ELSEVIER

Journal of Financial Markets 5 (2002) 169–195

Journal of FINANCIAL MARKETS

www.elsevier.com/locate/econbase

Depth improvement and adjusted price improvement on the New York stock exchange ☆

Jeffrey M. Bacidore[a], Robert H. Battalio[b], Robert H. Jennings[c,*]

[a] *Goldman Sachs & Co., Derivatives and Trading Research, 1 New York Plaza, 42nd Floor, New York, NY 10004, USA*

[b] *Department of Finance, Mendoza College of Business, University of Notre Dame, Notre Dame, IN 46556, USA*

[c] *Department of Finance, Kelley School of Business, Indiana University, 1309 E. 10th Street, Bloomington, IN 47405, USA*

Abstract

Traditional price improvement improperly assesses large orders' execution quality by ignoring additional liquidity depth-exceeding orders receive at the quoted price and viewing orders that "walk the book" as "disimproved". Ignoring this additional liquidity is particularly problematic when assessing execution quality in markets with significant non-displayed liquidity. To correct this deficiency, we modify the price benchmark used to determine whether an order is price improved by making the benchmark a function of the order's size relative to the quoted depth. We document that the differences between conventional price improvement and our

☆ The authors thank Katharine Ross and Susan Farkas for their help in data collection. The authors also thank an anonymous referee, Michael Goldstein, Craig Holden, Paul Irvine, Mark Lipson, Gideon Saar, George Sofianos, Avanidhar Subrahmanyam, workshop participants at Baruch College, Indiana University, the New York Stock Exchange, and the Securities and Exchange Commission, and participants at the 2000 Western Finance Association meeting for comments on earlier versions of this paper. Any remaining errors are the responsibility of the authors. The comments and opinions expressed in this paper are the authors' and do not necessarily reflect those of the directors or officers of Goldman Sachs & Co.

*Corresponding author. Tel.: 812-855-2696; fax: 812-855-5875.

E-mail address: jennings@indiana.edu (R.H. Jennings).

PII: S1386-4181(01)00026-X

measure, adjusted price improvement, can be dramatic and show that the difference depends on trading volume, stock price, and volatility.

JEL classificaion: G0

Keywords: Execution quality; Price improvement; Liquidity

1. Introduction

Many academic studies document the frequency and amount of "price improvement", which is defined to occur when orders execute at prices better than the quoted prices. These studies focus on quantifying execution quality (e.g., determining which market provides the best executions among the venues trading a security) and on characterizing the behavior of liquidity suppliers such as floor brokers and specialists who augment the quoted liquidity schedule.[1] Prior research, however, emphasizes the relationship between execution prices and the quoted bid/offer prices.[2] For orders with less-than-quoted size (small orders), this focus is reasonable because trading venues must, according to Securities and Exchange Commission Rule 11Ac1-1, provide execution prices no worse than the quoted price. For orders with more size than the quoted size (large orders), however, the quoted price is not the worst price an investor can expect. Indeed, unless additional depth is supplied, a large order will exhaust the quote and "walk the book". Conventional measures of execution quality view large orders executing at prices worse than the quote as "disimproved", even if these orders simply execute against the quote and standing orders in the limit order book. Furthermore, suppose a large order executes entirely at the quoted price. Such an order fared better than if it simply walked the book, yet conventional execution-quality measures view this order as neither improved nor disimproved. Clearly, this order benefits from "depth improvement", i.e., additional depth at the quoted price,

[1] Examples of execution quality studies include Blume and Goldstein (1992), Lee (1993), Petersen and Fialkowski (1994), Easley et al. (1996), Huang and Stoll (1996), Ross et al. (1996), Bessembinder and Kaufman (1997), SEC (1997), and Battalio et al. (1998). Studies focusing on characterizing floor participant behavior include Angel (1997) and Ready (1999).

