

Island Goes Dark: Transparency, Fragmentation, Liquidity Externalities, and Multimarket Regulation

TERRENCE HENDERSHOTT

and

CHARLES M. JONES*

First Version: March 14, 2003

This Version: October 29, 2003

* Hendershott is at the Haas School of Business, University of California at Berkeley and Jones is at the Graduate School of Business, Columbia University. We thank Maureen O'Hara, two anonymous referees, Mike Barclay, Ekkehart Boehmer, Larry Harris, Marc Lipson, Gideon Saar, Patrik Sandås, and seminar participants at the Courant Institute Mathematical Finance Seminar, HEC Montreal, the New York Stock Exchange, the University of California at Irvine, the University of Southern California, and the University of Wisconsin for helpful comments and suggestions. This research began while Jones was the visiting economist at the New York Stock Exchange. The comments and opinions expressed in this paper are the authors' and do not necessarily reflect those of the directors, members, or officers of the New York Stock Exchange, Inc. Hendershott gratefully acknowledges financial support from the National Science Foundation. Any errors or omissions are the responsibility of the authors alone.

Island Goes Dark: Transparency, Fragmentation, Liquidity Externalities, and Multimarket Regulation

ABSTRACT

In response to a regulatory enforcement, the Island electronic communications network recently stopped displaying its limit order book in the three most active exchange-traded funds (ETFs). As a result of this reduction in pre-trade transparency, this dominant venue's share of trading activity and price discovery falls, fragmenting the market. ETF prices adjust more slowly when Island goes dark, and substantial price discovery moves from the ETF market to the futures market. Island's effective and realized spreads increase, while effective and realized spreads fall in other markets. The increase in spreads on Island more than offsets the reduction in trading costs in other markets, increasing overall trading costs. Overall ETF market quality falls.

As they have for decades, financial market regulators and researchers are currently wrestling with a fundamental issue: how to facilitate competition between securities markets without destroying the liquidity externalities that arise when traders are able to come together in space and time. With the rise of electronic stock markets, the U.S. Securities and Exchange Commission (SEC) in 1998 tried to maintain this balance by issuing complex rules (Regulation ATS) governing these alternative trading systems. Among other things, Regulation ATS was intended to provide information to the public about the terms of trade (mainly quoted prices) at various venues, as well as integrate these new automated markets into the existing linkages between venues.

On September 23, 2002, the SEC enforced Regulation ATS by requiring that the Island electronic communications network (ECN) comply with ATS display requirements in several of the highest-volume exchange-traded funds (ETFs) where Island dominated trading and price discovery. Island chose to comply by going “dark,” no longer displaying its automated limit order book to any market participant. Trading and price discovery continue on Island after it goes dark, but at much lower levels, and the repercussions are observable across other markets. The regulatory enforcement provides an excellent experiment to examine quote transparency,¹ liquidity externalities, market fragmentation, and the adequacy of theoretical models in describing actual multimarket trading.

Our results show that this regulatory enforcement makes the ETF market worse in terms of price discovery and trading costs. When Island goes dark, ETF prices become less efficient and are slower to adjust to new information than the corresponding futures prices. This has implications well beyond ETFs. To our knowledge, this paper demonstrates for the first time that a reduction in transparency and/or the migration of order flow away from the dominant market worsens overall price discovery. The main challenge is disentangling these two effects, because both result from the same event. However, it is important to try to disentangle these two, because the reduction in ETF trading concentration raises the possibility that this regulatory enforcement harms the ETF market by reducing liquidity externalities.

¹ The word transparency has a number of meanings. It sometimes refers to announcing trading intentions in advance (Admati and Pfleiderer (1991)), but the term usually refers either to pre-trade transparency, the availability of quotes and other information about trading interest, or to post-trade transparency, defined as the disclosure of executed trades. In this paper, the term transparency alone refers to quote transparency.

Liquidity externalities arise when traders are able to come together either physically or virtually in space and time, reducing trading/search costs and enhancing price discovery. This in turn attracts more traders, further reducing trading costs and improving price discovery. The mechanism underpinning liquidity externalities is often somewhat vague, but competition between liquidity providers is one natural explanation. If competition across markets is not perfect, then a fragmented (less concentrated) trading arrangement may result in weaker competition among liquidity providers. Lack of transparency and integration are the main reasons for imperfect competition among securities markets. If traders cannot see the relative state of different markets and instantly choose the best one, it matters where trading occurs. However, less concentrated trading need not result in lower market quality. For example, if each venue is sufficiently liquid, and competition for liquidity provision is intense at each market center, the externalities lost from splitting trading across venues may be minimal.

Transparency can also affect market quality directly. For example, Milgrom and Weber (1982) show that bid transparency increases competition in a static common value auction. Glosten (1999) writes that “transparency should lead to greater commonality of information. Greater commonality of information means that adverse selection becomes less of an issue.” We find evidence to support both of these ideas. Island trading costs and overall ETF trading costs both rise. Based on a simple Glosten (1987) spread decomposition, we find that Island liquidity providers earn more post dark, and Island trades move prices more after the event. The former indicates a reduction in competition within Island, while the latter is consistent with greater adverse selection there.

While liquidity externalities are one of the most important outstanding issues in market design (Madhavan (2000)) and market regulation (Macey and O'Hara (1999)), it has been difficult to document their importance in practice.² Here we find evidence that liquidity externalities are important. To disentangle the effect of liquidity externalities vs. transparency, we focus on venues other than Island. If markets are completely segmented, then a change in one market's transparency should not by itself directly affect price discovery and trading costs in other markets. In fact, if order flow is segmented but other markets rely on information in

² See Barclay and Hendershott (2002) for evidence on liquidity externalities arising from traders consolidating their activity in time.

Island's book, a reduction in transparency on Island could harm price discovery and increase trading costs elsewhere. In contrast, if liquidity externalities are important, then these other markets should become more competitive as they gain order flow. We find that trading costs (effective spreads) in all ETFs decline in non-Island markets that pick up market share, which are the other electronic systems (mainly ARCA and Instinet ECNs). Realized spreads, which proxy for revenue to liquidity providers, decline off Island, indicating that these other markets become more competitive as they obtain more order flow. Thus, order flow migration and liquidity externalities must be an important part of the story.

The multimarket models of Pagano (1989a, 1989b) are also helpful for interpreting our results. Pagano (1989b) shows that markets with greater order flow have more efficient price discovery, while Pagano (1989a) can explain the relationship between order flow and trading costs. In both of these models, competition between markets is based solely on price, so either trading concentrates in a single market or all markets are identical. In reality, securities markets compete along other dimensions as well. For example, ECNs and floor-based exchanges differ in their speed of execution, anonymity, and transparency. Surprisingly, we find that about half of Island's trading volume remains after it goes dark. This suggests that there are different clienteles with different preferences, and we should not expect order flow to concentrate in a single market when competition is multidimensional and traders are heterogeneous.³ Perhaps future research will develop a richer model of multimarket competition to help sort out these issues.

As a practical matter, the presence of such multidimensionality and heterogeneity among traders and trading systems makes the regulator's task much harder. The Island event highlights the real-world difficulties of facilitating multimarket competition without damaging liquidity externalities. We believe that improved linkages between markets are an important part of the solution, and we provide specific suggestions at the end of the paper.

The paper proceeds as follows. Section 1 reviews some related literature. Section 2 provides an overview of ETFs, the Island ECN, and Island's decision to go dark. Section 3

³ In fact, multiple equilibria are possible in such models because traders face a coordination problem, even if competition is multidimensional and traders are heterogeneous. Thus, our results may depend on the particular equilibrium the market was in before and after the event, though we think it unlikely that this explains our results.

describes our data. Section 4 explores the effect going dark had on Island's market share and role in price discovery. Section 5 studies the impact on the efficiency of price discovery within the ETF market. Section 6 analyzes price discovery across the ETF and futures markets. Section 7 investigates the effect of Island going dark on trading costs across markets. Section 8 examines the effect of Island going dark on private information on Island and other markets. Section 9 concludes.

1. Related Literature

There is a great deal of related research on multimarket regulation, multimarket trading, and transparency. For example, nearly every regulatory change that increases competition and facilitates integration across markets improves market quality. Recent examples include the Nasdaq order handling rules (Barclay et al. (1999)) and the cross-listing of equity options (de Fontnouvelle, Fishe, and Harris (2003), Battalio, Hatch, and Jennings (2003)). Island's going dark represents a substantially different event, in that robust competition exists before the regulatory change, and the change harms market quality. Thus, this paper provides evidence that care and caution are needed in crafting a national market system.

Theoretical research on multimarket trading typically assumes similar market structures to focus on traders' choice of venue, e.g., Pagano (1989a, 1989b) and Chowdhry and Nanda (1991). Similarly, Glosten (1994) examines whether competing liquidity suppliers in multiple markets can coexist. However, there are a few models of different competing market structures, including Parlour and Seppi (2003), who study competition between a specialist/limit order market and a pure limit order market, and Hendershott and Mendelson (2000), who analyze the effect of a crossing network on a dealer market.

There are a number of models that assess the impact of revealing or suppressing security price information. Most pre-trade transparency papers contrast quote disclosure with negotiation in a multiple dealer market (Bloomfield and O'Hara (1999), Flood, Huisman, Koedijk, and Mahieu (1999), and de Frutos and Manzano (2002)). Madhavan (1992, 1995), Biais (1993), Pagano and Roell (1996), and Rindi (2002) study the disclosure of trade information following trades under different market structures.

There is also theoretical work on the display of limit orders. Harris (1996) notes that exposing a limit order may reveal private information and this may allow others to front-run the order. Madhavan, Porter, and Weaver (2000) develop a model where liquidity traders use limit orders, and they conclude that greater order book transparency can reduce liquidity. In contrast, Baruch (2002) develops a model where liquidity is supplied by smart limit order traders and a specialist. An open limit order book improves liquidity in his model.

In terms of the empirical literature on limit order display, two directly relevant papers come to opposite conclusions. Boehmer, Saar, and Yu (2003) study the 2002 introduction of OpenBook, a real-time NYSE limit order book data feed, and find that the resulting increase in quote transparency is associated with lower trading costs and more aggressive limit order submission. In contrast, Madhavan, Porter, and Weaver (2000) examine the April 1990 decision by the Toronto Stock Exchange to provide the top five prices on either side rather than just the top of the book. They find wider spreads post-event. Both of these papers examine decisions by markets to increase transparency to off-exchange traders, rather than the regulatory shock that results in a complete loss of transparency that we study here. We find no support for the Madhavan, Porter, and Weaver (2000) results, although the Island event is not directly comparable because it involves the complete loss of transparency at one of many trading venues, and significant order flow migration occurs. Foucault, Moinas, and Theissen (2002) also study a decline in quote transparency that is less closely related to the Island event. They examine Euronext's decision to stop displaying the identities of limit order traders. They find that average quoted spreads decline significantly while the average quoted depth increases significantly.

There are also several papers that focus on the impact of ECNs on Nasdaq trading. Barclay, et al. (1999) analyze the impact of Nasdaq's order-handling rules and show that ECN quotes play an important role in reducing trading costs. Simaan, Weaver, and Whitcomb (2003) show that ECN quotes are more likely to be on an odd tick than market-maker quotes and how quote anonymity affects market-makers' quotes. Huang (2002) finds that ECN quote updates are more informative than market-maker quote updates. Barclay, Hendershott, and McCormick (2002) find that more informed traders choose to trade on ECNs and are willing to pay higher effective spreads to do so. Conrad, Johnson, and Wahal (2002) compare trading costs for institutional trades on crossing networks and Instinet to those with traditional brokers. Weston

(2001) uses ECNs' monthly activity to examine their evolving competitive role in Nasdaq. There are two other papers that focus specifically on Island. Hasbrouck and Saar (2001) examine the impact of volatility on the use of limit orders on Island. Biais, Bisiere, and Spatt (2002) study quote competition between Nasdaq and Island and the profitability of limit order strategies.

Finally, there is related literature on ETFs. Hasbrouck (2002) examines the intraday price discovery of indexes across the ETF and index futures markets. Boehmer and Boehmer (2003) show that prior to August 2001 the ETF market did not appear competitive, and the NYSE's entry at that time increases competition and reduces trading costs. Also, Elton, Gruber, Comer, and Li (2002) compare ETF returns to changes in the underlying securities. Poterba and Shoven (2002) analyze ETF investing from a tax perspective. Engle and Sarkar (2002) study the dynamic relationship of ETF prices to the underlying securities.

2. An Overview of Island's Decision to Go "Dark"

Before discussing Island and its decision to go dark, some background on ETFs is useful. ETFs are registered investment companies under the Investment Company Act of 1940. Like open-end index funds, they hold portfolios designed to track a particular stock or bond index. However, they differ from open-end index funds in two important ways. First, ETF shares can be traded continuously on major stock markets during the day. Second, at the end of every trading day, ETF shares can be created or redeemed by exchanging ETF shares (typically in multiples of 50,000 shares) for the appropriate number of shares of the underlying stocks. This ability to conduct arbitrage ensures that the ETF closely tracks its corresponding index. The first ETF, the S&P 500 Trust Series I (SPY), listed on AMEX in 1993. The success of SPY has led to related products for the Nasdaq-100 (QQQ), Dow Jones Industrial Average (DIA), sector-specific indexes, and value and growth-oriented indexes.⁴ ETFs typically charge a management fee that is lower than that of a comparable open-end index mutual fund, though investors must also pay brokerage commissions to buy or sell ETFs.

