THE JOURNAL OF FINANCE • VOL. LVIII, NO. 6 • DECEMBER 2003

Intraday Price Formation in U.S. Equity Index Markets

JOEL HASBROUCK*

ABSTRACT

The market for U.S. equity indexes presently comprises floor-traded index futures contracts, exchange-traded funds (ETFs), electronically traded, small-denomination futures contracts (E-minis), and sector ETFs that decompose the S&P 500 index into component industry portfolios. This paper empirically investigates price discovery in this environment. For the S&P 500 and Nasdaq-100 indexes, most of the price discovery occurs in the E-mini market. For the S&P 400 MidCap index, price discovery is shared between the regular futures contract and the ETF. The S&P 500 ETF contributes markedly to price discovery in the sector ETFs, but there are only minor effects in the reverse direction.

THIS PAPER IS AN EMPIRICAL STUDY of the short-run price dynamics in three important U.S. equity index markets (S&P 500, S&P MidCap 400, and Nasdaq-100). Traditionally these markets have been viewed as principally comprising floor-traded index futures contracts and the underlying spot or cash markets, consisting of the individual markets for the component stocks. Against this backdrop, the present paper is motivated by the relatively recent emergence of:

- Exchange-traded index mutual funds (ETFs), which permit the direct purchase or sale of an index portfolio at any time.
- Electronically-traded index futures contracts.
- Exchange-traded S&P 500 index sector funds, which decompose the overall index into industry components.

At first glance these new securities would seem to be highly redundant and incapable of expanding investors' opportunity sets in any significant way. Nevertheless, in terms of trading volume, they are highly successful. The natural question, which this paper seeks to address, is the extent to which these new securities have altered the dynamics of index price formation, particularly in respect to price discovery and leadership.

An exchange-traded fund (ETF) is most closely related to a closed-end mutual fund in that it may traded continuously (during regular trading hours) and

*Stern School of Business, New York University. For helpful discussions and comments I am indebted to Ned Elton, Rob Engle, Richard Hogan of Spear Leeds and Kellog, Gideon Saar, Allan Zebedee, and seminar participants at the University of Chicago, the University of Iowa, and the University of Michigan. All errors are my own responsibility.

allows no cash investments or redemptions. Creations and redemptions are allowed, however, in kind: Units of the fund are swapped for portfolios of stocks and a small cash component.[1] The first ETF, the S&P 500 fund (SPDR or "spider"), began trading in 1993. Elton et al. (2002) 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 by off-floor traders. This system is similar to electronic communication networks (ECNs) in U.S. equity markets.

The third development is the introduction of nine exchange-traded S&P 500 sector funds. These 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 index funds. 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. Within this enriched set of index securities, the present paper attempts to quantify the informational contribution of each to determination of the index price.

The paper's key results may be summarized as follows. Numerous preexisting studies generally find that the floor-traded index futures contract "leads" the cash stock index. The present study finds that for the S&P 500 and Nasdaq-100 indexes, the largest informational contributions presently arise from the electronically

[1] The creation and redemption process for the first ETF, the S&P 500 fund (SPDR or "spider") is typical. Creation or redemption must occur in multiples of 50,000 SPDR shares (roughly \$7.5 million during the sample period). The SPDR's trustee determines the exact composition of the exchange portfolio (numbers of shares of each component stock and size of the cash component, which represents accumulated dividends). A party who wishes to create SPDRs must deliver the exchange portfolio to the Depository Trust Company at the close of business. A party who wishes to redeem SPDRs pays cash and receives the exchange portfolio. The creation and redemption fees start at \$3,000.

traded futures contracts, with the ETFs playing a smaller role. The pattern is different for the S&P 400 index, however, where no E-mini contract existed over the sample period. For this index, the ETF provides substantial price discovery. Finally, the information contribution of the sector ETFs to S&P 500 pricing is very small (about 10% by the paper's metric). This modest role is surprising because some of these funds (such as the technology ETF) are traded actively, 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 first section reviews the literature related to the paper's principal economic themes. Section II presents a preliminary look at the data. Sections III to V describe the overall method, specification of the price vector, and estimation. Section VI presents results for the index ETFs and their futures contracts; Section VII, the sector ETFs. A brief summary concludes the paper in Section VIII.

I. 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 versus 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 versus Cash

As noted in the introduction, the cash market for stock indices has traditionally consisted primarily of the markets for the individual component stocks. This has motivated theoretical models that feature multiple stocks and basket securities (Subrahmanyam (1991) and Kumar and Seppi (1994), among others). From an empirical perspective, a substantial academic and practitioner literature explores the dynamics of U.S. stock index and index futures prices with the aim of determining which market is dominant. The methodologies of these studies vary widely, but the consensus finding is that the dominant influence runs from the futures market to the cash, and weaker (though still measurable) effects occur in the reverse direction.[2]

B. Electronic versus 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 secu-

[2] 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), and 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.

rities 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. Coppejans and Domowitz (1999) compare the nighttime Globex market and daytime floor markets for regular contracts and conclude that the Globex system performs well during a period (the night) when the flow of liquidity traders is likely to be relatively low. However, Venkataraman (2001) finds trading costs to be lower on the NYSE (a floor market) than on the electronic limit order book of the Paris Bourse (also see Madhavan (2001)).

C. Baskets versus Individual Securities

This paper's analysis of the joint dynamics of the SPDR and the S&P sector ETFs is most strongly connected to the models of Gorton and Pennacchi (1993) and Subrahmanyam (1991), which incorporate trading in baskets and securities in the presence of asymmetric information. The equilibria 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.

