Intraday Price Formation in US Equity Index Markets

Joel Hasbrouck

Department of Finance
Stern School of Business
New York University
44 West 4th St. Suite 9-190
New York, NJ 10012-1126
212.998.0310

jhasbrou@stern.nyu.edu

October, 2000
This draft: November 16, 2000

Preliminary draft
Not for attribution
Comments welcome

Current drafts of this paper and the associated computer programs and data will be posted to my web site: http://www.stern.nyu.edu/~jhasbrou.

For helpful discussions and/or comments I am indebted to Ned Elton, Rob Engle, Richard Hogan of Spear Leeds and Kellog, and seminar participants at the University of Michigan.

All errors are my own responsibility.

Intraday Price Formation in US Equity Index Markets
Abstract / Summary

The market for US equity indexes has traditionally comprised floor-traded index futures contracts and the individual markets for the component stocks. This picture has been altered by the advent of exchange-traded funds (ETFs) that mirror the indexes, electronically-traded, small-denomination ("E-mini") futures contracts, and (for the S&P 500) a family of sector ETFs that break the index into nine components. This paper empirically investigates price discovery (price leadership) in this new environment. The specifications are estimated at very fine (up to one second) time resolution. The principal findings are as follows.

- For the S&P 500 and Nasdaq-100 indexes, the index market comprises an ETF, a regular floor-traded futures contract and an E-mini futures contract. The paper finds that most of the price discovery for both indexes occurs in the E-mini markets: price changes in the E-mini futures prices generally lead those in the regular futures contracts and the ETFs.
- The market in the S&P 400 MidCap index comprises only the ETF and the regular futures contract. Both securities contribute substantially to price discovery, but the ETF appears to dominate.
- The sector ETFs can closely replicate the S&P 500 ETF. Nevertheless, trading activity across the sector ETFs varies considerably. The most actively traded sector (technology) contributes a modest amount to price discovery in the overall index. The other sector ETFs play only a minor role.

1. Introduction

This paper is an empirical study of the short-run price dynamics in three important US equity index markets (S&P 500, S&P MidCap 400, and Nasdaq-100). The statistical analysis attempts to quantify the contribution of each index-related security to the evolution of an implicit "optimal forecast" index price. From an economic perspective, these contributions measure the information production in the various markets, which may in turn be related to security and market characteristics.

The market for US equity indices has traditionally been viewed as comprising the futures contract for the Standard and Poor's 500 index (traded on the floor of the Chicago Mercantile Exchange) and the underlying spot or cash market, consisting of the individual markets for the component stocks. This definition could arguably have been broadened to encompass the markets for futures options and index options (both exchange-traded and over the counter), but in terms of visible trading activity, the spot and futures markets have stood as preeminent. In this context, the paper is motivated by the relatively recent emergence of:

- Exchange-traded index mutual funds, which permit the direct purchase or sale of a portfolio in a single trade at any time during regular exchange trading sessions.
- Electronic trading in index futures contracts, which facilitates market participation by off-floor traders.
- Exchange-traded S&P 500 index sector funds, which facilitate trading in a small number of S&P 500 industry-aligned components (technology, financial, energy, etc.).

Each of these developments has carried the possibility of altering pre-existing patterns of information- and liquidity-motivated trading.

An exchange-traded fund (ETF) is most closely related to a closed-end mutual fund in that it is exchange traded continuously (during regular trading hours) and allows no cash investments or redemptions. Creations and redemptions only occur in kind, and involve swapping fund units and index stock portfolios. The first of these, the S&P 500

fund, began trading in 1993. This is the Standard and Poor's Depositary Receipt (SPDR or "spider"), traded on the American Stock Exchange. Elton et al. (2000) discuss the investment and tracking performance of this fund. Following on the SPDR's success, similarly-structured basket products were introduced for the NASDAQ-100 and S&P MidCap 400 indexes, both of which are also analyzed in this paper.

Each of the three indexes analyzed here underlies a futures contract traded on the floor of the Chicago Mercantile Exchange. The S&P 500 and Nasdaq-100 indexes underlie, in addition, "E-mini" contracts. These are sized to be one-fifth of the value of the regular contracts, making them more accessible to traders of modest capital. More importantly for present purposes, they are traded on an open electronic limit order book system (Globex) that is accessible in real time worldwide. A number of brokers have constructed interface systems to the Globex system. As a result, outside customers can obtain electronic access to the market for purposes of viewing the best bid and offer, recent trade prices and order entry, similar to the screens that an online equity trader might possess.

The third development is the introduction of exchange-traded S&P 500 sector funds. There are nine of these funds. They are constructed to reflect broad industrial sectors, and can closely replicate the S&P 500 index. Like the SPDR, they can be traded intraday and are exchangeable in kind. A natural clientele for these funds would consist of liquidity traders motivated by special diversification needs. Given the existence of liquidity trade, one might expect other traders to find it advantageous to produce and trade on sector-related private information, as suggested by the models of Admati and Pfleiderer (1988), Foster and Viswanathan (1990), and Subrahmanyam (1991).

All of these developments are interesting by virtue of their close connection to the index price formation process. Additional importance arises, however, because each development is in some sense representative or prototypical. The SPDR was the prototype for what is now a large class of exchange-traded funds based on existing indexes. The shift of electronic trading in futures contracts will be seen to mirror in many respects a transition that has occurred in other markets. The sector index products

represent an attempt to complete a market by offering an enriched set of spanning possibilities.

This paper assesses the importance of these developments by a detailed examination of price discovery in the index markets. The methodology attempts to quantify price dominance in related markets. Pre-existing studies, for example, have generally found that the regular (floor-traded) index futures contract is the price leader. To what extent has this role shifted to the electronically-traded futures contract or the exchange-traded fund? Furthermore, in view of the expanded range of trading possibilities in S&P 500 components, to what extent are sector fund prices informative about movements in the broad index?

Briefly, the results suggest that for the S&P 500 and Nasdaq-100 indexes, the electronically-traded futures contracts (the E-minis) are the new price leaders. The corresponding ETFs play a smaller informational role. The pattern is different for the S&P 400 index, however, where no E-mini contract trades. For this index, the ETF provides substantial price discovery. The contribution of the sector funds to S&P 500 pricing, however, is very small (about 10% by the paper's metric). This modest role is surprising because trading activity in some of these funds (typically technology) is substantial, and because media reports often depict market moves as being led or driven by a particular sector (typically technology).

The paper is organized as follows. The next section reviews the literature related to the paper's principal economic themes. Section 3 presents a preliminary look at the data. The following two sections describe the econometric framework of the analysis. Section 4 discusses the time series representation of the data and the derived measures of price discovery. The models used here are vector error correction specifications. The information measures are those proposed in Hasbrouck (1995), and are based on Stock and Watson (1988) common-trend (random walk) components. Section 5 discusses estimation and the modifications and extensions to standard procedures that are employed to capture more closely the economic features of the data. The paper then turns to the estimation results. Section 6 presents estimations for systems consisting of the index

basket, the pit-traded futures contract, and (when available) the electronically-traded futures contract. Section 7 presents estimates for the S&P 500 and sector indexes. A brief summary concludes the paper in Section 8.

2. Related work

The present paper is related to three principal themes in the literature. The first theme loosely falls under the rubric "stock index futures vs. cash"; the second concerns the competition between floor- and electronically-based trading systems; the third deals broadly with the relationship between liquidity trading and information production in individual and basket securities.

a. Stock index futures vs. cash

As noted in the introduction, the cash market for stock indices has traditionally consisted primarily of the markets for the individual component stocks. Most economic analyses of trading and arbitrage motivated by this market structure, therefore, feature multiple stocks and a single basket security. The model of Subrahmanyam (1991) features a one-factor structure of returns and private information. Kumar and Seppi (1994) consider the cross-market information differentials induced by lagged dissemination of information. Both of these factors may be present in the present paper's setting. From a trading perspective, however, the electronically-traded futures contract and the exchange-traded basket security are much closer substitutes for the regular futures contract than the individual stocks. Accordingly, models in which a single security trades in multiple markets are also relevant (Chowdhry and Nanda (1991)).

From an empirical perspective, a substantial academic and practitioner literature explores the dynamics of US stock index and index futures prices with the aim of determining which market is dominant, characterizing and quantifying this dominance, and (in some cases) apportioning blame for market crashes. The methodologies of these studies vary widely, reflecting the diverse characteristics of the data and aims of the

analyses. The conclusions of these studies are mostly similar and are not highly sensitive to methodology. The consensus finding is that the dominant influence runs from the futures market to the cash, and weaker (though still measurable) effects in the reverse direction.[1]

Simple inferences about futures/cash price dynamics may be drawn from the serial and cross correlations of the two return series. Most studies, however, employ some sort of time series modeling. The early papers in this area employ lead-lag regressions in which each return (in separate specifications) is regressed against leads and lags of the other. These specifications support tests against the null hypothesis of "no causality". They can roughly characterize the time horizon of the dependencies. They do not, however, provide an obvious metric for assessing relative importance of the two markets.

Later studies tend to characterize the two return series by joint models, which treat the two securities in a unified, symmetric fashion and clarify the role of innovations. Cash and futures prices are generally viewed as impounding a single common random walk component. This implies that the price difference (the basis) is stationary, which suggests a cointegrated specification, typically a vector error correction model (VECM).

There remain, however, unresolved issues concerning the specification and interpretation of cointegrated models in these applications. Different approaches to model specification arise from differing views of the sense in which the basis is considered "stationary". Alternative interpretations arise because VECMs are inherently reduced-form models, and they usually feature a very large number of parameters. To

---

[1] Representative studies include: Kawaller, Koch, and Koch (1987), MacKinlay and Ramaswamy (1988), Stoll and Whalley (1990), Wahab and Lashgari (1993), Harris, Sofianos, and Shapiro (1994), Choi and Subrahmanyam (1994), Fleming, Ostdiek, and Whalley (1996), Chu, Hsieh, and Tse (1999). The last study also examines the SPDR, concluding that it is dominated by the regular index futures contract. The Chu et al. methodology also differs from that in the present paper.

map these parameters into a manageably small summary set of price discovery measures, the analyst must identify structural features of the model that are economically meaningful. The basis stationarity and identification issues are discussed more fully later in the paper.

b. Electronic vs. floor-based markets

In the S&P 500 and Nasdaq-100 index markets, there exist simultaneous floor and electronic markets for substantially similar contracts. A number of empirical studies analyze the competition that arises when closely substitutable securities trade on electronic and floor-based markets. Surveying studies of comparative market quality, Domowitz and Steil (1999) find that electronic markets tend to offer liquidity similar to that of floor markets, but at lower cost. These cost advantages suggest an eventual displacement of floor markets.[2]

At all times, the CME's E-mini contracts trade only on the Globex system. During normal trading hours, the regular index futures contracts are traded only on the floor of the CME. At all other times, they may be traded on Globex. Coppejans and Domowitz (1999) compare the nighttime Globex market and daytime floor markets for regular contracts. They find that the Globex system performs well during a period (the night) when the flow of liquidity traders is likely to be relatively low. Their comparisons are across disjoint time periods as well as across the two market systems. The present paper takes the perspective that the regular and E-mini contracts are close substitutes, and

---

[2] Of particular noteworthiness is the recent experience of competing markets in German government bond futures contracts. Similar contracts were traded on the London (LIFFE) floor and the German (DTB) electronic book. There was a prolonged period when the two markets coexisted. At least one commentator suggested that this coexistence constituted an equilibrium. In 1997, however, the market dramatically "tipped" in favor of the DTB electronic book, leading London to abandon expansion plans and eventually close its trading floor.

analyzes the joint dynamics during normal trading hours, when the former are floor-traded and the latter are electronically traded. It is emphasized that the dynamics studied here are joint, i.e., conditional on both markets being simultaneously open. The analysis does not support formal inferences about how either mechanism might behave if it were the sole operating market (the focus of Coppejans and Domowitz).

c. Baskets vs. individual securities

This paper's analysis of the joint dynamics of the SPDR and the S&P sector ETFs is connected to several themes in the literature. Most directly relevant are Gorton and Pennacchi (1993) and Subrahmanyam (1991), who present basket/multiple security models with asymmetric information. The equilbria in these models generally exhibit both informed and uninformed trading in both the basket and individual securities. A key feature of these models is that private information is diversified in the basket security. This mitigates adverse selection in the basket market. The resulting increase in basket liquidity induces uninformed agents to trade this security. Subrahmanyam concludes that this diversification effect is strong. It supports trade in the basket even when informed traders possess private information about the common factor, and even in the absence of nondiscretionary liquidity traders who are required to trade the basket.

The present situation differs from these models in an important respect. Specifically, while sector portfolios are baskets of individual securities, they are also themselves components of the broader index. The sector portfolios might therefore be considered to offer an intermediate level of aggregation. It might be conjectured that the sector baskets provide a private-information diversification effect similar to (but smaller than) that of the market basket, and that this effect would support liquidity trading in the sector baskets.

Within such a model, however, discretionary liquidity traders (in the Subrahmanyam model) in the sector baskets would view as substitutes not only the individual component securities, but also the overall market basket. The diversification in the sector baskets is by design, however, incomplete, since they are constructed to load

on a single "sector" factor. A discretionary trader desiring broad diversification, therefore, would naturally prefer the market basket. This suggests that nondiscretionary liquidity traders might be more important for the viability of the sector basket markets than they are for the existence of trade in the overall basket.

