Technical Analysis and Liquidity Provision

Kenneth A. Kavajecz
University of Wisconsin–Madison

Elizabeth R. Odders-White
University of Wisconsin–Madison

The apparent conflict between the level of resources dedicated to technical analysis by practitioners and academic theories of market efficiency is a long-standing puzzle. We explore a previously unexamined feature of technical analysis — namely its relation to liquidity provision. We demonstrate that support and resistance levels coincide with peaks in depth on the limit order book and moving average forecasts reveal information about the relative position of depth on the book. Furthermore, we show that these relationships stem from technical rules locating depth already in place on the limit order book.

There is an ongoing belief in the investment community that technical analysis, which involves the study of past price and volume data, can be used to infer the direction of future prices. A substantial segment of the investments industry is dedicated to this form of analysis, with virtually all investment banks and trading firms employing some technical trading strategies.[1] Many different types of technical indicators are used in practice, including candlestick charts, moving average forecasts, point and figure charts, and support and resistance levels. Technical analysts view these measures as a way to capitalize on identifiable price patterns that stem from investors' repetitive behaviors.

Despite its popularity among practitioners, academics have historically dismissed technical analysis. They argue that it is inconsistent with the theory of market efficiency, which states that all available information must be reflected in security prices. Much of the finance literature relies on the assumption of market efficiency because, in its absence, investors could earn excess profits while assuming little or no risk, a state that is not sustainable in equilibrium.

The traditional academic wisdom and investment community practices appear to be at odds. In hopes of resolving this conflict, many

We gratefully acknowledge helpful comments from seminar participants at the Society for Financial Studies' Conference on Investments in Imperfect Capital Markets, Rutgers University, University of Georgia, and University of Pennsylvania as well as Simon Gervais, David Hsieh, Pete Kyle (discussant), Craig MacKinlay, Maureen O'Hara (editor), Carol Osler, Mark Ready, and an anonymous referee. All remaining errors are our own. Address correspondence to Kenneth A. Kavajecz, University of Wisconsin–Madison, 975 University Ave., Madison, WI 53706, or e-mail: kkavajecz@bus.wisc.edu.

[1] See also the International Federation of Technical Analysts (www.ifta.org) and the affiliated country associations.

doi:10.1093/rfs/hhg057 Advance Access publication March 26, 2004

researchers have undertaken studies of technical analysis. Some have found results consistent with the practitioners' view by providing evidence that technical analysis can predict price movements or by developing models of (inefficient) markets in which investors benefit from conditioning on historical information. For example, Neftci and Policano (1984), Neftci (1991), Brock, Lakonishok, and LeBaron (1992), Neely, Weller and Dittmar (1997), and Lo, Mamaysky, and Wang (2000), among others, test different technical trading rules and find evidence consistent with technical analysis providing information beyond that already incorporated into the current price. In theoretical work, Treynor and Ferguson (1985) and Brown and Jennings (1989) examine settings in which privately informed investors use past prices to determine whether their information has been revealed to the market and to learn about the private signals of other traders, respectively. Similarly Blume, Easley, and O'Hara (1994) demonstrate that volume may provide relevant information if prices do not react immediately to new information. Furthermore, there is a growing literature examining the possibility that common biases in human judgment lead to market inefficiencies [see, e.g., DeBondt and Thaler (1985) and Barberis, Shleifer, and Vishny (1998)].

Others provide results consistent with the traditional academic argument by demonstrating empirically that technical analysis does not predict future prices. For example, Allen and Karjalainen (1999) and Ratner and Leal (1999) find little support for the technical trading rules they examine, and Fama and Blume (1966), Bessembinder and Chan (1998), and Ready (2002) show that transaction costs offset the benefits of technical analysis found by Alexander (1961) and Brock, Lakonishok, and LeBaron (1992). Furthermore, Jensen and Bennington (1970), Sullivan, Timmermann, and White (1999), and Jegadeesh (2000) warn that data-snooping and survivorship biases can be severe when evaluating technical rules, which can lead researchers to falsely conclude that technical trading strategies can predict future price movements. Given this mixed evidence, it is perhaps not surprising that academics have yet to agree on one of the two explanations discussed above.

In this article we take a different approach and explore a previously unexamined feature of technical analysis — namely its relation to liquidity provision. Rather than attempting to link technical trading rules directly to future transaction prices, as is typically done in the literature, we claim that technical analysis measures capture changes in the state of the limit order book. While transaction prices are determined by the interplay between those supplying liquidity and those demanding immediacy, here we focus solely on one side of this relation — the supply of liquidity through limit orders. The difference between liquidity provision and transaction prices may partially explain the tenuous conclusions in the previous literature on technical analysis.

We test two specific hypotheses regarding the link between common technical trading strategies and limit order book liquidity. The first hypothesis exploits the natural similarities between support and resistance levels and limit prices corresponding to high depth on the book. Limit order books, which organize outstanding limit orders according to price/time priority, are an important source of liquidity within many markets. Kavajecz (1999) has shown that these supply and demand schedules are rarely smooth. Instead, there are often a few limit prices that contain a disproportionate number of shares or “peaks” of liquidity. When there are many shares standing ready at a particular limit price to satiate liquidity demanders’ trading needs, this price is likely to become an obstacle to achieving more extreme prices. Likewise, support (resistance) levels occur when substantial selling (buying) pressure is required to breach a given price level. Consequently we hypothesize that technical analysis support and resistance levels are related to peaks in liquidity on the limit order book.

The second hypothesis states that moving average indicators reveal information about the relative position of depth on the limit order book. Moving average indicators measure momentum by comparing a long-run and short-run moving average of prices. If the short-run moving average rises above the long-run average, for example, then technicians believe that prices are moving into a new, higher trading range. In this case we would expect the quoted midpoint to grow closer to the peak in liquidity on the sell side of the book as prices continue to rise. At the same time, the distance between the quoted midpoint and the peak on the buy side of the limit order book would increase as the quotes move higher. In other words, the technical analysis moving average indicator should provide a sense of the “skewness” of liquidity on the two sides of the book, relative to the quoted prices.

Consistent with the first hypothesis, our results show that technical analysis support and resistance levels have a significant ability to explain limit prices with high cumulative depth even after controlling for the state of the trading environment. We investigate the underlying source of this relationship and conclude that it is tied to traders’ order placement strategies. Specifically, the results of Granger-causality tests and analyses of limit order placement activity suggest that technical analysis measures discover pockets of depth already in place on the limit order book. These clusters of orders may have been placed by informed traders who have strong opinions as to the asset’s value, or by liquidity traders who congregate on particular prices when trying to manage adverse selection and execution risk. In either case, the findings suggest that the connection between technical analysis and liquidity stems from technical levels detecting changes in the limit order book.

In support of the second hypothesis, we demonstrate that moving average indicators explain changes in the relative position of cumulative depth on the limit order book, even after controlling for the trading environment factors. As expected, the extreme buying pressure that triggers a positive moving average signal causes the quotes to "walk up" the book, as lower-priced sell orders are either executed or canceled. This result provides additional evidence of the strong connection between technical analysis and liquidity provision.

This study represents a crucial step toward deepening our understanding of technical analysis by demonstrating its connection to liquidity provision. The identification of this link not only offers insight into the foundation of technical analysis, but also provides information that may help traders reduce transaction costs. While mounting evidence suggests that market participants can benefit from knowing the composition of the limit order book, this information is often unavailable or difficult to convey in real time. Our results show that technical analysis provides a method of obtaining information about the limit order book using readily available data. Although we focus solely on the New York Stock Exchange (NYSE), we believe that this connection extends beyond limit order book markets to other trading structures, for example, foreign currency markets.[2] Whether liquidity is provided explicitly through limit orders or implicitly through other order placement strategies, we believe technical analysis is related to the provision of liquidity.

The remainder of the article is organized as follows. Section 1 describes the data used in the analysis, as well as the limit order book estimation technique. Section 2 describes the limit order book liquidity measures we employ. Section 3 defines the technical analysis rules we consider. Sections 4 and 5 present and interpret our results. Section 6 concludes.

1. Data and Limit Order Book Estimation

We construct estimates of limit order books using NYSE SuperDOT data, provided by the NYSE. These data contain quote, transaction, and order information for 110 NYSE stocks from July 1997 through September 1997.[3]

Estimates of the limit order books are constructed from the order data using the technique described in Kavajecz (1999). The principle behind the limit order book estimation is that at any instant in time, the limit order

[2] In a related study, Osler (2002) examines the relationship between order placement activity — specifically the clustering of stop-loss and take-profit orders — and the success of technical analysis in the foreign exchange market.

[3] Our dataset consists of the 110 of the 144 original TORQ stocks still traded on the NYSE as of November 1996. We thank the NYSE for the use of these data. For more information on the TORQ dataset, see Hasbrouck (1992).

book should reflect those orders remaining after netting all prior execution and cancellation records. The first step in the estimation is to search the entire dataset for records that have order arrival dates prior to July 1, 1997. We use this set of limit orders to form an estimate of the initial limit order book (or "prebook") at the start of the NYSE order dataset. After the prebook is constructed, current records are processed. To estimate the limit order book for a given date and time, all records with a date and time stamp prior to the chosen date and time are selected, new orders are added to the prebook, and execution and cancellation records are matched to existing orders on the limit order book, where matched orders are eliminated. The remaining or residual set of orders becomes our estimate of the limit order book for the chosen date and time. By sequentially updating the limit order book estimates, we create a sequence of "snapshots" of the limit order books for each stock. These estimates are calculated at the opening quote for each stock and then at half-hour intervals thereafter, generating 14 limit order book estimates each day for each stock.