II. A First Look at the Data

The sample period considered here is March 1, 2000 through May 31, 2000, chosen as the most recent available three-month period when the study was commenced. The primary data sources are the NYSE's TAQ database (for the ETF data) and the volume-price files from the CME's web site. In the TAQ data, only regular AMEX quotes were used. All price data were filtered.[3]

[3] TAQ quotes were screened to remove zero and negative spreads, and spreads greater than one dollar. TAQ trades were screened to remove exchange-identified erroneous trades and trades with nonstandard settlement. Both trades and quotes were screened for outliers using a reversal filter that removed prices that differed by more than 50 cents from a centered moving average over the nearest 10 prices. CME trades were screened visually, resulting in the removal of five extreme outliers.

A. Description of the Instruments

Table I 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.[4] 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. Trading volume in the S&P 500 sector funds is substantially lower than that in the overall index fund.

Table II 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 asset, 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 to an outside counterparty. An electronic limit order book, in contrast, favors direct interaction of customer orders without intermediation.

B. The Basis

In stock index futures studies, 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 spot value of the underlying index. This definition is also adequate in the case where S_t is the price of the index ETF, subject to two qualifications. First, ETFs are generally scaled so that prices per share are similar to those of other stocks (see Table I). Thus, S_t is to be interpreted as the ETF price scaled up by the appropriate factor. The second qualification is that the ETF includes a cash component, which 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. Thus, the no-arbitrage condition (forward-spot parity) states that $F_t - S_t = c_t$, where c_t is the cost-of-carry less the cash component of the ETF.

Model specification and interpretation depend crucially on the time series properties of c_t. To illustrate the behavior of c_t over a typically quarterly cycle, Figure 1 depicts the daily averages of the log bases (constructed at 1-minute intervals) for the three indexes. The time paths of these daily averages are dominated by rough declining trends reflecting forward-spot convergence. There are two breaks in each path. The first (March 9) reflects rollover from the March contract to the June contract. The second break (March 17)

[4] In July 2001 (subsequent to the present sample), the NYSE began trading the SPDR. It now trades all of the ETFs studied in this paper.

Table I

Summary Statistics for Exchange-traded 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		Avg. daily volume (1M sh)	Avg. AMEX vol. share	Average daily trades	Avg. 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 II

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

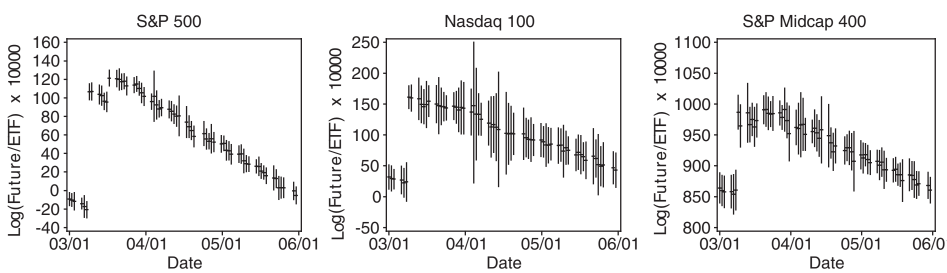

Figure 1. Average daily futures/ETF bases. The basis is defined here as $\log(F_t/S_t) \times$ 10,000 where F_t is the futures price and S_t is the price of the corresponding exchange-traded fund, as of the end of minute t. The E-mini futures contract is used for the S&P 500 and Nasdaq 100 indexes; the regular floor contract is used for the S&P Midcap 400. The sample is March 1, 2000 through May 31, 2000.

occurs when the ETF goes ex dividend.[5] The figure also demarcates (with vertical bars) a two-standard-deviation range. All graphs convey strong evidence of interday nonstationarity.[6]

Intraday, however, the bases are approximately stationary. Figure 2 depicts the bases for the three indexes at 1-minute intervals on a representative day (the first day of the sample). Each series appears to be rapidly mean reverting, with no obvious trend. 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

[5] For the first 6 days of the sample, prior to the rollover and the ex-dividend day, the S&P 500 basis is negative, that is, the spot price lies above the futures price. This is a consequence of accrued dividends in the cash component of the ETF.

[6] 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 whereas the S&P 500 had 8. Since the liquidity of the component stocks is relatively low, frequent index changes would lead to high (and variable) trading costs.

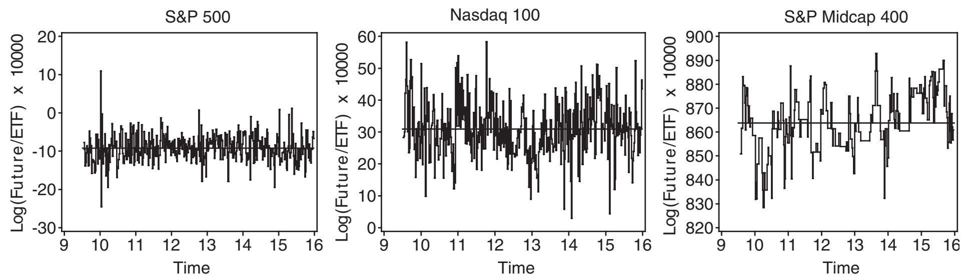

Figure 2. Intraday futures/ETF bases on March 1, 2000. The basis is defined here as $\log(F_t/S_t) \times 10{,}000$ where F_t is the futures price and S_t is the price of the corresponding exchange-traded fund, as of the end of minute t. The E-mini futures contract is used for the S&P 500 and Nasdaq 100 indexes; the regular floor contract is used for the S&P Midcap 400.