Whence might sector-specific nondiscretionary uninformed trading needs arise? In the Gorton and Pennachi model, uninformed trading arises from endowment shocks. Along these lines, there might exist agents who are subject to shocks in their endowments of nonmarketable assets. A person newly employed in the finance sector, for example, might prefer might prefer a reduced exposure to that sector's factor risk in her marketable portfolio. To the extent, however, that this human capital is associated with proximity to the firm or sector (e.g., managerial equity stakes or compensation), such traders are unlikely to be viewed as uninformed.

One source of liquidity trading in the overall market basket arises from portfolio managers who pursue indexed strategies. The S&P 500 index is the most commonly accepted index for this purpose, and is furthermore a widely used benchmark for performance evaluation. The sector indexes, in contrast, do not appear to fulfill target or benchmark roles. Liquidity trading can arise in any market, of course, from noise traders, irrational traders or traders who incorrectly believe that they possess information. The characteristics that induce such agents to concentrate on one security rather than another, however, are not known. There is accordingly no reason to suspect that noise traders might play particularly influential roles in the sector markets.

Security-design considerations suggest another perspective. From the viewpoint of the sponsor, an exchange-traded fund is like any other mutual fund in that the revenue derives from management fees. From the viewpoint of the listing exchange, however, the exchange-traded fund is like any other listing in that revenue derives from trading activity. Both objectives are important in the present case due to the close relationship between the exchange and the sponsor. Essentially, the American Stock Exchange *is* the sponsor of the S&P 500 and S&P 400 trusts. The sponsor of the Nasdaq-100 trust is Nasdaq, but this trust was initiated when the Amex and Nasdaq were merged entities.

The sector basket securities are legally constituted as mutual funds; the Amex has no sponsorship role. The Amex computes, however, the indexes on which these funds are based.

These remarks carry no suggestion of conflict of interest or impropriety. They are instead advanced in support of the view that the economic position of the Amex is similar to that of a futures exchange designing a new contract, wherein trading volume enters strongly into the objective. Duffie and Jackson (1989) model the introduction of futures contracts by a volume-maximizing exchange. The desire for trade in their model is driven by endowment differentials. The optimal contracts are those that facilitate exchange by being maximally correlated with the differential, or by being maximally correlated with the component of the differential that is not spanned by previously introduced contracts. Corkish, Holland, and Vila (1997) provide some empirical evidence consistent with this model.

3. A first look at the data

a. Description of the instruments

Table 1 reports summary statistics for the ETF index products. (Here and henceforth, the SPDR is identified by its ticker symbol *SPY*.) Although the American Stock Exchange is the primary listing venue for these securities, substantial trading activity occurs away from the Amex, on the regional exchanges and over-the-counter. The pattern is similar to that found in NYSE-listed securities in that the listing exchange (Amex) accounts for the preponderance of activity measured by share volume, but much less activity measured by number of trades. Like other listed securities, payment for order flow and other inducements appear to divert smaller (typically retail) trades away from the Amex. Furthermore, trading volume in the S&P 500 sector funds is substantially lower than that in the overall index fund.

Table 2 reports summary statistics for the index futures contracts. Of particular note here is the fact that contract volume in the electronically-traded contracts is roughly

comparable to that in the corresponding regular contracts. The regular contracts are five times as large as the E-minis, however, so in terms of dollar-value of the underlying, the regular contracts clearly dominate.

That noted, it must be emphasized that since the pit and electronic markets are fundamentally different, trading volume figures may not be directly comparable. The pit resembles a dealer market, wherein one outside customer order may lead to multiple trades as the order gets passed to multiple dealers and eventually an outside counterparty. An electronic limit order book, in contrast, favors direct interaction of customer orders without intermediation.

b. Trading costs

Table 3 provides some rough estimates of one-way trading costs in the index instruments. The smallest regular trade is one contract for the futures products and 100 shares for the baskets. In terms of dollar value of underlying the baskets are by far the smallest. A round lot trade in *SPY*, for example, corresponds to roughly $14,000; the corresponding E-mini, $71,000; the regular contract $355,000.

In estimating the trading cost, the half-spread is assumed to be one-half of the minimum tick for the futures contract, and one-half of the estimated average spread for the baskets. In addition, a fixed commission of $10 per trade is assumed. (This is an estimate of the commission charge for online equity and futures brokers.) In absolute terms, the one-way costs are roughly comparable. The differences in relative trading costs are therefore primarily reflect the differences in underlying dollar values. By this measure, the trading cost of the basket product is an order of magnitude larger than that of the futures contracts.

Note, however, that the net trading costs over a given investment horizon are not necessarily twice the estimated one-way cost. The near-maturity futures contract lives for at most three months. An investor wishing to maintain the position must then roll over the contract, incurring further trading costs. The roll-over may also lead to

unavoidable taxation on short term capital gains. Most of the capital gains taxes on the ETF, in contrast, are realized when the position is liquidated.

c. The basis

The basis is generally defined as the difference, $F_t - S_t$, where F_t is the price of the index futures contract and S_t is the price of the underlying index. This definition is adequate for analysis of the futures and ETFs, subject to two qualifications. First, ETFs are generally scaled so that their per share prices are similar to those of other stocks (see Table 1). Thus, S_t is to be interpreted as the ETF price scaled up by the appropriate factor.

The second qualification is that the portfolio underlying the ETF includes not only the stock, but also a cash component. The latter consists of the dividends that have accrued on the portfolio since the ETF's most recent ex dividend date, less trading costs of the fund and management fees. It is therefore convenient to write

$$S_t = S_t^{stock} + C_t \tag{1}$$

where C_t is the cash component, and S_t^{stock} is the implied market value of the stock portfolio.

The principle of forward-spot parity normally asserts the equivalence of the basis and the cost of carry. Applying this to the stock component of the ETF gives:

$$F_t = S_t^{stock} + c_t \tag{2}$$

where c_t is the cost of carrying the stock position until maturity of the futures contract (interest expense less dividends anticipated between *t* and maturity). Equations (1) and (2) imply that forward-spot parity, when expressed in terms of the ETF's market price, yields:

$$F_t - S_t = c_t - C_t . \tag{3}$$

That is, the cash component of the ETF can be viewed as an offset to the cost of carry. For the remainder of this paper, "basis" refers to the left-hand side of (3), and "cost of carry" refers to the right-hand side.

Model specification and interpretation in this paper depend crucially on whether (and in what sense) the basis is stationary. Assuming that arbitrage drives the basis toward the cost of carry, then, stationarity of the basis is equivalent to stationarity of the cost of carry.

Starting at its highest value when the contract is introduced, the cost of carry normally declines, reaching zero at maturity due to forward-spot convergence. For any given contract, therefore, the time path of the basis approximates a declining trend, which is nonstationary. The approximation arises because quite apart from pure time-to-maturity effects, the expected cost of carry might also change due to new information regarding prospective changes in the index composition, expected dividends over the remaining life of the contract, and expected interest expense.

Figure 1 depicts the daily averages of the basis (constructed at one-minute intervals) for the S&P 500 (price of the regular futures contract less the price of *SPY*). The time path of these daily averages is dominated by the declining trend. There are two breaks in this path. The first (March 9) reflects roll-over from the March contract to the June contract. The second break (March 17) occurs when the SDPR goes ex dividend.[3] The figure also demarcates (with vertical bars) two-standard-error bounds on the daily means. (For the S&P 500, these bars are generally so small as to be obscured by the plotting symbol.)

Figure 2 graphs the basis for the Nasdaq-100 index. The pattern resembles that of the S&P 500. It is somewhat noisier, however, evincing greater variation around the

---

[3] For the first six days of the sample, prior to the roll-over and the ex-dividend day, the basis is negative, i.e., the spot price lies above the futures price. This is a consequence of accrued dividends in the cash component of the ETF.

trend. Figure 3 depicts the basis for the S&P 400 index, which is even more irregular.[4] In all cases, however, the distinguishing feature is the downward trend predicted by forward-spot convergence.

The behavior of the basis is somewhat different, however, at a finer level of time resolution. Figure 4 depicts the bases for the three indexes at one-minute intervals on a representative day (the first day of the sample). Each series appears to be mean reverting, with no obvious trend. The figures are typical of most days in the sample. Unit root tests applied to each day's basis path uniformly and overwhelmingly rejected the null (unit root) hypothesis. The basis paths in Figure 4 furthermore exhibit both substantial variation and rapid reversion to the mean.

Intraday stationarity of the basis is consistent with the fact that the main components of the cost of carry (dividends and interest) do not accrue intraday. That is, interest is paid only on overnight positions, and dividends are paid only to the holder of record at the end of the day prior to the ex-dividend day. As noted above, trading costs on the ETF are also impounded into the cash component (and the cost of carry). The cash component is only reported once per day, however, shortly before the market opening.

In principle, of course, forward-spot parity reflects the expected cost of carry over the life of the contract. These expectations could change within the day, even though the accruals were overnight. Such changes might reflect shifts in interest rates or dividend

---

[4] The relative volatility in the S&P 400 basis may be due in part to the greater turnover of its components. During the sample period (March 1, 2000 through May 31, 2000), the S&P 400 had 19 replacements; the S&P 500, 8. Since the liquidity of the component stocks is relatively low, frequent index changes would lead to high (and variable) trading costs.

expectations. The results of the unit root tests, however, suggest that these changes do not constitute a major component.[5]

4. Methodology: Representation

Most of this paper's claims and attributions of contributions to price discovery are based on an innovations representation for price changes. Implicit in such a representation are forecasts for each of the security prices. The changes to a given price forecast are linear in the innovations, i.e., forecast updates are linear in the newly-arrived information. To the extent that an innovation can be identified with a particular security or market, the contribution of this innovation to the forecast price update measures price discovery attributable to that security or market. The approach is based on random-walk decompositions of integrated time series first developed and applied to macroeconomic modeling (Beveridge and Nelson (1981), Watson (1986), Stock and Watson (1988)), and later applied to microstructure data (Hasbrouck (1991), Hasbrouck (1995)). This section discusses model specification and representation. Estimation is the subject of the next section.

a. The innovations representation

Let p_t denote a column vector of n security prices. At the present level of generality, p_t might comprise bids, asks and/or trade prices for one or more securities. An innovations representation for the prices, also called the vector moving average (VMA), may be written:

$$\Delta p_t = u_t + A_1 u_{t-1} + A_2 u_{t-2} + \ldots \qquad (4)$$

---

[5] One likely source of changes in expected dividends is changes in index composition. Standard and Poor generally announces revisions, however, after normal trading hours, and with a one-week lead time.

where u_t is the vector of zero-mean innovations with $Var(u_t) = \Omega$ and $Eu_t u_s' = 0$ for $t \neq s$. The A_i are $n \times n$ coefficient matrices.

This representation implies that the total (cumulative) impact of innovation u_t on prices through k periods ahead is the cumulative impulse response function given by:

$$D_t(k) \equiv E\left[\Delta p_t + \Delta p_{t+1} + \Delta p_{t+2} + \cdots + \Delta p_{t+k} \middle| u_t\right] = \left(\sum_{i=0}^{k} A_i\right) u_t \tag{5}$$

where A_0 is understood to be the identify matrix. Plots of these functions are useful in summarizing price dynamics.

Denote the limiting cumulative price impact by

$$w_t = \lim_{k\to\infty} D_t(k) = \left(\sum_{i=0}^{\infty} A_i\right) u_t = A u_t \text{ where } A = \sum_{i=0}^{\infty} A_i \tag{6}$$

This quantity possesses a useful interpretation. Suppose that traders are primarily interested in the model's predictions about the conditional expectation of prices in the infinitely distant future, $\lim_{k\to\infty} E\left[p_{t+k} \middle| p_t, p_{t-1}, \ldots\right]$. This is analogous to the conditional expected end-of-trading security value that underlies most theoretical microstructure models. When new information arrives at time t in the form of the innovation u_t, the model identifies the cumulative impact on p_t as Au_t. This quantity, then, must be the amount by which the conditional expectation is revised. Since the u_t are serially uncorrelated, so too are the w_t.

b. Analysis of $Var(w_t)$

The variance in the revision of conditional expectations can be decomposed into constituents attributable to the individual components of the innovation vector. To illustrate this, consider the bivariate case $p_t = \begin{bmatrix} p_{1t} & p_{2t} \end{bmatrix}'$. Formally,

$$Var(w_t) = Var\left(\begin{bmatrix} w_{1t} \\ w_{2t} \end{bmatrix}\right) = A\Omega A' = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{bmatrix} \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix}', \tag{7}$$

where A and $\Omega = Var(u_t)$ are presented in full form. The elements of $Var(w_t)$ are the variances and covariances of the revisions for the two prices. Now consider the variance of the revision for the first security:

$$Var\left(w_{1t}\right) = \begin{bmatrix} a_{11} & a_{12} \end{bmatrix} \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{12} & \sigma_2^2 \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \end{bmatrix}. \tag{8}$$

If the innovation covariance matrix is diagonal $\left(\sigma_{12} = 0\right)$, then equation (8) implies a clean decomposition of the long-run variance into components explained by or attributable to each of innovations:

$$Var\left(w_{1t}\right) = a_{11}^2 \sigma_1^2 + a_{12}^2 \sigma_2^2 \tag{9}$$

The relative size of these contributions indicates the importance of the series. As in Hasbrouck (1995), the information share of security or market *j* in $Var\left(w_{it}\right)$ is defined as

$$I_{ij} = \frac{\sigma_{ij}}{Var\left(w_{it}\right)} \tag{10}$$

If the innovation covariance matrix is not diagonal, the information share is not exactly identified. In this case, one can examine alternative factor rotations for the innovations that either minimize or maximize the contribution of an innovation. This permits computation of bounds for the information share. This approach is used in De Jong, Mahieu, and Schotman (1998b), Huang (2000), Martens and Kofman (1998), and Martens (1998).

c. Cointegration

Certain characteristics of the joint price dynamics have important implications for model structure. For example, in the bivariate case the long-term dynamics of p_{1t} and p_{2t} are governed by their random walk components. Two random walks will generally diverge over time, even if their increments are correlated. It may happen, however, that the economic structure and logic dictates that two random walks are in fact identical. This would be expected to occur if the two prices represented the bid and ask on the same security, or a stock index and an index futures contract (over a sample where the cost of carry is stationary). If there exists a linear combination of the two prices that is stationary, then the two prices are said to be cointegrated (Engle and Granger (1987)).