[2] Of course, there are additional dimensions of execution quality beyond price and depth improvement. Macey and O'Hara (1997) discuss and Bacidore et al. (1999) and Battalio et al. (1999) empirically examine several of these additional dimensions of execution quality.

yet this benefit is ignored by conventional measures.[3] By not considering the size of the order relative to the quoted depth, conventional metrics misclassify large orders.[4]

We show that ignoring the quoted depth when measuring price improvement can be misleading. We find that 16 percent of NYSE system market orders have an order size greater than the quoted depth, which implies that conventional price improvement may misclassify as many as one in six NYSE system market orders. Furthermore, 70 percent of these orders are depth improved. Investors, therefore, have many opportunities to receive depth improvement and do so frequently. This result is similar to that of Handa et al. (1999) who document price and depth improvement on the American Stock Exchange. Our results suggest that conventional execution-quality metrics underestimate the number of superior executions and overestimate the number of inferior and "neutral-quality" trades.

Although depth improvement and price improvement together provide a more complete picture of how execution prices can differ from quoted prices (i.e., how the effective liquidity supply schedule differs from the posted schedule), simultaneous use of both metrics can lead to ambiguous results. For example, suppose market A provides more price improvement than market B, but market B provides more depth improvement. Which market provides better executions? Similarly, suppose market makers in stock A provide more price improvement, but those in stock B provide more depth improvement. Which market makers provide the most additional liquidity? This problem can be rectified if one notes that depth improvement is a form of price improvement. To the extent that depth-improved orders do not walk the book, the shares in excess of the quoted depth receive price improvement relative to the prices market makers are required to provide. In this sense, price improvement can be used to measure execution quality for all orders, so long as the benchmark price is a function of the relative order size.

[3] Others use the term "liquidity enhancement" instead of depth improvement. We use depth improvement because it is similar to the term price improvement, which is widely used in the academic literature. Handa et al. (1999) documents the frequency with which investors can trade more than the quoted number of shares at the quoted price, a phenomenon they refer to as "quantity improvement".

[4] Some studies present market-quality statistics by order size (e.g., Lee, 1993; Petersen and Fialkowski, 1994). We show, however, that the important conditioning variable is order size relative to the quoted depth, not unconditional order size. Although some studies control for relative order size (e.g., Lee, 1993; Ready, 1999), they do not fully correct the bias because they ignore the fact that relatively large orders may receive superior executions via depth improvement. Irvine et al. (2000) develop a measure that captures the liquidity in the limit order book beyond the best quoted price, thereby taking into account the relative order size when measuring the displayed or "committed" liquidity in the market.

We address this issue by developing a new concept, adjusted price improvement (API), which combines conventional price improvement and depth improvement. API compares the volume-weighted average trade price to a volume-weighted benchmark price to assess whether an order receives a favorable execution and how favorable that execution is. We compute both conventional and adjusted price improvement and show that ignoring depth improvement when measuring execution quality can be significant, especially with share-weighted statistics. We also find that the difference between conventional and adjusted price improvement depends on trading volume, price, volatility, and the percent of time the stock trades in minimum variation markets.[5]

Our work has implications for quantifying execution quality as well as for gauging the liquidity floor traders (i.e., floor brokers and specialists) provide. Ignoring quoted depth when measuring execution quality can lead to incorrect inferences regarding which market center is most likely to offer superior executions if one market receives more depth-exceeding orders than the others. Simply computing conventional price improvement and depth improvement separately can produce ambiguous results. Adjusted price improvement, however, consolidates conventional price and depth improvement into a single measure. Furthermore, floor traders can improve trade quality by bettering quoted prices (which may be necessary for them to participate) or by adding depth for large orders.[6] Although conventional price improvement captures the former, it ignores the latter. By ignoring liquidity added at the quoted price, researchers fail to capture those cases where floor traders supplement quoted depth. The proper measurement of price improvement, regardless of the purpose, should incorporate depth improvement because conventional price improvement and depth improvement are conceptually equivalent.

Finally, decimalization (decreasing securities' minimum price variation to a penny) will increase the importance of incorporating depth into execution

[5] Price improvement is but one common execution-quality statistic traditionally using the quoted price as a benchmark. Another is the effective spread, which often is reported conditional on the prevailing quoted spread (see, e.g. Blume and Goldstein, 1992; Battalio et al., 1998; Lightfoot et al., 1998). Effective spread and price improvement are related in that price improvement is a discrete comparison of effective spreads to the quoted spread. Our contribution regarding effective spreads is to advocate that the benchmark spread become a function of the size of the order relative to the quoted size (see Section 5).

