

Comments Welcome

Technical Analysis and Liquidity Provision

Kenneth A. Kavajecz* and Elizabeth R. Odders-White**

May 16, 2002

*Assistant Professor of Finance
The Wharton School
University of Pennsylvania
Philadelphia, PA 19104-6367
Email: kavajecz@wharton.upenn.edu
Phone: 215-898-7543
Fax: 215-898-6200

**Assistant Professor of Finance
University of Wisconsin - Madison
975 University Avenue
Madison, WI 53706
Email: ewhite@bus.wisc.edu
Phone: 608-263-1254
Fax: 608-265-4195

We gratefully acknowledge helpful comments from seminar participants at the Society for Financial Studies' Conference on Investments in Imperfect Capital Markets and the University of Pennsylvania as well as Simon Gervais, David Hsieh, Pete Kyle (discussant), Craig MacKinlay, Carol Osler, and Mark Ready. All remaining errors are our own.

Technical Analysis and Liquidity Provision

Abstract

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 offer an alternative explanation for the value of technical analysis that is consistent with market efficiency – specifically, that it reveals information about liquidity provision. We find evidence consistent with 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. These results demonstrate that technical analysis can have value even in an efficient market, and provide a practical method for estimating the level of liquidity on the book.

1. Introduction

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 because it is inconsistent with one of the most fundamental theories in traditional finance – the theory of market efficiency. The theory of (semi-strong form) market efficiency states that all publicly 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 excessive profits while assuming little or no risk, a state that is not sustainable in equilibrium. If markets are indeed efficient, then technical analysis, which relies heavily on publicly available past price and volume data, cannot predict future prices despite analysts' claims.

The traditional academic wisdom and investment community practices are clearly at odds. If practitioners' use of technical analysis is well-founded, then it appears that markets must be inefficient. Alternatively, if markets are efficient, then it seems that the investment community is wasting an enormous amount of resources on technical analysis.

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

In light of this conflict, many 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 (1991), Brock, Lakonishok, and LeBaron (1992), Neftci and Policano (1984), 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 incremental 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 or 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.²

Others provide results consistent with market efficiency 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), Ready (1997), and Bessembinder and Chan (1998) show that transaction costs offset the benefits of technical analysis found by Alexander (1961) and Brock et al. (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. In addition, countless tests of market

² See, for example, DeBondt and Thaler (1985) and Barberis et al. (1998).

efficiency have been conducted and many find strong support for efficiency.³ 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 paper, we propose a previously unexamined third explanation in which the value of technical analysis is not in direct conflict with market efficiency. Rather than attempting to link technical trading rules to future transaction prices as is typically done in the literature, we claim that technical analysis is related to the level of liquidity provision – although we do *not* claim that technicians are necessarily aware of this relationship. The fact that technical analysis and market efficiency can coexist hinges on the subtle, yet critical, difference between liquidity provision and future prices. While 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. We claim that the difference between liquidity provision and transaction prices explains the previous literature’s tenuous conclusions regarding the value of technical analysis, and 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, summarize supply and demand schedules and hence 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’

³ See, for example, Fama (1991) for a survey of past articles on market efficiency.

trading needs, this price is likely to become an obstacle to achieving more extreme (higher ask prices and lower bid prices) transaction prices. Likewise, support (resistance) levels occur when substantial selling (buying) pressure is required to breach a given price level. Consequently, we hypothesize that 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 are defined by the relation between a long-run and short-run moving average of quoted midpoint prices. If the short-run moving average rises above (falls below) the long-run average, then this serves as an indication that prices are going to move into a new, higher (lower) trading range. Given this definition, we would expect sell-side limit order book prices to be closer to (farther from) the quoted midpoint and buy-side limit order prices to be farther away (closer) when the moving average indicator signals that prices are rising (falling). In other words, the 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 are cointegrated with limit prices having high cumulative depth. Moreover, after controlling for both the cointegrating relations and the state of trading environment, we find that these technical analysis support and resistance levels still have significant ability to explain limit prices with high cumulative depth. 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 aforementioned cointegrating relations, as well as other trading environment factors.

This study represents a crucial step towards reconciling the long-standing disagreement between practitioners and academics about the value of technical analysis by demonstrating that technical analysis need not be inconsistent with market efficiency. In an efficient market, the value of technical analysis lies in its ability to minimize transaction costs. For example, market participants can use the results of this study to determine where depth is likely to be positioned on the limit order book, which will allow them to better estimate and minimize potential transaction costs by strategically placing limit orders or by splitting orders to reduce price impact.

While we investigate the explicit link between technical analysis and liquidity provision on the limit order book, we view this connection much more broadly. We note that technical analysis is often used in markets that do not have limit order books, e.g., foreign currency markets.⁴ Whether limit orders are placed explicitly or implicitly, through unobservable order placement strategies, we believe that technical analysis is related to the provision of liquidity. Therefore, despite the move to make limit order books more transparent (e.g., Openbook at the NYSE and SuperMontage on Nasdaq) the ability of market participants to work an order through a floor broker or dealer clouds the transparency of liquidity provision and validates the applicability and importance of this link.

It is important to emphasize that we do not test market efficiency, nor do we examine the ability of technical analysis to predict returns. Consequently, we do not claim that markets are either efficient or inefficient, nor do we claim that technical analysis can be used *only* to locate liquidity on the book, but rather that it is one source of value. Furthermore, we do not argue that technicians currently implement technical trading strategies with the goal of minimizing

⁴ 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.

transaction costs. To the contrary, both anecdotal evidence and work by Fung and Hsieh (2001), for example, suggest that many technical traders are “trend followers,” who bet on large price changes. We suspect that many technicians are unaware that technical analysis can be used to locate liquidity.

The remainder of the paper is organized as follows. Section 2 describes the data used in the analysis, as well as the limit order book estimation technique. Section 3 describes the limit order book liquidity measures we employ, and section 4 defines the technical analysis rules we consider. Section 5 discusses our results along with some robustness checks and also presents our causality investigation. Section 6 concludes.

2. Data and Limit Order Book Estimation

We construct estimates of limit order books using NYSE SuperDOT order and quote data, provided by the NYSE. These data contain quote, transaction, and order information for 110 NYSE stocks from July 1997 through September 1997.⁵

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 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 data set 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 TORQ data set. 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

⁵ 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).

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.

3. Limit Order Book Liquidity Measures

While many measures of limit order book liquidity can be constructed, our focus is on cumulative depth. Cumulative depth provides an estimate of the stated demand and supply functions for the stock. The level and slope of these functions succinctly measures the relative position of depth on the limit order book. Therefore, cumulative depth is a powerful and convenient way to summarize the state of liquidity on a limit order book.

We consider two specific measures of cumulative depth on the limit order book. The first is the limit price on each side of the limit order book at which the largest number of shares is placed (i.e., the buy and sell limit price mode). The limit prices with the largest increases in cumulative depth on each side of the limit order book implicitly represent the greatest obstacles or, in other words, the most significant support or resistance levels. 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) as the mode on the buy side (sell side) of the limit order book. This is done to eliminate noise in the limit order book mode estimates that come from large, stale orders that are far from the current quoted bid and ask.

Our second measure of cumulative depth reflects the relation between trading volume and limit order book depth. This measure is defined as the highest (lowest) limit buy (sell) price with cumulative depth of at least 15% of the average daily trading volume for the given stock. As such, it reflects the last share price that an investor would receive (pay) when selling (buying) 15% of the average daily volume in that stock from the limit order book. We choose 15% because orders submitted through SuperDOT constitute approximately 30% of total order flow by volume.⁶ If we assume that this activity is evenly split between the buy and sell sides of the limit order book, we might expect that on a typical day, each side of the limit order book would handle 15% of average daily trading volume. Therefore, our second measure can be thought of as the buy-side and sell-side limit prices containing enough cumulative depth to satisfy a typical day's liquidity demand.

Figure 1 provides an illustration of the relation among quoted prices and limit order book prices with high cumulative depths. The chart covers a one-week period August 25-29, 1997 for Dresser Industries. Panel A displays our first measure of cumulative depth – the limit prices on each side of the market with the most number of shares (thick lines) – along with the quoted bid and ask prices. 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 our second measure of cumulative depth – the limit prices containing cumulative depth equal to 15% of average daily trading volume – with the quoted bid and ask prices. It is important to note that the volume-based cumulative depth measures are not always

⁶ Although this percentage was computed for the time period covered by the original TORQ (1990-1991) sample, it is the best available estimate.

defined. In other words, it is sometimes impossible to sell or purchase 15% of average daily trading volume from the limit order book.