More generally it is assumed that all components of p_t are integrated, but that $m<n$ linear combinations of p_t given by $\alpha' p_t$ (where α is $n \times m$) are stationary.

In the bivariate case, where both prices refer to the "same" underlying security (e.g., bid and ask or index and index futures prices), economic logic suggests that the combination be formed as the difference between the two prices, i.e., $\alpha' = \begin{bmatrix} 1 & -1 \end{bmatrix}$. In the general case, when all *n* prices refer to the same security, $\alpha' = \begin{bmatrix} \iota_{n-1} & -I_{n-1} \end{bmatrix}$, where ι_{n-1} is a column unit vector of size *n*-1 and I_{n-1} is the identity matrix of order *n*-1. Microstructure applications of cointegration have exclusively (to my knowledge) involved situations of this "same-security" type. The present analyses that involve futures contracts and their corresponding exchange-traded funds are of this form.

Same-security cointegration implies a regularity in the model. In equation (6), it implies that all rows of *A* are identical, and therefore that all components of w_t are identical. It further implies that information share calculations are identical for any price.

The cointegration specification used in same-security models is not appropriate in the joint specification of the S&P 500 ETF and its nine sector portfolios. Here, cointegration arises from replication. Denote the price vector by $\begin{bmatrix} p_t^{SPDR} \\ p_t^{Sector} \end{bmatrix}$ where p_t^{Sector} is the (9×1) vector of sector portfolio prices. If "on average" $p_t^{SPDR} = a\, p_t^{Sector}$, where *a* is a vector of replication weights, then the appropriate stationary linear combination is $p_t^{SPDR} - a\, p_t^{Sector}$, i.e., $\alpha' = \begin{bmatrix} 1 & -a \end{bmatrix}$.

d. Representations useful for estimation

The innovations representation given in (4) is useful for interpretation. An alternative form that is more amenable to estimation is the vector error correction model (VECM) of order *M*:

$$\Delta p_t = B_1 \Delta p_{t-1} + B_2 \Delta p_{t-2} + B_3 \Delta p_{t-3} + \cdots + B_M \Delta p_{t-M} + \gamma \left(\alpha' p_{t-1} - E\alpha' p_t \right) + u_t \qquad (11)$$

where the term in parentheses is the deviation ("error") between $\alpha' p_{t-1}$ and its long-run value $E\alpha' p_t$. There is no intercept term in the specification, which essentially forces the

drift to zero. The γ matrix contains the error correction coefficients. There are a number of procedures for obtaining (4) from (11). Perhaps the easiest and most intuitive approach involves stepping the system forward in response to a unit innovation (Hamilton (1994), pp. 318-323).

e. Alternative interpretations of the VECM

The information share is but one of the quantities derived from the VECM that have been proposed to measure price discovery. One group of studies focuses on the error correction coefficients, the γ_i in equation (11). These can be interpreted as measuring as the rates of adjustment or convergence in the price differences occurring in one step (e.g. from t to $t+1$). The γ_i cannot by themselves, however, be reworked to give meaningful multi-step adjustment rates, because the latter depend on the full set of VECM coefficients. Studies that feature the error-correction coefficients include He (1997), Kofman and Moser (1977), and Pizzi, Economopoulos, and O'Neill (1998).

The "target" price series underlying the information share measures is an implicit optimal forecast of the future price. It has also been suggested, however, that the target price series be constructed as a weighted average of current market prices, and that price discovery be measured by the size of these weights. (This weighted average is identified by the additional restriction that differences between this average and each of the component prices are orthogonal to changes in the average.) Studies adopting this perspective include Ding and et al (1999).

5. Estimation

The last section discussed models and information shares, the functions of model parameters that are used here to assess price discovery. The present section discusses estimation. Estimation is mostly standard and straightforward, employing ordinary least squares applied to specification (11). There nevertheless arise in the present application certain noteworthy details and issues.

a. Cointegration in futures/spot specifications

The analysis of the time series properties of the index bases in section 3.c concluded that the bases are trending (and therefore nonstationary) between days, but stationary within days. In terms of specification (11), this implies that $E\alpha' p_t$ is constant only intraday. This in turn suggests estimating specification (11) separately for each day in the sample, using the daily average of $\alpha' p_t$ as an estimate of $E\alpha' p_t$. This paper follows this procedure. Summary statistics for the whole sample are formed from the full set of daily estimates.

This practice represents a departure from the procedure used in other studies, however, and so warrants further discussion. In most studies, the sample consists of intraday spot ("S_t") and near-maturity futures prices ("F_t"), usually spanning several expiration cycles. The cointegrating vector is taken to be $F_t - S_t$. This is included in a VECM, together with a constant term, and the specification is estimated over the full sample (see Pizzi, Economopoulos, and O'Neill (1998), Chu, Hsieh, and Tse (1999), Booth, So, and Tse (1999)).

The price dynamics implied by this specification, however, force $F_t - S_t$ to converge to its overall sample mean. There is certainly a sense in which this occurs. In a sample that consists of multiple near-maturity rollovers, a plot of $F_t - S_t$ will resemble a saw tooth, centered at the sample mean, jumping at the rollover points, and trending downwards between rollover points (cf. Figures 1, 2 or 3).[6] The adjustment of $F_t - S_t$ towards its sample mean, however, is driven essentially by the time trend. It does not

[6] These remarks apply to specifications of futures prices and contemporaneous spot prices that motivated by forward-spot parity. Cointegration also arises in models of futures prices and spot prices as of contract maturity. In these studies, cointegration arises from some form of the expectations hypothesis. See Chow (1998) and the references therein.

capture or reflect the adjustment due to arbitrage. Arbitrage drives the basis toward its current fair value, not the fair value averaged over the life of the contract.[7,8]

b. Statistical properties of model estimates

Assuming covariance stationary price changes, model residuals are homoskedastic and serially uncorrelated. This ensures the asymptotic consistency of OLS coefficient and residual covariance estimates. The distribution of these estimates, however, is not easily characterized. The asymptotic sample distributions of these parameters are known only when the model disturbances are i.i.d. normal. Given price discreteness and the irregular timing of price updates, normality is unlikely. Furthermore, the statistics of principal interest here are impulse response functions and information shares, which are nonlinear functions of the model parameters.

As noted above, nonstationarity of the futures-spot bases across days motivates estimation of the model in daily subsamples. From a statistical viewpoint, this practice conveys an additional benefit in that the properties of the daily estimates (e.g., mean, standard deviation, etc.) may be easily computed from the full sample of daily estimates. For example, let $\hat{I}_{ij}^{d}$ denote the information share of *j* to *i*, i.e., the contribution of price *j*'s innovations to the long run variance of price *i*, implied by the model estimated on day *d*'s

---

[7] All of the cited papers report statistical tests for $F_t - S_t$ that invariably reject the unit root null hypothesis. It is not difficult to see why these tests fail to indicate that anything is amiss. When a series possesses a unit root, it tends to diverge without bound as the sample size increases. The divergences in a repeating saw-tooth pattern, however, are only local. Nor are tests for global trend stationarity likely to be more revealing because the basis more closely resembles a step function. Trend stationarity is only likely to be detected in daily-averaged data spanning a single expiration cycle.

[8] In connection with Booth, So, and Tse (1999), Yiuman Tse writes (personal communication), "We check that if we use residuals from OLS regressions (with trend included), the results are qualitatively the same as reported in the paper."

sample. The overall sample estimate of I_{ij} is the simply the mean $\hat{I}_{ij} = \left(\sum_d \hat{I}_{ij}^d\right) \Big/ N^d$ where N^d is the number of days in the sample $\left(N^d = 64\right)$. Assuming independence across days, the standard error of this mean may be computed in the usual fashion. Estimates of impulse response functions are formed in a like manner. The general approach here, construction of an overall sample estimate as the mean of subsample estimates, was advocated in time series analysis by Bartlett (1950).

c. Choice of price variables

The securities traded in this paper are traded in electronic, futures pit and equity exchange markets. Different trading protocols give rise to different data, both in real time and for purposes of historical analysis. These differences force the analysis to accommodate various kinds of prices.

The data are most comprehensive for the exchange-traded basket securities, comprising last sale prices, and bid and offer quotes. While the econometric specification should incorporate all of these prices in principle, computational considerations dictate otherwise. For these securities, I rely on the quote midpoint. Quotes are used in preference to trade prices because they are more timely. Furthermore, the quotes should reflect the information contained in the trades. The E-mini contracts are traded on the electronic Globex system. Anyone can view a real-time data feed that contains the bid, ask, sizes at the bid and ask, and recent trade prices. The available historical data, however, are somewhat more limited, containing only trade prices.

d. "Nonsynchronous" prices

The specifications are estimated for time series in which *t* indexes intervals of wall-clock time that are (in different estimations) one-, five- or ten-seconds wide. When p_t is composed of prices described in the preceding section, therefore, it will often include prices that were set prior to the start of interval *t*. For example, p_t for *t*=10:05:01 might include a quote midpoint based on a bid and ask that were first set at 10:01 (and are still prevailing as of 10:05:01), the price of an E-mini trade that occurred at 10:02 and a tick

price on the futures contract that occurred at 10:03. Thus, p_t reflects prices that were not set contemporaneously and are in a sense "stale" as of t=10:05:01.

This practice is common in empirical studies where t indexes intervals of a day or longer. CRSP daily closing prices are widely used, for example, irrespective when during the day the last trade actually occurred. The practice is less orthodox in microstructure studies, however, and so some discussion is in order.

Use of non-contemporaneous last-sale or last-tick prices can be justified on both informal and formal grounds. The informal argument rests on the point that even though a last-tick price time-stamped somewhat prior to t is in a sense stale, it still represents the most recent datum available to market participants. A VECM may accordingly be motivated as an ad hoc forecasting device that represents the best linear prediction based on current data.

More formally, the statistical justification for the existence of a VECM representation usually rests on the premise that the data are covariance stationary.[9] A last-sale price will satisfy this requirement under some fairly reasonable and straightforward assumptions.

Specifically, suppose that there is a "true" unobserved price process, $\{p_t\}$ and an associated observation process $\{q_t\}$, where $q_t \in \{0,1\}$, with "1" corresponding to the event that p_t is observed ("a trade occurs"). The most recently observed ("last sale") price is given recursively by $p_t^* = q_t\, p_t + (1-q_t)\, p_{t-1}^*$. Appendix A shows that sufficient conditions for the covariance stationarity of p_t^* are that $\{q_t\}$ is stationary and that $\{p_t\}$ is conditionally covariance stationary. The first condition admits the possibility that q_t is serially correlated, e.g., that trades occur in "waves". The second condition admits the

---

[9] By the Wold theorem, a non-deterministic covariance-stationary stochastic process possesses a VMA representation corresponding to (4) (see, e.g., Hamilton (1994), p. 108). To obtain the VECM representation in (11), an invertibility assumption is also necessary.

possibility that trades are more frequent when price changes are large. The principal substantive restriction is that covariance terms not depend explicitly on *t*.

Other studies have employed a number of approaches to the problem of nonsynchronous prices. One approach is to simply use long intervals of time aggregation. For example, if it can be assumed that a security trades at least once per minute, then the last-sale price in any sixty-minute interval is likely to be "close" to the end of the interval. The drawback, of course, is that time aggregation over long intervals obscures the short-term dynamics that are of primary interest.

Other approaches may be broadly classified as sample-thinning or statistical adjustment procedures. Sample-thinning aims at approximate synchronization by deleting data points (times) when a trade did not occur in any of the markets under consideration, where non-occurrence is operationally defined relative to some window (e.g., one minute). Thinning is applied in Harris et al. (1995), Chu, Hsieh, and Tse (1999) and Booth, So, and Tse (1999).

Like any censoring, however, thinning reduces the information in the sample. Furthermore, the effects are not uniform across markets. A market that posts frequent, promptly updated prices is forced to mimic the same revision pattern as the slowest market in the sample. From an economic viewpoint, prices that are reported during times when any other market is not reporting are viewed as having no informational value (Hasbrouck (2000)).

Statistical adjustment procedures adopt a perspective similar to the present one. The observed price series is viewed as arising from a latent price process, which is filtered through a stochastic observation mechanism to generate observed prices or returns. Approaches along this line include Lo and MacKinlay (1988), Stoll and Whalley (1990), de Jong and Nijman (1997), de Jong, Mahieu, and Schotman (1998a), and de Jong and Donders (1998). These papers have as their primary goal the characterization of the latent price process. This end generally requires, however, making the strong assumption that observations (trade occurrences) are independent of the latent price

process. This rules out, among other things, dependency between the magnitude of a price change and the likelihood of a trade.