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. The specifications used to assess price discovery assume that the basis is stationary within the day. This reflects an economic assumption that the fair-value (no-arbitrage) basis is constant within the day.

III. Methodology

The analysis follows the information share approach of Hasbrouck (1995). An overview of this and alternative procedures is presented in a collection of papers: Baillie et al. (2002), de Jong (2002), Harris, McInish, and Wood (2002), Hasbrouck (2002), and Lehmann (2002). This method focuses on the random-walk components in a set of security prices, that is, the remainder after transient (presumably microstructure-related) effects have been removed. In the case where there is a single underlying security or index, the random-walk component is the same for all prices. The random-walk innovation variance is decomposed into components that can be attributed to innovations in each price. The relative contribution of a given price to this variance is defined as that price's information share.

The approach assumes that the prices are linked by one or more linear arbitrage relationships. For example, in the bivariate case where p_{1t} is the spot price and p_{2t} is the futures price, the quantity $p_{1t} - p_{2t}$ does not diverge over time. Formally, the prices are said to be cointegrated. Define $\mu = E(p_{1t} - p_{2t})$ as the mean deviation. In the parlance of time series analysis, $(p_{1t} - p_{2t}) - \mu$ is called the "error." Here, μ corresponds to the cost of carry, and the error corresponds to the arbitrage profit (before transaction costs) of buying security 1 and selling security 2. More generally, if there are n securities, there are $n-1$ linearly independent differences, and it is convenient to define these as

$$z_t = [(p_{1t} - p_{2t})\ (p_{1t} - p_{3t}) \cdots (p_{1t} - p_{nt})]'.$$

The reduced-form econometric specification underlying the analysis is a vector error correction model (VECM) of order M:

$$\Delta p_t = B_1 \Delta p_{t-1} + B_2 \Delta p_{t-2} + \cdots + B_M \Delta P_{t-M} + \gamma(z_{t-1} - \mu) + u_t, \tag{1}$$

where p_t is the column vector of prices, the B_i matrices are the autoregressive coefficients, $\gamma(z_{t-1} - \mu)$ is the error correction term, and γ is an adjustment coefficient. The VECM is relatively easy to estimate. As a reduced-form specification, however, the parameters may be difficult to interpret. It is often more useful to examine the impulse response functions, which characterize the joint price dynamics subsequent to initial shocks to the individual component prices, and the information shares.

The information shares are defined as follows. Each price in the system contains a latent random walk component (the "efficient" price). Although this component is unobservable without further identification restrictions, its innovations have the property that they are linear in the disturbances. In the bivariate case,

$$w_t = \begin{bmatrix} w_{1t} \\ w_{2t} \end{bmatrix} = A\, u_t = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{bmatrix} u_{1t} \\ u_{2t} \end{bmatrix}, \tag{2}$$

where w_{it} is the random walk innovation in the ith price, and the a_{ij} are determined from the VECM parameters. In the present case, the rows in the coefficient matrix may be shown to be identical, which implies that the random-walk innovations are identical. Intuitively, both prices reflect the same efficient price. It therefore suffices to concentrate on either one. The innovation variance for the first price is

$$Var(w_{1t}) = \begin{bmatrix} a_{11} & a_{12} \end{bmatrix} \begin{bmatrix} \sigma_1^2 & \sigma_{12} \\ \sigma_{21} & \sigma_2^2 \end{bmatrix} \begin{bmatrix} a_{11} \\ a_{12} \end{bmatrix}. \tag{3}$$

If the innovation covariance matrix is diagonal ($\sigma_{12} = 0$), then equation (3) implies a clean decomposition of the long-run variance into components explained by or attributable to each of innovations: $\sigma_w^2 = Var(w_{1t}) = a_{11}^2\, \sigma_1^2 + a_{12}^2\, \sigma_2^2$. The relative size of these contributions indicates the importance of the series. As in Hasbrouck (1995), the information share of security or market i in σ_w^2 is defined as

$$I_i = a_{ii}^2 \sigma_i^2 / \sigma_w^2. \tag{4}$$

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 upper and lower bounds for the information share. The bounds tighten as we approach the diagonal case, that is, as the magnitude of the off-diagonal covariance σ_{12} decreases. Intuitively, if the price innovations are highly correlated, it is not possible to assign explanatory power with any precision. In practical terms, this suggests employing a very fine time resolution, to avoid introducing correlation by time aggregation.

Equation (1) may readily be estimated by nonlinear least squares. The estimated parameters may then be transformed to give estimates of the impulse response functions and information shares. Various practical issues arise in the implementation, however. The next section discusses specification of the price vector. Estimation details are covered in Section V.

IV. Specification and Construction of the Price Vector

From a statistical viewpoint, the prices in equation (1) are simply the variables of interest in a prediction problem. From an economic perspective, however, the prices impound information, the price vector p_t defines the information set, and the conclusions about price discovery depend on this set. This section discusses this specification, with particular emphasis on the diverse nature of the price series.

The institutional particulars are as follows. The Globex (E-mini) and stock (ETF) markets report in real time the price and volume of each trade. The pit markets do not generally report the occurrence or volume of a trade. The price is reported only if it differs from that of the immediately preceding trade. The information for the floor-traded futures contract is therefore coarser than the E-mini/ETF information. The present analysis incorporates last-sale prices for all three securities. Thus, the volumes and times of trades (at a given price) for the ETFs and E-minis are not used.