4. Technical Analysis Rules

Much of technical analysis involves pattern recognition using specific frequency (intra-day, daily, weekly) charts that display opening, high, low, and closing prices, as well as trading volume in some form. While there are many trading strategies under the rubric of technical analysis, we focus on two rules, support and resistance levels and moving average forecasts. We choose these rules because both represent standard technical analysis techniques that are widely used in practice. In addition, many alternative technical rules are variations on those we consider. For example, formations such as head-and-shoulders, pennants, round-tops, etc., are all variations on the principles underlying support and resistance levels. It is worth reiterating that we do not test whether these rules can successfully predict future price movements (i.e., whether they should truly be used as *trading* rules), but rather whether they can be used to uncover information about liquidity on the limit order book.

The fact that technical analysis remains more of an art than a science creates methodological difficulties.⁷ Using the same data, different technicians may arrive at different forecasts because there are no definitive rules for their construction. The fact that technicians often incorporate their own subjective views into their forecasts exacerbates this problem. Given the absence of explicitly defined 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 specific technicians, the existence of a relation between our rules and liquidity on the limit order

⁷ 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 Besignor (2000).

book would demonstrate the ability of past prices to relay information about the current liquidity profile.

At each half-hour interval, we define the support level as the lowest bid price that has been attained at least twice during the preceding week (i.e., during the previous 70 half-hour observations). This definition reflects the fact that technicians commonly use recent lows to determine support levels. If no price level meets the specified criterion at a given time, the support level is not defined. Resistance levels are defined analogously (see Equation 1 below).

$$\text{Support Level}_t = P_t^S = \begin{cases} \min(\text{Bid}_{t-w}, \dots, \text{Bid}_t) \text{ provided } P_t^S = \text{Bid}_i = \text{Bid}_j \text{ for some } i, j \text{ in } (t-w, t) \\ \text{undefined otherwise} \end{cases} \quad (1)$$

$$\text{Resistance Level}_t = P_t^R = \begin{cases} \max(\text{Ask}_{t-w}, \dots, \text{Ask}_t) \text{ provided } P_t^R = \text{Ask}_i = \text{Ask}_j \text{ for some } i, j \text{ in } (t-w, t) \\ \text{undefined otherwise} \end{cases}$$

Our moving average indicator is defined by the relation between the long-run and short-run moving averages of quoted midpoint prices. The long-run moving average is defined over the past two weeks, while the short-run moving average is defined over the past 2½ days. If the short-run moving average exceeds (falls below) a band around the long-run moving average, then the observation is assigned the value 1 (-1), otherwise the observation is assigned a neutral value of zero. As explained by Brock et al. (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.

$$\text{Moving Average Indicator} = \begin{cases} +1 \text{ if } MAS_t > MAL_t + B \\ -1 \text{ if } MAS_t < MAL_t - B \\ 0 \text{ otherwise} \end{cases} \quad (2)$$

$$\text{where } MAS_t = \frac{\sum_{k=1}^S P_{t-k}}{S},$$

$$MAL_t = \frac{\sum_{k=1}^L P_{t-k}}{L},$$

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.

Figure 2 provides an illustration of the technical analysis measures that we investigate, again for Dresser Industries. The technical analysis support and resistance levels and the quoted bid and ask prices for the same one-week period as in Figure 1 are displayed in Panel A. 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 that are common in traditional technical forecasts. First, the technical analysis measures tend to move with the quoted bid-ask spread (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, they are not always defined. 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. Recall that our indicator variable is non-zero whenever the short-run average is more than $\pm 3\%$ from the long-run average.

The limit order book variables in Figure 1 and the technical analysis measures in Figure 2 display a number of similar characteristics. Figure 3, which contains both the limit order book depth and the technical analysis prices, allows for a more direct comparison of these series. The graph in Panel A demonstrates the way the technical resistance level and the corresponding limit

order book liquidity series move together as prices increase over the period. The close relationship between the limit order book mode and the technical resistance level is particularly striking. Panel B displays the analogous series on the support side. Despite the volatility of the limit order book mode series, we still observe strong relationships between limit order book liquidity and the technical analysis rule.

5. Results

5.1 General Properties of Technical Analysis Rules and Liquidity Provision Measures

Table 1 begins to formalize the relations displayed in Figures 1 through 3 for our entire sample. Panel A displays statistics on the support/resistance 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 (resistance) levels are defined for 50.4% (50.6%) of the observations in our sample. When they exist, resistance prices exceed the quoted midpoint by \$0.77 on average, with a median difference of \$0.34. Similarly, the quoted midpoint price exceeds the support level by \$0.84 on average, with a median difference of \$0.44. 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, consistent with Figure 3. Median differences between quoted prices and the limit order book mode measure are very similar to those reported for the support and resistance levels, as are average differences using the 15% trading volume measure. A notable characteristic of all these measures is the tendency to cluster disproportionately on whole dollar prices, then half-dollar prices, then quarter-dollar increments, etc.⁸ The limit order book measures tend to cluster at half-dollar increments much

⁸ The tendency of prices to cluster has been documented elsewhere in the literature. Niederhoffer and Osborne (1966) were perhaps the first to discuss this feature of price data. More recently, Hasbrouck (1999) examines the

more than our technical rules, however. Between 38 and 46% of limit order book mode observations occur on even half-dollars, while only 23-27% of the technical analysis prices fall on these increments. This difference is probably partially responsible for the median (mean) absolute deviations between technical prices and limit order book modes of \$0.25 (\$0.77) or more.

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), however, the summary statistics suggest that these non-neutral periods last for approximately 4 days and are associated with negative (positive) median returns. This suggests that the moving average indicator may be associated with price reversals of quote midpoints.

5.2 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 conduct further statistical analysis of these series. Because each of the series is specified in terms of price levels, the existence of unit roots becomes a concern. Time-series with unit roots present an econometric challenge because inference, both within and between series, is altered. Table 2 displays the results of unit root tests on the limit order book and technical analysis rules on a stock by stock basis. The results using both Dickey-Fuller and Phillips-Perron tests, which allow for serial correlation in the error terms, show overwhelmingly that the series in question fail to reject the null hypothesis of a unit root. Therefore, any investigation of a relation between these limit order book and technical analysis series must take into account the possible presence of unit roots.

clustering of discrete transaction price series, and Osler (2002) documents clustering in stop-loss and take-profit orders in currency markets.

We also recognize that these series are likely to be cointegrated, since there are bounds on how much the prices on the limit order book will realistically differ from rules based on quoted prices. If the series are cointegrated, then imposing the cointegrating relations creates covariance stationary series, thereby eliminating problems associated with the unit roots. We test for cointegration using two different tests, the results of which are contained in Table 3. The Dickey-Fuller test imposes the cointegrating vector (1,-1), which assumes that the simple difference between the two series is covariance stationary. In contrast, the Phillips-Ouliaris test is based on an estimated, rather than predetermined, cointegrating vector.

The results of both tests confirm that the majority of the technical price and limit order book liquidity series are cointegrated. Dickey-Fuller tests (Panel A) indicate that between 57 and 84% of these series are cointegrated at a 5% significance level. Phillips-Ouliaris tests (Panel B) show that between 51 and 70% of the series are cointegrated. The results in Table 3 also demonstrate that the limit order book and quoted midpoint series are cointegrated, as are the technical analysis levels and the quoted midpoint.⁹

Econometrically, the fact that these series are cointegrated identifies a long-run relation between them, which is consistent with our model. However, we admit that these results may not be surprising given the nature of the limit order book. In particular, the possibility exists that these results may simply be a reflection of a mechanistic connection between the limit order book and the quoted prices. To account for this possibility, we test whether the technical analysis support and resistance levels are related to the limit order book depth variables *after*

⁹ The results for the technical price series are weaker, but still suggest that these series are cointegrated, at least for a subset of the stocks in the sample. Dickey-Fuller tests that include a non-zero drift term (final columns of Panel A) provide stronger support for cointegration between the quoted midpoint and the support and resistance levels. This may stem from the fact that these series can slowly drift apart, with deviations persisting for long periods of time (relative to our window).

imposing these cointegrating relations. Specifically, by taking the difference between our series of interest (limit order book depth measures and technical analysis measures) and the midpoint of the quoted spread we explicitly account for these cointegrating relations, thereby constructing covariance stationary series that can be used in ordinary least-squares regression analysis. If these differences are significantly related in a regression framework it would imply that the relation between technical analysis and liquidity provision extends beyond a simple connection to the quoted prices.