[6] Angel (1997) and Ready (1999), for example, focus on the liquidity supplied by floor traders beyond that reflected in the quote. Angel investigates who receives and supplies (conventional) price improvement, while Ready investigates how specialists use price improvement in the context of guaranteed orders as a means of sampling future order flow before committing to provide price improvement. In both cases, the authors consider only those cases where the execution price receives a price favorable to the quoted price, ignoring depth improvement entirely.

quality measures.[7] Several studies find narrower spreads and less quoted depth after the minimum price variation is reduced (e.g., Ahn et al., 1996, 1998; Bacidore, 1997; Goldstein and Kavajecz, 2000a,b, Porter and Weaver, 1997; Ronen and Weaver, 1998; Chakravarty et al., 2001).[8] Unless investors alter their order submission strategies, less quoted depth suggests that a greater fraction of orders will exceed the quoted depth in decimal trading. As a result, more orders will be incorrectly classified with conventional price improvement.

The remainder of the paper is organized as follows. We discuss our data in the next section. Section 3 provides an empirical analysis of the frequency of depth improvement, including a cross-sectional analysis. In Section 4, we document the interaction between depth improvement and conventional price improvement. We introduce adjusted price improvement, document significant differences between API and conventional price improvement, and examine how that difference varies as a function of stock characteristics in Section 5. Section 6 concludes.

2. Data

We use data from the NYSE's system order database (SOD) for August 1999. We focus on system orders because we lack comparable order-level data for orders handled by floor brokers and cannot reliably infer this information from trade data.[9] Nevertheless, examining system orders is useful because it allows us to ignore issues associated with differential floor broker commissions, the relationships between floor traders, and the information sharing among floor traders that may affect the probability of price/depth improvement.[10] It is worth noting that we study a sample of orders less likely to be eligible for depth improvement than the average NYSE order because system orders tend to be for fewer shares than floor orders.

Ignoring closed-end mutual funds, there are 2,416 common stock issues trading on the NYSE throughout August 1999. We eliminate issues with an

[7] See Wall Street Journal, "NYSE Seeks to Start Decimal Stock Trade Beginning on August 28", July 26, 2000.

[8] This finding is consistent with Lee et al. (1993) who note that the quoted spread-depth combination reflects only one point on the underlying liquidity supply function. If the smaller tick size results in a tighter spread, then quoted depth also may decline as a new point on the function is exposed. This is true even if the introduction has no impact on the shape of the underling liquidity supply function.

[9] The NYSE recently proposed a change to Rule 123 that would require that all orders be recorded electronically prior to representation or execution.

[10] See Beneviste et al. (1992) for a discussion of the role of floor broker/specialist relationships in mitigating the adverse selection problem. See also Sofianos and Werner (1997) for an analysis of floor broker activity.

average share price less than \$3.00 per share or greater than \$500.00 per share. We also exclude stocks that split during our sample period because the mix of orders and the stock's liquidity tend to change significantly following splits (e.g., Lipson, 1999; Easley et al., 2000; Schultz, 2000). Finally, we focus on issues with a round lot of 100 shares. This provides a sample of 2,128 stocks. We delete limit orders, odd lots, tick-sensitive orders, and orders for which a valid quote is unavailable. Limit orders actually may supply liquidity, so price improvement may not be an appropriate execution-quality metric.[11] Odd lot orders are excluded because they automatically execute against the dealer's inventory at the prevailing quoted price with no opportunity for improvement (see Bacidore et al., 1999 for more detail). A valid quote is required to determine if an order receives depth/price improvement. As a result, orders participating in the open and market-on-close orders are ignored, as are orders arriving when the bid price is equal to or greater than the offer price.[12]

To determine whether an order receives price/depth improvement, we must determine the National Best Bid and Offer (NBBO) quote. This is done with SOD's quote-companion file, which contains the NYSE's best bid and offer prices and sizes and the best off-NYSE bid and offer prices and sizes existing when the order is displayed to the specialist (DBTIME in SOD). The NYSE uses the Consolidated Quotation System to align the order with the quote prevailing at that time. The NBB(O) price is computed as the highest bid (lowest offer) price across all markets quoting the stock. The size associated with the bid (offer) price is defined as the size posted by the market with the best bid (offer). If more than one market is posting the best price, then we use the size of the market posting the greatest depth.[13]