Bid and ask quotes in the different markets are subject to various interpretations. In the Globex and ETF markets, bids and offers are promptly and widely disseminated, and remain valid until they are explicitly canceled. The analysis incorporates the midpoint of the prevailing bid and ask for the ETF, but historical bid and offer data for the Globex products are not available. In the pit, bids and offers are good "only as long as the breath is warm." Indicative (nonfirm) quotes are occasionally disseminated, when there is no trading activity and the floor reporter wishes to show an indicative price.

Wireless devices are generally prohibited on both the AMEX and CME trading floors (with the exception of exchange-provided devices for transmitting customer orders). Thus, the comprehensive collection of all current price data (bid, ask, and trade prices, oral and transcribed, both on the floors and off) defines a set that no single agent possesses, obviously a poor proxy for common public information.

In summary, then, the price vector comprises last-sale prices for the floor contract, the E-mini contract and the ETF, and the bid-ask midpoint for the ETF. Each 1-second observation reflects the current or most recent realization of each price. For example, p_t for $t = 10{:}05{:}01$ might include an ETF quote midpoint based on a bid and ask that were first set at 10:01 (but are still prevailing), an ETF trade price set at 10:10, the price of an E-mini trade posted at 10:02, and a price on the pit-traded futures contract posted at 10:03.

This set of prices captures key summary data that any market participant might easily possess. It is obviously not comprehensive. Each of the markets produces information supplementary to the price. To the extent that these data are incrementally informative about future price movements, the information shares computed from a broader information set might differ from the present estimates.

The criteria for including variables in the information set are availability and relevance. In particular, there is no presumption that market participants limit themselves to one type of price data (say, quotes). Furthermore, efficient use of the information would not normally require that diverse price data be transformed to make the price series in some sense comparable. As a trivial example, although price increments differ on the various markets, one would not normally restate prices to artificially impose a uniform tick size.

More substantive concerns might arise in connection with the use of last-sale and quote midpoint prices in the same specification. The former presumably include a bid-ask bounce component not present in the latter. It would obviously not be appropriate to render the series more alike by adding a simulated bid-ask bounce to the quote midpoint, as this would degrade the information set. If any transformation is to be attempted, it would be in the direction of prefiltering the last-sale price to purge it of estimated bid-ask bounce. Univariate prefilters of this sort have been used by Stoll and Whalley (1990), Grunbichler, Longstaff, and Schwartz (1994), Fleming et al. (1996), and Pizzi, Economopoulos, and O'Neill (1998), among others.

A useful starting point for evaluating the impact of such a procedure is to consider the filtering already performed in the VECM specification equation (1). Implicit in this specification is a forecast of the future security price that is optimal within the class of linear functions of current and past price data. That is, within this class the specification extracts the maximal useful information from each of the component series. A univariate prefiltering aimed at removing estimated bid-ask bounce that is implemented as a linear function of current and past prices can at best leave unchanged (and at worst degrade) the information in the series.

The restriction of the prefiltering to current and past data is a substantive one. In the Roll (1984) model of the bid-ask spread, for example, the estimate of the underlying efficient price at a given time is made more precise by conditioning on future trades (in addition to those current and past). Incorporation of these future trades into the analysis, however, would enrich the information set beyond what agents would normally possess. The restriction to the class of linear models may also be substantive. There are many structural models for which linear forecasts are suboptimal.[7] These models are generally stylized ones, however, and the forecast improvement is contingent on correct specification. The linear VECM specification, in contrast, makes minimal statistical assumptions on the data.

[7] In Garbade and Lieber (1977), for example, the underlying security price follows a normal diffusion, trade prices contain Gaussian bid-ask bounce, and trade occurrence is modeled as an independent Poisson arrival process. The optimal forecast here is given by a continuous-time Kalman filter.

V. Estimation Details

A. Estimation of the Mean Cointegrating Vector, μ

In the present application, μ in equation (1) corresponds to the no-arbitrage basis. Section II.B noted that for economic and institutional reasons, this should be essentially constant within the day. If the intraday basis exhibits rapid reversion to μ in-sample, then its intraday sample average is a good estimate of μ. This is generally the case (see Figure 2, for example). Thus, the principal approach is simply to estimate (1) separately for each day, and report summary statistics of these daily estimates. As discussed below, this also facilitates computation of distributional results.

The standard deviations plotted in Figure 1, though, suggest that on certain days the basis is quite volatile. On such days, the sample mean may be a poor estimate of the no-arbitrage basis. As an alternative to the separate daily estimates, therefore, a single estimation was made over the full sample in which μ was parameterized as a trending variable subject to contract rollover and dividend payment dummies. Specifically for day d, $\mu_d = \alpha_0 + \alpha_1$ *Dummy* (March 9, 2000) $+ \alpha_2$ *Dummy* (March 17, 2000), for $d = 1, \ldots, 63$. This model produces smoother estimates for the μs. Since it imposes more structure on the data, however, misspecification is a larger concern. Among other things, breaks in the trend might result from index reconstitutions. In any event, the point estimates of the information shares in this second approach are similar to those in the first.[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

[8] In the VECM stock index futures literature, one common practice involves estimating μ as the average basis over the entire multiday sample (Pizzi et al. (1998), Booth, So, and Tse (1999), and Chu et al. (1999)). On any given day, however, this average has no particular connection to the no-arbitrage condition. It is therefore difficult to see why this should be a plausible target of price adjustment. In connection with Booth et al. (1999), Yiuman Tse writes (October 2, 2002, 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 of daily estimates. For example, let $\hat{I}_i^d$ denote the information share of i, that is, the contribution of price i's innovations to the long-run variance of price i, implied by the model estimated on day d's sample. The overall sample estimate of I_{ij} is simply the mean $\hat{I}_{ij} = \left(\sum_d \hat{I}_{ij}^d\right)/N^d$ where N^d is the number of days in the sample ($N^d = 64$). 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. Construction of an overall sample estimate as the mean of subsample estimates is an approach familiar from Bartlett (1950) and Fama and MacBeth (1973).