Table 4 shows the regressions relating the transformed series. Several specifications are shown for each level. Differences between each limit order book liquidity variable and the quoted midpoint are regressed on the technical analysis levels independently in univariate regressions, as well as conditioning on current market characteristics. 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 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.¹⁰

The control variables are important for a number of reasons. First, they summarize 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 all trading on one side of the market. Lastly, work by Mingelgrin (2000) and Gervais, Kaniel and Mingelgrin (2001)

¹⁰ 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.

show that volume and net orderflow measures contain information useful for predicting future returns both at the intra-daily and daily levels.

The univariate regression results (models 1, 3, 5, and 7) demonstrate positive and significant relationships between the technical and limit order book liquidity levels. The multivariate regression results (models 2, 4, 6, and 8) 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 49-73% of the stocks in our sample when the limit order book mode is used as the dependent variable. When the volume-based liquidity measure is used, the technical levels are statistically significant for 28-58% of the stocks, which 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 4 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 limit price corresponding to 15% of average daily trading volume.

Figure 4 provides a graphical illustration of the correspondence between technical support and resistance levels and peaks in depth on the book. Panel A contains graphs of the distribution of cumulative limit order book depth on the ask side of the market for three different series. These series are created by grouping observations according to the distance between the technical resistance level and the quoted ask price, where observations in which the resistance

level exceeded the ask by $\$1/8$ are placed into one group, by $\$1/4$ into another group, and by $\$1/2$ into the final group.¹¹ The depth at each successive limit price is aggregated across all observations in each group and the empirical distribution of cumulative depth is computed. The process is then repeated using the technical support level and the quoted bid price (Panel B).

Consistent with the findings above, technical levels tend to lie at limit prices with large increases in depth. Consider, for example, the series in Panel A that contains observations for which the technical resistance level exceeds the quoted ask by $\$1/8$. The steepest increase in depth for this group occurs at a price $\$1/8$ higher than the ask – exactly where the technical level lies (see bolded segment of distribution). The graphs of the other series in Figure 4 display analogous patterns. As in the regression analysis, this relationship is strongest on the ask side of the market.

Figure 4 offers strong visual evidence that technical analysis can be used to locate liquidity, but there are limitations to the inferences that should be drawn. Because the empirical distributions were created using all observations in each category, stocks with more shares on their books receive more weight. Accordingly, Figure 4 tells us little about the accuracy with which technical support and resistance levels locate liquidity for individual stocks in the sample.

The summary statistics in Table 5 provide a more detailed description of the magnitude of differences between technical levels and limit order book modes. The results in Table 5 suggest that the graphs in Figure 4 are fairly representative of all stocks in the sample. A substantial fraction of the technical analysis levels are within $\$1/8^{\text{th}}$ of the limit order book mode (Panel A). Furthermore, the results show that even when the technical price and the limit order

¹¹ To standardize the relation between the quoted and limit order book prices, we condition on the specialist reflecting the best price on the limit order book in the quotes. This resulted in a 23% decrease in the sample size. In addition, we created 20 different series representing differences between $\$1/16$ and $\$20/16$. The graphs of all series are consistent with those presented in Figure 4 and are available from the authors upon request.

book mode differ in dollar terms, there is often no difference in the number of shares at these prices.¹² Therefore, on average, the technical analysis levels 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).

5.3 Moving Average Indicator

Given the nature of the moving average indicator variable, it does not have a unit root and therefore cannot be part of a cointegrating relation. However, the methodology we employed for the support/resistance levels can also be used for the moving average indicator. Again exploiting the natural relation between the limit order book and quoted prices, we difference the limit order book depth variables with the quoted midpoint, imposing the cointegrating relation as we did above. 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 order prices to be closer to the quoted midpoint and the buy-side prices to be farther away. In other words, the moving average indicator may reveal information about the “skewness” of liquidity on the two sides of the book.

¹² This commonly occurs when the technical level falls into a gap between prices on the limit order book. For example, suppose that the buy-side limit order book mode of \$30 offers a cumulative depth of 8000 shares and that the next buy order on the book has a limit price of \$29.75. If the technical support level lies between the two orders, at say \$29.81, then the cumulative depth available at the support level is equal to that available at the mode, in spite of the price difference.

The results of the regressions are contained in Table 6. As in Table 4, several different model specifications, both univariate and multivariate, are presented. In all cases, the moving average indicator variable is negative and statistically significant for a substantial fraction (31-58%) of the stocks in the sample. The negative coefficient confirms our hypothesis that a positive (negative) moving average indicator corresponds to a shift in quoted prices toward (away from) the sell-side liquidity levels and away from (toward) the buy-side levels.¹³ As in Table 4, the relationships between the technical analysis prices and the limit order book mode appear stronger than those between the technical levels and the volume-based depth measure.

5.4 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 to 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. This redefinition is likely to generate more extreme (i.e., farther from the quoted prices) technical analysis measures, as there is a larger pool of prices to consider.

We indeed find that the technical analysis levels constructed using the longer window tend to be farther from both the quoted prices and the limit order book depth measures, but the

¹³ Differences between buy-side limit order book variables and the quoted midpoint are negative, so a negative coefficient implies an increase in this distance.

relationship between the technical and limit order book levels remains strong. The results of the unit root and cointegration tests are similar to those in Tables 2 and 3 and, with the exception of slightly larger coefficient estimates on the technical analysis variables and constant terms, the regression results are indistinguishable from those reported in Table 4.

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. Using this definition, the moving average indicator variable assumes a non-zero value only 10% of the time, compared to 23% from Table 1, and the mean and median duration of these values is shortened as well. This is perhaps not surprising because a very large one-day price change is required to exceed (or fall below) the 3% band around the long-run average. Consistent with the strength of this signal, the economic and statistical significance of the moving average variables in Table 6 increases when the shorter window is used, while the other coefficients remain unchanged.

Third, we reduce the band surrounding the long-run moving average from 3% to 1%. This change makes it easier for the short-run average to deviate sufficiently from the long-run average to trigger the moving average indicator variable. Accordingly, the indicator assumes non-zero values almost 55% of the time, much more frequently than in Table 1, which results in slightly higher R^2 for the moving average regressions. The results are otherwise identical to those presented in Table 6.

The robustness checks described above demonstrate 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. Recall that we use two complementary measures of cumulative depth on the limit order book. The limit order book mode 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 price on each side of the market. Taking the closer of two modes eliminates the noise that stems from large, stale limit orders that are far from the current quoted bid and ask. The daily volume measure is the highest (lowest) limit price on the buy (sell) side of the book having cumulative depth that equals or exceeds 15% of average daily trading volume.

We rerun the analysis using two 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 – and second, we use the limit price associated with 10% rather than 15% of average daily trading volume. In both cases, the results are similar to those presented here. The only notable difference is that the true mode tends to be farther from the technical levels and the quotes than the “closer” mode used in the tables, which occurs by construction. The results using the 10% measure are almost identical to those using the 15% measure.

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 re-run our entire analysis on the original TORQ data, which cover 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 $1/16^{\text{th}}$ rather than $1/8^{\text{th}}$. 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.

Table 7 demonstrates that the relationships we document in our 1997 sample exist in the original TORQ sample period as well. A comparison across periods reveals that the results are both qualitatively and quantitatively similar, which demonstrates that our findings are robust to time and substantial changes in the trading environment.

5.5 Causality Tests

The results thus far demonstrate a strong relationship between technical analysis measures and liquidity provision, but the question of causality naturally arises: Does technical analysis identify prices with depth already in place on the limit order book, or do liquidity providers submit limit orders at the place where technicians forecast support and resistance levels?

We address this question using two complementary approaches. 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 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 limit order book.

The Granger-causality tests analyze whether 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. The regression is then reversed to see if the limit order book depth measures have explanatory power for the support and resistance levels after controlling for lagged technical levels.