2. Limit Order Book Liquidity Measures

We consider four specific measures of cumulative depth on each side of the limit order book. Each measure captures somewhat different aspects of liquidity provision. The first two measures reflect the absolute location of depth on the book, while the second two measures reflect depth relative to average transaction volume. The first measure, called the "mode," captures the price on each side of the limit order book at which there is a large increase in cumulative depth. Specifically, this measure is created by considering the two limit prices with the most shares on each side of the book and selecting the limit price that is closer to the quoted bid (ask) on the buy side (sell side) as the limit order book mode. Selecting the price that is closer to the quote on each side ensures that our liquidity measure reflects current trader interest rather than picking up large, stale orders that are sitting far from the current quoted bid and ask.[4] The limit order book mode measure implicitly represents the first great obstacle or, in other words, the most significant support or resistance level on the book.

Our second measure of limit order book liquidity, labeled "near depth," captures how concentrated depth is near the quoted prices. For each stock, we create a band around the quoted midpoint, where the bandwidth is the product of the standard deviation of half-hour returns and the quoted midpoint, multiplied by 10.[5] On the buy side of the limit order

[4] We conduct robustness checks using the actual mode rather than the closer of two, and the results are similar. (See the appendix for a more detailed discussion of all of our robustness checks.)

[5] We also conducted the analysis using five as the multiplier, and the results are similar.

book, our near depth measure is computed as the total depth within one-half of the bandwidth below the quoted bid price (i.e., between bid and bid-$\frac{1}{2}$bandwidth) divided by the total depth within a full bandwidth of the quoted bid price. The sell-side measure is computed analogously. Values close to one indicate that depth is concentrated near the quotes, while measures near zero imply that depth is dispersed over a wider range of limit prices. One possible interpretation of this variable is the degree of consensus among liquidity providers concerning the appropriate buy and sell limit prices.

Measures three and four reflect cumulative depth on the limit order book relative to trading volume. Our third measure is defined as the highest limit buy price and the lowest limit sell price with a cumulative depth of at least 5% of the average daily trading volume for the given stock. As such, it reflects the last share price that an investor would pay when buying (receive when selling) 5% of the average daily volume in that stock from the limit order book.[6] We refer to this measure as "5% ADV."

Our fourth measure of limit order book liquidity, called "90th ATP," is also related to transaction volume, but differs from the third in two ways. First, it uses the 90th percentile of the distribution of trade sizes for each stock rather than 5% of average daily volume. Second, it is computed as the average price paid for the shares rather than the price of the final share.[7]

Figure 1 provides an illustration of the relationships among the quoted midpoint and our four limit order book liquidity measures. The chart covers a one-week period, August 25–29, 1997, for Dresser Industries. Panel A displays our first measure, the mode, along with the quoted midpoint. Notice that the limit order book mode variables are often close to the quoted prices, but also have discrete jumps. As a result, the difference between the current quote and the position of the limit order book mode sometimes becomes very large.

Panel B displays the near depth ratio. Notice that this is the only measure that is specified as a ratio instead of a price. While by construction this variable is almost always defined, it also has the potential to be extremely volatile, as displayed by Dresser's ratios.

Panel C displays the 5% ADV measure of limit order book liquidity. It is important to note that the volume-based cumulative depth measures are not always defined. In other words, the limit order book sometimes

[6] One way to interpret the choice of 5% is that, according to the best available estimates [TORQ (1990–1991)], SuperDOT orders constitute approximately 30% of the total order flow by volume. If we assume that this activity is evenly split between the buy and sell sides of the limit order book, then 5% of average daily volume represents one-sixth of a typical day's order flow on each side of the limit order book (assuming volume is one-half of total order flow). We also use 15% of average daily trading volume and we obtain qualitatively and quantitatively similar results.

[7] We thank an anonymous referee for suggesting the near depth and 90th ATP measures.

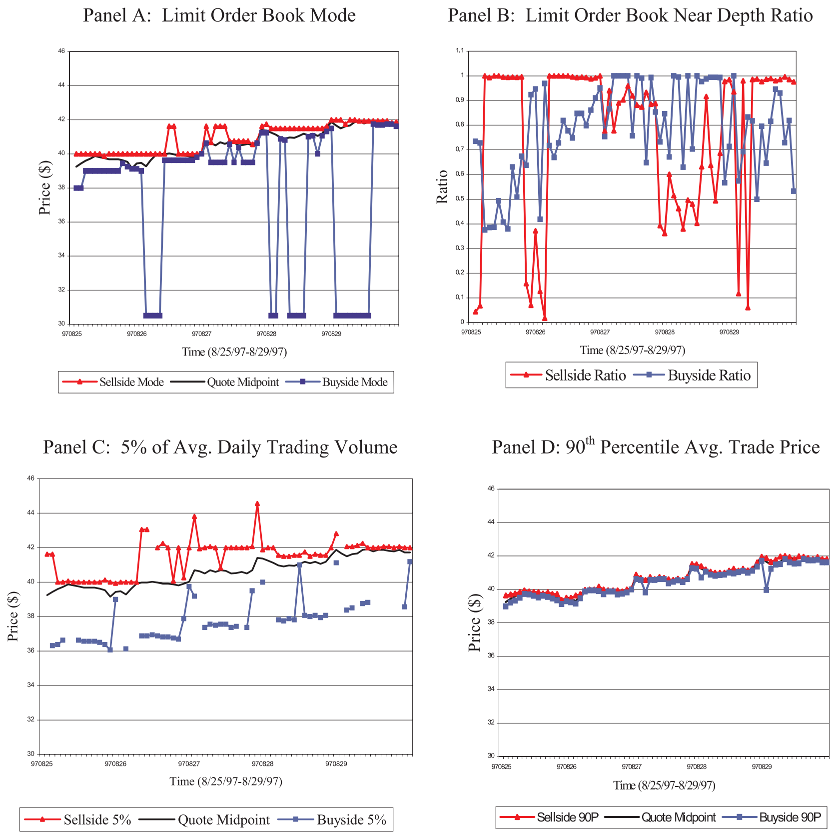

Figure 1
Position of limit order book depth measures for Dresser Industries
Buy-side and sell-side modes (panel A) represent the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. The buy-side (sell-side) near depth ratio (panel B) represents the ratio of buy-side depth within one-half of a stock-specific band below (above) the quoted midpoint to total buy-side (sell-side) depth within the stock-specific band. Five percent of average daily trading volume (panel C) is defined as the last share price from the limit order book that an investor would pay when buying or receive when selling a quantity equal to 5% of the average daily trading volume. The 90th percentile trade size is the average share price of executing a trade equal to the 90th percentile trade size for that stock.

contains so few shares that it is impossible to sell or purchase 5% of average daily trading volume from the book.

The 90th ATP measure is presented in panel D. Although there are similarities between the third and fourth measures, the graphs highlight key differences between them. First, unlike the 5% ADV measure, the 90th percentile average price is almost always defined. Second, the 90th ATP measure tracks the quotes much more closely than the 5% ADV because traders typically submit orders smaller than the quoted depth, and the quoted depth is a very small fraction of average daily volume.

3. Technical Analysis Rules

Much of technical analysis involves pattern recognition using specific frequency (intraday, daily, weekly) charts that display opening, high, low, and closing prices, as well as trading volume in some form. We focus on two types of rules that are widely used: support/resistance levels and moving average forecasts. Alternative technical indicators, such as head-and-shoulders, pennants, and round-tops, are all variations on these simple rules.

The fact that technical analysis remains more of an art than a science creates methodological difficulties.[8] Using the same data, different technicians may arrive at different forecasts because there are no definitive guidelines for the rules' construction. In addition, technicians often incorporate their own subjective views into their forecasts, exacerbating this problem. Without explicitly stated strategies to guide us, we define specific rules that capture the spirit of the technical trading strategies. Although these rules may differ from those used by particular technicians, the existence of a relation between our rules and liquidity on the limit order book would demonstrate the ability of past prices to relay information about the current liquidity profile.