⁴ SPY, QQQ, and DIA are sponsored by an affiliate of the American Stock Exchange. Other major ETF sponsors include Merrill Lynch, which sponsors Holding Company Depository Receipts (HOLDERS), and Barclays Global Investors, which markets and manages iShares.

The three most active ETFs – DIA, SPY, and QQQ – are all listed on the American Stock Exchange, but they trade many other places as well. They are traded at the NYSE, most of the regional stock exchanges, 3rd market dealers such as Madoff and Knight, and also at most ECNs. In fact, by the middle of 2002, the leading venue for the major ETFs was the Island ECN.

Island essentially operates a very fast, automated limit order book. This agency market accepts only priced limit orders. An order at a limit price that meets or crosses the best opposing price (a marketable limit order) executes immediately. If an incoming limit order does not cross an existing order, it goes into the book until a matching order is received or the limit order is canceled. For most securities, order submitters can choose to have an unmatched order displayed in the limit order book or remain hidden from view. At the end of each day all outstanding limit orders in the book expire.

Island matches orders based on strict price-display-time priority. This means that displayed orders have priority over undisplayed orders for orders at the same price. Island automatically executes matching orders on its system without regard to quotes posted by other market participants. Hence, Island does not guarantee the best price in the market, and the responsibility of ensuring best execution falls on the subscriber.

Island subscribers submit limit orders without charge. To encourage depth in the limit order book, if a limit order on Island does not execute immediately and subsequently executes against an incoming order, the subscriber receives a 0.11 cent rebate per executed share during our sample period. The incoming order executing against Island's book is charged 0.19 cent per executed share.

Like other ECNs, Island displays the price and the number of shares available to the market, but not the identities of the investors or the brokers. Hence, in addition to immediate execution, Island and other ECNs offer anonymity. Both of these appeal to investors who may have short-lived information about future price changes. Furthermore, Island subscribers are able to suppress the display of their limit orders in whole or in part. These hidden limit orders are an important fraction of the total liquidity available on Island. Together these differences suggest that Island may appeal to a clientele of investors that are not well-served by a slower, less anonymous floor-based exchange.

Island subscribers pay a fee for a complete real-time data feed, but for each security the top 15 orders on each side of the book are continuously displayed on Island's web site. For Nasdaq stocks, the top of Island's book is displayed on and is accessible via Supermontage. For NYSE and AMEX stocks, including the ETFs that are of interest here, Island transmits no quote information other than the website and subscriber data feeds, and only subscribers can trade these listed stocks on Island.

In December 1998, the SEC adopted Regulation ATS to establish a regulatory framework for alternative trading systems to more fully integrate them into the national market system.⁵ The goals of Regulation ATS are to provide investors with access to the best prices, facilitate competition, provide a complete audit trail and surveillance on alternative trading systems, and reduce the potential for market disruption due to system outages. An important section of Regulation ATS requires that if an alternative trading system executes more than five percent of the total volume in a given security in at least four of the past six months, it must do two things. First, if the alternative trading system displays quotes, it must display those quotes in the national market system (the Consolidated Quote System for stocks listed on the NYSE or AMEX). Second, it must allow participants at other venues to trade with these quotes. For listed stocks, this means that the alternative trading system must join the Intermarket Trading System (ITS).

ITS was developed in the late 1970's and early 1980's to electronically link together the NYSE, AMEX, Nasdaq, and regional stock exchanges. Its central purpose is to create a national market in NYSE and AMEX stocks, so that market participants can see and access the best bid and offer prices across all venues. A key requirement is that ITS members cannot trade through another venue. If a better price is offered in another market, the order must be sent to the other market via ITS. The market on the receiving end of an ITS order has up to 30 seconds to respond.

In effect, this means that the receiving market potentially has a 30-second option to execute the ITS order. Prices on the largest ETFs move so rapidly, and the securities have such

⁵ Most ECNs, including Island, are subject to Regulation ATS. Archipelago is partially exempt from Regulation ATS because it is regulated as an exchange following the October 25, 2001 SEC approval to launch the Archipelago Exchange as a regulated facility of the Pacific Exchange.

narrow spreads, that this option is quite valuable to an ETF market-maker. This value to the market-maker is a cost to the investor whose order is sent over ITS. Thus, ITS has the potential to increase investor trading costs, a standard measure of market quality. Island was also concerned that the trade-through rules would greatly inhibit the speed of its ETF executions.

The SEC recognized this and actually tried to address these concerns about ITS by issuing an exemption to the trade-through rule on August 28, 2002, effective September 4, 2002. The exemption applied only to the three largest ETFs (DIA, QQQ, and SPY), and it allowed firms and trading venues to trade through the consolidated quote's national best bid-offer (NBBO) by at most three cents.⁶ Many market observers, including the SEC, expected Island to work to integrate its quote and orders into ITS. Beginning on September 23, 2002, the SEC required that Island comply with the display requirements of Regulation ATS. Island's market share was over the five percent threshold in DIA, QQQ, and SPY.⁷

However, Island surprised the market and found another way to comply with Regulation ATS. If an electronic trading system does not display quotes to subscribers at all, it is not required to post quotes or send and receive orders through ITS. Rather than post its quotes and participate in ITS, Island chose to go dark. It stopped displaying its limit order book completely in the affected ETFs. It did this by converting all affected orders to undisplayed status. Figure 1 shows the state of Island's visible limit order book in QQQ after going dark and an example of a non-dark book (in MSFT). In an e-mail to its subscribers on September 19, 2002, Island explained, "We do not believe it would be possible for Island simultaneously to comply with these requirements and to maintain the system performance our subscribers expect."

⁶ While it is possible that the trade-through exemption could affect the examination of Island going dark, we did not find significant changes occurring around September 4th.

⁷ Island also exceeded the five percent threshold and stops displaying its quotes on September 23, 2003 in two other ETFs: MKH (Merrill Lynch Market 2000+ HOLDERS Index) and SMH (Merrill Lynch Semiconductor HOLDERS Index). MKH trades infrequently – fewer than 25 trades per day – and trading had left Island prior to doing dark, with an average of only four trades per day on Island prior to going dark and none afterwards. SMH is more active, with a few thousand trades per day. However, Island is far from the dominant market in SMH, with a market share of only 8% pre-event and 2% post, so we would expect little change in SMH market quality when Island goes dark. In results not reported, we confirm that there is little change in overall SMH effective and realized spreads when Island goes dark. However, Island's effective and realized spreads increase significantly.

Because Island's SMH market share fell below the 5% threshold after it went dark, on March 1, 2003 Island began displaying its limit order book in SMH again. With the reinstatement in March, Island's market share increases to a little less than 5%. Overall effective and realized spreads show little change, but Island's effective and realized spreads fall when its quotes return.

[Insert Figure 1 Here]

ETF trading continued on Island, but with the limit order book not visible to any market participant. Island describes this as a continuous call market. Island knew this would hurt its trading volume in the affected securities. In an attempt to stem the expected loss of order flow, Island reduced its net fees for the first week after going dark. It increased its rebate for adding liquidity from 11 to 19 cents per hundred shares executed. The fee for removing liquidity (executing against an existing limit order) remained at 19 cents per hundred shares, so Island received no net revenue from trading in the affected ETFs for that week.

Why Island chose to comply with the display requirements of Regulation ATS in such a novel way is difficult to know for sure. It may be due to its stated reason that participating in the Intermarket Trading System makes it unable to “maintain the system performance that our subscribers expect,” although other electronic markets, such as Archipelago, function as part of ITS. For example, even now Island could join ITS and simply reject all marketable orders that would trade through the existing quote by more than three cents. This would allow Island to maintain its status as a pure destination market; Island would accept incoming ITS orders, but would not need to route out any orders to other markets. In going dark instead, Island presumably knew both that it would lose some order flow, but, given that about half of Island’s ETFs trades execute against undisplayed limit orders prior to going dark, Island may have guessed that the clientele using hidden limit orders in the pre-dark period would remain when all limit orders became hidden.

More broadly, Island’s actions can be viewed as a strenuous objection to regulations that force ECNs to accommodate slower traditional floor-based exchanges, e.g., the 30-second option/delay for ITS orders. Once Island joined ITS, it probably felt that its business model would be permanently compromised, and perhaps avoiding that outcome was more valuable than the lost ETF market share. Given that Island’s obligation is to its shareholders and not the market as a whole, it makes sense that it would not consider the market-wide negative consequences of going dark.

In contrast, the SEC’s objective is to maximize the quality of the market as a whole. Given the current regulatory framework and the SEC’s attempts to address Island’s concerns

with the three cent trade-through exemption discussed above, the SEC had little choice but to require Island to cease its clear violation of the display requirements of Regulations ATS. Unfortunately, we show that the application of existing regulation made the ETF market worse in terms of price discovery and trading costs. After presenting our results in detail, we conclude by proposing a more flexible regulatory approach to integrate slower and faster markets.

3. Data and Descriptive Statistics

Our data consist of all the trades and quotes for three ETFs that went dark – DIA, QQQ, and SPY – during the normal trading hours from 9:30 a.m. to 4:00 p.m. from August 16, 2002 to October 31, 2002 (54 trading days, 25 in the pre period and 29 in the post period).⁸ Data for non-Island trades and quotes are gathered from the NYSE's TAQ database. The TAQ data identify several different trading and quoting venues, including the NYSE, AMEX, Boston, Cincinnati, Midwest, Pacific, and Philadelphia exchanges. The Pacific exchange trades using Archipelago's ECN technology, so it is essentially an ECN. We group together the other regionals (Boston, Midwest, and Philadelphia, as well as non-Island Cincinnati trades) because their combined market share is fairly small, and their ETF business is fairly homogeneous. Other market-makers, such as Knight Securities and Madoff, and ECNs such as Instinet, print their trades on the so-called 3rd market. The TAQ data do not distinguish between 3rd-market participants, so all 3rd-market trades other than those matched to Island trades are treated as a single venue.⁹

The ITCH data from Island are identical to those supplied in real time to Island subscribers. The data consist of time-sequenced messages of trade and book activity. An arriving order executing against an existing order in the book generates an Order Execution message. If any part of the order does not execute, the system sends an Add Order message. Add Order messages have a unique identification number and describe whether the order is a buy or sell order, a number of shares, a limit price, a display condition (normal, subscriber-only, or invisible), and a buy-sell indicator. When an order executes, an Order Execution message reports its identification number. If an order is cancelled or reduced in size, a Cancel Order message is

⁸ August 16, 2002 was chosen as the beginning date because ARCA-Pacific began trading the three ETFs that day.

⁹ Rule 11Ac1-5 reports show that about 2/3 of 3rd-market trading is Instinet, so we refer to this group as Instinet-3rd Market.

generated. We construct the visible book at each point in time by aggregating these messages from the start of the day up to that point in time. Although the arrival time and quantity of invisible orders are not available, special trade messages signal their execution.

To identify Island trades we match ITCH trades with TAQ trades in the markets where Island reports ETF executions (Cincinnati and the 3rd market). Due to timing delays in communicating across systems, the trade time in TAQ is usually slightly later than the ITCH trade time. We match trades on price and size using a sequence of time lags (2, 3, 1, 4, and 5 seconds). This approach identifies 99.9 percent of the Island trades in TAQ.

Because Island quotes are not in TAQ, it is more difficult to merge the ITCH data with the TAQ quotes. In constructing a consolidated NBBO our goal is to ensure accuracy, consistency, and comparability across the pre- and post-dark periods. There are two issues: there are no quotes in the ITCH data in the post-dark period, and there is substantial post-dark order migration from Island to venues that participate in the TAQ consolidated quote.

Because there are no Island quotes in the post-dark period, one approach is to exclude the ITCH quote data in the pre-dark period. Unfortunately, this introduces several possible problems. First, it turns out that excluding the Island quotes could bias Island trading costs downward and non-Island trading costs upward.¹⁰ Because this bias would appear only in the pre period, it would also bias estimates of changes in trading costs when Island goes dark. Specifically, this bias would falsely lead us to find an increase in Island spreads and a decrease in non-Island spreads when Island goes dark. This is exactly the result we expect from a reduction in transparency and the resulting order flow migration, and we want to ensure that this result is not artificially generated.