Table 8 displays the results of these bivariate Granger-Causality tests. Tests were conducted using two and five half-hour lags. In both cases, the limit order book liquidity measures are more likely to Granger-cause support or resistance levels (LOB to Tech) than the reverse (Tech to LOB), although no statistically significant causality can be documented for over half the sample. In addition, Osler (2002) claims that technical trading and order placement strategies may be jointly determined. This is the case for between 5 and 9% of the stocks in our sample. Overall, the findings in Table 8 demonstrate that when the direction of the causality can be determined, it is more likely to flow from the limit order book to the technical rules than vice-versa. This suggests that the relation is more likely to be characterized by technical analysis identifying limit prices with high depth than by traders placing limit orders where support and resistance levels have been identified.

The Granger-causality tests may be affected by the inherent persistence of some orders on the limit order book. Because many of the distant limit orders on the book are good-until-canceled, they remain a permanent fixture on the book until executed or canceled. This creates pockets of depth on the book that change relatively slowly. As a result, we conduct a complementary analysis of the *flow* of orders to the book, by analyzing the distribution of newly placed orders, rather than the number of shares already sitting on the book. Our metric is the weighted dispersion of an order, defined as the dollar difference between the limit price of a new order and the relevant reference price (quote, technical analysis level, or limit order book mode), weighted by the size of the order in shares. Negative dispersion values represent orders placed

below the relevant reference price, and positive values represent orders placed above the reference price.

When analyzing the placement of orders around technical analysis levels, we must account for the fact that limit orders are often positioned close to the quoted bid or ask prices. As shown in Panel A of Table 9, approximately 60% of all newly submitted limit orders are placed within \$0.25 of the quoted prices. One way to disentangle the impact of the technical analysis level from the tendency to place orders near the quotes is to separate the orders into groups based on the proximity of the technical level to the quotes. Technical or limit order book prices are classified as “near” (“far”) if they are less (more) than \$0.50 from the respective quote.

Panel B summarizes the dispersion of new limit orders relative to the technical levels and the limit order book modes. Not surprisingly, when the technical level is near the quotes, new orders tend to cluster around the technical level. In contrast, when the level is far from the quotes, limit orders do not cluster around the technical price as evidenced by only 7.6% (9.2%) of orders being placed within \$0.25 of the technical support (resistance) level. This suggests that the clustering that occurs around the technical level in the “near” group is likely due to the tendency to place orders near the quotes, not to the technical level. An analogous result holds for the dispersion of limit orders relative to the limit order book mode.

The lack of any significant relation between order positioning and technical analysis levels in Panel B is at first surprising, but may simply result from treating every technical analysis observation identically. Suppose that for most observations, investors place orders with little regard for the technical level, but they concentrate orders around new technical analysis prices when they *change*. In this case, a study of limit order placement across the entire sample is likely to yield few results.

To investigate this possibility, we focus on periods surrounding changes in technical analysis levels. To avoid “fleeting” levels, we identify cases in which the new technical analysis level remains constant for at least one day. We then analyze the placement of limit orders relative to this new level in the one-day window preceding and the one-day window following the technical analysis change. Given the findings in Panel B, we condition on the distance between the quoted prices and *both* the technical level and limit order book mode. Specifically, technical levels are considered “far” if both the technical analysis and limit order book mode prices are more than \$0.50 from the same side quote. The results are presented in panel C.

Again we find a concentration of limit orders close to the technical level when the level is near the quotes. However, unlike in Panel B, the “far” sample displays distinct clustering around both the support and resistance levels, particularly in the pre-period. Most striking are the large fractions of limit orders (18.20% on the support side and 21.97% on the resistance side) placed just ahead of (i.e., closer to the quotes than) the upcoming technical analysis levels in the period preceding the change. Differences between each of these fractions and their unconditional counterparts are statistically significant at the 1% level. The post-period results show that smaller, but still statistically significant, fractions of limit orders are placed ahead of the support level (11.41%) and behind the resistance level (12.05%).

The results in Panel C 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. The results are consistent with technical analysis being able to identify where limit orders have been placed.

In summary, the analysis of limit order placement suggests that limit orders are placed before the creation of a new technical analysis level and are positioned slightly ahead of (closer

to the quotes) the forthcoming technical level. Limit order placement after the technical level is much more diffuse. These findings reinforce the earlier findings of the Granger-causality tests suggesting that the connection between technical analysis and limit order book depth is driven by technical analysis being able to identify locations with high cumulative depth already in place on the limit order book.

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 in the finance literature. If technical analysis is in fact valuable, then it appears that markets must be inefficient. Alternatively, if markets are efficient, then it seems that practitioners' claims are unfounded. Despite numerous studies, this conflict has yet to be resolved.

In this paper, we make no claims about the efficiency of the market or about the ability of technical analysis to predict future returns. Instead, we propose a previously unexamined explanation in which the value of technical analysis is not in direct conflict with market efficiency. Specifically, we document a link between technical analysis measures and the amount of liquidity on limit order books. The fact that technical analysis and market efficiency can coexist in this setting hinges on the subtle, yet critical, difference between liquidity provision and future prices. While prices are determined by the interplay between the amount of liquidity supplied *and* the amount of liquidity demanded in the market, here we focus solely on one side of this relation – liquidity supplied on the limit order book.

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. Consistent with our first hypothesis, we find that support and resistance levels are cointegrated with limit prices containing high cumulative depth. Furthermore, after controlling for these cointegrating relations and the state of trading environment, we find that support and resistance levels still have significant ability to explain these limit order book prices. Granger-causality tests show that, when causality can be determined, the limit order book depth variables tend to cause the technical analysis variables more than the reverse. This implies that technicians can use their forecasts to find depth in place on the limit order book, although we believe that they may be largely unaware of this benefit. We find further support for the ability of technical analysis to locate liquidity in our investigation of order placement activity surrounding technical analysis changes. 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 aforementioned cointegrating relation, as well as other factors.

The results of this study show that the use of technical analysis need not be at odds with academic theories of market efficiency by demonstrating a link between technical trading rules and the level and nature of liquidity provision. Although we do not necessarily believe that practitioners perform technical analysis in hopes of locating liquidity, the value of this relationship is immediate and obvious. Armed with this information, market participants can better estimate potential transaction costs as well as implement trading strategies, like limit order submission or order splitting strategies, to minimize transaction costs.

Our findings are also important for regulators charged with managing the future of our trading venues. Understanding the limit order book supply and demand schedules as revealed through technical analysis will be an important consideration when comparing the costs and

benefits to liquidity suppliers and liquidity demanders of limit order book transparency or exchange consolidation.

References

- Alexander, Sidney S. 1961. Price movements in speculative markets: Trends or random walks. *Industrial Management Review* II: 7-26.
- Allen, Franklin, and Risto Karjalainen. 1999. Using genetic algorithms to find technical trading rules. *Journal of Financial Economics* 51: 245 – 271.
- Bensignor, Rick. 2000. *New thinking in technical analysis: Trading Models from the Masters*. Bloomberg Press, Princeton, New Jersey.
- Bessembinder, Hendrick and Kalok Chan. 1998. Market efficiency and the returns to technical analysis. *Financial Management* 27: 5-17.
- Blume, Lawrence, David Easley and Maureen O’Hara. 1994. Market statistics and technical analysis: The role of volume. *Journal of Finance* 49: 153-181.
- Brock, William, Josef Lakonishok and Blake LeBaron. 1992. Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance* 47: 1731-1764.
- Brown, David and Robert H. Jennings. 1989. On technical analysis. *Review of Financial Studies* 2: 527-551.
- Edwards, Robert D and John Magee. 1998. *Technical analysis of stock trends*. John Magee Inc., Boston, Massachusetts.
- Fama, Eugene. 1991. Efficient Capital Markets: II. *Journal of Finance* 46: 1575-1617.
- Fama, Eugene and Marshall Blume. 1966. Filter rules and stock market trading. *Journal of Business* 39: 226-241.
- Gervais, Simon, Ron Kaniel and Daniel Mingelgrin. 2001. The high-volume return premium. *Journal of Finance* 56: 877-919.
- Goldstein, Michael A. and Kenneth A. Kavajecz. 2000. Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the NYSE. *Journal of Financial Economics* 56: 125-149.
- Hasbrouck, Joel. 1992. Using the TORQ Database. Working paper. Stern School of Business, New York University.
- Hasbrouck, Joel. 1999. Security bid/ask dynamics with discreteness and clustering. *Journal of Financial Markets* 2: 1-28.
- Fung, William, and David A. Hsieh. 2001. The risk in hedge fund strategies: Theory and evidence from trend followers. *Review of Financial Studies* 14: 313-341.