The specifications in the present paper only attempt to use observed prices in a linearly optimal forecasting rule. No characterization of the latent price is claimed or attempted. The assumptions, therefore, are weaker.[10]

e. Reducing the size of the parameter set using polynomial distributed lags

The VECM specification (11) contains coefficient matrices B_i for each lag in the model. If the interval width is small and a modest span of wall clock time is required, the number of coefficients is extremely large. For example, if there are three prices, t indexes intervals of one second, and lagged terms up to five minutes are included, the number of coefficients in the model is roughly $3 \times 3 \times 300 = 2,700$. This is computationally unacceptable, so some pruning is necessary. The expedient used here (and in Hasbrouck (1995)) relies on polynomial distributed lags. This procedure is conventionally used when the analyst possesses a prior that the coefficients should lie on a smooth curve (Greene (1993)). The details are given below.

f. Lag lengths

In the futures/ETF analyses, maximal lag lengths were determined by analyses of the bases. Examination of autoregressive residuals suggested that five minute lags were

---

[10] The analysis of Stoll and Whalley (1990) reduces to a linear time series model which they approximate by a finite order autoregressive-moving average (ARMA) model. Other authors have followed this practice, using an ARMA model to prefilter the data (Grunbichler, Longstaff, and Schwartz (1994), Fleming, Ostdiek, and Whalley (1996), Pizzi, Economopoulos, and O'Neill (1998)). This filter, however, can easily be impounded within the VAR or VECM specifications, obviating the need for separate estimation.

adequate for the S&P 500 and Nasdaq 100 models, and that a ten minute lag sufficed for the S&P 400 analysis.

g. Summary

In alternative analyses, the time subscript *t* indexes intervals of one, five or ten seconds. The general form of the estimated VECM is:

$$\Delta p_t = B_1 \Delta p_{t-1} + B_2 \Delta p_{t-2} + \cdots + B_M \Delta p_{t-M} + \gamma \left(\alpha p_{t-1} - \overline{\alpha p_t}^{day} \right) + u_t \qquad (12)$$

where p_t is the price vector. The cointegrating vectors (more properly, a basis for the cointegrating vectors) is given in α.

The B_i in (12) are coefficient matrices. Each set of coefficients is constrained to lie on a function that is piecewise quadratic. Specifically, the coefficients of $\Delta p_{j,t}$ in the $\Delta p_{i,t}$ equation may be written as the $M \times 1$ column vector $b_{ij} = \left[B_1^{ij} \quad B_2^{ij} \quad \cdots \quad B_1^{ij} \right]'$. Let C be the $M \times m$ basis matrix for the piecewise quadratic function. Then $\underset{M\times 1}{b_{ij}} = \underset{M\times m}{C} \; \underset{m\times 1}{\beta_{ij}}$.

Parsimony is achieved by letting $m << M$. The same *C* matrix is used for all coefficient sets.[11]

---

[11] For example, if the coefficients were constrained to lie on only one quadratic segment,

$$\underset{M\times 1}{b_{ij}} = \underset{M\times 3}{C} \; \underset{3\times 1}{\beta_{ij}} = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 2 & 4 \\ & \vdots & \\ 1 & M & M^2 \end{bmatrix} \underset{3\times 1}{\beta_{ij}} .$$

In the specifications actually estimated the segments for the one-second analyses are defined by the breakpoints 5, 10, 30, 60, 90, 120, 150, 180, and 300 (i.e., lags up to five minutes); the five-second analyses used breakpoints 6, 12, 30, 60, 90 and 120 (ten minutes).

6. Results: The index ETFs and their futures contracts

This section discusses the estimated dynamics of the S&P 500, Nasdaq-100 and S&P 400 ETFs and their associated futures contracts.

a. Price discovery in the S&P 500 index

The price set modeled for the S&P 500 consists of the quote midpoint prevailing at the end of the interval for the ETF (symbol *SPY*), the most recently reported tick on the regular futures contract (*SP*), and the most recently reported trade on the E-mini contract (*ES*). Specification (12) was estimated at a one-second level of resolution, with lags through five minutes.

The estimated VECM parameters are not easily interpretable, and for the sake of brevity are not reported. Instead, model properties are more easily seen from the implied dynamics. Figure 5 depicts the cumulative impulse response functions, i.e., cumulative price impacts implied by initial perturbations (cf. equation (5)). Following the paper's general practice, each point is an average of daily estimates.

The cumulative impulse response functions for the S&P 500 are in the first row (panels A, B and C). Each column corresponds to an initial shock in a different security (*SPY*, *SP* and *ES*, left to right). In a given graph, the three lines track the implied cumulative price changes for these securities.

Panel A tracks the cumulative impact of a unit shock in *SPY* (the ETF). By construction at t=0, the impact is unity in the *SPY*, and zero in the other two securities. In the long run (fifteen minutes), the impact is essentially identical for all securities. This long run convergence is a consequence of cointegration. The transient behavior suggests rapid reversion in the *SPY* price, and slow adjustment in the other securities. The impacts of a shock to the regular futures contract (*SP*) are similar to those of the *SPY* innovation (panel B). The effect of a unit innovation in the E-mini contract price (*ES*), however, is stronger (panel C). The initial mean reversion in the *ES* equation is relatively small. The impacts on the other securities are large.

Table 4 presents estimates relating to the information shares. Panel A describes the correlations of the disturbances. At one-second resolution, the off-diagonal correlations are small. Panel B presents the long-run impact coefficient matrix, corresponding to any of the identical rows of *A* in equation (6)). These are identical to the amplitudes of the long-run cumulative price impacts shown in Figure 5.

Panel C in Table 4 shows the properties of the information shares. Medians and means (across days) are reported, along with the standard errors of the means, which are relatively small. Due to the presence of non-zero off-diagonal correlations in the innovations, the information shares can be estimated only up to rotations. Because the correlations are small, however, the minima and maxima are close together. The E-mini contract possesses by far the dominant information share, accounting for roughly 90% of the price discovery. The other two securities share the remainder, with the regular futures contract (*SP*) accounting for slightly more than the ETF.

b. Price discovery in the Nasdaq 100 index

An identical specification was estimated for the Nasdaq-100 securities: the basket (symbol *QQQ*), the regular contract (*ND*), and the E-mini contract (*NQ*). The impulse response functions (graphed in the middle row of Figure 5, panels D, E and F) are similar to those of the S&P 500 set. Innovations in the ETF (*QQQ*) and the regular futures contract (*ND*) are swiftly and substantially reversed, while an innovation in the E-mini contract (*NQ*) substantially persists.

Table 5 reports statistics related to the information shares for the Nasdaq-100 index. As with the S&P 500, the E-mini contract dominates, providing over 80% of the price discovery. In contrast to the S&P 500, however, the contributions of the ETF (*QQQ*) and regular futures contract (*ND*) are roughly equal.

c. Price discovery in the S&P 400 index

The specification for the S&P 400 index securities includes only the ETF (symbol *MDY*) and the regular futures contract (*MD*). It is estimated with five-second resolution

and encompasses lags through ten minutes. The increase in lag length (relative to the one-second estimates) is motivated by the somewhat slower decay in the basis. The decrease in resolution (five seconds is coarse relative to one second) is necessitated by computational limitations.

The results for the S&P 400 index differ markedly from those of the other two indexes. The bottom row of Figure 5 (panels G and H) depicts the cumulative impulse response functions. These show that unit shocks in both the ETF (symbol *MDY*) and the regular futures contract (*MD*) have substantial long-run impacts. The impact of the ETF shock (panel G) is somewhat higher, however.

Table 6 reports statistics related to the information shares. Both securities contribute substantially to the price discovery, but the ETF dominates with an information share of slightly over 60%. These results stand in marked contrast to those of the other two indices, for which the contribution of the ETF was relatively small.

d. Discussion

The results of this section can be summarized as follows. In the two index markets where E-mini contracts exist, these contracts now provide the bulk of the price discovery. This is consistent with the view of many observers (such as Domowitz and Steil (1999) that the lower costs offered by electronic systems portend their eventual displacement of floor markets. The analysis in Section 3.b, however, suggests that the obvious costs (spread and commission) do not work in favor of the electronic market. Transaction cost estimation is subject to error, of course. Consideration of other price components, such as price impacts and the human cost of order entry, might well favor the electronic market.

It should also be recalled, however, that relative to the E-minis the regular futures contracts have larger dollar trading volume and open interest. This suggests that many market participants still find the regular contract to be the cheapest vehicle. In view of this, the informational dominance of the E-minis is even more striking.

Not only do the E-minis eclipse the regular futures contracts, but they also offer more price discovery than the ETFs. It therefore appears that although the ETF is a superior index trading vehicle relative to the individual stocks, futures contracts are still preferred. This conclusion does not imply that futures contracts are superior investment vehicles. Elton et al. (2000) document the substantial investment advantages of the ETF.

The S&P 400 midcap index market differs substantially from other two index markets. There is no E-mini contract here, so price discovery is left to the ETF and the regular futures contract. In this market, the ETF appears to provide most of the price discovery. The reasons for this are unclear. Since the component firms are smaller and have liquidity that is likely to be markedly inferior to S&P 500 firms, the ETF is likely to be much cheaper to trade than the individual stocks. It is not immediately obvious why this would also confer an advantage relative to the futures contract. One conjecture is that the ETF's advantage relative to individual stocks serves to draw more liquidity trading, which supports more informational trading.

Finally, in addition to the usual caveats for an empirical analysis (sample size, model specification, etc.), it is worthwhile to note some overarching limitations of the analysis. First, the attributions of price discovery are contingent on a sample in which all markets are operating jointly. The observed dominance of the E-mini markets, for example, does not imply that price discovery would not be materially affected if the other two markets (ETF and regular futures contract) did not exist.

Second, the present attributions of price discovery are based on subsets of information produced by the markets. The Globex system on which the E-minis trade, for example, produces real-time bids, asks and depths. This is a somewhat richer set of information than the last-trade prices used here. Presumably, incorporation of these data would only increase the E-minis' advantage. Similarly, the ETF markets report recent trades, volumes and depths. Finally, the floor markets (the futures pits, and to an extent the Amex floor) also produce less quantifiable information arising from floor interactions that precede reported trades and quote revisions. The Amex permits cellular phones on the floor, which would presumably aid in the dissemination of this information. Cellular

phones are not presently allowed, however, on the floor of the Chicago Mercantile Exchange.

e. Robustness of the results to timing errors

Given the fine time resolution used in the present specifications, accuracy of the time stamps is important. Synchronization in the present analysis aims in principle at replicating the real-time public market data stream. The data are not, however, collected in real time. Instead, synchronization relies on exchange-reported time stamps. The S&P 500 and Nasdaq-100 analyses comprise three securities, with market activity reported on three different computer systems. Although market personnel generally attempt to set accurate system clocks, it must be acknowledged that any discrepancies would not be obvious. Incorrect synchronization might also arise from differences in reporting system latencies.

To investigate the sensitivity of the results to timing, time-stamp errors were artificially introduced in the following manner. The salient feature of the S&P 500 and Nasdaq-100 analyses is the informational dominance of the E-mini futures contracts. It therefore makes sense to assess how this result is weakened if the E-mini price reports are "delayed" by some fixed time. To investigate this, specifications analogous to those discussed above were estimated, but at a five-second (rather than one-second) resolution. For each index market, four estimations were conducted in which E-mini price reports were delayed by 0, 5, 10 and 15 seconds.

Table 7 reports the results of this analysis. For brevity, the table contains only the means (across days) of the information share minima and maxima. These bounds are somewhat wider than those reported in the one-second analyses (Table 4 and Table 5). This reflects an increase in the magnitudes of off-diagonal innovation correlations, which implies greater sensitivity to rotation.

Intuitively, a reporting delay should penalize the informativeness of a market. Stale prices are less valuable. The results in Table 7 are consistent with this conjecture. The E-mini information shares drop as the delay increases.

For the S&P 500 index market, a five-second delay weakens, but does not eliminate the E-mini's dominance. The E-mini information share is substantial even with a ten-second delay. This share only becomes overshadowed by the others under a fifteen-second delay. Examination of information share shifts in the other two securities suggests that the bulk of the price discovery lost from the E-mini is picked up by the regular futures contract (symbol *SP*) rather than the ETF (*SPY*).

The E-mini dominance result for the Nasdaq 100 index is even less sensitive to induced reporting delays. The information share for the E-mini contract (*NQ*) predominates even with a ten-second delay (and remains substantial with a fifteen-second delay). As with the S&P 500 index, most of the lost price discovery shifts to the regular futures contract.

In summary, while it must be acknowledged that large report timing errors would affect the results, the importance of the E-mini contracts does not seem highly sensitive to minor (five or ten seconds) errors.

7. Results: The S&P 500 and sector ETFs

The paper turns now to price discovery among the S&P 500 index and its sector indexes. The econometric specification in this section is a time series model of the full set of price variables, i.e., the S&P 500 ETF and the exchange-traded sector portfolios. The previous section established that the dominant security in the S&P 500 index market was the E-mini futures contract. For purposes of the present analysis, however, the ETF was used as the summary price index measure due to its strong structural similarity to the sector funds. The sector funds were designed to replicate the S&P 500 ETF, and so have the same dividend payment schedule. Furthermore, the S&P 500 ETF is traded in the same venue as the sector funds (the Amex). Clock synchronization across different exchanges is therefore less of a concern.