¹⁰ A simple example illustrates this effect. Assume that the best non-Island bid is 100 and ask is 101. Buys and sells are equally likely. If Island is not at the inside, the issue of Island quotes is moot. Also, assume that if Island is tied at the inside of either quote, then the probability that Island gets the trade is 0.5. Finally, assume if Island is alone at the inside the probability of a trade occurring on Island is greater than 0.5, say 0.75. If Island's ask price is 100.9, then using a consolidated NBBO Island's trading cost as measured by the effective (half) spread – the absolute difference between the transaction price and the bid-ask midpoint – is 0.45. The non-Island trades will have an effective spread of 0.4833 (one-third times 0.55 plus two-thirds times 0.45). If the Island quotes are excluded then the Island effective spread is 0.44 (two-fifths times 0.5 plus three-fifths times 0.4) and the non-Island 0.5. Thus, the exclusion of the Island quotes biases the Island effective spread downward and the non-Island effective spread upward. If Island is alone at the inside at both the bid and the ask this same bias exists as long as Island's improvement of the quotes is not symmetric and the quote that offers greater improvement over the non-Island quote is more likely to attract a trade.

The second problem involves the migration of order flow when Island goes dark. If displayed liquidity moves from Island to, say, ARCA, then that liquidity is in the consolidated quote post-event.¹¹ For consistency, that liquidity should also be in the consolidated quote pre-event. Based on these issues, we think it best to include the Island quotes in a consolidated NBBO for the pre-dark period. However, the paper's results on overall trading costs and non-Island trading costs are unaffected by the exclusion of the Island quotes from the NBBO in the pre period. When studying price discovery we present the pre-event period analysis both with and without the Island quotes and discuss the differences.

For the pre-dark period, we construct an Island inside quote of at least 100 shares from ITCH and merge that with quotes from TAQ to construct an NBBO. Because the differences in ITCH and TAQ times affect the construction of the NBBO, we perform the subsequent analysis using lags similar to those found between the ITCH and TAQ trade time. Using these different lags does not affect the results qualitatively. For example, an NBBO with Island quotes lagged two seconds provides somewhat stronger results than those in the paper, e.g., the increase in trading costs overall and on Island is larger. We also delete quotes which lock or cross the NBBO, although the paper's results are not sensitive to the exclusion of these quotes. Finally, to ensure that our results are not sensitive to the rounding of quotes in TAQ to the nearest penny, we test (without reporting the details) that rounding the pre-period ITCH quotes does not affect the paper's results.

Using the Lee and Ready (1991) algorithm, non-Island trades are classified as buyer-initiated if the trade price is greater than the contemporaneous quote midpoint and seller-initiated if the trade price is less than the quote midpoint. Trades executed at the midpoint are classified with the tick rule: midpoint trades on an up-tick are classified as buyer initiated, and midpoint trades on a downtick are classified as seller initiated. We represent the buy/sell indicator with the variable $x = +1$ for a buy order and $x = -1$ for a sell order. For Island trades we use the buy-sell indicator in the ITCH data. Using the ITCH buy-sell indicator only affects Island trades, and its primary benefit is to increase the precision of the estimates of Island's realized spreads and the

¹¹ This turns out to be the case and causes the quote midpoints in the pre period to be less accurate. This mainly affects Island's trading costs by increasing the effective spreads and decreasing realized spreads in the pre period.

price impact of Island trades in the pre period. The qualitative results are unchanged if we use the Lee and Ready (1991) algorithm for Island trades as well.

The data include almost 5 million trades for \$331 billion and a similar number of quote revisions. Table 1 provides daily trading activity, quoted spreads, and market share in terms of shares and trades across markets for the three ETFs in both the pre and post-dark periods. All three ETFs are very active, averaging two trades in three seconds for DIA and almost two trades per second for QQQ. Quoted spreads are narrow, ranging from 1.07 cents for QQQ up to 2.65 cents for DIA. Before going dark, Island is the dominant venue in share volume for DIA and QQQ. For SPY Island executes the majority of trades, but its smaller trade size results in it having only 21 percent of the share volume. In SPY both AMEX and Instinet-3rd market have greater share volume. One possible effect of Island going dark is some traders exiting from the market, but trading volume increases substantially in all three ETFs: 50 percent in DIA and SPY and 10 percent in QQQ.

[Insert Table 1 Here]

Going dark reduces Island's market share by 40 to 55 percent in both share and trade terms. In the post period, Island does not have the highest share volume in any of the ETFs, although it still has the greatest number of trades. Most of the volume leaving Island moves to other electronic venues, with ARCA-Pacific and Instinet-3rd market gaining the majority of Island's lost market share (Rule 11Ac1-5 reports show that over half of the growth in 3rd market volume is due to Instinet).

A standard way of measuring market concentration is a Herfindahl index. For example, Ellis, Michaely, and O'Hara (2002) and Schultz (2003) use the Herfindahl index to measure concentration of trading in Nasdaq stocks. The Herfindahl index is given by the sum of the squares of the percentage market share of each market. All trading being in a single market gives a Herfindahl index equal to one, while equally sized markets yield an index of one divided by the number of markets. Table 1 also provides the Herfindahl index in both share and trade terms. Note that due to data limitations, the third market is considered a single venue for this purpose. The average daily Herfindahl index shows that trading concentration decreases in the post-dark period. This change is significant at the 0.01 level, except for market share in share volume in

SPY where the change is positive but not statistically significant. There is no single correct measure of concentration in financial markets, so later when we calculate a measure of price discovery across markets – information share – we will calculate another Herfindahl index measuring the concentration of price discovery. The price discovery concentration index exhibits larger declines than the trading concentration index presented here.

Undisplayed limit orders are often an important feature of automated trading systems, but hidden limit orders are particularly prominent in Island's ETF trading. Table 1 shows that prior to going dark, between 45 and 50 percent of Island's share volume executed against undisplayed liquidity. When Island goes dark and all executed volume is by definition against undisplayed liquidity, Island's overall market share also declines on the order of 50 percent. Thus, to a first approximation, the volume of hidden order executions on Island remains constant post-event. One possible explanation for Island's continued significant market share is that certain market participants value a hidden limit order book market, and by extension value the absence of quote transparency. However, it does not necessarily follow that market quality improves when the dominant market becomes non-transparent, since the lack of transparency may impose costs on others. We return to trading cost measurement later, but first we examine price discovery, beginning with the information shares of various venues before and after Island goes dark.

4. Island Going Dark: Market Share and Information Share

The effects of Island going dark are immediate. Figure 2 provides the daily average market share for Island by share volume (Panel A) and number of trades (Panel B). Island's percentage of share volume in DIA and QQQ is about 35 percent through September 20, 2002 and then drops to about 20 percent the day it goes dark. For SPY, Island's percentage of share volume drops from over 20 percent until September 20 to 10 percent on September 23. Island's market share of trades is much higher, making the drop more pronounced. Island's trade share is between 55 and 60 percent in all three ETFs prior to going dark. This drops to below 40 percent in QQQ and just about 30 percent in DIA and SPY. Island's trade share continues to decline somewhat in all three.

[Insert Figure 2 Here]

Prior research shows that electronic markets' prices and quotes are informative beyond their fraction of trading (e.g., Huang (2002) and Hasbrouck (2002)), and trading on ECNs is associated with more informed trading (Barclay, Hendershott, and McCormick (2002)). Thus, the impact of Island going dark on the price formation process might be even greater than its precipitous decline in market share suggests. To examine this issue a framework for the price formation process across multiple markets must be laid out.

4.1 Information Share

Our analysis follows the information share approach of Hasbrouck (1995, 2002), Huang (2002), and others. A more detailed overview of this and alternative procedures is in Baillie et al. (2002), de Jong (2002), Harris, McNish, and Wood (2002), Hasbrouck (2002), and Lehmann (2002). These methods focus on the random-walk component of security prices across markets – the latent efficient price. The random-walk innovation variance is decomposed into components that can be attributed to innovations in each market's price. The relative contribution of a given price to the overall random-walk variance is defined as that price's information share.

Let p_t be a $(n \times 1)$ vector of transaction prices from n markets. If the prices are nonstationary, their future time path depends on past effects. Although the prices are individually nonstationary, they are for the same underlying asset, so a linear combination of the prices may be stationary. If all prices follow a random walk, they are integrated of order one and Δp_t is a stationary process. Prices are cointegrated if there exists a nonzero cointegrating vector b such that $b'p_t = e_t$ is stationary. The vector e_t can be interpreted as a deviation from long-run equilibrium. When e_t is stationary, pricing errors are corrected over time.

Cointegrated variables can be equivalently represented as a vector autoregression (VAR) in levels, as a vector error correction model (VECM), or as a vector moving average (VMA) representation. Suppose p_t can be expressed as a VAR:

$$p_t = \phi_0 + \sum_{j=1}^q \phi_j p_{t-j} + \varepsilon_t, \quad (1)$$

where $E(\varepsilon_t) = 0$ and $E(\varepsilon_t \varepsilon_s') = \Omega$ if $s=t$ and 0 otherwise. Then there exists a VECM:

$$\Delta p_t = \theta_0 p_{t-1} + \sum_{j=1}^{q-1} \theta_j \Delta p_{t-j} + \varepsilon_t, \quad (2)$$

where $\theta_j = -\sum_{i=j+1}^q \phi_j$ for $j = 1, \dots, q-1$ and $\theta_0 = -\left(I_n + \sum_{i=1}^q \phi_j\right)$.

Alternatively, the common trends representation proposed by Stock and Watson (1988) is:

$$p_t = p_0 + \Lambda(1) \sum_{i=1}^t \varepsilon_i + \Lambda^*(L) \varepsilon_t, \quad (3)$$

where $\Lambda^*(L)$ is a polynomial matrix with elements which are scalar polynomials in the argument L , p_0 is the initial price, and the last term of (3) is a stationary process. The middle right-hand term is the stochastic common trend to prices in all markets and captures the permanent impact of new information on prices. This impact excludes all transient price effects such as inventory adjustment and bid-ask bounce. Hasbrouck (1995) provides a measure of a market's contribution to price discovery using the total variance of the common stochastic trend, $\Lambda\Omega\Lambda'$, where Λ is a row in $\Lambda(1)$, and Ω is the $(n \times n)$ covariance matrix of the innovations. The share of this total variance attributed to market j is

$$I_j \equiv \frac{\Lambda_j^2 \Omega_{jj}}{\Lambda\Omega\Lambda'}, \quad (4)$$

If the innovation covariance matrix (Ω) is not diagonal, the information share is not exactly identified. To bound the contribution of an innovation, the matrix is typically diagonalized by finding the Cholesky factorization of Ω : $\Omega=FF'$, where F is a triangular matrix. Using F so that $\varepsilon_t =F u_t$, the innovations (u_t) are orthogonal to the preceding innovations in the ordering. Thus, the proportion of the information attributable to a market depends on the ordering. The upper and lower bounds of the information share of each market are calculated over all possible orderings of the markets.

In this case, with a vector of prices in six markets, the VECM specification (2) requires 36 estimated coefficients per lag. If prices are measured every second, then a two-minute

maximum lag results requires estimation of $36 \times 120 = 4,320$ coefficients. To reduce the number of estimated coefficients, we follow Hasbrouck (1995, 2002) and use polynomial distributed lags. The 6×6 coefficient matrices θ_j are constrained to be quadratic in the lag j . That is,

$$\theta_j = A_0 + A_1j + A_2j^2, \quad (5)$$

where the A_i are 6×6 estimated matrices. Separate quadratic polynomials are fit over various lag ranges, and at longer lags, the coefficient matrix θ_j is constrained to be constant over a pre-specified lag range.¹²

Figure 3 shows Island's maximum (Panel A) and minimum (Panel B) information share each day. Island's maximum information share is between 94 and 97 percent prior to going dark and falls to 74, 91, and 88 percent for DIA, QQQ, and SPY, respectively. Its maximum information share continues to decline through the post period finishing at 60, 90, and 85 percent for DIA, QQQ, and SPY. The minimum information shares follow a similar pattern starting between 90 and 95 percent and declining to 55 to 80 percent.¹³ Note that the fragmentation of price discovery and the fragmentation of market share in the post-dark period are significantly different. Even though Island's share of trades falls to 30 to 40 percent in the post-dark period, its information share remains at 60 to 90 percent. This demonstrates that Island's higher price discovery is not simply generated by its higher trading activity.

[Insert Figure 3 Here]

Just as trading activity shifts from Island to the other electronic markets, so does price discovery. Table 2 presents the daily averages of each market's maximum and minimum information share for the three ETFs in both the pre- and post-dark periods, and inference is conducted using the time series of daily information shares. The floor-based markets see some increases in price discovery that are significant at the 0.01 level, but these are small in economic terms, with none of AMEX, NYSE, or the regionals having a maximum information share of five

¹² Prices are measured at one-second intervals, and the maximum lag is two minutes, for a total of 120 lags. Three separate second-degree PDL's are imposed on lags 1-10, lags 11-20, and lags 21-30. VECM coefficients are constant for all lags 31-60. Similarly, the coefficient matrix is restricted to be constant over lags 61-90, and constant over lags 91-120.

¹³ For robustness, we also calculate information shares including non-Island quote midpoints in the vector of prices. This has almost no effect on Island's dominance.

percent in any of the ETFs. In contrast, all of the changes in information share for ARCA-Pacific and Instinet-3rd market are statistically significant. Island's going dark results in ARCA-Pacific's maximum information share increasing from about one percent to over ten percent in all three ETFs. Instinet-3rd market's information share in DIA increases from less than five percent to over 20 percent when Island goes dark.