C. Reducing the Size of the Parameter Set

The VECM specification (1) 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 1 second, and lagged terms up to 5 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 relies on polynomial distributed lags (Greene (1993) and Hasbrouck (1995)). Using this technique, the coefficients of lagged price changes $\Delta p_{i,t-k}$ in the equation for $\Delta p_{i,t}$ are constrained to lie on segments that are polynomials in the lag k.[9]

VI. 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, last-sale price for the ETF, last-sale price for regular floor-traded futures contract, and the last-sale price for the E-mini. The specification was estimated at a 1-second level of resolution, with lags through 5 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 3 depicts the cumulative impulse response functions, that is, cumulative price impacts implied by initial perturbations. Following the paper's general practice, each point is an average of daily estimates.

[9] For the S&P 500 and Nasdaq-100 estimations, second-degree (quadratic) polynomial distributed lags (PDLs) were used over lags 1–10 (seconds), 11–20, and 21–30 and zero-degree (constant) coefficients were imposed on lags 31–60, 61–120, and 121–300. For the S&P 400 estimations (5-second intervals), second-degree PDL's were imposed on lags 1–6, 7–12, 13–30, 31–60, 61–90, and 91–120; constant coefficients were imposed on lags 121–150, 151–180, and 181–300.

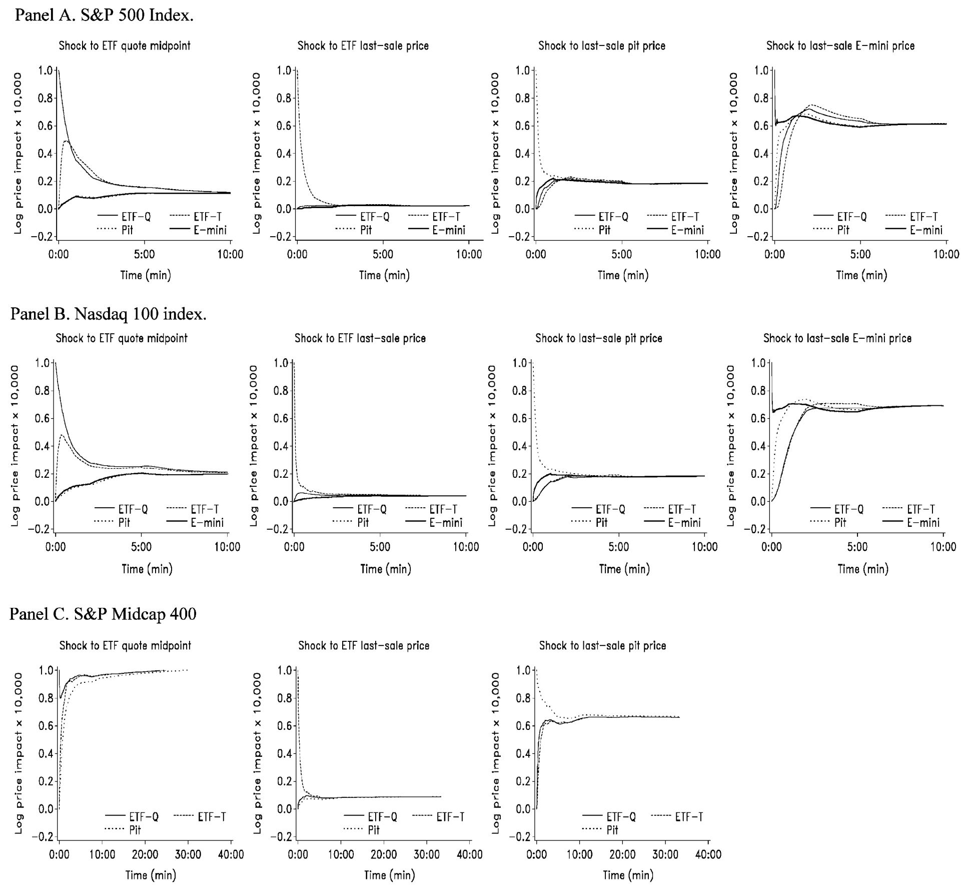

Figure 3. Cumulative impulse response (price impact) functions. The figures plot impulse response functions for prices of index securities. The estimates are based on vector error correction models of ETF Quote Midpoints (ETF-Q), ETF last-sale transaction prices (ETF-T), and last-sale prices for pit-traded and E-mini futures contracts. The models are estimated with 1-second resolution for S&P 500 and Nasdaq 100 indexes; and with 5-second resolution for the S&P 400 index.