- Jegadeesh, Narasimhan. 2000. Discussion of Foundations of Technical Analysis, *Journal of Finance* 55: 1765-1770.
- Kavajecz, Kenneth A. 1999. The specialist's quoted depth and the limit order book. *Journal of Finance* 52: 747-771
- Lee, Charles and Mark Ready. 1991. Inferring trade direction from intraday data. *Journal of Finance* 46, 733-746.
- Lo, Andrew W., Harry Mamaysky, and Jiang Wang. 2000. Foundations of Technical Analysis: Computations Algorithms, Statistical Inference, and Empirical Implementation. *Journal of Finance* 55: 1705 – 1765.
- Mingelgrin, Daniel H. 2000. The informational content of high-frequency trading activity. Working paper, The Wharton School, University of Pennsylvania.
- Murphy, John J., 1999. *Technical analysis of the financial markets*. New York Institute of Finance, Paramus, New Jersey.
- Neely, Christopher, Peter Weller, and Robert 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, Salih N. 1991. Naïve trading rules in financial markets and weiner-kolomgorov prediction theory: A study of technical analysis. *Journal of Business* 64: 549-571.
- Neftci, Salih N. and Andrew Policano. 1984. Can chartists outperform the market? Market efficiency tests for "technical analyst." *Journal of Futures Markets* 4: 465 - 478.
- Niederhoffer, Victor, and M. F. M. Osborne. 1966. Market making and reversal on the stock exchange. *Journal of the American Statistical Association* 61: 897-916.
- Osler, Carol L., 2002, Currency orders and exchange-rate dynamics: An Explanation for the Predictive Success of Technical Analysis, *Journal of Finance*, forthcoming.
- Pring, Martin J., 1991. *Technical analysis explained*. McGraw-Hill Inc. New York, New York.
- Ratner, Mitchell and Ricardo P.C. 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, Mark J. 1997. Profits from technical trading rules. Working Paper, University of Wisconsin-Madison.
- Sullivan, Ryan, Allan Timmermann and Halbert White. 1999. Data-snooping, technical trading rule performance, and the bootstrap. *Journal of Finance* 54: 1647-1691.

Treynor, Jack L. and Robert Ferguson. 1985. In defense of technical analysis. *Journal of Finance* 40: 757-775.

Table 1
Properties of Technical Analysis Rules

Descriptive statistics on the properties of the technical trading rules and limit order book liquidity measures are contained below. 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. “LOB mode” represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. “LOB 15% ADV” is defined as the last share price from the limit order book that an investor would pay (receive) when buying (selling) a quantity equal to 15% of the average daily trading volume.

Panel A: Support and Resistance Levels

	Support Levels			Resistance Levels		
	Technical Rule	LOB Mode	LOB 15% ADV	Technical Rule	LOB Mode	LOB 15% ADV
% Time Defined	50.4%	99.0%	86.9%	50.6%	99.5%	87.9%
Distance to Qtd Mdpt in \$						
Mean	-0.837	-1.428	-0.804	0.767	0.849	0.627
Median	-0.438	-0.313	-0.219	0.344	0.313	0.219
% on Half (Whole) Dollar	23.4 (6.1)	38.3 (26.4)	26.7 (16.0)	27.1 (8.4)	46.1 (32.9)	32.3 (21.6)
Tech – LOB Level						
Mean		1.129	0.717		0.771	0.563
Median		0.375	0.313		0.250	0.250

Panel B: Moving Average Indicator

	+1	0	-1
Frequency	13.9	77.1	9.0
Duration			
Mean	4.0 days	12.8 days	3.5 days
Median	3.8 days	7.0 days	3.6 days
Total Return			
Mean	1.8%	22.3%	2.0%
Median	-0.6%	3.2%	0.6%

Table 2
Unit Root Tests

Unit root tests were conducted on a stock-by-stock basis for each series using Dickey-Fuller parameter-based statistics (DF Stat), Dickey-Fuller T-statistics (DF T-stat), and Phillips-Perron rho (PP Rho) and tau (PP Tau) statistics. Phillips-Perron tests allow for serial correlation in error terms. Technical support/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. “LOB 15% ADV” is defined as the last share price from the limit order book that an investor would pay (receive) when buying (selling) a quantity equal to 15% of the average daily trading volume. For all tests, a 5% critical level was used to assess significance. Dickey-Fuller T-statistics are undefined for the technical levels in Panel B because the parameters were estimated without error in all cases (because support and resistance levels are piecewise constant where they are defined).

	Median Coefficient	Median Statistics				% Stocks w/Unit Root			
		DF Stat	DF T-stat	PP Rho	PP Tau	DF Test	DF TTest	PP Rho	PP Tau
Panel A: Limit Order Book									
Bid Side									
Mode	1.001	0.047	0.612	-0.048	-0.108	100.0	100.0	100.0	100.0
15% of ADV	1.001	0.105	0.988	0.011	0.050	100.0	99.0	100.0	100.0
Sell Side									
Mode	1.000	-0.058	-0.957	0.006	0.020	100.0	93.1	100.0	100.0
15% of ADV	0.999	-0.089	-0.983	0.012	0.088	100.0	85.3	100.0	100.0
Quoted Spread Midpoint	1.000	0.000	0.026	0.030	0.357	100.0	99.0	100.0	100.0
Panel B: Technical Trading Rules									
Support	1.000	0.000	undef.	0.035	0.681	100.0	N/A	100.0	99.1
Resistance	1.000	0.000	undef.	0.027	0.460	100.0	N/A	100.0	99.1

Table 3
Cointegration Tests

Cointegration tests were conducted on a stock-by-stock basis using Dickey-Fuller T-tests (Panel A) and Phillips-Ouliaris rho and tau test statistics (Panel B). The tests in Panel A assume a cointegrating vector equal to (1,-1) and were conducted both with and without a constant term in the regression. Technical support/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. "LOB 15% ADV" is defined as the last share price from the limit order book that an investor would pay (receive) when buying (selling) a quantity equal to 15% of the average daily trading volume. For all tests, a 5% critical level was used to assess significance.

Panel A: Dickey-Fuller tests assuming cointegrating vector = (1,-1)

Series 1	Series 2	Median Statistics: Dickey Fuller T-test		% Cointegrated	
		No Constant	Constant	No Constant	Constant
Support Level	LOB Mode (Buy-side)	-2.365	-2.929	57.8	72.9
Support Level	LOB 15% ADV (Buy)	-2.378	-3.413	57.8	83.8
Support Level	Quoted Midpoint	-0.896	-1.751	10.8	57.1
Resistance Level	LOB Mode (Sell-side)	-2.363	-2.858	56.9	73.1
Resistance Level	LOB 15% ADV (Sell)	-2.410	-3.273	58.8	78.6
Resistance Level	Quoted Midpoint	-1.089	-1.901	22.5	55.1
LOB Mode (Buy-side)	Quoted Midpoint	-2.551	-3.345	68.6	84.0
LOB 15% ADV (Buy)	Quoted Midpoint	-4.384	-6.007	83.3	88.7
LOB Mode (Sell-side)	Quoted Midpoint	-2.768	-3.576	64.7	83.8
LOB 15% ADV (Sell)	Quoted Midpoint	-4.044	-5.500	81.4	89.8

Panel B: Phillips-Ouliaris tests using estimated cointegrating vectors

Dependent Variable	Independent Variable	Median Estimates		Median Statistics		% Cointegrated	
		Intercept	Coeff	Rho	Tau	Rho	Tau
Support Level	LOB Mode (Buy)	12.036	0.488	-20.75	-3.417	50.9	51.9
Support Level	LOB 15% ADV (B)	7.911	0.683	-24.97	-3.857	63.6	61.8
Support Level	Quoted Midpoint	3.512	0.805	-15.24	-2.900	23.6	22.7
Resistance Level	LOB Mode (Sell)	9.085	0.647	-32.80	-4.207	68.2	68.2
Resistance Level	LOB 15% ADV (S)	6.212	0.807	-29.88	-4.205	70.0	69.1
Resistance Level	Quoted Midpoint	4.410	0.845	-17.83	-3.055	40.0	36.4
LOB Mode (Buy)	Quoted Midpoint	5.063	0.782	-257.8	-12.37	98.1	97.2
LOB 15% ADV (B)	Quoted Midpoint	0.520	0.965	-299.3	-13.87	99.1	99.1
LOB Mode (Sell)	Quoted Midpoint	4.420	0.879	-227.3	-11.59	96.4	96.4
LOB 15% ADV (S)	Quoted Midpoint	1.276	0.970	-307.8	-14.03	99.1	99.1

Table 4
Support and Resistance Level Regressions

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. “LOB mode” represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. “LOB 15% ADV” is defined as the last share price from the limit order book that an investor would pay (receive) when buying (selling) a quantity equal to 15% of the average daily trading volume. Differences between each variable and the quoted midpoint are taken to ensure stationarity. LOB buy- (sell-) side regressions use support (resistance) level technical measures.