[Insert Table 2 Here]

The final columns of Table 2 calculate a Herfindahl index for the concentration of price discovery across markets. As noted earlier, Instinet-3rd market is considered a single venue because TAQ does not provide data about the actual trading location within the third market. Because the maximum information shares add up to more than one and the minimum to less than one, each day we calculate the sum of the maximum and minimum information shares across markets and then divide each market's information share by the total. Another approach to deal with the contemporaneous component across markets would be to choose a particular orthogonalization. Since no orthogonalization is clearly preferable, we opt for our simpler approach. Given how tight the bounds on the information shares are, it is unlikely that the exact approach is important. The information share Herfindahl index decreases by roughly 40, 20, and 30 percent for DIA, QQQ, and SPY respectively. These decreases are much bigger than the Herfindahl index decreases based on trading volume, indicating that price discovery fragments more than trading volume.

5. The Effect of Island Going Dark on the Efficiency of Price Discovery

An important benefit of concentrating trading in a single market is more efficient price discovery (Pagano (1989b)). Thus, Island's going dark may hamper the incorporation of information into prices. This could occur for several reasons. First, having a dominant market for price discovery could simplify the process of determining the efficient price. Prior to going dark, the information share analysis in Table 2 suggests that traders could simply look at Island to obtain the efficient price. Once Island goes dark, traders may need look at Island, ARCA-Pacific, Instinet, and AMEX and decide how to weigh each venue. This process alone could slow price discovery over short horizons. Another possibility is that if trading costs increase when Island

goes dark (and they do; see below), information that was profitable before might no longer yield a profitable trade. As a result, prices may adjust more slowly. To measure effects on the efficiency of price discovery beyond the increase in bid-ask bounce, we examine quote midpoints and measure the autocorrelation of their changes.

Table 3 shows the average daily autocorrelation of quote midpoint returns over 15-, 30- and 60-second intervals by ETF for the pre and post periods. In the pre-dark period, we calculate the NBBO two ways: using all market center quotes, and using only non-Island quotes. We do this to determine whether changes in price discovery are driven by the change in transparency on Island or the resulting order flow migration. However, we first focus on overall price efficiency using all market center quotes.

[Insert Table 3 Here]

The autocorrelations in Table 3 are all positive, signifying lagged adjustment to information. The quote change autocorrelations increase from the pre period to the post period in all ETFs for all time intervals, indicating a reduction in the efficiency of price discovery in the post-dark period. Seven of the nine increases in quote autocorrelation are significant at the 0.01 level. For example, autocorrelations for one-minute quote returns in QQQ increase markedly from 0.04 before September 23 to 0.11 afterward. Autocorrelations over five-minute and longer intervals are not reported, but are indistinguishable from zero in both the pre and post periods.

Given that Island's going dark is a change in transparency that leads to order flow migration, it is difficult to completely disentangle the two effects. However, examining the non-Island markets provides some insight. If order flow moving from Island to other markets results in improved price discovery in these markets, the fragmentation results in a more efficient price on these markets. Table 3 shows that this appears to be the case for SPY at all intervals.

Alternatively, given Island's dominant role in price discovery in the pre period, it is possible that traders off Island watch and benefit from the Island quotes. If this effect is stronger than the benefits from increased order flow, it could result in less efficient prices off Island in the post period. This doesn't seem to be the case for SPY, where off-Island quotes are more efficient when Island goes dark. However, the transparency effect appears stronger for DIA, as non-Island

quote autocorrelations increase in the post period at all intervals. For QQQ, the transparency and order flow effects appear to cancel out, with none of the changes in the quote autocorrelation statistically significant. Together these results indicate that the externalities of order flow and the ability to observe prices in a dominant market center both contribute to the efficiency of price discovery.

For all three ETFs studied here, there are also futures traded on the same underlying securities. Changes in the ETF market may also affect the linkages between the ETF and futures markets. Hasbrouck (2002) shows that almost all price discovery in the Nasdaq 100 and S&P 500 futures markets occurs in the E-minis. E-minis are variants of the standard index futures contracts with NQ being the Nasdaq-100 contract and ES being the S&P 500 contract. They have a smaller contract size and are traded electronically rather than using open outcry. To look at price discovery in the ETF and futures markets, we calculate first-order cross-correlations of returns in the two markets. We use E-mini transaction price data and ETF quote midpoints.¹⁴ If price discovery is fully efficient across the ETF and futures markets, these cross-correlations should be zero.

[Insert Table 4 Here]

Table 4 shows the average daily first-order cross autocorrelation of ETF quote midpoint returns and E-mini futures transaction price returns for 15-, 30-, and 60-second intervals. Again, we calculate the ETF quote returns in the pre-dark period both using our NBBO of all market centers' quotes and using an NBBO of all market centers except Island. The cross-autocorrelation of ETF NBBO returns with lagged E-mini returns is positive, while the lagged ETF returns provide little predictive power for current E-mini returns. These imply that ETF prices lag futures prices. This is consistent with Hasbrouck (2002), who finds that price discovery in these indexes is concentrated in the E-mini futures market.¹⁵ When Island goes

¹⁴ Historical quotes for Globex trading are not available from the Chicago Mercantile Exchange, so transaction prices are used. ETF transaction prices provide the same qualitative results, but the bid-ask bounce increases the variance of ETF returns, reducing the autocorrelations toward zero in both the pre and post periods.

¹⁵ This could also be due to difficulties in reconciling the times across both markets. The discussion of merging the TAQ and ITCH data notes the several second lags between the ITCH time stamps and the TAQ time stamps. The changes between the pre- and post-dark periods in the cross auto-correlations here are not sensitive to this – similar results are found when ETF returns are calculated using ETF transaction prices and when ETF returns are calculated

dark, the lag in ETF prices increases and the cross-autocorrelation of ETF returns with lagged E-mini returns goes up significantly. This is true for both the Nasdaq-100 and the S&P 500, and it is true for 15-, 30-, and 60-second returns. For example, the cross-autocorrelation in 15-second Nasdaq-100 returns goes from 0.12 pre- to 0.41 post-event.

If we perform the same analysis using only non-Island quotes in the pre-dark period, the results are consistent with the autocorrelation results of Table 3. For the Nasdaq-100, the cross-autocorrelations using non-Island quotes in the pre period are similar to the post period, with the cross-autocorrelations higher in post period at 15 seconds, similar at 30 seconds, and lower at 60 seconds. For the S&P 500, E-mini futures lead the non-Island ETF quotes less in the post period than in the pre period.

Another way to assess price efficiency is to look at the intraday variability in the basis between futures and ETFs, as discussed in Hasbrouck (2002). While the basis varies over time depending on interest rates, dividend yields, and accrued but unpaid dividends on the ETF, the basis should be constant within a trading day if the ETF perfectly tracks the E-mini or vice versa. Thus, for each trading day we examine deviations from the mean basis for that day. We use all ETF and E-mini transaction prices at 60-second intervals based on TAQ timestamps, but the results are virtually identical for other sampling intervals. When Island goes dark, the two markets become less tightly linked. For the Nasdaq-100, the average absolute deviation rises from 5.52 basis points pre-dark to 5.91 basis points afterward. Similarly, when Island goes dark, the average absolute deviation for the S&P 500 rises from 1.69 to 2.54 basis points. Both of these changes are statistically significant at the 0.01 level, where inference is conducted using the daily time series of average absolute deviations.

Thus, quote autocorrelations, quote cross-autocorrelations, and the variance of deviations between ETF and futures prices all yield similar results. When Island goes dark, prices become less efficient both within the ETF market and across markets for related securities.

using Island transaction prices with the ITCH time stamps, but how much the futures prices lead the ETF prices may be sensitive to these timing issues. We explore this in the next section.

6. The Effect of Island Going Dark on Price Discovery Across Securities

Hasbrouck (2002) finds that the E-mini futures prices dominate intraday price discovery in the Nasdaq-100 and S&P 500 indices in March through May of 2000. However, two years later spreads on the ETF have fallen by an order of magnitude – spreads are one to two cents in our sample versus twenty cents in Hasbrouck (2002). Also, ETF trading volume has increased substantially to 40,000 and 27,000 daily trades in QQQ and SPY versus 10,000 and 1,700 daily trades in Hasbrouck’s 2000 sample. Both of these statistics suggest that ETFs’ role in index price discovery may have increased since 2000. If this is the case, then Island going dark may have important effects on price discovery between ETFs and futures.

To measure price discovery between the ETFs and futures we return to the information share analysis from Section 4 (and in Hasbrouck (2002)). We estimate each VECM using ETF and associated E-mini prices at one-second resolutions. Table 5 presents the daily averages of the ETF and E-mini maximum and minimum information share for the Nasdaq-100 and S&P 500 in both the pre- and post-dark periods. Using all ETF trade prices with the TAQ time stamps shows E-mini dominance similar to Hasbrouck (2002), although the 14 to 17 percent information share for ETFs in the Nasdaq-100 is higher than the five to seven percent in Hasbrouck (2002). The post-dark increase in the E-mini information share and decrease in the ETF information share is significant in the Nasdaq-100, with the ETF losing almost its entire role in price discovery. For the S&P 500, the ETF information share is only three or four percent in the pre period and declines to one percent or less after Island goes dark.

[Insert Table 5 Here]

From matching Island trades to their corresponding TAQ prints, we know that there is a lag of at least a second or two between the ITCH timestamp and the TAQ timestamp. This lag may bias downward the price discovery attributable to ETFs. To test this, we redo the analysis using just the Island trades and the ITCH time stamps to calculate the ETF information share. Table 5 shows that this makes a significant difference, with ETFs commanding sixty to eighty percent of the price discovery in the pre period in the Nasdaq-100 and 30 to 50 percent in the S&P 500. These numbers indicate that Island is responsible for a substantial fraction of price

discovery pre-event. ETF information shares drop by about one-half to two-thirds in the post period; price discovery is moving to the futures.

When ITCH data are used, there is a large difference between maximum and minimum information shares. This implies that significant contemporaneous price discovery is occurring across the E-mini and ETF markets. This also explains why using the less informative non-Island ETF trades reduces the ETF information share. Calculating the information shares with only Island trades (details not presented), but with their TAQ time stamps shows that about one-third of the increase in ETF information shares when using the ITCH data is due to using only Island trades and the other two-thirds is due to the lagged time stamps in TAQ.¹⁶

[Insert Figure 4 Here]

Figure 4 shows the daily maximum information shares for ETF and E-mini futures prices using only the ITCH ETF transactions prices. As with the other Figures in the paper, Figure 4 shows that the date Island went dark, September 23, 2002, was a shift in regime. When Island goes dark, index price discovery migrates substantially from ETFs on Island to E-mini futures on Globex.

7. The Effect of Island Going Dark on Trading Costs

We know that the elimination of quote transparency on Island reduces the information available to limit order providers. This may result in higher trading costs on Island due to greater adverse selection and/or less aggressive competition in liquidity provision. However, a change in one market's transparency should not by itself directly affect trading costs in other markets. Instead, changes in other markets are most likely due to order flow migration. If liquidity externalities are important, then fragmentation of the order flow away from Island should increase the liquidity externalities in the markets that pick up order flow and decrease trading costs in those markets.

¹⁶ There is a significant difference between our sample period and that of Hasbrouck (2002). He shows that E-minis continue to dominate price discovery even if they are lagged by 10-15 seconds, while we find a more developed ETF market where price discovery findings are sensitive to much smaller changes in time stamps.

7.1 Effective and Realized Spreads on Island

Panel A of Figure 5 shows Island's average daily effective half-spreads – measured as the absolute difference between the transaction price and the contemporaneous bid-ask midpoint of our consolidated NBBO – by ETF over the sample period.¹⁷ Island's effective spreads are low, less than two cents for all three major ETFs and close to half a cent in QQQ. Island's effective spreads double in DIA and QQQ and go up by 50 percent in SPY after going dark.

[Insert Figure 5 Here]

If the increase in Island's effective spreads is that reduction in transparency on Island reduced competition for liquidity provision, then profits from providing liquidity on Island should go up post-event. Realized spreads—the difference between the transaction price and the quote midpoint five minutes after the trade signed by the direction of the trade—provide an *ex post* measure of trading cost to the demander of liquidity and an *ex post* measure of profitability for the provider of liquidity. Panel B of Figure 5 shows Island's average daily realized half-spread by ETF over the sample period. The realized spreads on Island are small both before and after the ECN goes dark, but they increase after Island goes dark, indicating that competition for liquidity provision on Island declines. However, the increase in realized spreads is not as obvious or large as the increase in effective spreads, suggesting that prices move more in the direction of a trade after Island goes dark. That is, going dark may have increased adverse selection on Island. We return to these issues in a moment, but first we examine the effect of Island's going dark on trading costs in other markets.

7.2. Effective and Realized Spreads across Markets

While reducing liquidity on Island raises trading costs there, the migration of liquidity should lower trading costs in the markets where the liquidity moves. The net overall effect of rising costs on Island and falling costs elsewhere is ambiguous. To confirm that any change in spreads is not due to a marketwide change in liquidity, Figure 6 compares overall ETF spreads to spreads in similarly active individual stocks. Effective half-spreads are calculated for the average across DIA, QQQ, and SPY, as well as two highly active stocks, General Electric (GE) and

¹⁷ All trading cost measures are calculated on a share-weighted basis throughout. Daily averages are reported and used for statistical inference, assuming independence over time.