Each half hour, we compute the technical analysis measures using the bid and ask prices from the prior week (i.e., the 70 most-recent observations, sampled at half-hour intervals). The support level is defined as the lowest bid price attained during the last week, provided this minimum was achieved at least twice during that period and is undefined otherwise. We repeat this process at every half-hour interval, using rolling one-week windows. The resistance level is defined analogously. (See Equation 1 below.) These definitions reflect the fact that technicians commonly use recent lows to determine support levels and recent highs to determine resistance levels.

$$\begin{aligned}
\text{Support level}_t &= P_t^s \\
&= \begin{cases} \min(Bid_{t-w}, \dots, Bid_t) \text{ provided } P_t^s = Bid_i = Bid_j \\ \quad \text{for some } i, j \text{ in } (t-w, t) \\ \text{undefined otherwise} \end{cases} \\
\text{Resistance level}_t &= P_t^R \\
&= \begin{cases} \max(Ask_{t-w}, \dots, Ask_t) \text{ provided } P_t^R = Ask_i = Ask_j \\ \quad \text{for some } i, j \text{ in } (t-w, t) \\ \text{undefined otherwise} \end{cases}
\end{aligned} \tag{1}$$

Our moving average indicator is defined by the relation between long-run and short-run moving averages of quoted midpoint prices. The

[8] For a general overview of technical analysis, as well as an explanation of the various technical analysis indicators, see Pring (1991), Edwards and Magee (1998), Murphy (1999), and Bensignor (2000).

long-run moving average is defined over the past two weeks, while the short-run moving average is defined over one-fourth of that time (2.5 days).[9] If the short-run moving average exceeds a band around the long-run moving average, then the observation is assigned the value 1. Similarly if the short-run average falls below the band, the indicator assumes the value -1. Otherwise the observation is assigned a neutral value of zero. As explained by Brock, Lakonishok, and LeBaron (1992), the use of a band around the long-run average eliminates "whiplash signals" and focuses the analysis on more substantial differences between short- and long-run averages. Equation (2) specifies the moving average rule we adopt.

$$\text{Moving average indicator} = \begin{cases} +1 & \text{if } MAS_t > MAL_t + B \\ -1 & \text{if } MAS_t < MAL_t - B, \\ 0 & \text{otherwise} \end{cases} \qquad (2)$$

where $MAS_t = \sum_{k=1}^{S} P_{t-k}/S$ and $MAL_t = \sum_{k=1}^{L} P_{t-k}/L$, and B represents an exogenously determined band (set at 3%), P represents the quoted midpoint price, S represents the number of half-hour periods in the short window, and L represents the number of half-hour periods in the long window.[10]

Figure 2 provides an illustration of the technical analysis measures we investigate, again for Dresser Industries. Panel A displays the technical analysis support and resistance levels and the quoted midpoint for the same one-week period shown in Figure 1. Panel B shows the relation between the long-run and short-run moving averages for the entire month of August 1997. As Figure 2 illustrates, our technical analysis measures display a number of characteristics common in traditional technical forecasts. First, the technical analysis measures tend to move with the quoted midpoint (panel A). Note that this is particularly true for the resistance level, probably because prices were increasing over this period. Second, while the support and resistance levels can remain constant for substantial periods of time, it is important to note that they are not always defined and can contain discrete jumps. These features are strikingly similar to those for the limit order book measures in Figure 1. In panel B, the short-run moving average displays behavior consistent with both momentum and mean reversion, as it drifts away from and then reverts back to the long-run average.

[9] Our choice of window lengths was guided by the desire to have windows that were long enough to provide a reasonable benchmark but short enough to be relevant (not stale). In addition, Murphy (1999, p. 203) states that two popular moving average specifications have short-run window lengths that are one-quarter to one-fifth of the long-run window length.

[10] Three percent is one of the bandwidths used by Brock, Lakonishok, and LeBaron (1992). We also repeat the analysis using a 1% band and, although the moving average indicator assumes nonzero values more frequently, the results are qualitatively unchanged.

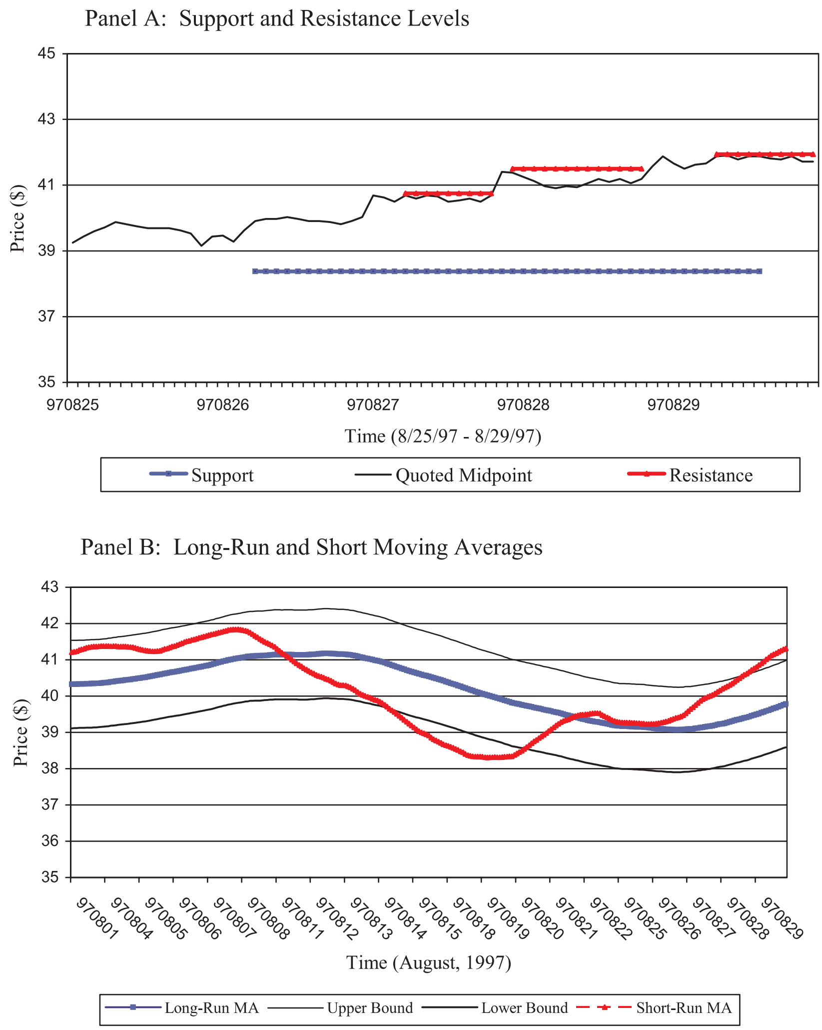

Figure 2
Position of technical analysis levels for Dresser Industries
Technical support (resistance) levels are based on the lowest (highest) bid (ask) price reached at least twice over the previous week. See Equation (1) for details. The long-run moving average is computed using data from the past 2 weeks, while the short-run moving average is based on data from the past 2.5 days. The upper and lower bounds represent a 3% band. See Equation (2) in the text.

4. Results

4.1 General properties of technical analysis rules and liquidity provision measures

Table 1 begins to formalize the relations displayed in Figures 1 and 2 for our entire sample. Panel A displays statistics on the support and resistance

Table 1
Properties of technical analysis and limit order book measures

Panel A: Support and resistance levels

	Technical rule	Limit order book measures: Mode	Depth near quotes	5% ADV	90th ATP
% Time defined	50.5%	99.3%	99.5%	90.4%	95.1%
Distance to qtd mdpt ($)					
Mean	0.802	1.139		0.740	0.216
Median	0.375	0.313		0.250	0.094
Distance to tech level ($)					
Mean		0.949		0.626	0.646
Median		0.375		0.250	0.250

Panel B: Moving average indicator

	+1	0	−1
Frequency	13.9	77.1	9.0
Duration			
Mean	4.3 days	13.8 days	3.8 days
Median	4.1 days	7.6 days	3.8 days
Total return			
Mean	0.3%	1.7%*	1.1%*
Median	−0.6%	3.6%	0.6%

Descriptive statistics on the properties of the technical trading rules and limit order book liquidity measures are contained above. Means, medians, and frequencies were computed across the entire sample unless otherwise noted. Technical support/resistance levels and the moving average indicator are defined in Equations (1) and (2) in the text. "Mode" represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. "Depth near quotes" is defined as the ratio of depth within a stock-specific interval around the quoted midpoint to the depth within twice the interval. "5% ADV" is defined as the last share price from the limit order book that an investor would pay when buying or receive when selling a quantity equal to 5% of the average daily trading volume. "90th ATP" is defined as the average price per share for a trade that executes against the limit order book where the trade size is the 90th percentile of the stock's distribution of trade size. "Distance to qtd mdpt" is the absolute value of the difference between the respective measure and the midpoint of the bid-ask spread. "Distance to tech level" is the absolute difference between the limit order book liquidity measure and the technical level on the same side of the market. Total return is computed from one change in the moving average indicator to the next (e.g., from the triggering of a positive signal to the return to a neutral signal) using quoted midpoints. All mean distances differ from one another at the 1% level.

levels we generate using Equation (1), as well as the corresponding limit order book liquidity measures, while panel B contains statistics on the moving average indicator. Support and resistance levels are defined for 50.5% of the observations in our sample. When they exist, support and resistance prices deviate from the quoted midpoint by $0.80 on average, with a median difference of $0.38. Analogous statistics for the limit order book variables support the hypothesis that these technical trading rules are related to the level of liquidity on the limit order book. Median differences between quoted prices and the limit order book mode and 5% ADV measures are similar to those reported for the support and resistance levels.

Panel B shows that the majority of the time (77%), the moving average indicator is neutral, meaning it assumes a value of zero. When the moving average indicator is positive (negative), the summary statistics demonstrate that these nonneutral periods last for approximately four days and are associated with negative (positive) median returns. This suggests that although these rules are designed to forecast short-term momentum, the moving average indicator may be associated with price reversals of quote midpoints in our sample.[11]

4.2 Support/resistance levels and limit order book liquidity

In order to directly test our hypothesis that technical analysis rules are related to liquidity provision on the limit order book, we regress the limit order book variables on the technical levels after subtracting the quoted midpoint from each series. Subtracting the midpoint of the bid-ask spread serves two purposes. First, it eliminates the nonstationarity that stems from the presence of unit roots in price series. Second, it removes the mechanistic part of the relationship that must exist between any functions of quoted prices and prices on the limit order book.[12]

Table 2 shows the regressions relating the differenced series. We regress the buy-side limit order book measures on the technical analysis support level and the sell-side limit order book measures on the technical analysis resistance level. In both cases we present results for univariate regressions and for multivariate regressions that control for current market conditions. The conditioning variables are those typically used in the literature to summarize the state of the market. We use the current quoted bid-ask spread and total quoted depth to control for the level of asymmetric information risk and inventory risk embedded into the quotes by liquidity providers. We also include the current (over the last 30 minutes) and total (over the entire one-week window) trading volume as well as the current and total net trade direction, which are defined as the difference between the number of buyer- and seller-initiated transactions in the given time interval.[13]

[11] To provide some perspective on these general statistics, we applied our technical analysis rules to a simulated pure random walk process, starting at \$50 per share and rounding to the nearest 16th. When the variance of the simulated series is set equal to the average variance in our sample, the moving average indicator is positive (negative) 5.8% (5.0%) of the time on average. We have to raise the volatility estimate to the 90th percentile of our sample distribution to generate moving average statistics close to those in Table 1.