Model	Dependent Variable	Median Coefficient Estimates [Median p-values] (% Significant at the 5% Level)								Medn Adj R ²
		Intercpt	Tech - QM	Quoted Spread	Tot Qtd Depth	Current Volume	Total Volume	Curr Buy-Sell Asym	Total Buy-Sell Asym	
1	LOB (Buy) Mode – QM	-0.5641 [0.0000] (83.6)	0.2321 [0.0393] (50.9)							0.03
2	LOB (Buy) Mode – QM	-0.5604 [0.0083] (59.1)	0.2210 [0.0581] (49.1)	-0.5655 [0.2067] (26.4)	1.7E-5 [0.0307] (57.3)	1.6E-6 [0.2355] (20.0)	9.3E-8 [0.0919] (43.6)	0.0056 [0.4326] (5.5)	-0.0131 [0.1438] (39.1)	0.16
3	LOB (Buy) 15% ADV – QM	-0.4975 [0.0000] (99.1)	0.0243 [0.1868] (34.5)							0.01
4	LOB (Buy) 15% ADV – QM	-0.5449 [0.0016] (78.2)	0.0351 [0.1698] (28.2)	-0.5680 [0.0620] (47.3)	1.2E-5 [0.0006] (80.9)	6.8E-7 [0.2918] (17.3)	4.2E-8 [0.2180] (30.9)	0.0039 [0.4007] (10.9)	-0.0011 [0.2420] (22.7)	0.18
5	LOB (Sell) Mode – QM	0.2967 [0.0000] (85.5)	0.3315 [0.0034] (72.7)							0.09
6	LOB (Sell) Mode – QM	0.3734 [0.0079] (60.9)	0.3235 [0.0044] (68.2)	0.3690 [0.2925] (21.8)	-1.0E-5 [0.0162] (60.9)	-2.9E-7 [0.3539] (12.7)	-5.5E-8 [0.1515] (34.5)	-0.0031 [0.4963] (6.4)	-0.0182 [0.1665] (37.3)	0.20
7	LOB (Sell) 15% ADV – QM	0.3663 [0.0000] (94.5)	0.1146 [0.0319] (58.2)							0.04
8	LOB (Sell) 15% ADV – QM	0.3243 [0.0047] (72.7)	0.1081 [0.0681] (43.6)	0.5164 [0.0440] (50.9)	-1.3E-5 [0.0001] (80.0)	-1.2E-7 [0.4243] (15.5)	7.8E-11 [0.2254] (28.2)	0.0030 [0.3956] (11.8)	-0.0017 [0.1474] (32.7)	0.24

Table 5
Distribution of Price and Share Differences between Technical Analysis and Limit Order Book Measures

The sample distribution of differences between the technical trading rules and the limit order book liquidity measures is contained below. 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.

	Percentiles				
	20 th	40 th	50 th	60 th	80 th
Panel A: Difference between Technical Levels and LOB Mode					
LOB Buy Mode Price – Tech Support Price in \$	-1.000	-0.125	0.000	0.063	0.500
LOB Sell Mode Price – Tech Resistance Price in \$	-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

Table 6
Moving Average Regressions

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). “LOB mode” represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. “LOB 15% ADV” is defined as the last share price from the limit order book that an investor would pay (receive) when buying (selling) a quantity equal to 15% of the average daily trading volume. Differences between the LOB variables and the quoted midpoint are taken to ensure stationarity.

Model	Dependent Variable	Median Coefficient Estimates [Median p-values] (% Significant at the 5% Level)								Medn Adj R ²
		Intercpt	Moving Average	Quoted Spread	Tot Qtd Depth	Current Volume	Total Volume	Curr Buy- Sell Asym	Total Buy- Sell Asym	
1	LOB Mode (Buy) – QM	-0.8382 [0.0000] (100.0)	-0.1200 [0.0425] (50.9)							0.01
2	LOB Mode (Buy) – QM	-0.6531 [0.0001] (75.5)	-0.0985 [0.0694] (47.3)	-0.6531 [0.1272] (35.5)	1.5E-5 [0.0013] (71.8)	8.9E-7 [0.3383] (9.1)	2.8E-9 [0.1275] (41.8)	0.0028 [0.3628] (14.5)	-0.0395 [0.1151] (40.0)	0.10
3	LOB 15% ADV (B) – QM	-0.6297 [0.0000] (100.0)	-0.0159 [0.1254] (36.4)							0.00
4	LOB 15% ADV (B) – QM	-0.6834 [0.0000] (86.4)	-0.0028 [0.1836] (30.9)	-0.6016 [0.0257] (59.1)	1.8E-5 [0.0000] (91.8)	7.6E-7 [0.1493] (25.5)	2.5E-8 [0.1349] (35.5)	0.0068 [0.3219] (19.1)	-0.0110 [0.1669] (35.5)	0.14
5	LOB Mode (Sell) – QM	0.6133 [0.0000] (100.0)	-0.1276 [0.0152] (58.2)							0.01
6	LOB Mode (Sell) – QM	0.6080 [0.0000] (90.9)	-0.0973 [0.0230] (54.5)	0.4819 [0.1254] (36.4)	-1.1E-5 [0.0008] (74.5)	-2.9E-7 [0.2967] (27.3)	-8.2E-9 [0.0679] (46.4)	-0.0201 [0.2133] (25.5)	-0.0357 [0.0654] (50.0)	0.11
7	LOB 15% ADV (S) – QM	0.5267 [0.0000] (100.0)	-0.0371 [0.0907] (42.7)							0.01
8	LOB 15% ADV (S) – QM	0.4621 [0.0000] (89.1)	-0.0100 [0.1065] (39.1)	0.6701 [0.0013] (66.4)	-1.6E-5 [0.0000] (90.9)	-3.4E-8 [0.4040] (17.3)	-5.5E-10 [0.2314] (35.5)	-0.0001 [0.4494] (12.7)	-0.0102 [0.1311] (41.8)	0.18

Table 7
Comparison of Results: 1997 versus 1990-1991

The table compares the current results with the results of the identical analysis using the TORQ sample (November 1990 through January 1991). Note only select results are presented. The full set of 1990-1991 results is available from the authors upon request.

		1997			1990-1991		
Panel A: Summary Statistics							
		Median Distance to Qtd Mdpt			Median Distance to Qtd Mdpt		
Support	Tech. Rule	0.438			0.375		
	Mode	0.313			0.313		
	15%	0.219			0.313		
Resistance	Tech. Rule	0.344			0.313		
	Mode	0.313			0.313		
	15%	0.219			0.313		
		Median Distance to Tech Level			Median Distance to Tech Level		
Support	Mode	0.375			0.250		
	15%	0.313			0.250		
Resistance	Mode	0.250			0.250		
	15%	0.250			0.234		
Panel B: Cointegration Phillips – Ouliaris Rho Test							
		Coef.	Rho Stats	%signif	Coef.	Rho Stats	%signif
Support	Tech to Mode	0.488	-20.74	50.9	0.489	-21.36	53.1
	Tech to 15%	0.683	-24.97	63.6	0.558	-24.48	56.6
Resistance	Tech to Mode	0.647	-32.80	68.2	0.493	-24.70	53.8
	Tech to 15%	0.807	-29.88	70.0	0.591	-29.57	62.2
Panel C: Multivariate Support/Resistance Regressions							
		Coef.	P-values	%signif	Coef.	P-values	%signif
Support	Mode	0.2210	0.0581	49.1	0.3433	0.0024	67.8
	15%	0.0351	0.1698	28.2	0.1020	0.0814	45.8
Resistance	Mode	0.3235	0.0044	68.2	0.2366	0.0194	60.1
	15%	0.1081	0.0681	43.6	0.0895	0.0576	46.9
Panel D: Multivariate Moving Average Regressions							
		Coef.	P-values	%signif	Coef.	P-values	%signif
Support	Mode	-0.0985	0.0694	47.3	-0.1077	0.0100	64.3
	15%	-0.0028	0.1836	30.9	-0.0303	0.1570	35.0
Resistance	Mode	-0.0973	0.0230	54.5	-0.0784	0.0413	51.7
	15%	-0.0100	0.1065	39.1	-0.0364	0.1370	40.6

Table 8
Summary of Granger Causality Tests

Columns contain the percentage of stocks in the sample with the given direction of Granger causality at the 5% level based on F-tests on two and five lags respectively. “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.