Microsoft (MSFT), which are listed on the NYSE and Nasdaq, respectively. Figure 6 shows that in the pre period, trading costs in cents per share are similar for the ETFs and the stocks. Because the average ETF share price of \$70 is higher – GE is about \$25 and MSFT is \$50 – the ETF costs are lower in percentage terms. GE and MSFT spreads stay roughly constant throughout the sample period, while the ETFs increase when Island goes dark. The increase in overall ETF effective spreads in Figure 6 is less pronounced than the increase in Island effective spreads in Figure 5, suggesting that effective spreads off Island fell when Island went dark.

[Insert Figure 6 Here]

To test the statistical significance of the increase in ETF effective spreads in Figure 6 and to examine the effects across markets, Table 6 reports average daily effective half-spreads by ETF and market for the pre and post periods. Because Island’s going dark resulted in an increase in average trade size, to examine market-wide costs we classify trades by size, 1,000 shares or less (small), 1,001 to 9,999 shares (medium), or 10,000 shares or more (large). The effective half-spreads are quite small, averaging between half a cent and six cents depending on trade size and trading venue. For DIA and QQQ, overall non-Island trading costs decrease, with most of the decreases statistically significant at the 0.01 level. In SPY, only ARCA-Pacific shows a statistically significant spread decline, and Instinet-3rd market effective spreads increase significantly. This increase for Instinet-3rd market is presumably due to Instinet’s October 1, 2002 decision to stop trading SPY, as its market share was also exceeding the five percent trigger for the display requirements of Regulation ATS.

[Insert Table 6 Here]

Prior to going dark, Island has the lowest effective spreads. Going dark raises spreads on Island and shifts trading to higher cost markets. Therefore, even though effective spreads generally fall in markets other than Island, the net effect is an increase in effective spreads. The increase is almost 10 percent for DIA and QQQ and over 20 percent in SPY. Therefore, we find no support for the Madhavan, Porter, and Weaver (2000) result that better transparency increases trading costs, although our event is not directly comparable because it is the complete loss of transparency at one of many trading venues and significant order flow migration occurs.

The evidence indicates that going dark increases Island's realized spreads. This implies a decrease in competition for providing liquidity on Island. If this decrease in competition is due to the loss of order flow on Island, there could be a corresponding increase in competition in the markets where liquidity moves. To explore this, Table 7 reports average daily realized half-spreads by ETF, trade size, and market for the pre and post periods. As noted earlier, realized spreads on Island increase and the increase is significant in DIA and QQQ. The realized spreads on ARCA-Pacific decrease sharply in all ETFs, indicating that the shift in volume increases competition in providing liquidity there. Aggregated across all markets, realized spreads increase and the effect is significant in DIA and SPY. As with effective spreads, the increase in SPY realized spreads in Instinet-3rd market is probably due to Instinet stopping trading in SPY in October.

[Insert Table 7 Here]

Tables 6 and 7 demonstrate that trading costs on Island increase when it goes dark, particularly relative to other markets. The fall in costs in other markets mitigates the overall impact, and it is possible that some other market-wide effect is responsible for the overall increase in effective spreads. To account for this, we include as controls both a volatility measure and a measure of spreads in other stocks. ETF spreads should be positively related to both of these. We use the contemporaneous day's transaction price range (the day's high minus the day's low) in the given ETF as the volatility measure. To measure changes in other stocks' bid-ask spreads, we construct a market-wide liquidity factor based on the dollar volume-weighted daily average effective spread for all NYSE-listed common stocks. Table 8 regresses average daily ETF percentage effective and realized spreads for all trades, Island trades, and non-Island trades on ETF dummies, a dummy for the dark period, trade size dummies, the dark dummy interacted with the trade size dummies, volatility, and the market liquidity factor.

[Insert Table 8 Here]

The first column of Table 8 shows that Island's going dark increases overall effective spreads by 0.60 basis points. The mean effective spread for the entire sample period in the regression is 4.8 basis points, implying that the effective spread increases by more than 10 percent. Similarly, overall realized spreads increase in both relative and absolute terms, by 0.79

basis points on an average of 3.7 basis points. That is, after controlling for differences in volatility and market-wide liquidity, overall trading costs rise and overall competitiveness of liquidity provision declines.

While we are attempting to measure changes in trading costs, it is important to acknowledge that the total expected trading costs include implicit, explicit, and opportunity costs. Because we do not have order submission data or a complete breakdown of commissions, we focus on the implicit trading costs as indicated by measures of the spread. One possible change in explicit trading costs is the liquidity fee that Island charges to demanders of liquidity that most other markets do not. Island's net fee is 0.08 cents per share, which corresponds to about 0.1 basis points for DIA and SPY and 0.3 basis points for QQQ. Table 1 shows that going dark reduces Island's market share (in shares) by 20%, 14%, and 9% respectively, so the overall increase in effective and realized spreads in Table 8 is more than 20 times the possible savings in liquidity fees for overall trading.

As in Tables 6 and 7, the increase in overall trading costs in Table 8 is driven by the increase in trading costs on Island, and the next two columns of Table 8 confirm the Island increase in the presence of volatility and market liquidity controls. The increase in effective spreads for Island trades is 1.12 basis points on an average of 2.5 basis points. Because realized spreads on Island are quite close to zero before going dark, the increase in realized spreads of 0.55 basis points is actually greater than the average realized spread of 0.3 basis points.

Tables 6 and 7 show a decrease in trading costs off Island, particularly on ARCA. The rightmost columns of Table 8 aggregate all non-Island trading together to test whether or not there is a significant decrease in trading costs in other markets when order flow leaves Island.¹⁸ The overall decrease in effective and realized spreads for non-Island trades is 0.46 and 0.30 basis points on means of 5.4 and 4.5 basis points. While the decrease in effective spreads is statistically significant, the fall in realized spreads for non-Island trades is not. There are several possible explanations. Realized spreads are noisier than effective spreads, which reduces statistical power. There is also a change in the composition of the order flow, as average trade

¹⁸ The cessation of SPY trading on Instinet on October 1, 2002 complicates the measuring of non-Island trading costs in SPY. In Table 8 non-Island spreads are calculated excluding all Instinet-3rd Market trades throughout the sample period. The analysis in Table 8 is qualitatively unchanged if SPY is excluded entirely.

sizes rise when Island goes dark. Finally, Island is dominated by relatively small trades, so we do not expect migration of these trades to affect realized spreads for large trades. To control for these, we disaggregate the trades by trade size and include trade size dummies. For both the small and medium trade sizes, the non-Island effective and realized spreads decline significantly. For large trades the changes in spreads are not significant.

Tables 6, 7, and 8 demonstrate that the reduction in transparency on Island diminishes the liquidity externalities and related competition for providing liquidity on Island and increases them in the markets where order flow moves. When Island goes dark, it costs more to trade there and liquidity providers earn more. When order flow moves elsewhere, the result in other markets is lower trading costs and lower profits for providing liquidity. Because Island is the dominant market with the lowest trading costs prior to going dark, the overall effect is a loss of liquidity externalities and higher trading costs. If Island were not the dominant market, the overall impact might be quite different.

Because Boehmer and Boehmer (2003) study the ETF market about a year prior to Island going dark, it is interesting to contrast our results to theirs. They find that while the NYSE's entry results in less concentrated trading, it significantly reduces spreads in all markets. Prior to the NYSE entry the ETF market appears relatively uncompetitive, with quoted spreads of 8.3, 2.7, and 6.4 cents for DIA, QQQ, and SPY respectively. These narrow by 20 to 50 percent with NYSE entry and narrow in all markets. These results are not inconsistent with ours. The NYSE entry improved competition, and this competitive effect most likely swamped any loss of liquidity externalities that might have occurred when trading became spread across more venues. By August 2002 quoted spreads decline even further, indicating a quite competitive market, and Island has captured significant volume.¹⁹ Thus, Island's going dark represents a substantially different event in that robust competition exists before the change.

¹⁹ Comparing our effective and realized spreads to those in Boehmer and Boehmer (2003) is somewhat difficult because they calculate spreads on a trade-weighted basis and in different trade size categories. However, it appears that effective spreads decrease similarly to quoted spreads from the end of their time period to the beginning of our time period. Boehmer and Boehmer (2003) report Nasdaq market share of 20 to 40 percent after NYSE entry whereas by our 2002 time period Island and Instinet-3rd Market represent about 60 percent of volume in all three ETFs.

An alternative approach to trading cost measurement is to use SEC Rule 11Ac1-5 (Dash5) filings by various market centers. Under Dash5, each market center must report monthly aggregate execution costs (including effective spreads and realized spreads) by stock for market orders and marketable limit orders of less than 10,000 shares. There are two main advantages to Dash5. First, there is no error in inferring whether a given trade was buyer or seller initiated. Second, execution costs are calculated based on the prevailing quotes at the time the order is received at the market center, not at the time the order is executed. In results not reported, we examine Dash5 filings from August 2002, the month before Island goes dark, and October 2002, the month after. We find roughly the same ordering and qualitative changes in trading costs for the non-ECN venues. In contrast, the ECN numbers are unreasonably large. For example, during August 2002 Island's average effective spread on marketable limits in DIA is more than 12 cents in every order size category. This is about four times the effective (whole) spread estimates in Table 6. Since we have the entire ITCH data feed, we have tried to reconstruct the Island Dash5 numbers but are unable to come anywhere close to these spread estimates. Others have encountered similar problems, and so many practitioners simply ignore the Dash5 numbers for Island and Instinet. In addition, Dash5 requires the use of the NBBO from the consolidated quote, thus excluding Island quotes in the pre-dark period. As we note above, since most price discovery occurs on Island, Dash5 trading costs are less useful if they are based on a less informative quote. Finally, Dash5 ignores the largest trades of more than 10,000 shares, which are quite common in these ETF markets. For all these reasons, we do not use Dash5 effective spread estimates in our main analysis.

8. The Effect of Island Going Dark on Private Information

Uninformed traders are likely to value transparency more than informed traders, and they would tend to migrate to transparent venues where they know the likely terms of trade in advance. If liquidity traders no longer congregate at the dominant market, this implies that those who remain on Island post-dark are more informed, resulting in higher costs of trading on Island. The expected effect on other markets is less clear. We know that the information share of non-Island markets goes up, which indicates that these other markets are picking up informed traders

as well as liquidity traders. Thus, the net effect on adverse selection in other markets is ambiguous.

The simplest measure of the amount of information in trades is the difference between the gross cost of the trade (the effective spread) and an *ex post* measure of net cost (the realized spread), or equivalently the change in the quote midpoint five minutes after the trade signed by the trade direction, often referred to as the price impact of the trade. The more information trades contain, the more prices will move in the direction of the trade (up following purchases and down following sales), reducing the *ex post* trading costs.

The difference between the effective spread in Table 6 and the realized spread in Table 7 is the average daily five-minute price impact by ETF, trade size, and market for the pre and post periods. Island price impacts generally increase. This implies that liquidity traders are leaving Island more than informed traders. With the exception of ARCA-Pacific, price impacts for non-Island markets do not show many significant changes. The significant increase in the price impact of trades on ARCA-Pacific suggests that the volume moving to it from Island is informative, whereas for the other markets it is less so.

The simple price impact of a trade is biased if trades are serially correlated, and in our data, the trade direction is persistent – buys are more likely to follow buys and sells are more likely to follow sells. But if a series of buys, for example, occurs in rapid succession, then each buy order is credited with the aggregate price impact of the entire series, and the price impact of any individual trade is overstated. Because trades are small and closer together on Island, the potential bias in the simple price impact is possibly significant.

Because the comparison of the effective and realized spreads in the pre and post periods show significant changes in the information content of trades occurs on Island, we construct a VAR with a quote equation, an Island trade equation, and a non-Island trade equation. Define x_t^n to be the sum of the signed non-Island trades (+1 for buys, -1 for sells) during the five-second interval t . Similarly, define x_t^I to be the sum of the signed Island trades during interval t , and

define r_t to be the percentage change (log return) in the quote midpoint during interval t . The VAR using 10 lags consists of three equations, two trade equations and a quote return equation:²⁰

$$x_t^I = \sum_{i=1}^{10} \phi_i r_{t-i} + \sum_{i=1}^{10} \kappa_i x_{t-i}^I + \sum_{i=1}^{10} \lambda_i x_{t-i}^n + \varepsilon_{3,t}, \quad (6)$$

$$x_t^n = \sum_{i=1}^{10} \delta_i r_{t-i} + \sum_{i=1}^{10} \phi_i x_{t-i}^I + \sum_{i=1}^{10} \gamma_i x_{t-i}^n + \varepsilon_{2,t}, \text{ and} \quad (7)$$

$$r_t = \sum_{i=1}^{10} \alpha_i r_{t-i} + \sum_{i=1}^{10} \beta_i x_{t-i}^I + \sum_{i=1}^{10} \zeta_i x_{t-i}^n + \varepsilon_{1,t}. \quad (8)$$

The VAR is inverted to get the VMA representation to focus on the impulse response functions that measure the permanent price impacts from a shock to each trade equation. As discussed in Hasbrouck (1991), this method is robust to price discreteness, lagged adjustment to information, and lagged adjustment to trades.