Can a portfolio of sector ETFs replicate the S&P 500 ETF? This is important from the investor's viewpoint because any cumulative discrepancies will affect her net wealth. It is important from a short-term trader's perspective because any discrepancies

in replication will give rise to arbitrage opportunities, and risk when the arbitrage is unwound. Finally, replication is important from an econometric perspective. If the S&P 500 ETF is spanned by the sector portfolios, then the specification must reflect this spanning (via a cointegration restriction).

Replication is to some extent hampered by differing expense ratios. The expense ratio of the S&P 500 ETF (currently and in the sample) is 0.12%. The expense ratios of the sector funds are 0.28%. One would therefore expect to see a cumulative shortfall in the replicating sector fund portfolio, even if its net holdings were identical to those of the S&P 500 ETF.

To assess tracking error, the S&P 500 ETF price was regressed against the set of sector fund prices. The prices used here were the daily closing quote midpoints. The relatively coarse time resolution is adequate for present purposes, because the immediate objective is characterization of any "long term" deviations. Table 8 reports the coefficients of this regression, together with the implied average weight (by dollars and relative value) of each sector in the total index. The R^2 in the regression was (to five decimal places) unity.[12] (The R^2 in the corresponding regression of price differences was 0.98, but in the presence of a cointegrating relationship, this regression is misspecified.) In summary, replication discrepancies resulting from cost differentials or other source, do not appear to be material.

Table 9 presents summary statistics. The table is presented as a matrix, with the entry in row *i* and column *j* measuring the contribution of fund *j* to the price discovery of fund *i*. For the sake of clarity, standard errors are not presented. The first row decomposes price discovery in the broad index (*SPY*). Approximately 80% of the price

---

[12] Normally, a coefficient of determination so close to unity would imply that the analyst had inadvertently estimated an identity, a linear specification that held exactly by virtue of some accounting relationship. This is not the case here. The variables are market prices of different securities. Furthermore, the sector fund prospectuses admit the possibility of tracking errors and caution against assuming exact replication.

discovery in *SPY* can be attributed to its own innovations. The sector funds share in the remainder, with the largest contribution originating from *XLK* (technology) at nearly 10%.

The second row decomposes price discovery in *XLB* (basic industries). Here, about 22% can be attributed to *XLB*'s own innovations, with *SPY* picking up the bulk of the remainder. This is typical of most rows in that most of the price discovery is shared by the *SPY* and a sector fund's own innovations. The *SPY*'s contribution is largest for technology fund (*XLK*, about 66%).

The strong contributions from *SPY* and *XLK*'s innovations to each others' long term variance is consistent with the prominence of *XLK*'s weight in the *SPY* (about 40%, from Table 8). In general, however, sector fund *SPY* weights are not closely related to price discovery either from the *SPY* to the sector fund or in the reverse direction.

Do price discovery contributions merely reflect the long-term return dependencies? Table 10 reports the bivariate correlations for the daily changes in end-of-day quote midpoints. The correlation between *SPY* and *XLK* is substantial (0.811), but the larger still is the correlation between *SPY* and *XLI* (0.902). Yet *XLI* contributes very little to the price discovery in *SPY* (about 1.9%, from the top row of Table 9).

These results are consistent with relatively low production of information at the sector level. Many factors might account for this. The cost of trading in these sectors might be a deterrent. Even though the average bid ask spreads reported in Table 1 are not especially high, trade volumes are often low. This cost is also determined by the extent of liquidity trading, which may also be low.

In light of mass media reports that over the sample period often identified the technology sector as the causal source of broad market movements, the *XLK* fund's contribution to *SPY* price discovery of 9.7% seems small. One possible explanation for this is the availability of better trading vehicles. The Nasdaq-100 index, for example, is particularly technology-heavy. However estimates based on a joint VAR analysis of the *SPY* and *QQQ* fund price changes found that latter's contribution to price discovery in the former was at most also about 10%.

8. Conclusions

Analyses of S&P 500 stock index and index futures dynamics have generally found that the latter lead the former, implying that price discovery takes place in the futures pit. The market characteristics have recently been changed, however, by the large volume of electronic trades in the futures market, the introduction of exchange traded index funds and the advent of exchange-traded sector funds (ETFs). This paper employs high-frequency time series analysis to re-examine index price discovery in the S&P 500, Nasdaq-100 and S&P 400 indexes in this new environment.

The paper's results suggest that for the S&P 500 and Nasdaq-100 indexes, price discovery is still dominated by futures trading. The contributions of the corresponding exchange traded funds, though statistically significant, are small. The results are less clear for the S&P MidCap 400 index. Here, the results suggest dominance of the ETF.

For the S&P 500 and Nasdaq-100 indexes, the Chicago Mercantile Exchange has introduced E-mini's. The E-minis have smaller size than the pit-traded contracts, and trade on the CME's GLOBEX electronic limit order book system. By underlying dollar volume and size of open interest, the E-minis are dominated by the regular contracts. Their trading volume is substantial, however, and the present analysis suggests that they now account for most of the price discovery in the markets for their respective indexes.

The S&P 500 sector funds are ETFs that are constructed on industry lines and can be used to replicate the overall index. Despite substantial trading activity in some of these funds, however, their contribution to price discovery in the overall index is modest. This suggests that information production is not occurring at the sector level.

9. Appendix: The time series properties of last-sale prices

It was asserted in section 5.d that if the first differences of a latent (implicit or unobserved) price are covariance stationary, then the first differences of last-sale prices are also covariance stationary under fairly mild regularity conditions. This appendix provides a formal analysis. The strategy is straightforward. The covariance structure of the observed price process is derived, and then examined to isolate the features that might cause a violation of covariance stationarity.

Suppose that we have a latent price process $\{p_t\}$ over the discrete times $t = \ldots, -1, 0, 1, \ldots$. The first differences are assumed to be zero-mean, covariance stationary: $E\,\Delta p_t = 0$; $E\,\Delta p_t\,\Delta p_{t-j} = \gamma_p(j)$. The observation process $\{q_t\}$ is defined by the indicator variable

$$q_t = \begin{cases} 1, \text{ if } p_t \text{ is observed at time } t \\ 0, \text{ otherwise} \end{cases} \tag{13}$$

It is assumed that $\{q_t\}$ is a stationary process. It may be autocorrelated (we may have "waves" of trading). The observed ("last sale") price is defined recursively as

$$p_t^* = q_t\, p_t + (1 - q_t)\, p_{t-1}^* \tag{14}$$

It is useful to define an indicator variable to mark the time of the most recent trade:

$$I_t(k) = \begin{cases} 1, \text{ if the most recent observation was at time } t-k \\ 0, \text{ otherwise} \end{cases}, \text{ for } k = 0, 1, \ldots \tag{15}$$

Alternatively,

$$I_t(k) = \bar{q}_t\, \bar{q}_{t-1} \ldots \bar{q}_{t-k+1}\, q_{t-k} \tag{16}$$

where $\bar{q}_t \equiv (1 - q_t)$, the "no observation" indicator variable. It is assumed that $\lim_{k\to\infty} I_t(k) = 1$, almost surely. (We can always find a trade by going back sufficiently far.) With these definitions, the observed price is:

$$p_t^* = \sum_{k=0}^{\infty} I_t(k)\, p_{t-k}\,. \tag{17}$$

The first difference of the observed price is

$$\Delta p_t^* = p_t^* - p_{t-1}^* = q_t\left(p_t - \sum_{k=0}^{\infty} I_{t-1}(k)\, p_{t-k-1} \right) = q_t\left(\sum_{k=0}^{\infty} I_{t-1}(k)\left(p_t - p_{t-k-1} \right) \right) \tag{18}$$

(The last equality follows because $\sum_{k=0}^{\infty} I_t(k) = 1$.)

In the case where the observation and latent price processes are independent, then

$$E\left[I_{t-1}(k)\left(p_{t-1} - p_{t-k-1} \right) \right] = E\left[I_{t-1}(k) \right] E\left[\left(p_{t-1} - p_{t-k-1} \right) \right] = 0\,, \tag{19}$$

implying $E\,\Delta p_t^* = 0$. The autocovariances of the observed prices are computed from

$$\begin{aligned} E\,\Delta p_t^*\,\Delta p_{t-j}^* &= E\left(q_t\left(\sum_{k=0}^{\infty} I_{t-1}(k)\left(p_t - p_{t-k-1} \right) \right) \right) \times \left(q_{t-j}\left(\sum_{l=0}^{\infty} I_{t-j-1}(l)\left(p_{t-j} - p_{t-j-l-1} \right) \right) \right) \\ &= E\left(\sum_{k=0}^{\infty}\sum_{l=0}^{\infty} \left(q_t\, q_{t-j}\, I_{t-1}(k)\, I_{t-j-1}(l) \right) \times \left(\left(p_t - p_{t-k-1} \right)\left(p_{t-j} - p_{t-j-l-1} \right) \right) \right) \\ &= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} E\left[\left(q_t\, q_{t-j}\, I_{t-1}(k)\, I_{t-j-1}(l) \right) \times \left(\left(p_t - p_{t-k-1} \right)\left(p_{t-j} - p_{t-j-l-1} \right) \right) \right] \\ &= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} E\left(q_t\, q_{t-j}\, I_{t-1}(k)\, I_{t-j-1}(l) \right) \times E\left(\left(p_t - p_{t-k-1} \right)\left(p_{t-j} - p_{t-j-l-1} \right) \right) \end{aligned} \tag{20}$$

From the stationarity of $\{q_t\}$ and the covariance stationarity of $\{p_t\}$, the two expectations in the summand depend only on $k, l,$ and j (not t).

The assumption of independence between $\{p_t\}$ and $\{q_t\}$ may be replaced by the weaker assumption that $\{p_t\}$ is conditionally covariance-stationary, in the following sense:

$$\begin{aligned} &E\left[\Delta p_t \,|\, q_t, q_{t-1}, \ldots \right] = 0 \\ &\text{and} \\ &E\left[\Delta p_t\, \Delta p_{t-j} \,|\, q_t, q_{t-1}, \ldots \right] = f\left(j, q_t, q_{t-1}, \ldots \right) \end{aligned} \tag{21}$$

The force of this assumption is simply to suppress explicit time dependence in the mean and autocovariances. It allows, however, for effects such as heightened volatility or stronger autocovariances during periods of elevated trading intensity. One might hypothesize, for example, $Var\left(\Delta p_t \,|\, q_t, q_{t-1}, \ldots \right) = E\left(\Delta p_t^2 \,|\, q_t, q_{t-1}, \ldots \right) = a_0 + a_1\, q_t + a_2\, q_{t-1}$, for

some positive coefficients a_0, a_1 and a_2. The important point here is that the conditional variance depends on *t* only through q_t and q_{t-1}.

The multiperiod price changes that appear in the summand of (20) can be written as the sum of one-period price changes. For example,

$$\left(p_t - p_{t-k-1}\right) = \left(\Delta p_t + \Delta p_{t-1} + \cdots + \Delta p_{t-k}\right). \tag{22}$$

The conditional expectation of the product of these multiperiod price changes can therefore be written as

$$E\left[\left(p_t - p_{t-k-1}\right)\left(p_{t-j} - p_{t-j-l-1}\right)\middle| q_t, q_{t-1}, \ldots\right] = F\left(j,k,l,q_t,q_{t-1},\ldots\right) \tag{23}$$

i.e., a function that does not depend on *t*. Re-examining the autocovariances of the observed price changes, we obtain

$$\begin{aligned} E\,\Delta p_t^*\,\Delta p_{t-j}^* &= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} E\left[\begin{array}{l}\left(q_t\, q_{t-j}\, I_{t-1}(k)\, I_{t-j-1}(l)\right) \\ \qquad \times E\left[\left(p_t - p_{t-k-1}\right)\left(p_{t-j} - p_{t-j-l-1}\right)\middle| q_t, q_{t-1}, \ldots\right]\end{array}\right] \\ &= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} E\left[\left(q_t\, q_{t-j}\, I_{t-1}(k)\, I_{t-j-1}(l)\right) \times F\left(j,k,l,q_t,q_{t-1},\ldots\right)\right] \\ &= \sum_{k=0}^{\infty}\sum_{l=0}^{\infty} G\left(j,k,l\right) \end{aligned} \tag{24}$$

for some function *G*, where (again) explicit time dependence vanishes.

10. References

Admati, A., Pfleiderer, P., 1988. A theory of intraday trading patterns: Volume and price variability. Review of Financial Studies 1, 3-40.

Bartlett, M. S., 1950. Periodogram analysis and continuous spectra. Biometrika 37, 1-16.

Beveridge, S., Nelson, C. R., 1981. A new approach to the decomposition of economic time series into permanent and transitory components with particular attention to the measurement of the 'business cycle'. Journal of Monetary Economics 7, 151-174.

Booth, G. G., So, R. W., Tse, Y., 1999. Price discovery in the German equity index derivatives markets. Journal of Futures Markets 19, 619-643.

Choi, H., Subrahmanyam, A., 1994. Using intraday data to test for effects of index futures on the underlying stock markets. Journal of Futures Markets 14, 293-322.