Tech Level	LOB Variable	Direction of Causality (%)					
		Two Lags			Five Lags		
		LOB to Tech	Tech to LOB	Both	LOB to Tech	Tech to LOB	Both
Support	Mode (Buy)	17.3	5.5	9.1	17.3	9.1	6.4
Support	15% ADV (Buy)	9.1	17.3	4.5	11.3	13.2	7.5
Resistance	Mode (Sell)	30.0	13.6	4.5	25.5	8.2	5.5
Resistance	15% ADV (Sell)	18.2	13.6	9.1	22.2	13.9	7.4

Table 9
Distribution of Limit Order Placement

Columns contain percentages of limit orders placed within the stated dispersion interval. Dispersion is the dollar difference between the limit price of a placed order and the relevant reference price (quote, technical analysis level, or limit order book mode) weighted by the size of the order in shares. Negative (positive) dispersion values represent orders placed below (above) the relevant reference point. The shaded percentages in Panel C were tested for statistical differences from their unconditional and pre/post counterparts. All differences were significant at the 1% level.

		Dispersion Intervals								
		(-inf,-1.0)	[-1.0,-0.5)	[-0.5,-0.25)	[-0.25,0.0)	[0.0, 0.0]	(0.0, 0.25]	(0.25, 0.5]	(0.5, 1.0]	(1.0, inf)
Panel A: Dispersion Surrounding the Quoted Prices										
Bid		12.25	10.78	13.31	39.60	13.71	10.35			
Ask					10.04	13.16	37.25	15.01	11.46	13.09
Panel B: Dispersion Conditioning by Distance to the Quote										
Technical Analysis Level										
Near	Support	4.56	3.37	4.25	17.59	15.80	39.43	13.99	1.01	0.00
	Resistance	0.01	0.79	12.94	35.05	14.66	21.73	5.37	4.37	5.14
Far	Support	2.53	2.30	0.88	2.11	0.92	4.53	13.30	33.03	40.39
	Resistance	38.36	33.40	11.95	4.24	1.76	3.20	1.11	1.90	4.09
Limit Order Book Mode										
Near	Support	5.07	5.25	8.13	31.33	14.66	28.92	6.23	0.41	0.00
	Resistance	0.02	0.49	6.58	28.74	17.18	27.69	8.48	5.82	5.00
Far	Support	3.27	1.58	1.38	2.52	2.75	2.95	6.53	16.84	62.17
	Resistance	57.80	19.75	7.68	2.77	3.99	1.51	1.21	1.62	3.67
Panel C: Dispersion Surrounding a Technical Analysis Change Conditioning by Distance to the Quoted Prices										
New Technical Level – Near										
Pre	Support	6.70	3.62	4.35	16.58	10.11	20.48	11.95	10.52	15.70
	Resistance	10.95	8.12	11.78	19.50	12.22	18.19	4.45	2.95	11.85
Post	Support	1.82	1.79	2.17	11.88	15.98	33.21	16.56	9.56	7.03
	Resistance	6.59	9.73	13.61	33.36	12.39	15.24	2.56	2.56	3.96
New Technical Level – Far										
Pre	Support	32.29	3.36	1.23	3.85	3.55	18.20	16.19	13.11	8.22
	Resistance	22.36	16.97	7.63	21.97	8.24	7.66	7.36	0.95	6.88
Post	Support	3.87	4.22	2.03	4.99	5.13	11.41	14.63	26.14	27.57
	Resistance	19.53	37.50	11.17	7.82	2.64	12.05	1.70	2.25	5.35

Panel A: Buyside and Sellside Modes



Panel B: 15% of Average Daily Trading Volume

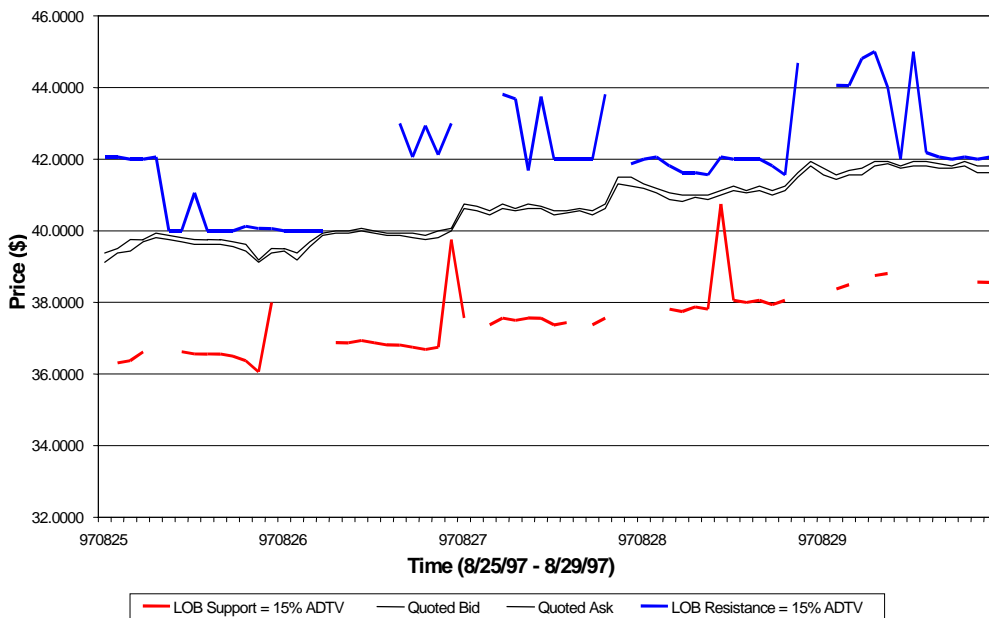
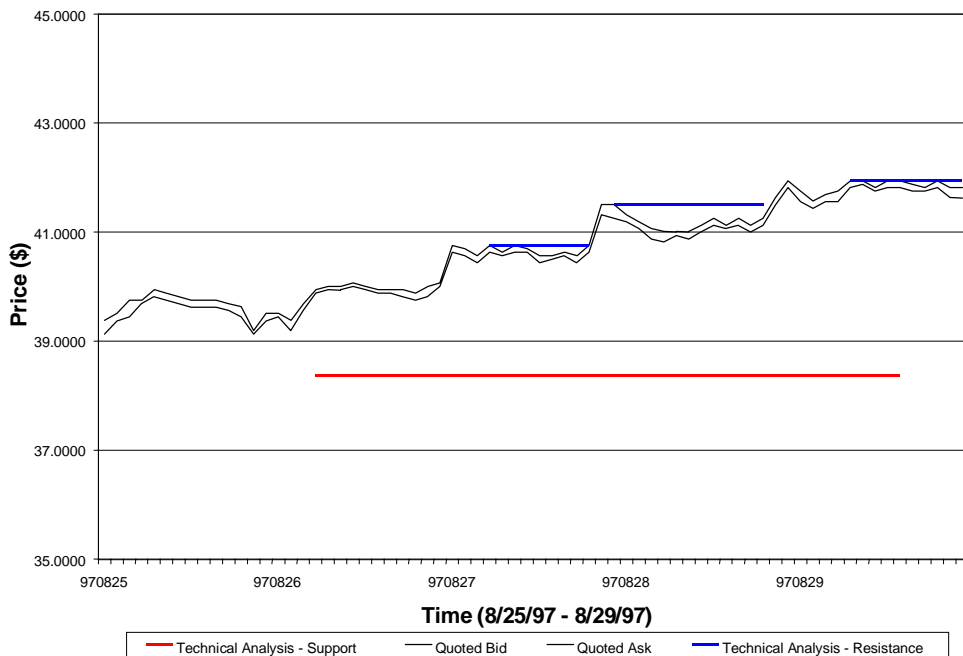


Figure 1: Position of Limit Order Book Depth Measures for Dresser Industries. Buyside and Sellside Modes (Panel A) represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. 15% of ADV (Panel B) is defined as the last share price from the limit order book that an investor would pay (receive) when buying (selling) a quantity equal to 15% of the average daily trading volume.