[Insert Table 9 Here]

Table 9 shows the average daily permanent price impacts of Island and non-Island trades by ETF for the pre and post periods. As with the simple price impacts, the impulse response shows a significant increase in the information in Island trades, but there is no such increase in information for non-Island trades. Before September 23, the impulse response to a trade on Island is between 100 and 550 percent greater than the response to a non-Island trade. When Island goes dark, the impulse response to an Island trade increases by 49, 66, and 40 percent in DIA, QQQ, and SPY respectively. These percentage increases are similar or somewhat higher than the simple price impact that can be calculated from the differences in effective and realized spreads in Tables 6 and 7. The price impacts in Table 9 suggest that Island appeals to traders with relatively better information about short-run price changes. Traders with this type of information are more likely to remain on Island after it goes dark.

Table 9 shows that adverse selection increases for both Island and non-Island trades in the post-dark period, although the increase in permanent price impact for non-Island trades is small and not statistically significant in DIA and SPY. To test if this is caused by a market-wide change in adverse selection in the post period, we estimate (details not reported) a VAR with trades from all markets aggregated in a single trade equation, and find little difference across

²⁰ Estimating the VAR with three to 20 second time intervals and five to 20 lags provides qualitatively similar results. VARs with orthogonalized residuals also provide qualitatively similar results.

periods. Similarly, the difference between the coefficient on the dark period dummy in the effective and realized spread regressions represents the change in the simple price impact. For all trades, the dark coefficient is larger in the realized spread regression, implying that the overall average simple price impact falls slightly in the dark period.

At first glance, it may seem odd that adverse selection increases on Island, does not change off Island, but does not increase overall. The explanation lies in the fact that adverse selection is significantly higher on Island. When some of the less informed traders leave Island, these traders in aggregate are still more informed than the average traders in the other markets. Therefore, adverse selection can increase on Island without an increase in adverse selection overall. To see this in our results look at SPY in Table 9. If Island had two equal-sized groups of traders in the pre period with price impacts of 0.06 and 0.16 respectively, then the average price impact for an Island trade in the pre period would be 0.11. If the group with the smaller price impact left Island when it went dark, then the price impact of an Island trade would be 0.16 in the post period. Because non Island trades in the pre period have a price impact of about 0.06 as well, the price impact of non Island trades would remain unchanged in the post period.

9. Conclusions

Recent experience with Nasdaq order handling rules (Barclay et al. (1999)), cross listing of options contracts (Battalio, Hatch, and Jennings (2003) and de Fontnouvelle, Fische, and Harris (2003)), and entry in ETF trading by the NYSE (Boehmer and Boehmer (2003)) demonstrate the substantial benefits of increasing competition and integration between trading venues. Through Regulation ATS the SEC is attempting to bring these benefits to investors. In response to the SEC enforcing the display requirements of Regulation ATS, the Island ECN stopped displaying its limit order book completely in several ETFs beginning on September 23, 2002. Island was the dominant market in these instruments, so this event had a significant effect on the ETF trading environment. Although the majority of price discovery still occurs on Island, and it remains the most active ETF trading venue, it loses substantial market share. ETF trading is significantly less concentrated than it was before Island went dark.

After Island goes dark, price discovery becomes less efficient both within the ETF market and between the ETF and futures markets. Substantial price discovery also moves from

the ETF market to the futures market. We also find that trading costs, measured by effective and realized spreads, increase on Island and decrease on other venues. The net effect is a substantial increase in overall effective and realized spreads. This indicates that competition on and off Island is significantly altered by this event, with the provision of liquidity becoming less competitive on Island and more competitive in the markets that pick up order flow. Overall, ETF market quality worsens when Island goes dark.

This paper also provides evidence on the importance of liquidity externalities and their relationship to competition within and among markets. A key finding is the decrease in realized spreads off Island. This indicates that these other markets become more competitive as they obtain more order flow. The conventional wisdom is that liquidity begets liquidity. The evidence here refines that virtuous cycle: liquidity (order flow) begets competition, which improves price discovery and reduces trading costs, which attracts additional order flow, which begets more competition. These results suggest that competition across markets is not perfect and a less concentrated trading arrangement results in weaker competition among liquidity providers. Competition among markets is not perfect because of a lack of complete transparency and integration.

Given the current situation, what would improve ETF market quality? Market quality should improve if the event is reversed, i.e., Island resumes display of its order book and regains its previous dominant position. However, the evidence here suggests other avenues. Market quality should also improve if another electronic limit order book is able to achieve the dominant position ceded by Island. Even so, a monopolistic trading arrangement is not necessarily optimal, and we would not advocate that trading be forced to concentrate on Island or any other single venue. If there is heterogeneity across traders, multiple venues can be optimal, and in fact competition across these venues is healthy and keeps them responsive to their customers. Thus, we would argue instead that improving the linkages between market centers may be the best approach to maintaining the benefits of competition across venues without destroying liquidity externalities.

Throughout the paper, we avoid attributing the decline in market quality to either the SEC or Island, but clearly the SEC bears final responsibility for its rules and their effect on market quality. This experiment could have been avoided if the SEC had held off enforcement of

Regulation ATS. However, this would have eroded the fundamental, Congressionally mandated principle of a national market system for trading stocks which has demonstrated its value in the past. This suggests that the details of the national market system should be reconsidered.

Some markets, including Island, have argued for complete repeal of the trade-through prohibitions of ITS. However, this could increase agency problems between brokers and customers. To address this concern, mandatory disclosure of trade-throughs could replace the trade-through prohibition. If the broker makes an order-routing decision, the broker would be required to disclose any trade-through to the customer. Such sunshine might work better than any other regulation in ensuring that brokers get their customers the best possible execution.

Another natural possibility is to modify ITS or the ITS rules. Currently, the receiving market has 30 seconds to either execute or reject an incoming ITS order. This delay imposes execution uncertainty on those seeking immediacy, because posted quotes are not firm quotes under existing ITS rules. In principle, it would be straightforward to require all markets to honor the quote in effect at the time an ITS order is received. This would force market-makers to be more diligent about keeping quotes up to date, and it would improve quote transparency, as all posted quotes would be firm. However, this may impose large costs on specialists and other market-makers. In fact, floor-based systems do not guarantee execution at the quote even for orders routed directly to them, so this suggestion may be impractical.

An alternative is to add a second tier to ITS for those seeking an immediate execution. This second tier would have a maximum delay of one to five seconds. With such a setup, market integration could be achieved for venues offering faster execution. Slower markets could choose to modify their internal systems to offer immediate execution for some orders or simply not participate in the faster tier of ITS. In submitting an order, investors could explicitly choose faster execution, or slower execution with the possibility of a better price. Trade-through rules would continue to apply, but orders designated for faster execution would only need to avoid trading through venues that meet the speed criteria. At first glance, this proposal might seem to disadvantage floor-based systems by relegating the NYSE and AMEX to the slower tier. However, the NYSE has an automated execution facility, DirectPlus, which would probably qualify for the faster ITS tier. Presumably AMEX could respond similarly. Overall, such a change to ITS would allow traders more choice along these two dimensions of execution quality,

and the tighter linkage between markets would encourage vigorous competition between venues without endangering liquidity externalities.

REFERENCES

- Admati, Anat, and Paul Pfleiderer, 1988, A theory of intraday patterns: volume and price variability, *Review of Financial Studies* 1, 3-40.
- Admati, Anat, and Paul Pfleiderer, 1991, Sunshine trading and financial market equilibrium, *Review of Financial Studies* 4, 443-481.
- Baillie, Richard, G. Geoffrey Booth, Yiuman Tse, and Tatyana Zobotina, 2002, Price discovery and common factor models, *Journal of Financial Markets* 5, 309-322.
- Barclay, Michael, William Christie, Jeffrey Harris, Eugene Kandel, and Paul Schultz, 1999, The effects of market reform on the trading costs and depths of Nasdaq stocks, *Journal of Finance* 54, 1-34.
- Barclay, Michael, and Terrence Hendershott, 2002, Liquidity externalities and adverse selection: Evidence from trading after hours, *Journal of Finance*, forthcoming.
- Barclay, Michael, Terrence Hendershott, and D. Timothy McCormick, 2002, Competition among trading venues: Information and trading on electronic communications networks, *Journal of Finance*, forthcoming.
- Baruch, Shmuel, 2002, Who benefits from an open limit-order book? Working paper, University of Utah.
- Battalio, Robert, Brian Hatch, and Robert Jennings, 2003, Toward a National Market System for the U.S. exchange-listed equity options, *Journal of Finance*, forthcoming.
- Biais, Bruno, 1993, Price formation and equilibrium liquidity in fragmented and centralized markets, *Journal of Finance* 48, 157-185.
- Biais, Bruno, Christophe Bisiere, and Chester Spatt, 2002, Imperfect competition in financial markets: Island vs Nasdaq, Working paper, Carnegie Mellon University.
- Bloomfield, Robert, and Maureen O'Hara, 1999, Market transparency: Who wins and who loses? *Review of Financial Studies* 12, 5-35.
- Boehmer, Beatrice, and Ekkehart Boehmer, 2003, Trading your neighbor's ETFs: Competition or fragmentation? *Journal of Banking and Finance*, forthcoming.
- Boehmer, Ekkehart, Gideon Saar, and Lei Yu, 2002, Lifting the veil: An analysis of pre-trade transparency at the NYSE, working paper, New York University.
- Chowdhry, Bhagwan and Vikram Nanda, 1991, Multimarket trading and market liquidity, *Review of Financial Studies* 4, 483-511.

- Conrad, Jennifer, Kevin Johnson, and Sunil Wahal, 2003, Institutional trading and alternative trading systems, *Journal of Financial Economics* 70, 99-134.
- de Fontnouvelle, Patrick, Raymond Fishe, and Jeffrey Harris, 2003, The behavior of bid-ask spreads and volume in options markets during the competition for listings in 1999, *Journal of Finance*, forthcoming.
- de Frutos, M. Angeles and Carolina Manzano, 2002, Risk aversion, transparency, and market performance, *Journal of Finance* 57, 959-984.
- de Jong, Frank, 2002, Measures of contributions to price discovery: A comparison, *Journal of Financial Markets* 5, 323-328.
- Ellis, Katrina, Roni Michaely, and Maureen O'Hara, 2002, The making of a dealer market: From entry to equilibrium in the trading of Nasdaq, *Journal of Finance* 57, 2289-2316.
- Foster, F. Douglas and S. Viswanathan, 1990, A theory of intraday variations in volume, variance, and trading costs in securities markets, *Review of Financial Studies* 3, 593-624.
- Foucault, Moinas, and Theissen, 2002, Does anonymity matter in electronic limit order markets? Working paper, HEC.
- Flood, Mark, Ronald Huisman, Kees Koedijk, and Ronald Mahieu, 1999, Quote disclosure and price discovery in multiple-dealer financial markets, *Review of Financial Studies* 12, 37-59.
- Glosten, Lawrence, 1987, Components of the bid-ask spread and the statistical properties of transaction prices, *Journal of Finance* 42, 1293-1307.
- Glosten, Lawrence, 1994, Is the electronic open limit order book inevitable? *Journal of Finance* 49, 1127-1161.
- Glosten, Lawrence, 1999, Introductory comments: Bloomfield and O'Hara, and Flood, Huisman, Koedijk, and Mahieu, *Review of Financial Studies* 12, 1-3.
- Harris, Frederick, Thomas McInish, and Robert Wood, 2002, Common factor components vs. information shares: A reply, *Journal of Financial Markets* 5, 341-348.
- Hasbrouck, Joel, 1991, Measuring the information content of stock trades, *Journal of Finance* 46, 179-207.
- Hasbrouck, Joel, 1995, One security, many markets: Determining the contributions to price discovery, *Journal of Finance* 50, 1175-1199.
- Hasbrouck, Joel, 2002, Intraday price formation in the U. S. equity index markets, *Journal of Finance*, forthcoming.