Chow, Y.-F., 1998. Regime switching and cointegration tests of the efficiency of futures markets. Journal of Futures Markets 18, 871-901.

Chowdhry, B., Nanda, V., 1991. Multimarket trading and market liquidity. Review of Financial Studies 4, 483-511.

Chu, Q. C., Hsieh, W. G., Tse, Y., 1999. Price discovery on the S&P 500 index markets: An analysis of spot index, index futures and SPDRs. International Review of Financial Analysis 8, 21-34.

Coppejans, M., Domowitz, I., 1999. Pricing behavior in an off-hours computerized market. Journal of Empirical Finance 6, 583-607.

Corkish, J., Holland, A., Vila, A. F., 1997. The determinants of successful financial innovation: an empirical analysis of futures innovation on LIFFE. Unpublished working paper. Bank of England, London.

de Jong, F., Donders, M., 1998. Intraday lead-lag relationships between the futures-options and stock market. European Financial Review 1, 337-359.

de Jong, F., Mahieu, R., Schotman, P., 1998a. Price discovery in the foreign exchange

market: An empirical analysis of the Yen/Dmark Rate. Journal of International Money and Finance 17, 5-27.

De Jong, F., Mahieu, R., Schotman, P., 1998b. Price discovery in the foreign exchange market: an empirical analysis of the Yen/Dmark rate. Journal of International Money and Finance 17, 5-27.

de Jong, F., Nijman, T., 1997. High frequency analysis of the lead-lag relationships between financial markets. Journal of Empirical Finance 4, 193-213.

Ding, D. K., et al, 1999. An investigation of price discovery in informationally-linked markets: Equity trading in Malaysia and Singapore. Journal of Multinational Financial Management 3, 317-329.

Domowitz, I., Steil, B., 1999. Automation, trading costs and the structures of the securities trading industry. In: Litan, R. E. and Santomero, A. M. (Eds.), Brookings-Wharton Papers on Financial Services (1999). Brookings Institution Press, Washington DC, 33-92.

Duffie, D., Jackson, M. O., 1989. Optimal innovation of futures contracts. Review of Financial Studies 2, 275-296.

Elton, E. J., Gruber, M. J., Comer, G., Li, K., 2000. Spiders: Where are the bugs. Unpublished working paper. Department of Finance, Stern School, New York University, New York.

Engle, R. F., Granger, C. W. J., 1987. Co-integration and error correction: representation, estimation and testing. Econometrica 55, 251-276.

Fleming, J., Ostdiek, B., Whalley, R. E., 1996. Trading costs and the relative rates of price discovery in stock, futures and options markets. Journal of Futures Markets 16, 353-387.

Foster, F. D., Viswanathan, S., 1990. A theory of the interday variations in volume, variance, and trading costs in securities markets. Review of Financial Studies 3, 593-624.

Gorton, G. B., Pennacchi, G. G., 1993. Security baskets and index-linked securities. Journal of Business 66, 1-27.

Greene, W.H., 1993. *Econometric Analysis.* 2nd ed., (New York: Macmillan).

Grunbichler, A., Longstaff, F. A., Schwartz, E. S., 1994. Electronic screen trading and the transmission of information: An empirical examination. Journal of Financial Intermediation 3, 166-187.

Hamilton, J.D., 1994. *Time Series Analysis.* (Princeton: Princeton University Press).

Harris, F. H. d., McInish, T. H., Shoesmith, G. L., Wood, R. A., 1995. Cointegration, error correction, and price discovery on informationally linked security markets. Journal of Financial and Quantitative Analysis 30, 563-579.

Harris, L. E., Sofianos, G., Shapiro, J. E., 1994. Program trading and intraday volatility. Review of Financial Studies 7, 653-85.

Hasbrouck, J., 1991. The summary informativeness of stock trades: An econometric analysis. Review of Financial Studies 4, 571-95.

Hasbrouck, J., 1995. One security, many markets: Determining the contributions to price discovery. Journal of Finance 50, 1175-99.

Hasbrouck, J., 2000. Stalking the efficient price in empirical microstructure specifications. Unpublished working paper. Department of Finance, Stern School, New York University, New York.

He, L. T., 1997. Price discovery in Hong Kong security markets: evidence from cointegration tests. Journal of International Financial Markets Institutions and Money 7, 157-169.

Huang, R., 2000. Price discovery by ECNs and Nasdaq market makers. Unpublished working paper. Owen School of Management, Vanderbilt University.

Kawaller, I. G., Koch, P. D., Koch, T. W., 1987. The temporal relationship between S&P 500 futures and the S&P 500 index. Journal of Finance 42, 1309-1329.

Kofman, P., Moser, J. T., 1977. Spreads, information flows and transparency across trading systems. Applied Financial Economics 7, 281-294.

Kumar, P., Seppi, D. J., 1994. Information and index arbitrage. Journal of Business 67, 481-509.

Lo, A. W., MacKinlay, A. C., 1988. Stock market prices do not follow random walks:

Evidence from a simple specification test. Review of Financial Studies 1, 41-66.

MacKinlay, A. C., Ramaswamy, K., 1988. Index-futures arbitrage and the behavior of stock index futures prices. Review of Financial Studies 1, 137-158.

Martens, M., Kofman, P., 1998. The inefficiency of Reuters foreign exchange quotes. Journal of Banking & Finance 22, 347-366.

Martens, M., 1998. Price discovery in high and low volatility periods: open outcry versus electronic trading. Journal of International Financial Markets, Institutions and Money 8, 243-260.

Pizzi, M. A., Economopoulos, A. J., O'Neill, H. M., 1998. An examination of the relationship between stock index cash and futures markets: A cointegration approach. Journal of Futures Markets 18, 297-305.

Stock, J. H., Watson, M. W., 1988. Variable trends in economic time series. Journal of Economic Perspectives 2, 147-174.

Stoll, H. R., Whalley, R. E., 1990. The dynamics of stock index and stock index futures returns. Journal of Financial and Quantitative Analysis 25, 441-468.

Subrahmanyam, A., 1991. A theory of trading in stock index futures. Review of Financial Studies 4, 17-51.

Wahab, M., Lashgari, M., 1993. Price dynamics and error correction in stock index and stock index futures markets: A cointegration approach. Journal of Futures Markets 13, 711-742.

Watson, M. W., 1986. Univariate detrending methods with stochastic trends. Journal of Monetary Economics 18, 49-75.

Table 1

Summary statistics for exchange-trade funds

All statistics except index factor are based on TAQ data from March 1, 2000 to May 31, 2000. The index factor indicates the scale of the exchange-traded fund share price. To obtain the corresponding index level, the share price is multiplied by the index factor.

Symbol		*Average daily volume (1M sh)*	*Average Amex vol share*	*Average daily trades*	*Average Amex share of trades*	*Avg Closing Price ($/share)*	*Avg Spread ($/share)*	*Avg Relative Spread*	*Tick size ($/share)*	*Index factor*	*Closing level of index (5/31/2000)*
Indexes											
SPY	S&P 500	9,888	0.76	1,735	0.62	144	0.194	0.0013	1/32	10	1,420.60
QQQ	Nasdaq-100	30,363	0.76	10,575	0.28	108	0.229	0.0023	1/32	40	3,324.08
MDY	S&P 400	745	0.84	245	0.56	87	0.218	0.0025	1/64	5	475.17
Sectors											
XLB	Basic industries	126	0.80	42	0.61	22	0.180	0.0081	1/64		
XLE	Energy	306	0.73	71	0.61	29	0.176	0.0060	1/64		
XLF	Financial	448	0.73	191	0.43	23	0.149	0.0064	1/64		
XLI	Industrial	113	0.75	24	0.66	29	0.169	0.0059	1/64		
XLK	Technologies	848	0.74	354	0.44	55	0.194	0.0036	1/64		
XLP	Consumer staples	274	0.72	87	0.54	22	0.169	0.0077	1/64		
XLU	Utilities	80	0.76	30	0.59	27	0.177	0.0065	1/64		
XLV	Consumer services	48	0.81	24	0.82	29	0.162	0.0055	1/64		
XLY	Cyclical/Transportation	96	0.78	27	0.64	28	0.163	0.0058	1/64		

Table 2

Summary Statistics for index futures contracts

All statistics are based on Chicago Mercantile Exchange data from March 1, 2000 to May 31, 2000.

Symbol	*Underlying*	*Type*	*Trading*	*Contract size (x index)*	*Tick size (index points)*	*Avg daily volume (contracts)*	*Avg end-of-month open interest (contracts)*
SP	S&P 500	Regular	Pit	250	0.100	98,324	384,476
ES	S&P 500	E-mini	Globex	50	0.250	73,779	32,331
ND	Nasdaq-100	Regular	Pit	100	0.500	24,096	36,189
NQ	Nasdaq-100	E-mini	Globex	20	0.500	30,991	18,247
MD	S&P 400	Regular	Pit	500	0.050	1,310	12,910

Table 3 Transaction cost analysis

Analysis of one-way transactions costs for the smallest regular trade for each type of index security. A fixed commission of $10 is assumed on all trades. ETF=Exchange-traded fund.

Index:	S&P 500			Nasdaq-100			S&P 400	
Symbol	*SP*	*ES*	*SPY*	*ND*	*NQ*	*QQQ*	*MD*	*MDY*
Contract/fund type	Regular	E-mini	ETF	Regular	E-mini	ETF	Regular	ETF
Smallest regular trade (contracts or shares)	1	1	100	1	1	100	1	100
Smallest regular trade ($)	355,150	71,030	14,206	332,408	66,482	16,620	237,585	9,503
Half-spread	12.5	6.25	18.47	25	5	11.45	12.5	21.8
Total (half-spread + $10)	22.5	16.25	28.47	35	15	21.45	22.5	31.8
Relative trade cost	0.0001	0.0002	0.0020	0.0001	0.0002	0.0013	0.0001	0.0033

Table 4

Price discovery in the S&P 500 index

Statistics are based on a vector error correction model of prices for the exchange-traded fund (symbol *SPY*), the regular futures contract (*SP*) and the E-mini (*ES*), estimated at a one-second resolution. The model is estimated for each day in the sample (March 1, 2000 through May 31, 200, 64 trading days). The table reports summary statistics for these daily estimates. Disturbance correlations are given in panel A. Panel B contains the coefficients of the disturbances in the long-run price revision. Information shares are reported in panel C. The entries in panels A and B are means of the daily estimates. The statistics in panel C are over the sample of daily estimates.

A. Disturbance correlation matrix

	SPY	*SP*	*ES*
SPY	1.000	0.003	0.002
SP	0.003	1.000	0.052
ES	0.002	0.052	1.000

B. Coefficients of efficient price

SPY	*SP*	*ES*
0.167	0.199	0.612

C. Information shares

	Source					
	SPY		*SP*		*ES*	
	Min	*Max*	*Min*	*Max*	*Min*	*Max*
Median	0.022	0.024	0.055	0.082	0.866	0.893
Mean	0.037	0.040	0.081	0.107	0.854	0.882
SEM	0.007	0.007	0.010	0.012	0.012	0.011
Std. Dev.	0.056	0.059	0.079	0.093	0.097	0.086

Table 5

Price discovery in the Nasdaq-100 index

Statistics are based on a vector error correction model of prices for the exchange-traded fund (symbol *QQQ*), the regular futures contract (*ND*) and the E-mini (*NQ*), estimated at a one-second resolution. The model is estimated for each day in the sample (March 1, 2000 through May 31, 200, 64 trading days). The table reports summary statistics for these daily estimates. Disturbance correlations are given in panel A. Panel B contains the coefficients of the disturbances in the long-run price revision. Information shares are reported in panel C. The entries in panels A and B are means of the daily estimates. The statistics in panel C are over the sample of daily estimates.

A. Disturbance correlation matrix

	QQQ	*ND*	*NQ*
QQQ	1.000	0.001	0.002
ND	0.001	1.000	0.031
NQ	0.002	0.031	1.000

B. Coefficients of efficient price

QQQ	*ND*	*NQ*
0.305	0.192	0.677

C. Information shares

	Source					
	QQQ		*ND*		*NQ*	
	Min	*Max*	*Min*	*Max*	*Min*	*Max*
Median	0.055	0.059	0.046	0.061	0.858	0.875
Mean	0.085	0.089	0.061	0.076	0.836	0.853
SEM	0.010	0.011	0.007	0.008	0.012	0.011
Std. Dev.	0.084	0.085	0.054	0.061	0.094	0.091

Table 6

Price discovery in the S&P 400 index

Statistics are based on a vector error correction model of prices for the exchange-traded fund (symbol *MDY*) and the regular futures contract (*MD*) estimated at a five-second resolution. The model is estimated for each day in the sample (March 1, 2000 through May 31, 200, 64 trading days). The table reports summary statistics for these daily estimates. Disturbance correlations are given in panel A. Panel B contains the coefficients of the disturbances in the long-run price revision. Information shares are reported in panel C. The entries in panels A and B are means of the daily estimates. The statistics in panel C are over the sample of daily estimates.