Panel A: Support and Resistance Levels



Panel B: Long-Run and Short-Run Moving Averages

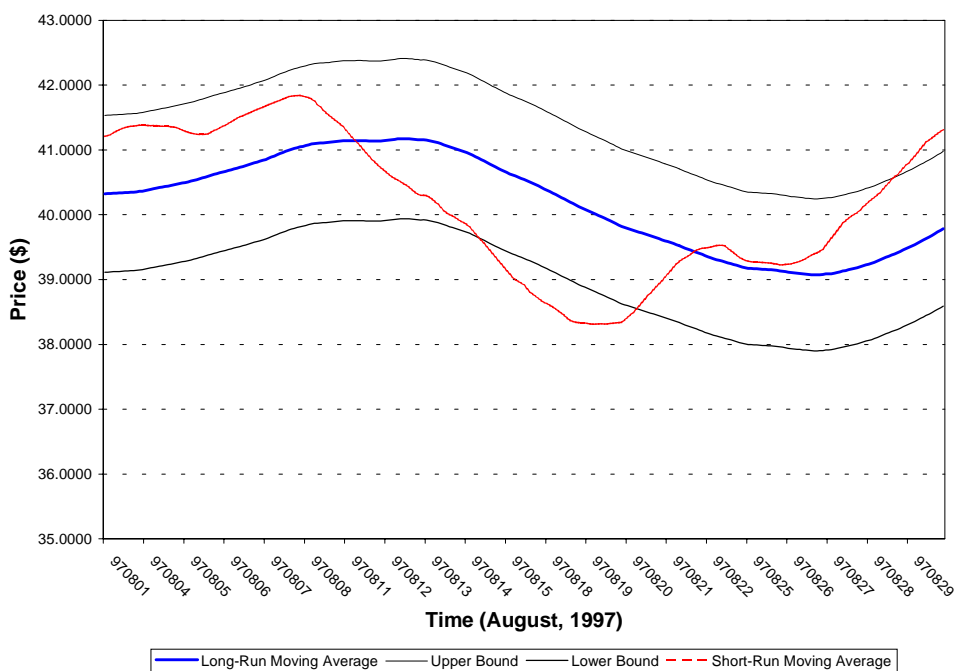
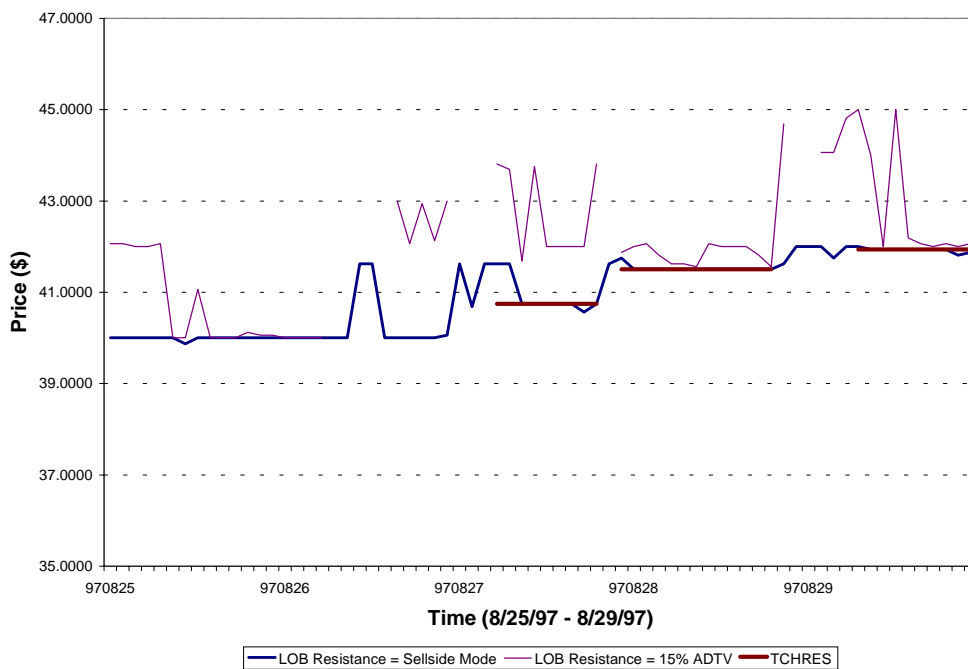


Figure 2: Position of Technical Analysis Levels for Dresser Industries. Technical support (resistance) levels (Panel A) are based on the lowest (highest) bid (ask) price reached at least twice over the previous week, see equation 1 for details. The long-run moving average is based on data over the past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

Panel A: Resistance Levels



Panel B: Support Levels


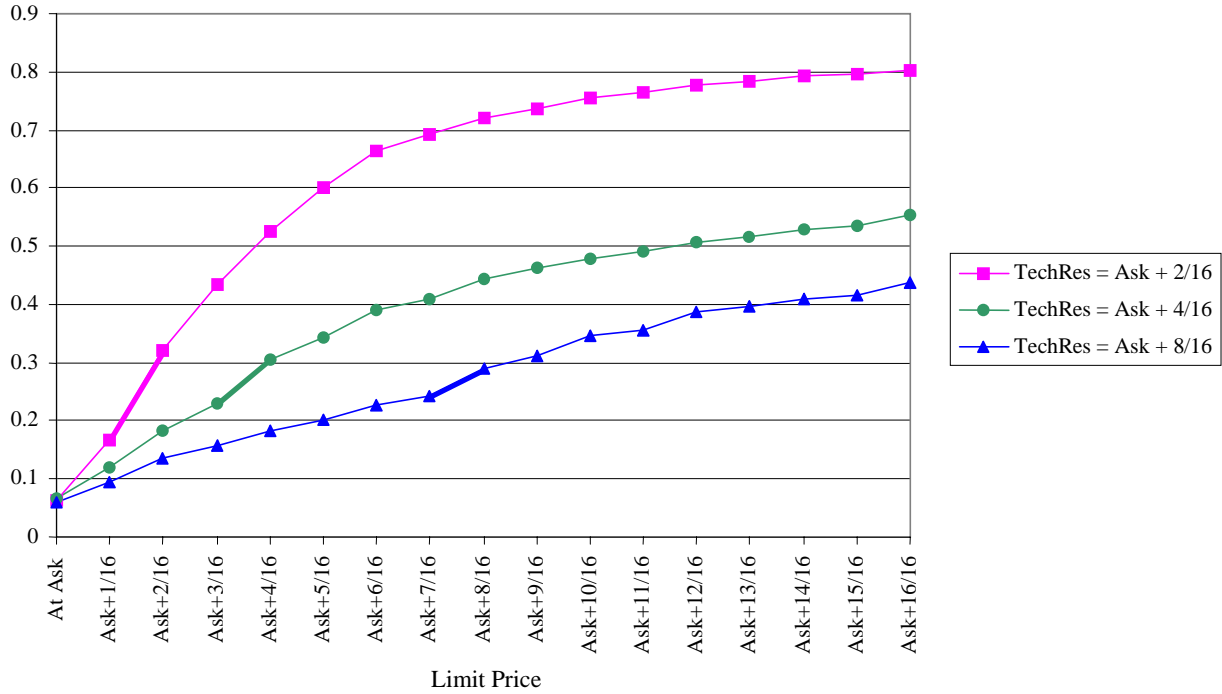


Figure 3: Limit Order Book Liquidity Measures and Technical Analysis Levels for Dresser Industries. Buyside and Sellside Modes represents the higher (lower) of the two prices on the buy (sell) side of the limit order book with the most shares. 15% of ADTV is defined as the last share price from the limit order book that an investor would pay (receive) when buying (selling) a quantity equal to 15% of the average daily trading volume. 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.

Panel A: Sell Side of the Limit Order Book



Panel B: Buy Side of the Limit Order Book




Figure 4: Distribution of Cumulative Depth on the Limit Order Book The empirical distribution of cumulative depth on the limit order book is computed by group, where observations are grouped according to the distance from the technical analysis level to the quoted price. The bold segments correspond to the position of the technical analysis level.