- Hasbrouck, Joel, and Gideon Saar, 2001, Limit orders and volatility in a hybrid market: The Island ECN, Working paper, New York University.
- Hendershott, Terrence, and Haim Mendelson, 2000, Crossing networks and dealer markets: Competition and performance, *Journal of Finance* 55, 2071-2115.
- Huang, Roger, 2002, The quality of ECN and Nasdaq market-maker quotes, *Journal of Finance* 57, 1285-1319.
- Lee, Charles, and Mark Ready, 1991, Inferring trade direction from intraday data, *Journal of Finance* 46, 733-747.
- Lehmann, Bruce, 2002, Some desiderata for the measurement of price discovery across markets, *Journal of Financial Markets* 5, 259-276.
- Macey, Jonathan, and Maureen O'Hara, 1999, Regulating exchanges and alternative trading systems: A law and economics perspective, *Journal of Legal Studies* 28, 17-54.
- Madhavan, Ananth, 1992, Trading mechanisms in securities markets, *Journal of Finance* 47, 607-642.
- Madhavan, Ananth, 1995, Consolidation, fragmentation, and the disclosure of trading information, *Review of Financial Studies* 8, 579-603.
- Madhavan, Ananth, 2000, Market microstructure: A survey, *Journal of Financial Markets* 3, 205-258.
- Madhavan, Ananth, David Porter, and Daniel Weaver, 2000, Should markets be transparent? Working paper, Baruch College.
- Milgrom, Paul and Robert Weber, 1982, A theory of auctions and competitive bidding, *Econometrica* 50, 1089-1122.
- Pagano, Marco, 1989a, Trading volume and asset liquidity, *Quarterly Journal of Economics* 104, 255-274.
- Pagano, Marco, 1989b, Endogenous market thinness and stock price volatility, *Review of Economic Studies* 56, 269-287.
- Pagano, Marco, and Ailsa Roell, 1996, Transparency and liquidity: A comparison of auction and dealer markets with informed trading, *Journal of Finance* 51, 579-611.
- Parlour, Christine, and Duane Seppi, 2003, Liquidity-based competition for order flow, *Review of Financial Studies* 16, 301-343.

Rindi, Barbara, 2002, Transparency, liquidity and price formation, Working paper, Bocconi University.

Schultz, Paul, 2003, Who makes markets, *Journal of Financial Markets* 6, 49-72.

Simaan, Yusif, Daniel Weaver, and David Whitcomb, 2003, Market-maker quotation behavior and pre-trade transparency, *Journal of Finance* 58, 1247-1268.

Stock, James, and Mark Watson, 1988, Testing for common trends, *Journal of the American Statistical Association* 83, 1079-1107.

Weston, James, 2001, Competition on the Nasdaq and the growth of electronic communication networks, Working paper, Rice University.

Table 1
Sample Descriptive Statistics

Descriptive statistics based on TAQ and ITCH data for three exchange-traded funds – DIA, QQQ, and SPY – from August 16, 2002 to October 31, 2002. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. Descriptive statistics are calculated for each ETF each day and then averaged across days.

Average Daily	DIA		QQQ		SPY	
	Pre	Post	Pre	Post	Pre	Post
Volume (M Shares)	9.20	13.33	77.43	85.42	35.46	51.71
Volume (000 Trades)	16.90	15.87	43.47	39.16	27.09	26.13
Price	85.72	79.90	23.59	22.38	91.22	85.18
Quoted Spread (cents)	2.62	2.65	1.07	1.23	2.15	2.18
Market Share (Shares)						
AMEX	0.25	0.27	0.18	0.19	0.27	0.27
Island	0.36	0.16	0.36	0.22	0.21	0.12
NYSE	0.05	0.10	0.07	0.08	0.08	0.09
ARCA-Pacific	0.03	0.08	0.02	0.06	0.02	0.05
Other Regionals	0.08	0.10	0.12	0.13	0.07	0.08
Instinet-3rd Market	0.21	0.29	0.26	0.33	0.35	0.40
Herfindahl Index	0.25	0.21	0.25	0.22	0.25	0.26
Market Share (Trades)						
AMEX	0.09	0.13	0.06	0.08	0.13	0.17
Island	0.61	0.29	0.59	0.37	0.56	0.30
NYSE	0.02	0.04	0.04	0.05	0.03	0.05
ARCA-Pacific	0.05	0.14	0.03	0.11	0.04	0.14
Other Regionals	0.08	0.12	0.15	0.19	0.11	0.16
Instinet-3rd Market	0.15	0.29	0.14	0.21	0.14	0.20
Herfindahl Index	0.41	0.22	0.39	0.24	0.36	0.20
Fraction of Island Volume Executed Against Undisplayed Orders						
Shares	0.45	1.00	0.50	1.00	0.50	1.00
Trades	0.54	1.00	0.55	1.00	0.56	1.00

Table 2
Information Shares within the ETF Market

Statistics are based on a vector error correction model (described in Section 3) of transaction prices estimated at a one-second resolution. Estimates are calculated for each ETF each day and then averaged across days. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. An asterisk (*) represents a statistically significant difference between the pre and post period at the 0.01 level, based on the daily time series of information share estimates.

	AMEX		Island		NYSE		ARCA-P		Regionals		Instinet-3rd		Herfindahl	
	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
DIA Max														
Median	0.009	0.040	0.940	0.687	0.003	0.004	0.009	0.095	0.002	0.008	0.045	0.225	0.841	0.441
Mean	0.014	0.043 *	0.941	0.679 *	0.007	0.007	0.014	0.106 *	0.006	0.012	0.045	0.257 *	0.842	0.485 *
Std Err	(0.002)	(0.004)	(0.005)	(0.026)	(0.002)	(0.001)	(0.003)	(0.015)	(0.002)	(0.003)	(0.004)	(0.026)	(0.010)	(0.018)
DIA Min														
Median	0.008	0.034	0.918	0.584	0.002	0.002	0.005	0.056	0.001	0.004	0.026	0.164	0.887	0.497
Mean	0.012	0.036 *	0.915	0.596 *	0.006	0.004	0.010	0.065 *	0.005	0.008	0.025	0.192 *	0.888	0.532 *
Std Err	(0.002)	(0.004)	(0.006)	(0.028)	(0.002)	(0.001)	(0.002)	(0.009)	(0.002)	(0.002)	(0.003)	(0.024)	(0.009)	(0.021)
QQQ Max														
Median	0.004	0.005	0.973	0.888	0.004	0.008	0.007	0.093	0.001	0.003	0.017	0.053	0.921	0.693
Mean	0.006	0.006	0.973	0.893 *	0.008	0.013	0.008	0.107 *	0.002	0.005 *	0.020	0.054 *	0.918	0.710 *
Std Err	(0.001)	(0.001)	(0.003)	(0.009)	(0.002)	(0.002)	(0.002)	(0.012)	(0.001)	(0.001)	(0.003)	(0.005)	(0.007)	(0.017)
QQQ Min														
Median	0.003	0.004	0.960	0.817	0.003	0.005	0.004	0.049	0.001	0.001	0.006	0.026	0.948	0.788
Mean	0.005	0.005	0.957	0.823 *	0.007	0.009	0.005	0.056 *	0.001	0.003 *	0.009	0.029 *	0.947	0.800 *
Std Err	(0.001)	(0.001)	(0.004)	(0.013)	(0.002)	(0.002)	(0.001)	(0.008)	(0.000)	(0.001)	(0.002)	(0.004)	(0.005)	(0.016)
SPY Max														
Median	0.005	0.028	0.971	0.822	0.003	0.015	0.005	0.168	0.002	0.008	0.025	0.051	0.903	0.593
Mean	0.007	0.034 *	0.968	0.809 *	0.006	0.020 *	0.011	0.159 *	0.004	0.013 *	0.026	0.066 *	0.901	0.592 *
Std Err	(0.001)	(0.006)	(0.004)	(0.021)	(0.001)	(0.004)	(0.003)	(0.021)	(0.001)	(0.003)	(0.003)	(0.009)	(0.009)	(0.023)
SPY Min														
Median	0.003	0.019	0.951	0.739	0.002	0.011	0.004	0.098	0.001	0.005	0.008	0.029	0.943	0.666
Mean	0.005	0.026 *	0.949	0.719 *	0.005	0.015 *	0.008	0.098 *	0.003	0.009 *	0.011	0.036 *	0.938	0.671 *
Std Err	(0.001)	(0.005)	(0.005)	(0.023)	(0.001)	(0.003)	(0.003)	(0.014)	(0.001)	(0.002)	(0.002)	(0.005)	(0.007)	(0.026)

Table 3
Quote Return Autocorrelations

The first order autocorrelation of quote midpoint returns is calculated over 15-, 30-, and 60-second intervals for each ETF each day and then averaged across days. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. The NBBO pre period uses quotes from all market centers and the non Island pre period uses quotes from all markets except Island. Standard errors are in parentheses. An asterisk (*) represents a statistically significant difference between the pre and post period at the 0.01 level, based on the daily time series. A dagger (†) represents a statistically significant difference between the non Island pre and post periods at the 0.01 level.

Interval	DIA			QQQ			SPY		
	Pre	Non Island Pre	Post	Pre	Non Island Pre	Post	Pre	Non Island Pre	Post
15 seconds	0.00 (0.01)	0.06 (0.01)	0.08 * (0.01)	0.01 (0.01)	0.08 (0.01)	0.09 * (0.01)	0.09 (0.01)	0.18 (0.01)	0.12 *† (0.01)
30 seconds	0.07 (0.01)	0.14 (0.01)	0.18 *† (0.01)	0.04 (0.01)	0.16 (0.01)	0.14 * (0.01)	0.10 (0.01)	0.19 (0.01)	0.14 † (0.02)
60 seconds	0.06 (0.02)	0.10 (0.01)	0.13 * (0.02)	0.04 (0.01)	0.14 (0.01)	0.11 * (0.01)	0.04 (0.01)	0.09 (0.01)	0.05 † (0.01)

Table 4
ETF and Futures Return Cross-Autocorrelations

The first-order cross-autocorrelation of ETF quote midpoint returns and E-mini future transaction price returns are calculated over 15-, 30-, and 60-second intervals for each index each day and then averaged across days. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. Standard errors are in parentheses. The pre period uses quotes from all market centers and the non Island pre period uses quotes from all markets except Island. An asterisk (*) represents a statistically significant difference between the pre and post periods at the 0.01 level, based on the daily time series. A dagger (†) represents a statistically significant difference between the non Island pre and post periods at the 0.01 level.

Panel A: ETF(t) - E-mini(t-1)							
Interval	Nasdaq-100			S&P 500			
	Non Island			Non Island			
	Pre	Pre	Post	Pre	Pre	Post	
15 seconds	0.12 (0.01)	0.34 (0.02)	0.41 *† (0.01)	0.18 (0.01)	0.38 (0.01)	0.35 *† (0.01)	
30 seconds	0.09 (0.01)	0.41 (0.02)	0.40 * (0.01)	0.16 (0.01)	0.35 (0.01)	0.28 *† (0.01)	
60 seconds	0.06 (0.01)	0.33 (0.02)	0.28 *† (0.02)	0.08 (0.01)	0.21 (0.01)	0.13 *† (0.02)	

Panel B: ETF(t-1) - E-mini(t)							
Interval	Nasdaq-100			S&P 500			
	Non Island			Non Island			
	Pre	Pre	Post	Pre	Pre	Post	
15 seconds	0.04 (0.01)	-0.02 (0.01)	0.00 * (0.01)	0.06 (0.01)	0.04 (0.01)	0.03 * (0.01)	
30 seconds	0.00 (0.01)	-0.02 (0.01)	-0.01 (0.01)	0.04 (0.01)	0.02 (0.01)	0.01 * (0.01)	
60 seconds	-0.01 (0.01)	-0.02 (0.01)	-0.03 (0.01)	-0.01 (0.01)	-0.03 (0.01)	-0.03 (0.01)	

Table 5
Information Shares Across the ETF and Futures Markets

Statistics are based on a vector error correction model (described in Section 3) of prices estimated at a one-second resolution. Estimates are calculated for each index each day and then averaged across days. Statistics are calculated using all ETF transaction prices with their TAQ time stamps and Island transaction prices with their ITCH time stamps. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. An asterisk (*) represents a statistically significant difference between the pre and post period at the 0.01 level, based on the daily time series.

	All ETF Trades (TAQ Times)				Island ETF Trades (ITCH Times)			
	E-mini		ETF		E-mini		ETF	
	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Nasdaq-100 Max								
Median	0.947	0.986	0.077	0.022	0.371	0.755	0.785	0.433
Mean	0.858	0.985 *	0.167	0.023 *	0.335	0.757 *	0.802	0.415 *
Std Err	(0.038)	(0.002)	(0.039)	(0.002)	(0.035)	(0.008)	(0.023)	(0.011)
Nasdaq-100 Min								
Median	0.923	0.978	0.053	0.014	0.215	0.567	0.629	0.245
Mean	0.833	0.977 *	0.142	0.015 *	0.198	0.585 *	0.665	0.243 *
Std Err	(0.039)	(0.002)	(0.038)	(0.002)	(0.023)	(0.011)	(0.035)	(0.008)
S&P 500 Max								
Median	0.989	0.997	0.019	0.006	0.730	0.922	0.463	0.191
Mean	0.970	0.994	0.040	0.010	0.710	0.918 *	0.481	0.183 *
Std Err	(0.018)	(0.001)	(0.020)	(0.002)	(0.028)	(0.009)	(0.023)	(0.017)
S&P 500 Min								
Median	0.981	0.994	0.011	0.003	0.537	0.809	0.270	0.078
Mean	0.960	0.990	0.030	0.006	0.519	0.817 *	0.290	0.082 *
Std Err	(0.020)	(0.002)	(0.018)	(0.001)	(0.023)	(0.017)	(0.028)	(0.009)

Table 6
Effective Half-Spreads

Effective spreads, the absolute difference between the transaction price and the quote midpoint, are calculated (in cents) for each ETF each day and then averaged across days. Daily averages are calculated on a per share basis. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. Trade-sizes are 100 to 1,000 shares (small), 1,001 to 9,999 shares (medium), and greater than or equal to 10,000 shares (large). An asterisk (*) represents a statistically significant difference between the pre and post period at the 0.01 level, based on the time series of daily averages.