A. Disturbance correlation matrix

	MDY	*MD*
MDY	1.000	0.014
MD	0.014	1.000

B. Coefficients of efficient price

MDY	*MD*
1.148	0.567

C. Information shares

	Source			
	MDY		*MD*	
	Min	*Max*	*Min*	*Max*
Median	0.661	0.680	0.320	0.339
Mean	0.616	0.629	0.371	0.384
SEM	0.033	0.033	0.033	0.033
Std. Dev.	0.263	0.262	0.262	0.263

Table 7

Price discovery in the equity indexes: Sensitivity to time shifts

Entries in the table are means (across days) of minimal and maximal information shares implied by vector error correction models estimated at five-second resolution. For the S&P 500 index, the model comprises the ETF (symbol *SPY*), the regular futures contract (*SP*) and the E-mini futures contract (*ES*); for the Nasdaq 100 index, the model comprises the ETF (*QQQ*), the regular futures contract (*ND*) and the E-mini contract (*NQ*). In these estimations, the reported prices from the E-mini markets (*ES* and *NQ*) are delayed by the indicated amounts.

S&P 500						
	SPY		*SP*		*ES*	
ES Delay (seconds)	*Min*	*Max*	*Min*	*Max*	*Min*	*Max*
0	0.049	0.060	0.077	0.213	0.730	0.874
5	0.049	0.092	0.197	0.472	0.452	0.752
10	0.066	0.150	0.449	0.684	0.214	0.469
15	0.096	0.182	0.631	0.770	0.114	0.225
Nasdaq 100						
	QQQ		*ND*		*NQ*	
NQ Delay (seconds)	*Min*	*Max*	*Min*	*Max*	*Min*	*Max*
0	0.117	0.131	0.042	0.103	0.769	0.839
5	0.124	0.142	0.094	0.232	0.632	0.779
10	0.131	0.153	0.225	0.406	0.450	0.638
15	0.140	0.163	0.392	0.553	0.298	0.460

Table 8 S&P 500 ETF sector replication and weights

Table is based on end-of-day TAQ closing quote midpoints for the S&P 500 ETF (*SPY*) and its component sector funds, March 1, 2000 through May 31, 2000. The *SPY* coefficients are the estimated a_i in the regression:

$$p_t^{SPY} = a_{XLB}\, p_t^{XLB} + a_{XLE}\, p_t^{XLE} + \cdots + a_{XLY}\, p_t^{XLY} + e_t$$

where the p_t^{XLB}, etc. are closing quote midpoints. (The R^2 of this regression was essentially unity.) The dollar components of the *SPY* are computed as $a_{XLB}\, \overline{p}^{XLB}$, etc., where $\overline{p}^{XLB}$ is the sample average price of the *XLB* fund. The percentage components of the *SPY* are computed as $a_{XLB}\, \overline{p}^{XLB} / \overline{p}^{SPY}$, etc.

Symbol	*Description*	*Average closing quote midpoint*	*Coefficient in SPY regression*	*Component of SPY ($)*	*Component of SPY (%)*
SPY	S&P 500	144.26			
XLB	Basic industries	22.37	0.070	1.56	0.011
XLE	Energy	29.27	0.256	7.48	0.052
XLF	Financial	23.42	0.861	20.16	0.140
XLI	Industrial	28.60	0.260	7.45	0.052
XLK	Technologies	55.07	1.053	58.01	0.402
XLP	Consumer staples	22.15	1.117	24.75	0.172
XLU	Utilities	27.31	0.261	7.14	0.049
XLV	Consumer services	29.37	0.136	3.99	0.028
XLY	Cyclical/Transportation	28.10	0.489	13.74	0.095
				144.26	1.000

Table 9

Price discovery between the S&P 500 and its component sector portfolios

Table entry $I_{i,j}$ (*i*th row and *j*th column) is the information share of security *j* in the long-term variance of security *i*. Due to residual correlations, these are not unambiguously determined. Table entries are information shares based on factor rotations when security *j* is assigned precedence. I.e, the entries are approximately the maximal information shares. The specification is a vector error correction model estimated with five-second resolution. The specification includes lags through five minutes.

	From:									
To:	*SPY*	*XLB*	*XLE*	*XLF*	*XLI*	*XLK*	*XLP*	*XLU*	*XLV*	*XLY*
SPY	0.801	0.020	0.014	0.020	0.019	0.097	0.021	0.016	0.014	0.018
XLB	0.216	0.661	0.014	0.020	0.018	0.029	0.023	0.019	0.019	0.017
XLE	0.176	0.027	0.674	0.020	0.021	0.041	0.023	0.020	0.017	0.018
XLF	0.447	0.024	0.015	0.417	0.018	0.050	0.018	0.019	0.017	0.020
XLI	0.401	0.028	0.016	0.021	0.465	0.034	0.016	0.021	0.019	0.024
XLK	0.658	0.018	0.016	0.020	0.016	0.244	0.023	0.019	0.014	0.018
XLP	0.316	0.019	0.016	0.021	0.013	0.079	0.521	0.019	0.017	0.020
XLU	0.316	0.021	0.018	0.020	0.031	0.041	0.020	0.541	0.017	0.017
XLV	0.430	0.019	0.013	0.026	0.021	0.040	0.018	0.014	0.446	0.021
XLY	0.427	0.024	0.015	0.024	0.022	0.044	0.017	0.015	0.030	0.429

Table 10

***SPY* and the sector portfolios: Correlations of daily price changes**

Correlations in daily quote-midpoint price changes for the *SPY* and its component sector funds, March 1, 2000 through May 31, 2000.

	SPY	*XLB*	*XLE*	*XLF*	*XLI*	*XLK*	*XLP*	*XLU*	*XLV*	*XLY*
SPY	1.000	0.492	0.184	0.727	0.902	0.811	0.500	0.536	0.804	0.661
XLB	0.492	1.000	0.243	0.648	0.654	0.048	0.684	0.522	0.335	0.703
XLE	0.184	0.243	1.000	0.286	0.177	-0.041	0.134	0.258	0.185	0.323
XLF	0.727	0.648	0.286	1.000	0.735	0.315	0.583	0.545	0.631	0.769
XLI	0.902	0.654	0.177	0.735	1.000	0.620	0.567	0.547	0.686	0.753
XLK	0.811	0.048	-0.041	0.315	0.620	1.000	0.030	0.180	0.656	0.204
XLP	0.500	0.684	0.134	0.583	0.567	0.030	1.000	0.563	0.260	0.661
XLU	0.536	0.522	0.258	0.545	0.547	0.180	0.563	1.000	0.389	0.599
XLV	0.804	0.335	0.185	0.631	0.686	0.656	0.260	0.389	1.000	0.626
XLY	0.661	0.703	0.323	0.769	0.753	0.204	0.661	0.599	0.626	1.000

Bivariate daily return correlations

Figure 1

S&P 500 Average daily basis

The basis is defined as $F_t - S_t$ where F_t is the price of the regular (pit-traded) index futures contract and S_t is the price of the corresponding exchange-traded fund (the "spot" price), as of the end of minute *t*. The sample is March 1, 2000 through May 31, 2000. The daily average is plotted as a dot. The vertical lines (sometimes completely obscured) mark the range, ± twice the standard error of the mean (corrected for autocorrelation).

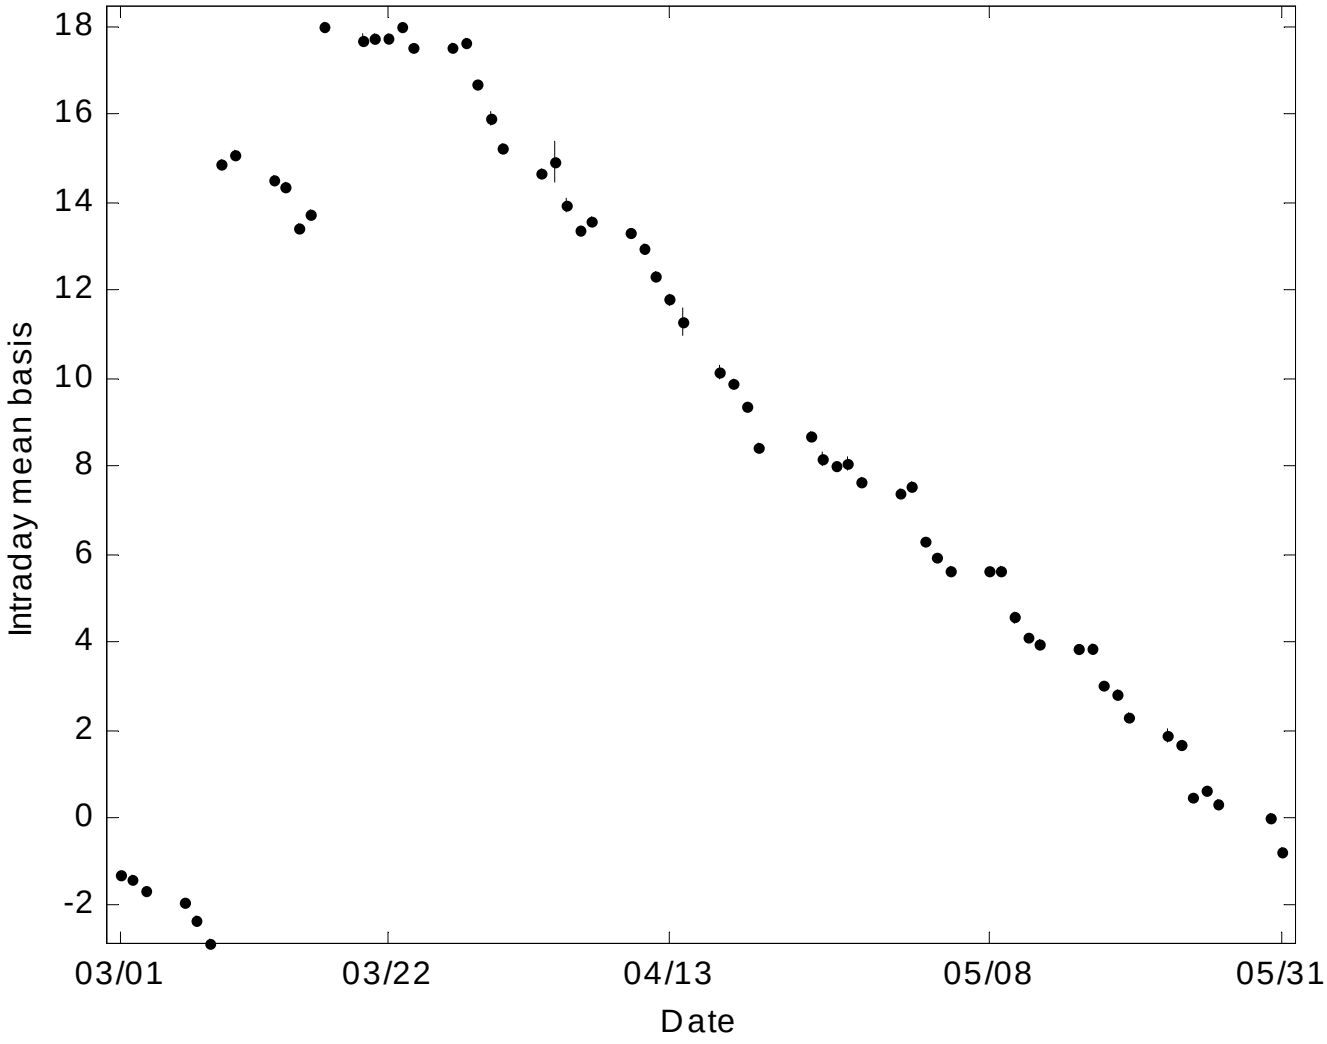

Figure 2

Nasdaq-100 Average daily basis

The basis is defined as $F_t - S_t$ where F_t is the price of the regular (pit-traded) index futures contract and S_t is the price of the corresponding exchange-traded fund (the "spot" price), as of the end of minute *t*. The sample is March 1, 2000 through May 31, 2000. The daily average is plotted as a dot. The vertical lines (sometimes completely obscured) mark the range, ± twice the standard error of the mean (corrected for autocorrelation).