Panel A of Figure 3 depicts the impulse response functions for the S&P 500. Each graph corresponds to an initial shock in a different security (indicated at the top of the graph). Within a given graph, the four lines track the implied cumulative price changes for these securities. The first graph tracks the cumulative impact of a unit shock in the ETF quote midpoint. By construction at $t = 0$, the impact is unity in the ETF quote midpoint, and zero in the other three securities. In the long run (10 minutes), the impact is essentially identical for all securities. This convergence is a consequence of cointegration imposed by the specification. The transient behavior indicates that the ETF quote shock is initially reflected most strongly in the ETF last-sale price, but both of the ETF price variables rapidly revert. Price adjustments in both futures contracts are small and protracted. A shock to the ETF last-sale price quickly dies out, with only a small persistent impact on the other prices (second graph).

The finding that ETF quotes appear to slightly lead the trade prices may reflect reporting delays for the transactions. The result implies that an off-floor observer can glean more timely information from the quotes than from the trades. This finding need not imply that that quotes cause trades in a functional sense.

A shock to the last-sale pit contract futures price (third graph) implies similar behavior, although with somewhat larger persistent impact. The most distinctive behavior arises subsequent to a shock in the E-mini last-sale price. The reversion

Table III

Price Discovery in the S&P 500 Index

Statistics are based on a vector error correction model of prices for index securities estimated at a 1-second resolution. The ETF trade price, pit contract price, and E-mini contract prices are last-sale prices. The model is estimated for each day in the sample (March 1, 2000 through May 31, 2000, 64 trading days). The table reports summary statistics for these daily estimates. 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.

Panel A: Disturbance Correlation Matrix

	ETF Quote Midpoint	ETF Trade Price	Pit Contract Price	E-mini Contract Price
ETF quote midpoint	1.000	0.010	0.003	0.002
ETF trade price	0.010	1.000	0.000	0.000
Pit contract price	0.003	0.000	1.000	0.052
E-mini contract price	0.002	0.000	0.052	1.000

Panel B: Coefficients of Efficient Price (Common Index Price Factor)

ETF Quote Midpoint	ETF Trade Price	Pit Contract Price	E-mini Contract Price
0.119	0.022	0.184	0.615

Panel C. Information Shares

	Source							
	ETF Quote Midpoint		ETF Trade Price		Pit Contract Price		E-mini Contract Price	
	Min	Max	Min	Max	Min	Max	Min	Max
Median	0.006	0.007	0.005	0.007	0.045	0.065	0.909	0.933
Mean	0.022	0.024	0.008	0.009	0.056	0.080	0.887	0.913
SEM	0.006	0.006	0.001	0.001	0.006	0.007	0.009	0.008
Std. dev.	0.048	0.050	0.009	0.010	0.046	0.056	0.069	0.062

is relatively small, and the persistent impact is much larger than those subsequent to shocks in the other prices.

Table III presents estimates relating to the information shares. Panel A describes the correlations of the disturbances. At 1-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 (2). These are identical to the amplitudes of the long-run cumulative price impacts shown in Figure 3.

Panel C in Table III 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 nonzero off-diagonal correlations in the innovations, the information shares can be estimated only up to linear 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 prices share the remainder, with the pit-traded futures contract accounting for slightly more than either ETF price series (or their total contribution).

B. Price Discovery in the Nasdaq-100 Index

An identical specification was estimated for the Nasdaq-100 securities. The impulse response functions (graphed in Panel B of Figure 3) are similar to those of the S&P 500 set. Innovations in the ETF quote midpoint, the ETF last-sale price, and the last-sale price of the pit contract are swiftly and substantially reversed, while most of an innovation in the E-mini last-sale price is permanent. Table IV reports statistics related to the information shares for the Nasdaq-100 index. As with the S&P 500, the E-mini last-sale price dominates, providing over 80% of the price discovery. In contrast to the S&P 500, however, the contributions of the ETF quote midpoint and last-sale price for the pit contract 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 quote midpoint, the ETF last-sale price, and the last-sale price for the pit contract. It is estimated with 5-second resolution and encompasses lags through 10 minutes. The increase in lag length (relative to the 1-second estimates) is motivated by the somewhat slower decay in the basis. The decrease in resolution (5 seconds is coarse relative to 1 second) is necessitated by computational limitations.

The results for the S&P 400 index differ markedly from those of the other two indexes. Panel C of Figure 3 depicts the cumulative impulse response functions. These show that unit shocks in both the ETF quote midpoint and the last-sale price of the pit contract have substantial long-run impacts, while a shock to the ETF last-sale price has only a minor effect. Table V reports statistics related to the information shares. The contribution from the ETF quote midpoint is the largest (approximately 50%). The sum of the quote-midpoint and last-sale information shares for the ETF (a rough indication of the total contribution from this

Table IV

Price Discovery in the Nasdaq-100 Index

Statistics are based on a vector error correction model of prices for index securities estimated at a 1-second resolution. The ETF trade price, pit contract price, and E-mini contract prices are last-sale prices. The model is estimated for each day in the sample (March 1, 2000 through May 31, 2000, 64 trading days). The table reports summary statistics for these daily estimates. 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.