ETF	Trade Size	Overall		AMEX		Island		NYSE		ARCA-P		Regionals		Instinet-3rd	
		Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
DIA	Small	1.95	2.23 *	3.61	3.35 *	1.36	1.86 *	3.30	2.47 *	1.55	1.37 *	3.15	3.04	2.11	2.12
	Medium	3.61	3.39	4.43	4.10 *	1.49	2.61 *	4.05	3.11 *	2.09	1.61	4.13	3.87	4.04	2.98 *
	Large	5.20	4.95	5.78	5.00 *	1.61	3.77 *	5.56	4.50 *	1.42	1.20	6.27	4.39 *	5.30	5.42
	Overall	3.16	3.39 *	4.72	4.30 *	1.39	2.07 *	4.43	3.63 *	1.63	1.42 *	3.66	3.42	4.04	3.72
QQQ	Small	0.84	1.06 *	1.35	1.20 *	0.54	0.98 *	1.21	1.16 *	0.72	0.74	1.45	1.21 *	1.11	1.11
	Medium	0.98	1.16 *	1.58	1.37 *	0.57	1.02 *	1.34	1.33	0.84	0.76 *	1.85	1.53 *	1.27	1.20 *
	Large	2.07	2.01	1.88	1.62 *	0.61	1.05 *	2.32	2.15	0.96	0.76	2.13	2.05	2.22	2.25
	Overall	1.41	1.53 *	1.77	1.52 *	0.56	1.01 *	1.97	1.85	0.78	0.75	1.81	1.58 *	1.95	1.91
SPY	Small	1.73	2.09 *	2.71	2.63	1.24	1.56 *	2.55	2.17 *	1.62	1.28 *	2.83	2.86	1.81	2.27 *
	Medium	2.49	2.55	3.18	3.05	1.36	1.74 *	3.07	2.84 *	1.78	1.34 *	3.32	3.41	2.37	2.79 *
	Large	4.32	5.10 *	3.82	4.13	1.46	1.83 *	3.48	3.32	1.18	1.70	5.35	5.28	4.76	6.00 *
	Overall	3.25	3.97 *	3.47	3.65	1.28	1.67 *	3.35	3.14	1.68	1.32 *	3.58	3.73	4.24	5.39 *

Table 7
Realized Half-Spreads

Realized spreads, the difference between the transaction price and the quote midpoint five minutes after the trade signed by the direction of the trade, are calculated (in cents) for each ETF each day and then averaged across days. Daily averages are calculated on a per share basis. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. Trade-sizes are 100 to 1,000 shares (small), 1,001 to 9,999 shares (medium), and greater than or equal to 10,000 shares (large). An asterisk (*) represents a statistically significant difference between the pre and post period at the 0.01 level, based on the time series of daily averages.

ETF	Trade Size	Overall		AMEX		Island		NYSE		ARCA-P		Regionals		Instinet-3rd	
		Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
DIA	Small	0.92	1.10	2.85	2.85	0.08	0.51 *	2.94	1.85 *	1.46	0.30 *	2.63	2.38	1.12	0.64 *
	Medium	2.92	2.71	4.10	3.56	0.48	1.84 *	2.92	2.57	1.92	0.77	3.51	3.62	2.70	1.94
	Large	4.74	4.74	3.86	4.97	-0.76	-1.00	4.44	4.29	2.22	2.71	5.82	4.80	6.15	4.85
	Overall	2.32	2.73 *	3.63	4.02	0.15	0.74 *	3.42	3.34	1.52	0.45 *	3.05	3.31	3.94	2.83
QQQ	Small	0.39	0.49 *	1.03	0.95	0.01	0.07	0.97	0.69 *	0.55	0.09 *	1.13	1.08	0.67	0.54
	Medium	0.53	0.56	1.31	1.17	-0.01	0.11 *	1.01	0.91	0.91	0.01 *	1.58	1.65	0.94	0.44 *
	Large	1.70	1.84	1.38	1.47	-0.17	0.01	2.28	1.51 *	-0.08	-0.20	2.44	2.03	1.79	2.21
	Overall	0.99	1.13	1.35	1.34	-0.01	0.11 *	1.85	1.28 *	0.70	0.00 *	1.64	1.58	1.53	1.67
SPY	Small	0.68	1.14 *	2.23	2.07	0.00	0.39 *	1.91	1.16 *	1.04	0.34 *	1.96	2.32	0.64	1.16 *
	Medium	1.70	1.57	2.73	2.46	0.12	0.19	2.61	1.50 *	0.60	0.53	2.70	2.93	1.42	1.68
	Large	3.41	4.57 *	2.64	3.76 *	0.68	0.63	2.35	2.22	-1.31	-3.50	5.42	3.77	4.04	5.71 *
	Overall	2.33	3.25 *	2.66	3.16 *	0.03	0.24	2.30	2.05	0.75	0.40	2.93	2.78	3.45	4.99 *

Table 8

Overall, Island, and Non Island Spread Regressions Controlling for Market-wide Changes in Liquidity

Effective spreads, the absolute difference between the transaction price and the quote midpoint, and realized spreads, the difference between the transaction price and the quote midpoint five minutes after the trade signed by the direction of the trade, are calculated (in basis points) for each ETF each day for all trades, Island, and non-Island trades. Daily averages are calculated on a per share basis. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. The dummy variable for the dark period is *Dark*. Dummy variables for the medium and large trade sizes are *Medium* and *Large*. The dark period dummy interacted with trade size dummies are *Dark*Small*, *Dark*Medium*, and *Dark*Large*. The three dummy variables for the ETFs are *DIA*, *QQQ*, and *SPY*. The daily difference between the highest and lowest transaction price in that ETF is *volatility*. The share-weighted daily average effective spread for all NYSE-listed stocks is *market liquidity*. *t*-statistics are in parentheses.

Variable	All Trades		Island Trades		Non Island Trades			
	Effective Spread	Realized Spread	Effective Spread	Realized Spread	Effective Spread		Realized Spread	
Dark	0.60 (4.93)	0.79 (3.53)	1.12 (10.31)	0.55 (3.06)	-0.46 (-3.31)		-0.30 (-1.07)	
Dark*Small					-0.50 (-3.43)		-0.82 (-2.71)	
Dark*Medium					-0.76 (-5.19)		-1.14 (-3.73)	
Dark*Large					-0.10 (-0.71)		0.39 (1.29)	
Medium					1.15 (8.22)		1.35 (4.63)	
Large					2.77 (19.75)		2.72 (9.32)	
DIA	1.95 (4.55)	1.52 (1.92)	0.86 (2.24)	0.50 (0.78)	3.08 (6.26)	1.88 (5.40)	2.55 (2.61)	1.14 (1.57)
QQQ	4.38 (10.21)	3.07 (3.86)	2.24 (5.80)	0.14 (0.22)	6.06 (12.29)	3.93 (11.28)	4.93 (5.04)	2.49 (3.44)
SPY	2.10 (4.89)	1.65 (2.08)	0.42 (1.08)	0.07 (0.12)	2.19 (4.44)	0.72 (2.07)	1.46 (1.49)	-0.15 (-0.21)
Volatility	1.54 (1.22)	0.33 (0.14)	1.46 (1.28)	0.12 (0.06)	2.01 (1.38)	1.70 (1.70)	1.44 (0.50)	1.38 (0.66)
Market Liquidity	8.62 (3.72)	5.90 (1.38)	3.27 (1.57)	-1.05 (-0.31)	9.29 (3.49)	8.66 (4.73)	8.16 (1.54)	8.83 (2.32)

Table 9
Impulse Response (Permanent Price Impact)

The permanent price impact of a trade (in basis points) is based on the VAR/VMA model (described in Section 8) of prices estimated with 10 lags of the signed net number of Island and non-Island trades over 10-second intervals. Estimates are calculated for each ETF each day and then averaged across days. The pre period is August 16, 2002 to September 20, 2002 and the post period is September 23, 2002 to October 31, 2002. An asterisk (*) represents a statistically significant difference between the pre and post period at the 0.01 level, based on the daily time series of the estimated permanent price impact.

	DIA				QQQ				SPY			
	Non Island		Island		Non Island		Island		Non Island		Island	
	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post	Pre	Post
Median	0.086	0.104	0.192	0.285	0.047	0.116	0.239	0.381	0.046	0.056	0.109	0.144
Mean	0.091	0.103	0.185	0.276 *	0.037	0.107 *	0.242	0.401 *	0.054	0.055	0.113	0.157 *
Std Err	(0.02)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)	(0.01)	(0.01)	(0.01)	(0.02)

**Figure 1
Island's Limit Order Book**

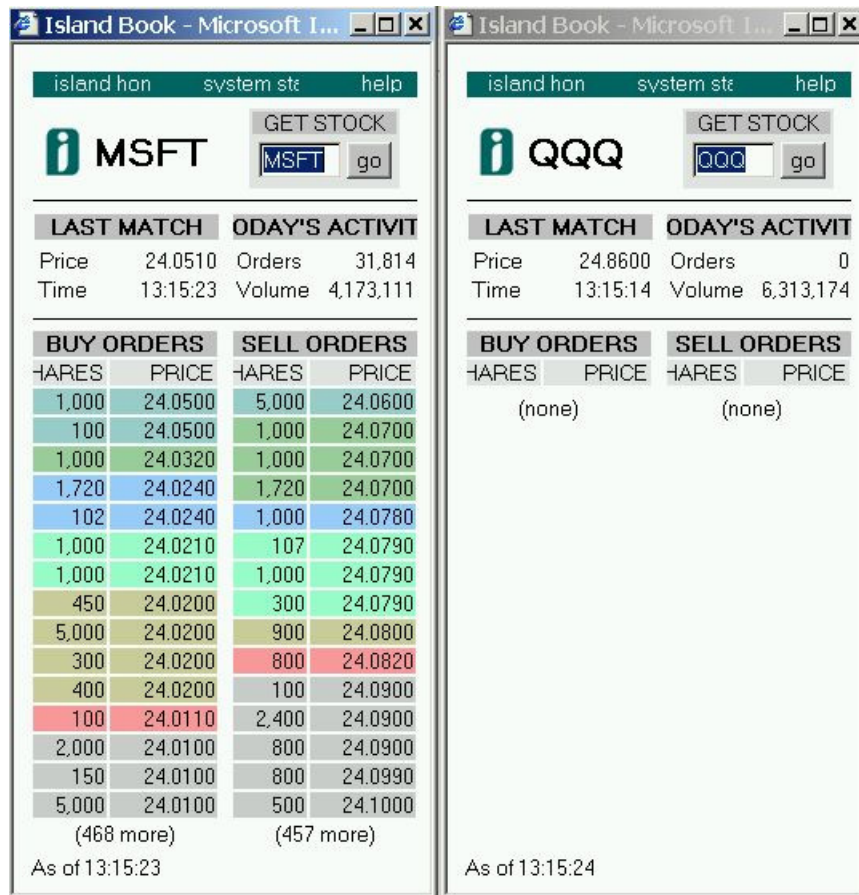



Figure 2
Island's Daily Market Share in Shares and Trades

For the three ETFs this chart graphs Island's daily market share as a fraction of share volume (Panel A) and trades (Panel B) from August 16, 2002 to October 31, 2002.

Panel A: Island Market Share (Shares)



Panel B: Island Market Share (Trades)





Figure 3
Island's Daily Information Share

This chart graphs the bounds on Island's daily information share with Panel A being the maximum and Panel B being the minimum from August 16, 2002 to October 31, 2002.

Panel A: Island's Maximum Information Share



Panel B: Island's Minimum Information Share





Figure 4
Daily Futures and ETF Information Shares

This chart graphs the maximum daily information share for ETF and E-mini transaction prices. ETF prices are Island transactions using ITCH time stamps. The Nasdaq-100 index is in Panel A, and the S&P 500 index is in Panel B. Data extend from August 16, 2002 to October 31, 2002.

Panel A: Nasdaq-100 ETF/E-mini Maximum Information Shares



Panel B: S&P 500 ETF/E-mini Maximum Information Shares





Figure 5
Daily Effective and Realized Half-Spreads on Island

For the three ETFs this charts graphs Island's daily average effective (Panel A) and realized half-spreads (Panel B) from August 16, 2002 to October 31, 2002. Daily averages are calculated on a per share basis.

Panel A: Island's Effective Spread



Panel B: Island's Realized Spread




Figure 6
Daily Effective Half-Spreads

This chart graphs daily average effective half-spreads from August 16, 2002 to October 31, 2002 for one of the most active NYSE-listed stocks (GE) and one of the most active Nasdaq stocks (MSFT), and for the average across the three ETFs. Daily averages are calculated on a per share basis.