[12] We have been careful to address the econometric concerns. Specifically we have tested each series for the presence of a unit root. Tests conducted using the Dickey-Fuller and Phillips-Perron test, which allows for serial correlation in the error terms, show overwhelmingly that the series in question fails to reject the null hypothesis of a unit root. We also tested for cointegrating relationships to ensure our difference variables were covariance stationary. Cointegration tests were done using the Dickey-Fuller and Phillips-Ouliaris tests. The results of both tests confirm not only that the differenced variables are covariance stationary, but that a large fraction of the stocks have technical analysis measures and limit order book variables that are cointegrated at the 5% level. The details of these tests are available from the authors upon request.

[13] Trades were signed using the Lee and Ready (1991) trade classification algorithm rather than determining the initiator directly using the order data to allow practitioners to apply our methods to other data.

Table 2
Support and resistance level regressions

Dependent variable	Univariate regn. coefficients: Intercept	Univariate regn. coefficients: Tech – qtd mdpt	Multivariate regression median coefficient estimates [Median p-values] and (% significant at the 5% level): Intercept	Tech – qtd mdpt	Quoted spread	Tot qtd depth	Current volume	Total volume	Curr buy-sell asym	Total buy-sell asym	Medn adj. R^2
Panel A: Buy-side limit order book measures regressed on technical analysis support levels											
Buy mode – qtd mdpt	−0.564	0.232	−0.560	0.221	−0.566	1.7E-5	1.6E-6	9.3E-8	0.006	−0.013	0.16
	[0.000]	[0.039]	[0.008]	[0.058]	[0.207]	[0.031]	[0.236]	[0.092]	[0.433]	[0.144]	
	(83.6)	(50.9)	(59.1)	(49.1)	(26.4)	(57.3)	(20.0)	(43.6)	(5.5)	(39.1)	
Buy depth near quotes	0.767	0.021	0.754	0.031	−0.163	3.0E-6	2.4E-7	1.7E-8	0.001	0.002	0.14
	[0.000]	[0.073]	[0.000]	[0.047]	[0.226]	[0.018]	[0.335]	[0.137]	[0.427]	[0.202]	
	(100.0)	(43.6)	(94.5)	(50.9)	(28.2)	(58.2)	(13.6)	(40.0)	(11.8)	(35.5)	
Buy 5% ADV – qtd mdpt	−0.513	0.039	−0.578	0.027	−0.522	1.5E-5	3.8E-7	1.0E-8	0.010	−0.002	0.18
	[0.000]	[0.112]	[0.001]	[0.234]	[0.078]	[0.001]	[0.333]	[0.213]	[0.450]	[0.295]	
	(96.4)	(39.1)	(72.7)	(35.5)	(46.4)	(76.4)	(20.0)	(31.8)	(6.4)	(29.1)	
Buy 90th ATP – qtd mdpt	−0.151	0.011	−0.114	0.016	−0.545	6.9E-6	2.0E-7	1.2E-8	0.002	−2.1E-5	0.21
	[0.000]	[0.167]	[0.002]	[0.233]	[0.000]	[0.000]	[0.428]	[0.158]	[0.363]	[0.274]	
	(99.1)	(33.6)	(76.4)	(27.3)	(82.7)	(92.7)	(15.5)	(33.6)	(8.2)	(20.9)	
Panel B: Sell-side limit order book measures regressed on technical analysis resistance levels											
Sell Mode – qtd mdpt	0.297	0.332	0.373	0.324	0.369	−1.0E-5	−2.9E-7	−5.5E-8	−0.003	−0.018	0.20
	[0.000]	[0.003]	[0.008]	[0.004]	[0.293]	[0.016]	[0.354]	[0.151]	[0.496]	[0.167]	
	(85.5)	(72.7)	(60.9)	(68.2)	(21.8)	(60.9)	(12.7)	(34.5)	(6.4)	(37.3)	
Sell depth near quotes	0.732	−0.031	0.730	−0.041	−0.196	5.5E-6	2.3E-8	5.5E-11	−0.000	−0.002	0.18
	[0.000]	[0.031]	[0.000]	[0.090]	[0.218]	[0.004]	[0.391]	[0.119]	[0.428]	[0.101]	
	(100.0)	(51.8)	(98.2)	(41.8)	(31.8)	(68.2)	(14.5)	(40.9)	(12.7)	(40.9)	
Sell 5% ADV – qtd mdpt	0.363	0.149	0.382	0.103	0.562	−1.8E-5	−1.7E-7	−2.3E-8	0.001	−0.004	0.25
	[0.000]	[0.034]	[0.002]	[0.074]	[0.074]	[0.000]	[0.349]	[0.191]	[0.425]	[0.163]	
	(95.5)	(55.5)	(71.8)	(43.6)	(55.5)	(80.9)	(17.3)	(30.9)	(7.3)	(37.3)	
Sell 90th ATP – qtd mdpt	0.107	0.035	0.067	0.028	0.533	−5.7E-6	−1.3E-8	−6.5E-10	0.002	−0.001	0.28
	[0.000]	[0.042]	[0.007]	[0.078]	[0.000]	[0.000]	[0.350]	[0.232]	[0.427]	[0.262]	
	(99.1)	(51.8)	(64.5)	(36.4)	(90.9)	(94.5)	(17.3)	(26.4)	(11.8)	(19.1)	

Median coefficient estimates [p-values] across stocks are presented, as well as the percentage of stocks for which the given explanatory variable is significant at the 5% level (in parentheses). Technical support/resistance levels are defined in Equation (1) in the text. "Mode" represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. "Depth near quotes" is defined as the ratio of depth within a stock-specific interval around the quoted midpoint to the depth within twice the interval. "5% ADV" is defined as the last share price from the limit order book that an investor would pay when buying or receive when selling a quantity equal to 5% of the average daily trading volume. "90th ATP" is defined as the average price per share for a trade that executes against the limit order book where the trade size is the 90th percentile of the stock's distribution of trade size. Differences between each variable and the quoted midpoint are taken to ensure stationarity. Limit order book buy-side (sell-side) regressions use support (resistance) level technical measures.

The control variables are important for a number of reasons. First, they are related to liquidity demanders' needs, which could have implications for the state of the limit order book. Second, they may capture momentum effects, if liquidity demanders are trading on one side of the market. Work by Mingelgrin (2000) and Gervais, Kaniel, and Mingelgrin (2001) shows that volume measures contain information useful for predicting future returns both at the intraday and daily levels. The number of net transactions has been used by Easley, Kiefer, and O'Hara (1997), among others, to measure adverse selection risk.

The univariate regression results demonstrate positive and significant relationships between the technical and limit order book price variables (i.e., mode, 5% ADV, and 90th ATP). For the near depth variable, the coefficient is positive in the buy-side regression and negative in the sell-side regression. Both coefficients are significant and imply that when the technical level is near the quotes, depth is concentrated near the quotes as well. Analogously, when the distance between the technical level and the quotes is large, a substantial fraction of depth lies far from the quotes, near the technical level.

The multivariate regression results show that the technical variables remain important even after controlling for other market characteristics. In fact, both significance levels and median coefficient estimates decline only slightly when the additional variables are added to the regression. The support and resistance variables are statistically significant for 42% to 73% of the stocks in our sample when the limit order book location measures (mode and near depth ratio) are used as the dependent variable. When the volume-based liquidity measures (5% ADV and 90th ATP) are used, the technical levels are statistically significant for 27% to 56% of the stocks. The limit order book mode results, which are statistically significant for 49% to 73% of the stocks, are particularly strong. This suggests that the mode may be a slightly better measure of liquidity in this context.

By demonstrating a specific relationship between technical analysis rules and liquidity provision on the limit order book, the results in Table 2 not only support our hypothesis, but also provide a practical method for estimating the location of depth on the book. Investors can use the technical analysis levels defined in Equation (1) to approximate the location of the limit order book mode and, to a lesser extent, the other limit order book liquidity measures.