Panel A: Disturbance Correlation Matrix

	ETF Quote Midpoint	ETF Trade Price	Pit Contract Price	E-mini Contract Price
ETF quote midpoint	1.000	0.027	0.001	0.002
ETF trade price	0.027	1.000	0.002	0.002
Pit contract price	0.001	0.002	1.000	0.031
E-mini contract price	0.002	0.002	0.031	1.000

Panel B: Coefficients of Efficient Price (Common Index Price Factor)

ETF Quote Midpoint	ETF Trade Price	Pit Contract Price	E-mini Contract Price
0.212	0.041	0.180	0.690

Panel C: Information Shares

	Source							
	ETF Quote Midpoint		ETF Trade Price		Pit Contract Price		E-mini Contract Price	
	Min	Max	Min	Max	Min	Max	Min	Max
Median	0.023	0.027	0.010	0.012	0.044	0.056	0.884	0.900
Mean	0.043	0.047	0.020	0.024	0.047	0.061	0.870	0.887
SEM	0.007	0.007	0.004	0.004	0.004	0.005	0.009	0.008
Std. dev.	0.053	0.056	0.030	0.033	0.034	0.038	0.069	0.067

security) is slightly above 50%. 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 provide the bulk of the price discovery in the sample. 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. In fact, the sums of commissions and half-spreads for a typical small trade in each security were similar

Table V

Price Discovery in the S&P 400 Index

Statistics are based on a vector error correction model of prices for index securities estimated at a 5-second resolution. The ETF trade price, pit contract price, and E-mini contract prices are last-sale prices. The model is estimated for each day in the sample (March 1, 2000 through May 31, 2000, 64 trading days). The table reports summary statistics for these daily estimates. 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.

Panel A: Disturbance Correlation Matrix

	ETF Quote Midpoint	ETF Trade Price	Pit Contract Price
ETF quote midpoint	1.000	0.038	0.015
ETF trade price	0.038	1.000	0.006
Pit contract price	0.015	0.006	1.000

Panel B: Coefficients of Efficient Price (Common Index Price Factor)

ETF Quote Midpoint	ETF Trade Price	Pit Contract Price
1.010	0.089	0.661

Panel C: Information Shares

	Source					
	ETF Quote Midpoint		ETF Trade Price		Pit Contract Price	
	Min	Max	Min	Max	Min	Max
Median	0.484	0.514	0.035	0.046	0.418	0.440
Mean	0.474	0.497	0.065	0.079	0.435	0.450
SEM	0.038	0.038	0.010	0.011	0.040	0.040
Std. Dev.	0.303	0.307	0.078	0.085	0.316	0.317

(about $20 for one floor contract, one E-mini contract, or 100 shares of the ETF). Relative to the value of the underlying asset, however, the floor contract (with its large size) is by far the cheapest to trade. Consideration of other trading cost components, however, 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.

It was noted earlier that the Globex electronic system on which the E-minis trade is functionally similar to the ECNs found in equity markets. The strong informational role of the E-mini found in the present study therefore parallels the finding of Huang (2002) that ECN quotes in Nasdaq stocks are more informative than the quotes posted by dealers. On both the ECNs and Globex system, speed of execution may enhance the price discovery function.

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. (2002)).

The S&P 400 Midcap index market differs substantially from other two index markets. There is no E-mini contract in the sample, 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. 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.

Third, the timing of an event in the analysis reflects when it entered the public information set, not when it (or a necessary precursor event) actually occurred. The time stamp on a last-sale price, for example, reflects when the sale was reported, not when it occurred, when the order was received by the exchange, or when the order was generated by the customer.

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 S&P 500 and Nasdaq-100 results to timing, time-stamp errors were artificially induced by lagging the E-mini price reports by 0, 5, 10, and 15 seconds. At a lag of 10 seconds, the S&P 500 E-mini lost dominance (to the regular floor contract). For the Nasdaq-100, the E-mini lost dominance only with a lag of 15 seconds. Thus, 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 (5 or 10 seconds) errors.

VII. The S&P 500 and Sector ETFs

The paper turns now to price discovery among the S&P 500 ETF and the sector ETFs. The overall econometric approach described in Section III is still appropriate, but requires modification. Specifically, the price vector has 10 components: $p_t = [p_t^{SPDR} \quad p_t^{Sector'}]'$, where p_t^{Sector} is the (9×1) vector of sector portfolio prices (quote midpoints). Replication implies that "on average" $p_t^{SPDR} = a\,p_t^{Sector}$, where a is a vector of replication weights. The system therefore has one cointegrating vector given by $p_t^{SPDR} - a\,p_t^{Sector}$. Moreover, there is no longer a single random-walk component common to all prices. Corresponding to equation (2), we have for this system a separate random-walk component for each security. In this system, the information share I_{ij} is defined as the proportion of the ith security's random-walk innovation variance $(Var(w_{it}))$ that is explained by the innovations in security j.

The econometric specification in this section is a time-series model of the full set of price variables, that is, 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

Table VI

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} and so forth are closing quote midpoints. (The R^2 of this regression was essentially unity.) The dollar components of the *SPY* are computed as $a_{XLB}\,\bar{p}^{XLB}$ and so forth, where $\bar{p}^{XLB}$ is the sample average price of the *XLB* fund. The percentage components of the *SPY* are computed as $a_{XLB}\,\bar{p}^{XLB}/\bar{p}^{SPY}$ and so forth. The R^2 in the regression is 1.000.

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

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 VI 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.[10] (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.

[10] 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.

Table VII

Price Discovery between the S&P 500 and Its Component Sector Portfolios

Table entry $I_{i,j}$ (ith row and jth 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, that is, the entries are approximately the maximal information shares. The specification is a vector error correction model estimated with 5-second resolution. The specification includes lags through 5 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 VII presents summary statistics on price discovery. 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. The first row decomposes price discovery in the broad index (*SPY*). Approximately 80% of the price 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 VI). 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.

To check whether the price discovery contributions merely reflected the long-term return dependencies, I examined the correlations across ETFs in daily quote-midpoint returns. The correlation between *SPY* and *XLK* is substantial (0.811), but larger still is the correlation between *SPY* and *XLI* (0.902). The *XLI* fund, however, contributes very little to the price discovery in *SPY* (about 1.9%, from the top row of Table VII).