The conditioning variables in the cross-sectional analysis come from two additional data sources, the NYSE Master file (MAST) and the Consolidated Trades Summary file (CTS). We use MAST to determine the number of shares outstanding as of July 31, 1999 for each sample stock. We multiply the shares outstanding by the closing price on July 31, 1999 from CTS to estimate each firm's market capitalization. CTS also provides the average daily NYSE trading volume for each stock during our sample period and the average percent difference between the intra-day high and low trade prices (our proxy for security price volatility).

[11] We exclude marketable limits, i.e., those that should execute immediately because the limit buy (sell) price is greater (less) than or equal to the quoted offer (bid) to avoid issues surrounding partial fills.

[12] Approximately, 52% of system orders are excluded because they are limit orders and 17% are excluded because they are not regular market orders. As a result, approximately 32% of system orders are included in our sample. See Bacidore et al. (1999) for additional background statistics associated with NYSE SOD data.

[13] See Bacidore et al. (1999) for a more detailed discussion of the issues surrounding calculation of the reference quote.

Table 1
Sample summary statistics

We compute the cross-sectional mean and median for each statistic below. Price is the mean closing trade price during the sample period. Market capitalization is the number of shares outstanding at the start of the sample period multiplied by the New York Stock Exchange closing price as of July 31, 1999. The daily dollar volume is the number of shares traded each day multiplied by the NYSE closing price for the day, and the number of trades equals the number of NYSE trades reported to the Consolidated Tape. Data on trading activity are taken from the Consolidated Tape Summary (CTS) database. The number of shares outstanding is taken from the NYSE Master file. System order data come from the NYSE system order data files for August 1999. All non-tick sensitive, non-opening, non-market-on-close orders are included. Daily trading volume in shares (i.e., system plus non-system volume) is taken from the CTS database.

	Mean	Median
Price	$28.83	$23.61
Market capitalization (in $ millions)	5,405	822
Daily dollar volume	$13,606,243	$1,779,115
Daily trading volume (in shares)	312,646	78,409
Number of trades per day	144	53
Daily system volume (eligible orders only)	64,354	14,637
Daily number of eligible system orders	88	21
Eligible system volume as % of twice trading volume	10.7%	10.1%
Average order size of eligible system orders (in shares)	715	678
Number of stocks in sample		2,128

For each stock, we compute the order-weighted and share-weighted conventional price improvement, depth improvement, and adjusted price improvement rates and effective spreads.[14] We then compute the cross-sectional means and medians of these statistics. Consequently, although we begin with over 4.5 million orders, our statistics are cross-sectional averages of 2,128 stocks. Table 1 contains cross-sectional summary statistics for our sample.

The mean (median) daily number of system market orders is 88 (21), or 64,354 (14,637) shares. This amounts to 10.7 percent (10.1 percent) of two times total daily trading volume.[15] The mean (median) system market order is 715 (678) shares.

[14] The effective spread is defined as twice the distance between the transaction price and the midpoint of the contemporaneous benchmark spread (see Huang and Stoll, 1996).

[15] Because the system order volume is the sum of system buy and sell volume, we divide by the sum of *total* buy and sell volume, which is essentially twice the reported trading volume.

3. Empirical analysis of depth improvement

3.1. Methodology

A sell (buy) order is eligible for depth improvement if the order's size exceeds the contemporaneous NBB(O) quoted depth. Similarly, we define the number of shares eligible for depth improvement as the number of shares by which the order size exceeds the relevant quoted depth. We define a stock's depth improvement rate as the sum of all orders (shares) receiving depth improvement divided by the number of orders (shares) eligible for depth improvement. An order receives depth improvement if the number of shares executing at or within the quote exceeds the number of shares quoted. The number of shares receiving depth improvement is defined as the number of shares executing at or within the quote less the number of shares quoted.