Figure 3 provides a graphical illustration of the correspondence between technical support and resistance levels and peaks in depth on the book. The figure presents average normalized depth at various points on the limit order book for those cases in which the support and resistance levels are more than \$1/8 from the respective quote. We first compute the depth at each limit price surrounding the support and resistance levels for each stock. We then scale by the total depth on the respective side of the

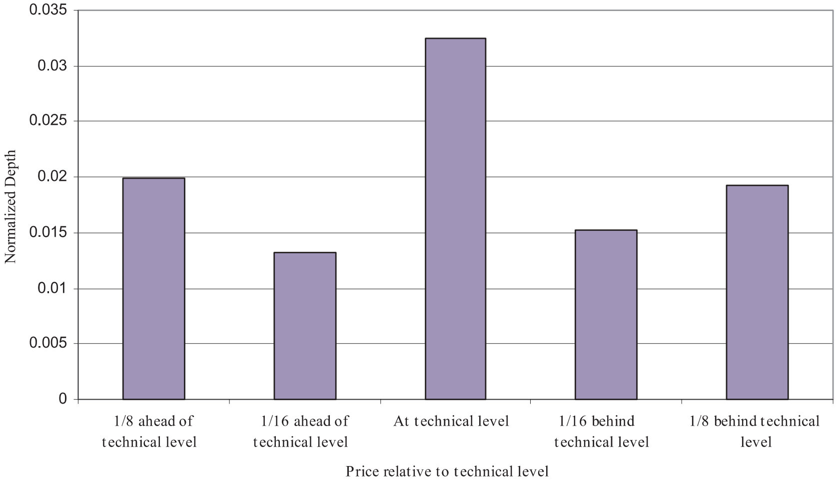

Figure 3
Normalized depth on the limit order book surrounding technical analysis levels
Using support and resistance levels that are more than \$1/8 from the respective quote, we compute the average normalized (i.e., scaled by total) depth at various limit prices surrounding the technical level across all stocks on both the buy and sell sides of the limit order book.

book to arrive at the normalized (i.e., fraction of total) depth for the stock. Finally, we average across all stocks over both sides of the book. The average normalized depth at selected prices near the technical levels is graphed in Figure 3. Consistent with the findings above, we see a large concentration of depth at the technical level.

Figure 3 offers strong visual evidence that technical analysis can be used to locate peaks of liquidity. The summary statistics in Table 3, which provide a more detailed description of the magnitude of differences between technical levels and limit order book mode variables, support the findings in Figure 3. A substantial fraction of the technical analysis levels are within \$1/8 of the limit order book mode (panel A). Furthermore, the results show that even when the technical price and the limit order book mode differ in dollar terms, there is often no difference in the number of shares at these prices. This commonly occurs when the technical level falls into a gap between prices on the limit order book. On average, the technical analysis levels within our sample locate cumulative depth of approximately 9,300 shares on the bid side and 8,500 shares on the ask side of the market (panel B). This observation suggests a new potential use for technical analysis, namely to provide a rough indication of the cost of trading approximately 9,000 shares. For example, if the technical level is close to the quotes, it may be easy to obtain a large block of stock, while if the technical level is far from the quotes, that same block of stock may be expensive to trade due to large price concessions.

Table 3
Distribution of price and share differences between technical analysis and limit order book measures

	Percentiles				
	20th	40th	50th	60th	80th
Panel A: Difference between technical levels and LOB mode					
LOB buy mode price – tech support price ($)	−1.000	−0.125	0.000	0.063	0.500
LOB sell mode price – tech resistance price ($)	−0.375	−0.063	0.000	0.000	0.500
Shares at tech support – shares at LOB buy mode	−15,000	0	0	0	4,600
Shares at LOB sell mode – shares at tech resistance	−3,774	0	0	0	13,650
Panel B: Shares at technical analysis levels					
Shares at support level	2,000	6,000	9,300	14,000	33,000
Shares at resistance level	1,900	5,500	8,500	12,300	31,000

The sample distribution of differences between the technical trading rules and the limit order book liquidity measures is contained above. Percentiles were computed using the entire sample. Technical support and resistance levels are defined in Equation (1) in the text. "LOB mode" represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares.

4.3 Moving average indicator

As above, we begin by subtracting the quoted midpoint from the limit order book depth variables. We then analyze whether the moving average indicator variable has any power to explain the magnitude of these differences. The motivation for this analysis is as follows. While technical analysis support and resistance levels are directly related to the position of liquidity on the limit order book, moving average indicators provide insight into the proximity of high cumulative depth to the quoted midpoint. For example, if moving average indicators signal that prices are going to move into a new, higher range (because the short-run average is above the long-run average), we would expect sell-side limit prices with high depth to be closer to the quoted midpoint and buy-side prices with high depth to be farther away. In other words, the moving average indicator may reveal information about the "skewness" of liquidity between the two sides of the book.[14]

The results of the regressions are contained in Table 4. As in Table 2, two different model specifications, both univariate and multivariate, are presented. For the limit order book mode and the 5% ADV measures, the moving average indicator variable is negative and statistically significant at the 5% level for a substantial fraction (27%–58%) of the stocks in the

[14] We also ran the moving average regressions using direct measures of asymmetry on the limit order book. Asymmetry measures were computed for the mode, 5% ADV, and 90th ATP variables by taking two times the difference between the sell-side limit order book price and the quoted midpoint divided by the difference between the sell- and buy-side prices. The results of the asymmetry regressions support the conclusions drawn here.

Table 4
Moving average regressions

Dependent variable	Univariate regn. coefficients: Intercept	Univariate regn. coefficients: Moving average	Multivariate regression median coefficient estimates [Median p-values] and (% significant at the 5% level): Intercept	Moving average	Quoted spread	Tot qtd depth	Current volume	Total volume	Curr buy-sell asym	Total buy-sell asym	Medn adj. R^2
Panel A: Buy-side limit order book measures regressed on technical analysis moving average indicator											
Buy mode – qtd mpt	−0.838 [0.000] (100.0)	−0.120 [0.043] (50.9)	−0.653 [0.000] (75.5)	−0.099 [0.069] (47.3)	−0.653 [0.127] (35.5)	1.5E-5 [0.001] (71.8)	8.9E-7 [0.338] (9.1)	2.8E-9 [0.128] (41.8)	−0.003 [0.363] (14.5)	−0.040 [0.115] (40.0)	0.10
Buy depth near quotes	0.739 [0.000] (100.0)	0.004 [0.097] (45.5)	0.735 [0.000] (98.2)	−5.3E-18 [0.124] (39.1)	−0.184 [0.143] (40.9)	3.6E-6 [0.001] (71.8)	6.4E-8 [0.329] (20.0)	4.1E-9 [0.130] (40.0)	−0.002 [0.354] (18.2)	−0.007 [0.174] (34.5)	0.09
Buy 5% ADV – qtd mdpt	−0.595 [0.000] (99.1)	−0.039 [0.099] (44.5)	−0.758 [0.000] (89.1)	−0.011 [0.173] (27.3)	−0.659 [0.030] (60.0)	1.9E-5 [0.000] (90.0)	8.5E-7 [0.296] (24.5)	1.3E-8 [0.171] (31.8)	0.013 [0.307] (20.0)	−0.005 [0.317] (31.8)	0.14
Buy 90th ATP – qtd mdpt	−0.168 [0.000] (100.0)	−0.002 [0.143] (34.5)	−0.123 [0.000] (90.9)	−0.001 [0.221] (23.6)	−0.573 [0.000] (92.7)	6.4E-6 [0.000] (97.3)	1.2E-7 [0.316] (21.8)	3.7E-9 [0.196] (35.5)	0.003 [0.275] (17.3)	−0.002 [0.279] (18.2)	0.17
Panel B: Sell-side limit order book measures regressed on technical analysis moving average indicator											
Sell mode – qtd mpt	0.613 [0.000] (100.0)	−0.128 [0.015] (58.2)	0.608 [0.000] (90.9)	−0.097 [0.023] (54.5)	0.482 [0.125] (36.4)	−1.1E-5 [0.001] (74.5)	−2.9E-7 [0.297] (27.3)	−8.2E-9 [0.068] (46.4)	−0.020 [0.213] (25.5)	−0.036 [0.065] (50.0)	0.11
Sell depth near quotes	0.694 [0.000] (100.0)	−2.9E-17 [0.045] (52.7)	0.696 [0.000] (99.1)	−0.002 [0.059] (48.2)	−0.246 [0.061] (45.5)	4.9E-6 [0.000] (81.8)	3.9E-8 [0.407] (22.7)	−1.3E-10 [0.074] (42.7)	0.003 [0.430] (12.7)	0.007 [0.090] (44.5)	0.12
Sell 5% ADV – qtd mdpt	0.524 [0.000] (100.0)	−0.044 [0.055] (48.2)	0.519 [0.000] (90.9)	−0.028 [0.110] (40.0)	0.737 [0.001] (70.9)	−1.9E-5 [0.000] (91.8)	−8.4E-8 [0.358] (25.5)	−1.1E-8 [0.194] (35.5)	−0.000 [0.398] (16.4)	−0.013 [0.134] (36.4)	0.19
Sell 90th ATP – qtd mdpt	0.146 [0.000] (100.0)	−0.001 [0.211] (31.8)	0.102 [0.000] (90.0)	−0.001 [0.218] (30.9)	0.573 [0.000] (95.5)	−5.8E-6 [0.000] (96.4)	−3.3E-9 [0.353] (19.1)	−5.1E-11 [0.273] (33.6)	1.2E-5 [0.371] (15.5)	−0.001 [0.270] (24.5)	0.25

Median coefficient estimates [p-values] across stocks are presented, as well as the percentage of stocks for which the given explanatory variable is significant at the 5% level (in parentheses). The moving average indicator variable assumes values of 1, 0, or −1 based on the relation between short- and long-run average prices [see Equation (2) in the text]. "Mode" represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. "Depth near quotes" is defined as the ratio of depth within a stock-specific band around the quoted midpoint to the depth within twice the band. "5% ADV" is defined as the last share price from the limit order book that an investor would pay when buying or receive when selling a quantity equal to 5% of the average daily trading volume. "90th ATP" is defined as the average price per share for a trade that executes against the limit order book where the trade size is the 90th percentile of the stock's distribution of trade size. Differences between each variable and the quoted midpoint are taken to ensure stationarity.

sample. Slightly weaker results obtain for the 90th ATP measure, with 24% to 35% of the stocks in the sample yielding statistically significant coefficients. Nonetheless, the combined negative coefficients of these three variables confirm our hypothesis that a positive moving average indicator corresponds to a shift in quoted prices toward the sell-side liquidity levels and away from the buy-side levels (and vice versa).[15]

The results for the near depth concentration ratio have a consistent and complimentary interpretation. The negative coefficient on the sell side indicates that depth is less concentrated near the ask when the short-run average exceeds the long-run average. This implies that as the quotes move higher (i.e., the moving average indicator is positive), sell-side orders near the quotes are either executed, canceled, or repositioned higher, thereby reducing the near depth ratio. At the same time, the positive median coefficient (in the univariate regression) on the buy side suggests that orders are placed near the quotes, thereby increasing the concentration of debt near the bid.