In light of mass media reports that over the sample period the technology sector was often identified 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 the estimates based on a joint VAR analysis of the *SPY* and *QQQ* fund price changes found that the latter's contribution to price discovery in the former was at most also about 10%.

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 I are not especially high, trade volumes are often low.

One hypothesis is that the sector funds may have few natural liquidity traders. Following the intuition of Subrahmanyam (1991), discretionary liquidity traders would view the sector funds as substitutes not only for the individual component securities, but also for 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.[11]

VIII. 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 reexamine 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.

[11] In comparing the strategies of holding the SPDR and alternatively replicating it with the sector ETFs, the latter strategy would incur higher trading costs. It would, however, offer tax advantages, since there would be more opportunities for tax-loss selling. 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 reliable information.

For the S&P 500 and Nasdaq-100 indexes, the Chicago Mercantile Exchange has introduced E-minis. The E-minis have a 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.

REFERENCES

Admati, Anat, and Paul Pfleiderer, 1988, A theory of intraday trading patterns: Volume and price variability, *Review of Financial Studies* 1, 3–40.

Baillie, Richard T., G. Geoffrey Booth, Yiuman Tse, and Tatyana Zabotina, 2002, Price discovery and common factor models, *Journal of Financial Markets* 5, 309–322.

Bartlett, Maurice S., 1950, Periodogram analysis and continuous spectra, *Biometrika* 37, 1–16.

Booth, G. Geoffrey, Raymond W. So, and Yiuman Tse, 1999, Price discovery in the German equity index derivatives markets, *Journal of Futures Markets* 19, 619–643.

Choi, Hong, and Avanidhar Subrahmanyam, 1994, Using intraday data to test for effects of index futures on the underlying stock markets, *Journal of Futures Markets* 14, 293–322.

Chu, Quentin C., Wen-liang Gordon Hsieh, and Yiuman Tse, 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, Mark, and Ian Domowitz, 1999, Pricing behavior in an off-hours computerized market, *Journal of Empirical Finance* 6, 583–607.

de Jong, Frank, 2002, Measures of contributions to price discovery: A comparison, *Journal of Financial Markets* 5, 323–328.

Domowitz, Ian, and Benn Steil, 1999, Automation, trading costs and the structures of the securities trading industry, in Robert E. Litan, Anthony M. Santomero, eds., *Brookings-Wharton Papers on Financial Services*, (Brookings Institution Press, Washington DC), 33–92.

Elton, Edwin J., Martin J. Gruber, George Comer, and Kai Li, 2002, Spiders: Where are the bugs? *Journal of Business* 75, 453–472.

Fama, Eugene F., and James D. MacBeth, 1973, Risk, return and equilibrium, *Journal of Political Economy* 71, 607–636.

Fleming, Jeff, Barbara Ostdiek, and Robert E. Whalley, 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. Douglas, and S. Viswanathan, 1990, A theory of the interday variations in volume, variance, and trading costs in securities markets, *Review of Financial Studies* 3, 593–624.

Garbade, Kenneth, and Zvi Lieber, 1977, On the independence of transactions on the New York Stock Exchange, *Journal of Banking and Finance* 1, 151–172.

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

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

Grunbichler, Andreas, Francis A. Longstaff, and Eduardo S. Schwartz, 1994, Electronic screen trading and the transmission of information: An empirical examination, *Journal of Financial Intermediation* 3, 166–187.

Harris, Frederick H. deB., Thomas H. McInish, and Robert A. Wood, 2002, Common factor components vs. information shares: A reply, *Journal of Financial Markets* 5, 341–348.

Harris, Lawrence E., George Sofianos, and James E. Shapiro, 1994, Program trading and intraday volatility, *Review of Financial Studies* 7, 653–685.

Hasbrouck, Joel, 1995, One security, many markets: Determining the contributions to price discovery, *Journal of Finance* 50, 1175–1199.

Hasbrouck, Joel, 2002, Stalking the efficient price in empirical microstructure specifications, *Journal of Financial Markets* 5, 329–339.

Huang, Roger D., 2002, The quality of ECN and market maker quotes, *Journal of Finance* 57, 1285–1319.

Kawaller, Ira G., Paul D. Koch, and Timothy W. Koch, 1987, The temporal relationship between S&P 500 futures and the S&P 500 index, *Journal of Finance* 42, 1309–1329.

Kumar, Praveen, and Duane J. Seppi, 1994, Information and index arbitrage, *Journal of Business* 67, 481–509.

Lehmann, Bruce N., 2002, Some desiderata for the measurement of price discovery across markets, *Journal of Financial Markets* 5, 259–276.

MacKinlay, A. Craig, and Krishna Ramaswamy, 1988, Index-futures arbitrage and the behavior of stock index futures prices, *Review of Financial Studies* 1, 137–158.

Madhavan, Ananth, 2001, Discussion of "Automated vs. floor trading", *Journal of Finance* 56, 1485–1488.

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

Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* 39, 1127–1139.

Stoll, Hans R., and Robert E. Whalley, 1990, The dynamics of stock index and stock index futures returns, *Journal of Financial and Quantitative Analysis* 25, 441–468.

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

Venkataraman, Kumar, 2001, Automated vs. floor trading: An analysis of execution costs on the Paris and New York Exchanges, *Journal of Finance* 56, 1445–1485.

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