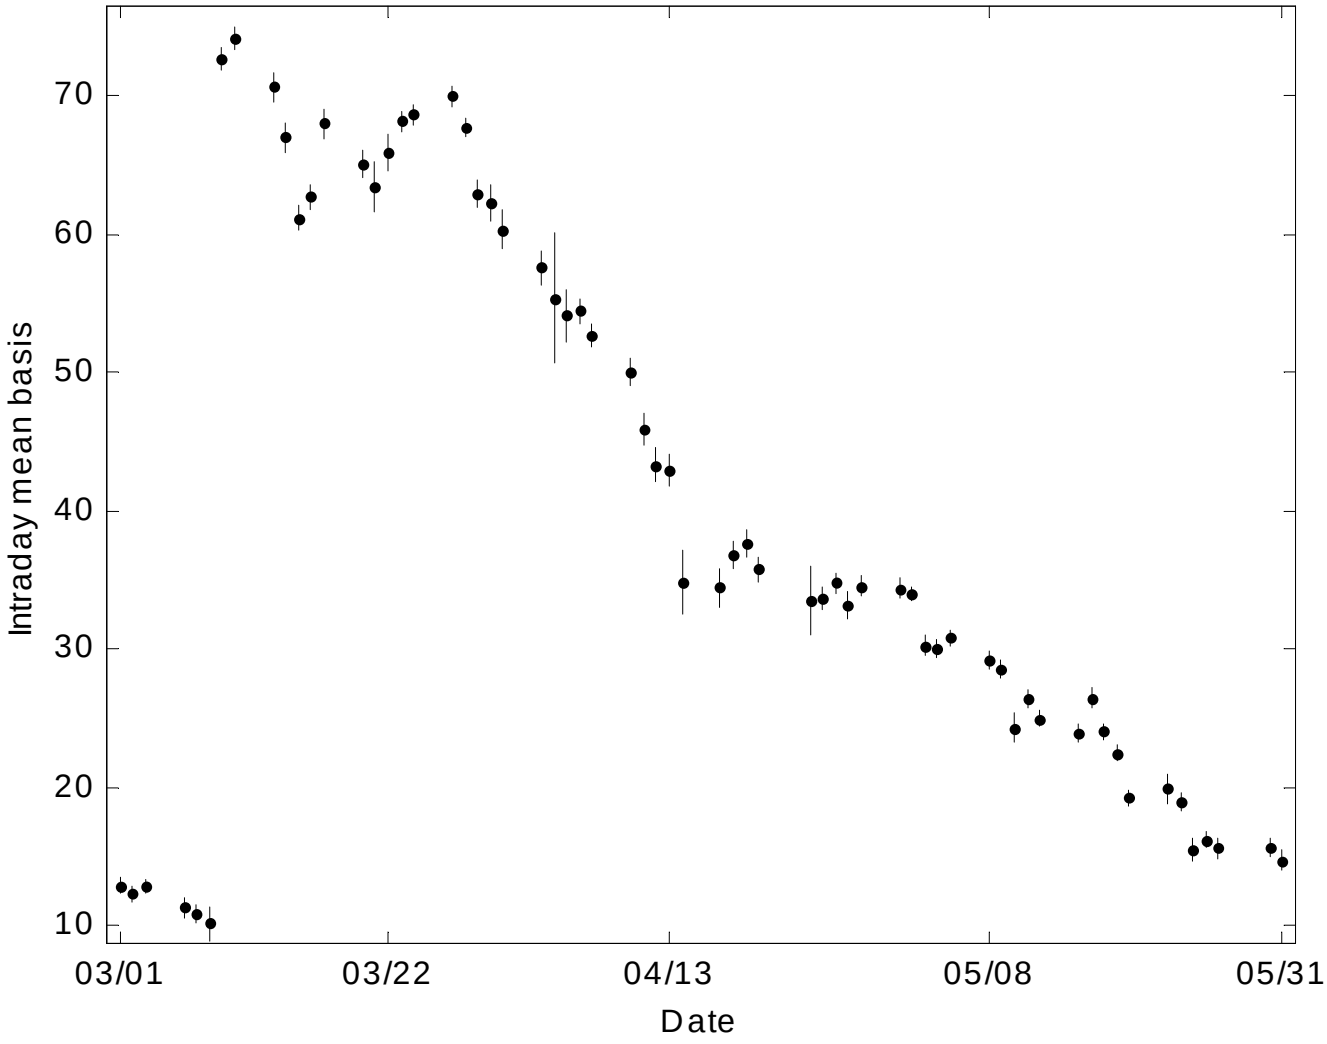

Figure 3

S&P 400 Average daily basis

The basis is defined as $F_t - S_t$ where F_t is the price of the regular (pit-traded) index futures contract and S_t is the price of the corresponding exchange-traded fund (the "spot" price), as of the end of minute *t*. The sample is March 1, 2000 through May 31, 2000. The daily average is plotted as a dot. The vertical lines (sometimes completely obscured) mark the range, ± twice the standard error of the mean (corrected for autocorrelation).

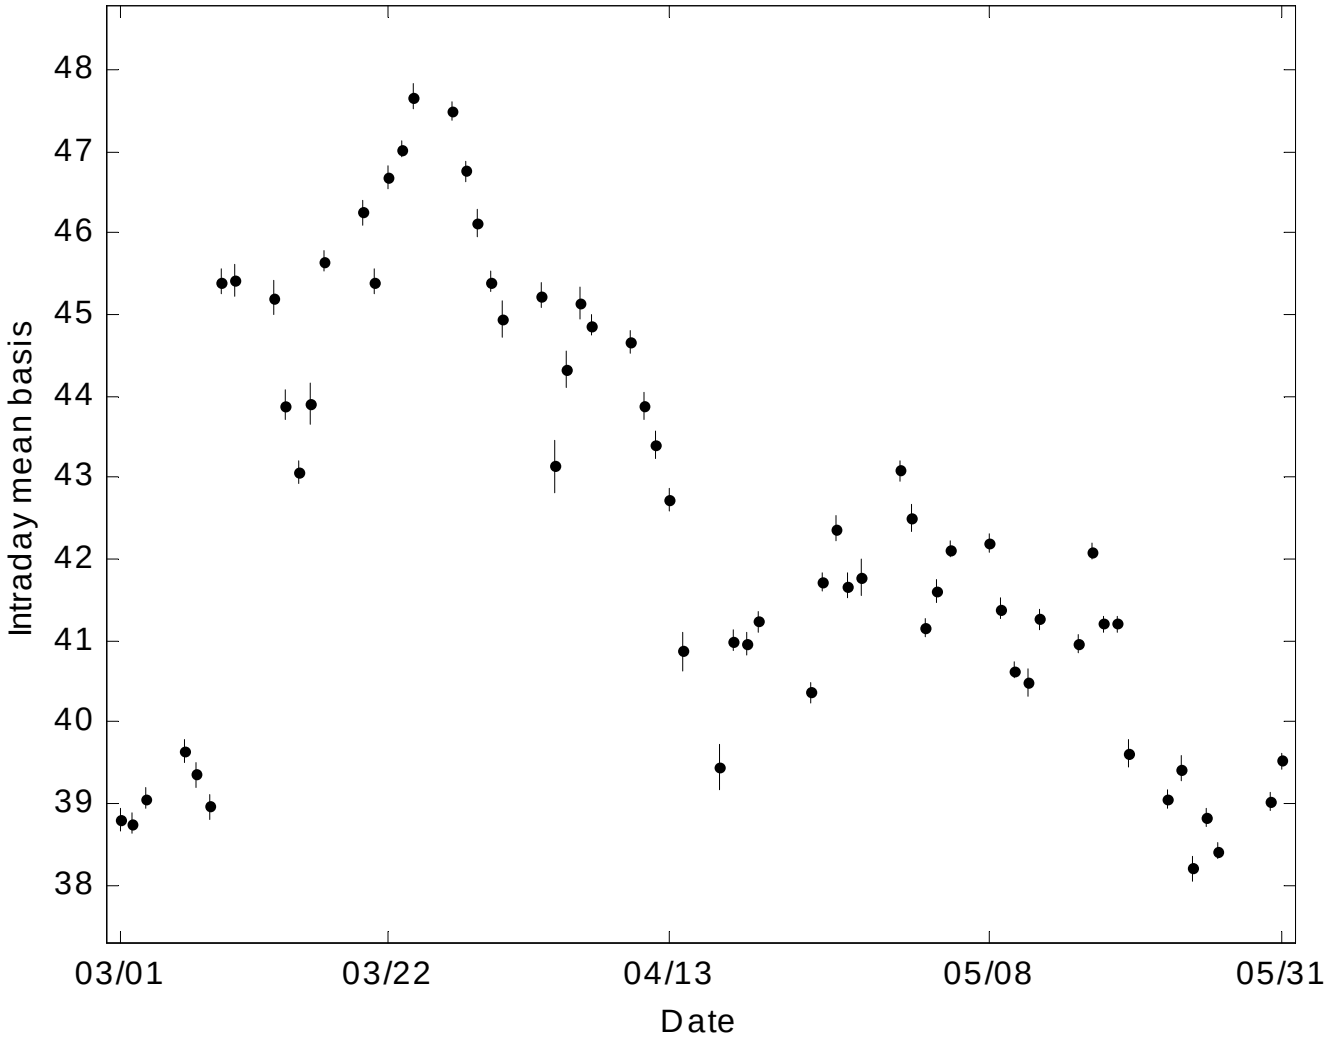

Figure 4

Intraday bases on March 1, 2000

The basis is defined as $F_t - S_t$ where F_t is the price of the regular (pit-traded) index futures contract and S_t is the price of the corresponding exchange-traded fund (the "spot" price), as of the end of minute *t*.

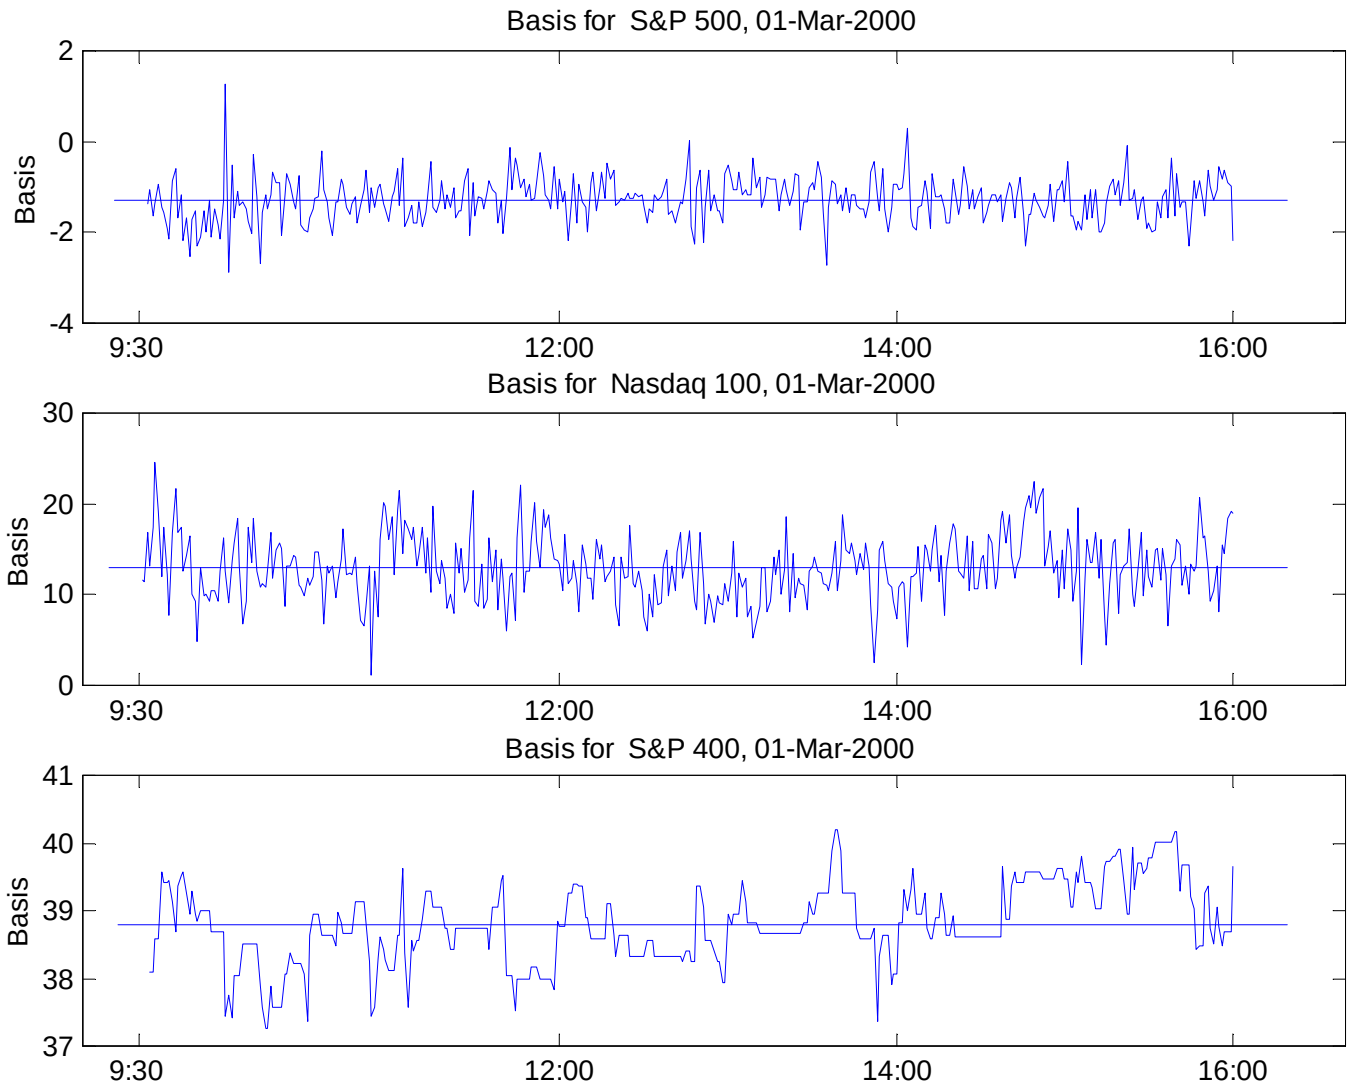

Figure 5

Cumulative impulse response (price impact) functions

Cumulative price impacts subsequent to initial unit shocks implied by vector error correction models. The models are estimated with one-second resolution for S&P 500 and Nasdaq 100 indexes; five-second resolution for the S&P 400 index. The top row (panels A, B and C) refers to the S&P 500 market, with symbols *SPY* (the exchange-traded fund); *SP* (the regular futures contract); and, *ES* (the E-mini futures contract). The middle row refers to the Nasdaq-100 market, with symbols *QQQ* (the exchange-traded fund); *ND* (the regular contract); and *NQ* (the E-mini). The bottom row refers to the S&P 400 market, with symbols *MDY* (the exchange-traded fund); and *MD* (the regular contract).