As in Table 2, the relationships between the technical analysis variable and the limit order book mode appear stronger than those between the technical analysis variable and the other limit order book depth measures. The strength of the mode results may be due to the "stickiness" of the limit order book. The other liquidity measures, which are influenced more by newly placed orders, are likely to respond more quickly to new order flow, thereby reducing the asymmetry in depth.

5. Interpretation

Our empirical results thus establish strong contemporaneous relationships between technical analysis indicators and liquidity on the limit order book. On a fundamental level, these relationships must arise from the strategies different market participants adopt in response to the trade-offs they face. In the case of the moving average indicator, the underlying strategies are clear; positive signals are triggered when strong buying pressure causes prices to rise. Likewise, negative signals are triggered by the price decrease that results from intense selling pressure. In contrast, the behaviors underlying the relationship between technical support/resistance levels and the limit order book are less obvious. The technical levels may respond to limit order placement behavior, or limit order placement may respond to the existence of technical support/resistance levels, or both. We focus the remainder of the discussion on interpretation of the support and resistance results, which are not only richer, but also capture the more natural link between technical analysis and liquidity

[15] Differences between buy-side limit order book variables and the quoted midpoint are negative, so a negative coefficient implies an increase in this distance.

provision. To begin to uncover the source of this link, we consider the objectives of various types of traders and the resulting characteristics of limit order books.

5.1 Theoretical and empirical background and hypotheses

A number of theoretical models offer insight into the characteristics of limit order books and provide benchmarks for empirical work. These models typically categorize traders into one of two groups — those who trade because they have private information about the value of the asset ("informed traders") and those who trade for reasons unrelated to information ("liquidity traders"). For example, Glosten (1994) and Rock (1996) present models in which uninformed, risk-neutral limit order traders compete to provide liquidity that is then consumed by (potentially) informed market order traders. In this setting, traders place limit orders on the book until the expected profit from supplying liquidity is driven to zero. While intuitively appealing, this is inconsistent with the empirical evidence gathered by Sandas (2001) and Hollifield, Miller, and Sandas (2003), who show that depth on the book is less than the competitive level suggested by these models.

Work by Kavajecz (1999) also demonstrates a number of limit order book characteristics that are either inconsistent with these models or outside their scope. For example, he shows that depth on limit order books is not distributed evenly. Instead, depth is concentrated near the quoted prices. In addition, some limit prices have no shares while others carry substantial depth. Another notable characteristic of limit order placement is the tendency of limit prices to cluster disproportionately on whole dollar prices, then half-dollar prices, then quarter-dollar increments, etc., as described first by Niederhoffer and Osborne (1966) and more recently by Hasbrouck (1999) in the equity markets and Osler (2002) in the foreign exchange market.

Not surprisingly, we see evidence of these features in our limit order book measures. For example, 42% of the limit order book mode observations occur on even half-dollars. In comparison, technical analysis levels are less concentrated, with only 25% of the technical prices falling on these increments (not reported). This tendency may drive a wedge between technical levels and limit prices with high depth. Indeed, it may partially explain the median (mean) absolute deviations between the technical prices and limit order book prices of up to \$0.38 (\$0.95) presented in Table 1. We also find that approximately 60% of all newly submitted limit orders in our sample are placed within \$0.25 of the quoted prices (Table 6 panel A). This feature will be important when analyzing the placement of orders around technical analysis levels.

Efforts to address these limit order book characteristics have led theoretical researchers, such as Angel (1995), Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2003), and Parlour and Seppi (2003), to focus on the strategic trade-offs facing liquidity providers and liquidity demanders in a dynamic setting. This work explicitly models the two basic risks traders face when they post a limit order — the risk of trading with someone who has superior information and the risk that the order goes unexecuted. The relative importance of these two risks is likely to differ between informed and uniformed traders. On the one hand, informed traders, believing they have superior information, are likely to be more concerned with execution risk than adverse selection risk. On the other hand, liquidity traders are likely to worry more about adverse selection than execution risk since they realize they do not have superior information. These strategic trade-offs may provide a fruitful avenue on which to investigate the source of the link between technical analysis and liquidity provision. With this in mind, we consider several possible explanations of this relationship.

First, suppose that a large pocket of depth exists on the limit order book for reasons unrelated to technical analysis. Informed traders, who are likely to have strong opinions as to the appropriate limit price, may have placed this collection of limit orders as in Bloomfield, O'Hara, and Saar (2003). Alternatively, a liquidity trader may have placed a large order at this price. In either case, the relationship between technical analysis and liquidity would then stem from the technical rules' ability to locate depth already in place on the book.

Perhaps instead, technical analysis reveals information unconnected to liquidity provision, but traders create this link by submitting limit orders near the technical levels. Traders may set their limit prices equal to the technical analysis levels in order to profit if prices remain in the current trading range. Traders may also choose to undercut the technical levels slightly. Undercutting, as described by Harris (1994), is a means of increasing the probability of execution by placing limit orders at slightly more aggressive prices, thereby gaining price priority over other orders. Under this hypothesis, the response of limit order traders to the technical signals establishes the link between technical analysis and the book.

Finally, it is possible that neither of these explanations is satisfactory in isolation, but that in fact, both occur. For example, technical analysis may locate prices where depth is already high, and liquidity providers may subsequently place orders just ahead of the technical levels in hopes of making a profit. In so doing, they reinforce the relationship between technical analysis and limit order book liquidity. This bidirectional hypothesis is consistent with work by Osler (2002), who claims that technical trading and order placement strategies may be jointly determined.

The contemporaneous results presented above are consistent with each of these explanations and thus do little to distinguish among them. We use two complementary approaches to determine which of the three hypotheses is best. First, we conduct Granger causality tests to analyze the lead-lag relation between the technical analysis support and resistance prices and the prices with high cumulative limit order book depth. These tests focus on how the *stock* of orders that sit on the limit order book relate to the technical analysis measures. In the second analysis, we investigate where new limit orders are placed relative to technical analysis support and resistance levels. In contrast to the first analysis, this investigation focuses on the *flow* of limit orders to the book.

5.2 Granger causality

The Granger causality tests allow us to determine whether (a) technical analysis support and resistance levels have any ability to predict the limit order book depth variables after controlling for lagged values of limit order book depth and/or (b) the limit order book depth measures have explanatory power for the support and resistance levels after controlling for lagged technical levels. Causality running from the book to the technical levels would provide support for the hypothesis that technicians locate depth already on the book. Causality from the technical levels to the book would suggest that limit orders are placed in response to technical indicators. Bidirectional causality would be consistent with the third explanation, in which technicians locate depth on the book and traders respond by placing orders near these prices.

Table 5 displays the results of the Granger causality tests. Tests are conducted using two and five half-hour lags. In both cases, the first three limit order book measures (i.e., the mode, near depth ratio, and 5% ADV) are more likely to Granger cause technical analysis support or resistance levels (LOB to tech) than the reverse (tech to LOB). The results for the 90th ATP are less conclusive, with causality running from technical levels to the book on the support side and causality running from the book to the technical levels on the resistance side. In addition, the results document bidirectional causality for as much as 16% or as little as 1% of the stocks in the sample, depending on the limit order book measure chosen and the number of lags included.

Overall the results in Table 5 provide stronger support for the hypothesis that technical analysis identifies limit prices with high depth than for the other hypotheses. The findings are not clear cut, however, and thus do not lead to a definitive conclusion as to the foundation of the relationship between technical analysis and limit order book liquidity.

5.3 Limit order book dynamics

We extend our investigation by taking a complementary approach that focuses on the flow of new limit orders surrounding changes in technical levels. This focus stems from our expectation that investors are most

Table 5
Lead-lag relationships between technical analysis levels and limit order book liquidity

	Direction of Granger causality (%)							
	Support				Resistance			
	Mode	Depth near quotes	5% ADV	90th ATP	Mode	Depth near quotes	5% ADV	90th ATP
Two lags								
LOB to tech	17.3	17.3	14.5	15.5	30.0	27.3	20.0	20.0
Tech to LOB	5.5	13.6	14.5	20.0	13.6	10.9	17.3	15.5
Both	9.1	6.4	7.3	8.2	4.5	0.9	10.0	15.5
Five lags								
LOB to tech	17.3	17.3	12.1	10.9	25.5	24.5	20.0	21.8
Tech to LOB	9.1	11.8	11.2	17.3	8.2	7.3	15.5	21.8
Both	6.4	5.5	12.1	10.9	5.5	4.5	8.2	10.0

This table contains the percentage of stocks in the sample with statistically significant Granger causality (at the 5% level) based on F-tests using two and five lags. "LOB to tech" refers to the limit order book depth variables Granger-causing the technical trading rules, "Tech to LOB" refers to the technical trading rules Granger-causing depth on the limit order book, and "both" refers to causality that runs in both directions.

conscious of technical analysis support and resistance prices when they change. The benefits of this approach are that, unlike the Granger causality tests, this investigation allows us to examine more directly the timing of changes in the limit order book relative to changes in technical measures and is not influenced by the persistence of the limit order book.