For example, suppose three 500-share buy orders arrive at different times throughout the day. When the first order arrives, the quoted offer price is \$20 with a corresponding depth of 300 shares. The first order has 200 shares eligible for depth improvement. Suppose this order fills entirely at \$20 so all 200 eligible shares receive depth improvement. Later, the second order arrives when the quote is \$20 1/16 with a size of 200 shares. Here, the number of eligible shares is 300. Suppose that 200 shares execute at \$20 1/16 and the remaining 300 execute at \$20 1/8. In this case, no shares receive depth improvement because the order simply exhausts the quote and walks up the book. The third order arrives when the quote is \$20 1/8 for 100 shares. Now, 400 shares are depth-improvement eligible. Suppose 300 shares trade at \$20 1/8 and the remaining 200 execute at \$20 3/16. Because only 100 shares are offered at \$20 1/8, 200 shares are depth-improved. The order-weighted depth improvement rate is 66.67 percent because two of the three depth-improvement-eligible orders receive depth improvement. In terms of shares, 1,500 shares are submitted, with 800 of those shares eligible for depth improvement. Of these 800 shares, 400 receive depth improvement, giving a share-weighted depth improvement rate of 50 percent.

3.2. Results

Table 2 reports the fraction of orders eligible for depth improvement and the sample's improvement rate.

About 16 percent of system orders are eligible for depth improvement. Of these orders, 70.4 percent are depth-improved. In terms of volume, 22.7 percent of system volume is eligible for depth improvement and 45.9 percent of these shares are depth-improved. Together these statistics imply that approximately 11 percent of orders and 10 percent of shares are

Table 2
Depth improvement summary statistics

System order data are taken from the New York Stock Exchange system order data (SOD) files from August 1999. Tick sensitive, opening, and market-on-close orders are excluded. Depth improvement-eligible orders are those with order sizes exceeding the size of the National Best Bid or Offer computed using the quote companion file to the SOD database. Depth-improvement-eligible volume is the number of shares exceeding the quoted size. Depth improvement is defined as the number of depth-improvement eligible shares that receive a price equal to or less (greater) than the prevailing offer (bid) for a market sell (buy) order. The cross-sectional means of the statistics are presented below.

	Mean (%)	Median (%)
% Of orders eligible for depth improvement		
Order-weighted	16.0	14.6
Share-weighted	22.7	21.2
% Of depth improvement-eligible orders which receive depth improvement		
Order-weighted	70.4	70.4
Share-weighted	45.9	45.2

depth-improved.[16] The disparity between our order-weighted and volume-weighted depth improvement rates suggests that, although only large orders are eligible for depth improvement, the smaller of these orders are more likely to be improved. This is consistent with the theoretical predictions of Easley and O'Hara (1987), Stoll (1978), and Ho and Stoll (1981). Easley and O'Hara posit that large orders tend to be most informative and, as such, the least likely to receive better-than-quoted prices. The latter two papers argue that risk-averse liquidity providers are more reluctant to execute large orders against their own inventory. Our findings also support with the empirical findings of Angel (1997), Handa et al. (1999), and Ready (1999) who find that the likelihood of receiving price improvement decreases in order size.[17]

To improve our understanding of the cross-sectional variation in depth improvement, we estimate a regression using variables extant research finds

[16] This compares to 5 percent of AMEX system orders and 11 percent of shares receive quantity improvement, as reported in Handa et al. (1999).

[17] We could conduct a cross-sectional analysis of the likelihood of a given order receiving depth improvement similar to Handa et al. (1999) and Ready (1999) regarding price improvement where we condition on market conditions. However, because our focus is on how ignoring depth improvement affects estimates of execution quality cross-sectionally, we specify our regressions in terms of depth improvement rates, and, in this sense, we integrate out market conditions. Also, we do not isolate stopped orders since, unlike during Ready's sample period where nearly 30 percent of system orders were stopped, stopped orders comprise only about 3 percent of system orders in our sample.

correlated with trading costs: trading volume, market capitalization, price, and volatility (see, e.g., Harris, 1994; Bessembinder and Kaufman, 1997). Specifically, we estimate the following regression equation:

$$DI_i = \beta_0 + \beta_1 \,\mathrm{Ln}(Volume_i) + \beta_2 \,\mathrm{Ln}(MktCap_i) + \beta_3 \, Volatility_i + \beta_4 \, InvPrice_i, \quad (1)$$

where *DI* is the order-weighted depth improvement rate, *Volume* is the average daily trading volume, *MktCap* is the company's market capitalization at the July 31, 1999 close, *Volatility* is the average percentage difference between the intra-day high and low prices, *InvPrice* is the average reciprocal of the quote midpoint price, and $\mathrm{Ln}(\cdot)$ denotes the natural logarithm.

