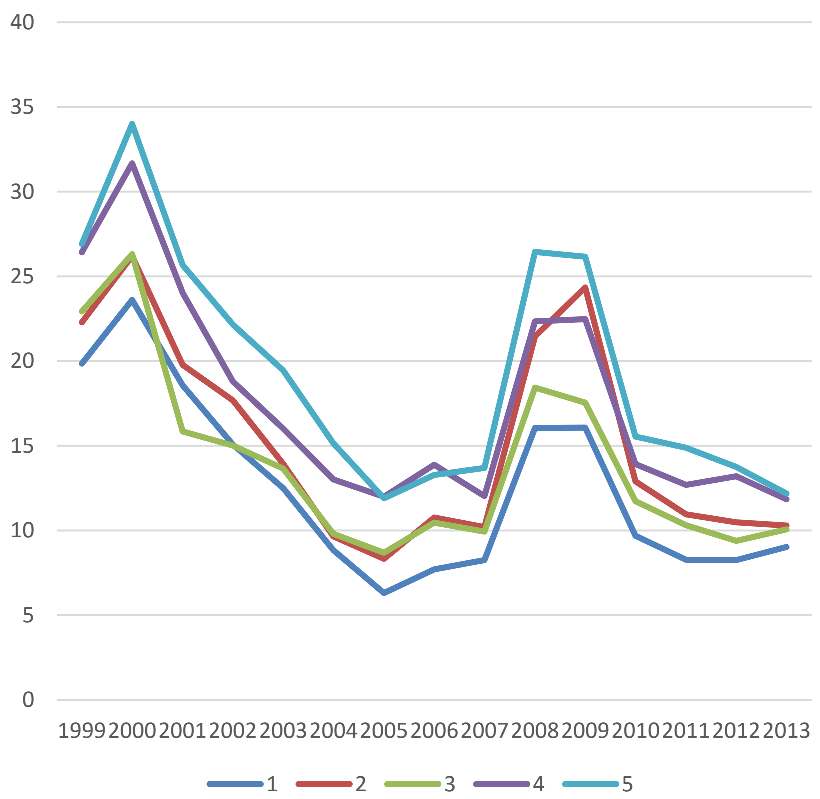

$2MM Notional

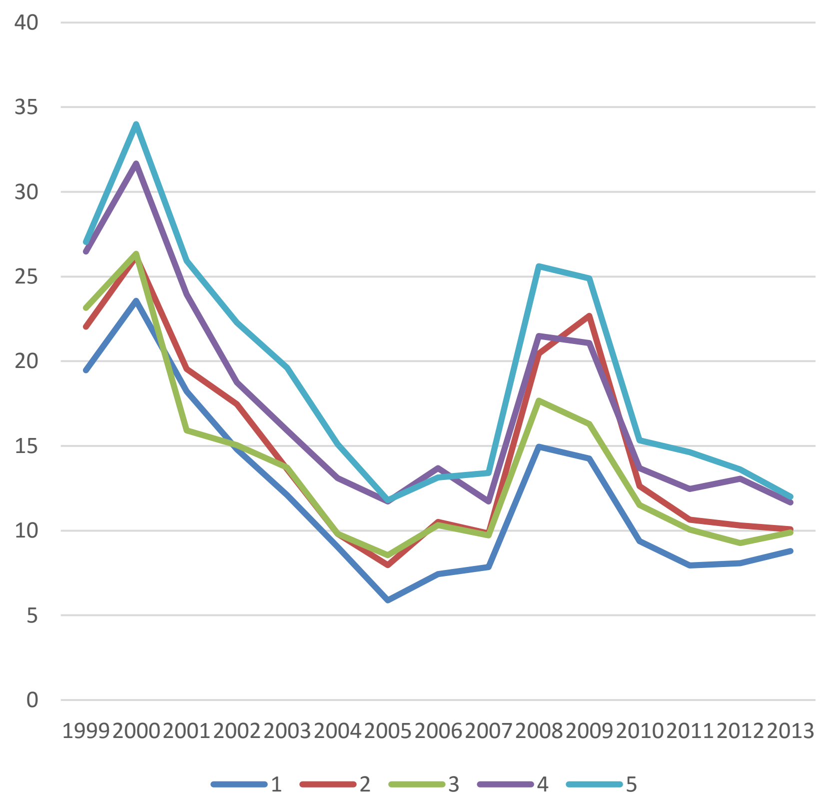

The Amihud Illiquidity Measure

Conclusions

- The Amihud ILLIQ measure seems to capture the overall trend in market liquidity well but not the magnitude
- It is more sensitive to overall market volume than realized IS, declining rapidly with increasing volume while return stays relatively constant
  - Perhaps missing the increasing fraction of market volume that originates from HFT
    - Less of the volume is “available” to institutional investors
  - Certainly missing the fact that fixed cost is the dominant component of total cost for Ancerno trades and fixed cost has declined far less rapidly than volume has increased

Conclusions