To disentangle the impact of the technical analysis level from the tendency to place orders near the quotes, we focus only on technical levels that are more than \$0.50 from the quoted price. For the same reason, we place an identical restriction on the limit order book mode. Furthermore, we analyze only those changes in which the new support or resistance level remains constant for at least one day in order to avoid "fleeting" changes.

For each newly placed order, we compute the dollar difference between the limit price and the technical analysis level. We then weight each order by its size in shares (relative to the average order size for that stock) to determine the fraction of new orders that are placed at, in front of, or behind the technical levels.

Panel B of Table 6 contains the distribution of newly placed limit orders relative to the technical levels for one day preceding and one day following the technical analysis change. The first columns of panel B demonstrate that a substantial fraction of orders are placed more than \$0.50 in front of the technical levels in both the pre- and postchange periods. This is consistent with the results in panel A, which document the tendency of new orders to concentrate around the quotes. (Recall that panel B focuses exclusively on technical levels that are more than \$0.50 from the quotes.)

Table 6
Limit order placement

Panel A: Distribution of newly placed limit orders by distance to the quoted prices

	In front of the quoted price					Behind the quoted price			
	>\$1	(\$0.50,\$1.00]	(\$0.25,\$0.50]	(\$0.00,\$0.25]	At the Quote [\$0.00]	(\$0.00,\$0.25]	(\$0.25,\$0.50]	(\$0.50,\$1.00]	>\$1
Limit buy orders				10.35	13.71	39.60	13.31	10.78	12.25
Limit sell orders				10.04	13.16	37.25	15.01	11.46	13.09

Panel B: Distribution of newly placed limit orders by distance to the new technical analysis prices

	In front of the technical analysis price					Behind the technical analysis price			
	>\$1	(\$0.50,\$1.00]	(\$0.25,\$0.50]	(\$0.00,\$0.25]	At tech price [\$0.00]	(\$0.00,\$0.25]	(\$0.25,\$0.50]	(\$0.50,\$1.00]	>\$1
Support side (limit buy orders)									
Preceding change	14.96	21.26	21.32	22.68	5.09	4.71	1.45	3.99	4.54
Following change	31.42	24.87	14.07	10.51	2.96	4.34	2.73	3.93	5.17
Resistance side (limit sell orders)									
Preceding change	22.62	15.53	7.74	22.57	9.83	8.11	7.82	1.03	4.76
Following change	23.06	34.95	10.61	7.50	3.16	10.16	1.92	2.85	5.79

The distribution of newly placed limit orders surrounding changes in technical levels (Panel B) was computed for all new support and resistance levels that persist for at least one day and are more than \$0.50 from the quoted price. When computing the distribution, we weighted each order by its size in shares scaled by the average order size for the stock. The one-day window preceding the change and the one-day window following the change are analyzed. “In front of” the quoted or technical price indicates that the order was placed closer to the spread midpoint and “behind” means the order was placed farther from the spread midpoint.

The results in panel B also reveal a concentration of orders just in front of the technical levels in the period before the levels are established. Since the technical levels in panel B are more than \$0.50 from the quotes, we can use the final columns of panel A as an approximate benchmark. Roughly 23% of new orders are placed just ahead of the new technical levels in the prechange period, which is substantially higher than the 11% to 13% indicated in panel A.

After the technical analysis levels have been established, we see a distinctly different pattern. In general, newly placed limit orders are posted far less frequently at, or just ahead of, the technical level. Specifically, the total fraction of orders posted at and just ahead of the technical level in the prechange period ranges from 28% to 32%, while in the postchange period those same frequencies range from 11% to 14%. Despite this difference, there appears to be a small concentration (approximately 10%) of orders placed just ahead of support levels and just behind resistance levels in the postchange period, which is perhaps due to an upward trend in prices over our sample period.

These results demonstrate that limit orders tend to be placed near a new technical analysis level prior to the level's creation more frequently than orders are placed near the new level after its creation. Limit orders placed before the creation of the new level tend to be positioned slightly ahead of (closer to the quotes) the forthcoming technical level, while limit order placement after the technical levels is much more diffuse. The fact that the concentration of orders before the change lies slightly in front of the technical analysis price suggests that traders may be trying to step ahead of a large position already in place on the book.

The results that limit orders are placed prior to the existence of the technical level is consistent with the hypothesis that technical analysis locates depth already in place on the book. These findings reinforce the results of the Granger causality tests, which suggest that the flow from the book to the technical levels is stronger than the flow from the technical levels to the book. Although there is some evidence of bidirectional causality, the results in aggregate imply that the connection between technical analysis and limit order book depth is driven by technical analysis being able to identify prices with high cumulative depth already in place on the limit order book.[16]

5.4 The limit order book as a principal source of information

Taken together, the causality results suggest that technical analysis measures capture changes to the limit order book — changes that result from

[16] An important caveat is in order. If the rules used by technical analysts determine support and resistance levels more quickly (or take longer) than our method, then the period we are labeling "pre" may actually include some postchange order flow or vice versa. The fact that the results are generally consistent with those from the Granger causality tests mitigates this concern.

the strategic behavior of various market participants. Accordingly, it is likely that limit order books, which synthesize the order placement strategies of both informed and uninformed traders into supply and demand functions for the stock, contain information relevant for future prices.

Several arguments support this idea. First, limit orders submitted by informed traders contain information that has not been priced into the asset by definition. Second, even uninformed orders may impact future prices, particularly in the short run, through the interplay between supply and demand. Finally, there is considerable empirical evidence that the limit order book contains information relevant for other aspects of the market. For example, Kavajecz and Odders-White (2001a, b) demonstrate that the book contains information that is not immediately reflected in the quotes. Furthermore, Irvine, Benston, and Kandel (2000) show that the slope of the limit order book can be used to predict trading activity. In addition, Harris and Panchapagesan (2002) estimate the option value of limit orders and relate it to future prices. Collectively this body of work demonstrates that the limit order book contains forward-looking information.

If order placement activity by strategic traders is indeed the source of the relationship between technical analysis and the book, as the prior results suggest, then the limit order book contains information that has not yet been revealed directly to the market and may therefore be relevant for future prices. If, on the other hand, the limit order book is merely responding to information that has already been revealed to the market in the form of announcements, past trading activity, or technical analysis measures, it should not be able to forecast future prices. We evaluate the predictive power of the limit order book by regressing one-week-ahead returns (computed using quoted midpoints) on each of our limit order book liquidity measures.[17] The results (not reported) demonstrate that price reversals are likely near points of high cumulative depth on the limit order book. Specifically, future returns are likely to be large and positive when the bid price approaches a limit buy price with high depth, and smaller when the ask price is close to a limit sell price with high depth. These relationships are statistically significant for between 29% and 57% of the stocks in our sample, with the strongest evidence of predictability coming from the limit order book mode. This suggests that limit prices with high depth are in fact difficult to penetrate. In summary, the results of the return regressions are consistent with the earlier finding that the link between technical analysis and liquidity stems from the ability of the technical rules to capture changes in the limit order book.

[17] As in the earlier regressions, we subtract the quoted midpoint from the dollar-based measures (i.e., the mode, 5% ADV, and 90th ATP measures).

6. Conclusion

The apparent conflict between the level of resources dedicated to technical analysis in the investment community and academic theories of market efficiency is a long-standing puzzle. Despite numerous studies, this conflict has yet to be resolved.

In this article, we explore a previously unexamined feature of technical analysis — its relationship to liquidity. Specifically, we test the hypotheses that support and resistance levels coincide with peaks in depth on the limit order book and that moving average forecasts reveal information about the relative position of depth on the book. After controlling for cointegrating relations and the state of the trading environment, we find that technical analysis support/resistance levels as well as moving average indicators are significantly related to the state of liquidity on the limit order book.

We consider several possible explanations of the relationship between technical analysis and the limit order book and conclude that it is tied to the strategic behavior of limit order traders. Specifically, the results of Granger causality tests and analyses of limit order placement activity reveal that technical analysis measures discover depth already in place on the limit order book. The fact that limit order book liquidity measures have some predictive power for future returns offers additional evidence that the limit order book is both a key source of forward-looking information and the origin of the relations that we highlight.

We believe that our results provide a novel way of thinking about both technical analysis and liquidity provision, and have lessons for both sides of the ongoing debate regarding technical analysis. For practitioners seeking information about the limit order book's composition, we show that technical analysis provides a convenient method of locating liquidity on the book. Using this information, traders can better estimate and reduce transaction costs and increase the probability of execution, by placing limit orders strategically. For academics interested in the price formation process, we demonstrate that the state of liquidity, as measured by the position of depth on the book, may be an important determinant of future prices.

Appendix A: Robustness Checks

We recognize that throughout our analysis, we make several assumptions that could influence our findings. For example, we rely on windows of particular lengths and on specific bandwidths when constructing the technical analysis measures. To investigate the dependence of the results on the particular parameters chosen, we conduct a number of robustness checks.

First, we examine the sensitivity of our results to the length of the window used to obtain the technical analysis measures. Recall that our support (resistance) levels are defined as the lowest (highest) quoted price achieved at least twice over the previous week. We rerun the analysis using a window length of two weeks instead of one week. Second, we investigate a similar redefinition of our moving average measure. While the original analysis defines the moving average variables using short- and long-run averages of 2.5 days and 2 weeks, respectively, we reconstruct averages using windows of one day for the short-run and one week for the long-run average. Third, we reduce the band surrounding the long-run moving average from 3% to 1%. In all cases, we find that our results are invariant to changes in the window lengths and bandwidth used to create the technical analysis levels.

We also recognize that our definitions of limit order book liquidity are an important part of the analysis, and we investigate the sensitivity of our results to these definitions. We rerun the analysis using four alternative definitions. First, we use the true limit order book mode — the price with the largest number of shares regardless of its position relative to the quotes. Second, we use a multipler of 5 instead of 10 for the near depth ratio. Third, we use the limit price associated with 15% rather than 5% of average daily trading volume. Last, we estimate the average price of a trade in the 75th percentile, rather than the 90th percentile, of a stock's distribution of trade sizes. Again, in all cases the results are similar to those presented here.

Finally, we examine the applicability of our results to another time period — one in which the nature of liquidity provision on the limit order book differs from that during our sample period. Specifically, we rerun our analysis on the original TORQ data, which covers the period from November 1990 through January 1991. This comparison is especially useful for several reasons. First, our sample period, which covers July–September 1997, was characterized by a minimum tick size of 1/16th rather than 1/8th. Goldstein and Kavajecz (2000) have shown that the reduction in minimum tick size resulted in a dramatic decrease in cumulative depth across all stocks in the market, which could certainly affect our results. Second, the regulatory environment changed substantially following the implementation of the Securities and Exchange Commission's Order Handling Rules, which began in January 1997. Third, trading volume increased dramatically between the two periods as well. A comparison across the 1997 and TORQ periods reveals that the results are both qualitatively and quantitatively similar, which demonstrates that our findings are robust to substantial changes in the trading environment. All robustness results are available from the authors upon request.

References

Alexander, S., 1961, "Price Movements in Speculative Markets: Trends or Random Walks," *Industrial Management Review*, II, 7–26.

Allen, F., and R. Karjalainen, 1999, "Using Genetic Algorithms to Find Technical Trading Rules," *Journal of Financial Economics*, 51, 245–271.

Angel, J., 1995, "Limit versus Market Orders," working paper, Georgetown University.

Barberis, N., A. Shleifer, and R. Vishny, 1998, "A Model of Investor Sentiment," *Journal of Financial Economics*, 49, 307–343.

Bensignor, R., 2000, *New Thinking in Technical Analysis: Trading Models from the Masters*, Bloomberg Press, Princeton, NJ.

Bessembinder, H., and K. Chan, 1998, "Market Efficiency and the Returns to Technical Analysis," *Financial Management*, 27, 5–17.

Bloomfield, R., M. O'Hara, and G. Saar, 2003, "The 'Make or Take' Decision in an Electronic Market: Evidence on the Evolution of Liquidity," working paper, Cornell University: forthcoming in *Journal of Financial Economics.*

Blume, L., D. Easley, and M. O'Hara, 1994, "Market Statistics and Technical Analysis: The Role of Volume," *Journal of Finance*, 49, 153–181.

Brock, W., J. Lakonishok, and B. LeBaron, 1992, "Simple Technical Trading Rules and the Stochastic Properties of Stock Returns," *Journal of Finance*, 47, 1731–1764.

Brown, D., and R. Jennings, 1989, "On Technical Analysis," *Review of Financial Studies*, 2, 527–551.

DeBondt, W., and R. Thaler, 1985, "Does the Stock Market Overreact?," *Journal of Finance*, 40, 793–805.

Easley, D., N. Kiefer, and M. O'Hara, 1997, "One Day in the Life of a Very Common Stock," *Review of Financial Studies*, 10, 805–835.

Edwards, R., and J. Magee, 1998, *Technical Analysis of Stock Trends*, John Magee Inc., Boston.

Fama, E., and M. Blume, 1966, "Filter Rules and Stock Market Trading," *Journal of Business*, 39, 226–241.

Foucault, T., 1999, "Order Flow Composition and Trading Costs in Dynamic Limit Order Markets," *Journal of Financial Markets*, 2, 99–134.

Foucault, T., O. Kadan, and E. Kandel, 2003, "Limit Order Book as a Market for Liquidity," working paper, HEC School of Management.

Gervais, S., R. Kaniel, and D. Mingelgrin, 2001, "The High-Volume Return Premium," *Journal of Finance*, 56, 877–919.

Glosten, L., 1994, "Is the Electronic Limit Order Book Inevitable?," *Journal of Finance*, 49, 1127–1161.

Goldstein, M., and K. Kavajecz, 2000, "Eighths, Sixteenths, and Market Depth: Changes in Tick Size and Liquidity Provision on the NYSE," *Journal of Financial Economics*, 56, 125–149.

Harris, L., 1994, "Minimum Price Variations, Discrete Bid-Ask Spreads and Quotation Sizes," *Review of Financial Studies*, 7, 149–178.

Harris, L., and V. Panchapagesan, 2002, "The Information-Content of the Limit Order Book: Evidence from Specialist Trading Decisions," working paper, Washington University.

Hasbrouck, J., 1992, "Using the TORQ Database," Working Paper 92-05, New York Stock Exchange.

Hasbrouck, J., 1999, "Security Bid/Ask Dynamics with Discreteness and Clustering: Simple Strategies for Modeling and Estimation," *Journal of Financial Markets*, 2, 1–28.

Hollifield, B., R. Miller, and P. Sandas, 2003, "Empirical Analysis of Limit Order Markets," working paper, University of Pennsylvania: forthcoming in *Review of Economic Studies*.

Irvine, P., G. Benston, and E. Kandel, 2000, "Liquidity Beyond the Inside Spread: Measuring and Using Information in the Limit Order Book," working paper, Hebrew University.

Jegadeesh, N., 2000, "Discussion of Foundations of Technical Analysis," *Journal of Finance*, 55, 1765–1770.

Jensen, M., and G. Bennington, 1970, "Random Walks and Technical Series: Some Additional Evidences," *Journal of Finance*, 25, 469–482.

Kavajecz, K., 1999, "The Specialist's Quoted Depth and the Limit Order Book," *Journal of Finance*, 52, 747–771.

Kavajecz, K., and E. Odders-White, 2001a, "An Examination of Changes in Specialists' Posted Price Schedules," *Review of Financial Studies*, 14, 681–704.

Kavajecz, K., and E. Odders-White, 2001b, "Volatility and Market Structure," *Journal of Financial Markets*, 4, 359–384.

Lee, C., and M. Ready, 1991, "Inferring Trade Direction from Intraday Data," *Journal of Finance*, 46, 733–746.

Lo, A., H. Mamaysky, and J. Wang, 2000, "Foundations of Technical Analysis: Computations Algorithms, Statistical Inference, and Empirical Implementation," *Journal of Finance*, 55, 1705–1765.

Mingelgrin, D., 2000, "The Informational Content of High-Frequency Trading Activity," working paper, University of Pennsylvania.

Murphy, J., 1999, *Technical Analysis of the Financial Markets*, New York Institute of Finance, Paramus, NJ.

Neely, C., P. Weller, and R. Dittmar, 1997, "Is Technical Analysis in the Foreign Exchange Market Profitable? A Genetic Programming Approach," *Journal of Financial and Quantitative Analysis*, 32, 405–426.

Neftci, S., 1991, "Naïve Trading Rules in Financial Markets and Weiner-Kolomgorov Prediction Theory: A Study of Technical Analysis," *Journal of Business*, 64, 549–571.

Neftci, S., and A. Policano, 1984, "Can Chartists Outperform the Market? Market Efficiency Tests for 'Technical Analyst'," *Journal of Futures Markets*, 4, 465–478.

Niederhoffer, V., and M. Osborne, 1966, "Market Making and Reversal on the Stock Exchange," *Journal of the American Statistical Association*, 61, 897–916.

Osler, C., 2003, "Currency Orders and Exchange-Rate Dynamics: An Explanation for the Predictive Success of Technical Analysis,"*Journal of Finance*, 58, 1791–1819.

Parlour, C., 1998, "Price Dynamics in Limit Order Markets," *Review of Financial Studies*, 11, 789–816.

Parlour, C., and D. Seppi, 2003, "Liquidity-Based Competition for Order Flow," *Review of Financial Studies*, 16, 301–343.

Pring, M., 1991, *Technical Analysis Explained*, McGraw-Hill, New York.

Ratner, M., and R. Leal, 1999, "Tests of Technical Trading Strategies in the Emerging Equity Markets of Latin America and Asia," *Journal of Banking and Finance*, 23, 1887–1905.

Ready, M., 2002, "Profits from Technical Trading Rules," *Financial Management*, Autumn, 43–61.

Rock, K., 1996, "The Specialist's Order Book and Price Anomalies," working paper, Harvard University.

Sandas, P., 2001, "Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market," *Review of Financial Studies*, 14, 705–734.

Sullivan, R., A. Timmermann, and H. White, 1999, "Data-Snooping, Technical Trading Rule Performance, and the Bootstrap," *Journal of Finance*, 54, 1647–1691.

Treynor, J., and R. Ferguson, 1985, "In Defense of Technical Analysis," *Journal of Finance*, 40, 757–775.