Trading volume often is viewed as a proxy for liquidity, while market capitalization can proxy for both liquidity and relative information asymmetry (see Bacidore, 1997; Bacidore and Sofianos, 2000; Madhavan and Sofianos, 1998 for example). Therefore, we predict that these variables' coefficient estimates are positive. Harris (1997) argues that relative tick size affects the willingness of professional traders (e.g., the NYSE specialist) to step ahead of the limit order book. Because the *dollar* tick size is fixed at $0.0625 for all sample stocks, the relative tick is determined by the stock price. Consequently, if Harris's argument is valid, we expect to find a positive relationship between share price and depth improvement (a negative relationship between inverse price and depth improvement). Finally, based on the arguments in Stoll (1978) and Ho and Stoll (1981) we expect that risk-averse liquidity providers may be more reluctant to step in front of the limit order book to supply depth improvement if the value of the resulting inventory position is less predictable. This suggests an inverse relationship between stock-price volatility and depth improvement.

Table 3 contains the estimated coefficients from this regression.

We find no significant relationship (at traditional significance levels) between the depth improvement rate and either market capitalization or trading volume.[18] As expected, we find a negative relationship between volatility and the depth improvement rate, consistent with the notion that liquidity suppliers are less likely to provide depth improvement when the stock price is difficult to predict. There is a positive relationship between the inverse of stock price and depth improvement rates, a finding opposite of that predicted by Harris (1997).

One possible explanation for the latter finding is that the percent of depth-improvement-eligible orders is determined endogenously. If liquidity providers are reluctant to place limit orders because other traders exploit the option value implicit in such orders, then the limit order book is thinner in stocks with large

[18] All tests involving regression coefficients are also done using a White (1980) heteroskedasticity-consistent covariance matrix. The results are qualitatively similar.

Table 3
Depth improvement rates as a function of trading activity, market capitalization, volatility, and price

Cross-sectional regressions are estimated with the order-weighted depth improvement rate as the dependent variable. Independent variables are the natural logarithm of average daily New York Stock Exchange (NYSE) volume, the natural logarithm of market capitalization, the average inverse quoted mid-point price, and the average percentage difference between the intraday high and low price (a proxy for intraday volatility). The depth improvement rates are calculated using NYSE system order data (SOD) from August 1999. Tick-sensitive, opening, and market-on-close orders are excluded. Depth improvement-eligible orders are those with order size's exceeding the size of the National Best Bid or Offer computed using the quote companion file to the SOD database. Data on the conditioning variables come from the NYSE Master file and the Consolidated Trade Summary file. Depth-improvement-eligible volume is the number of shares that exceed the quoted size. An order receives depth improvement when the order receives more shares than the amount quoted at a price equal to (or better than) quoted price. *P*-values are reported in parentheses.

	Independent variables					
Dependent variable	Intercept	Ln (average daily trading volume)	Ln (market cap.)	Intraday volatility	Inverse price	R^2
% Of depth-improvement eligible orders that receive depth improvement	78.90 (0.000)	0.426 (0.155)	−0.426 (0.197)	−3.869 (0.000)	64.204 (0.000)	0.085
Probability that slope coefficients jointly equal zero (*F*-test)	$P=0.000$					

relative ticks (low prices). Knowing that the underlying liquidity supply function is steeper in such stocks, market order traders may decrease their order size, ceteris paribus, leading to a reduction in the number of depth-improvement-eligible orders. To account for this, we estimate the same model using the depth improvement rate as a percent of *all* orders as the dependent variable (results not reported). Here, the relationship between relative tick size and depth improvement is significantly *negative*, a result consistent with Harris (1997).

To summarize, our results show that a significant number of orders are eligible for and receive depth improvement. Ignoring depth improvement when assessing the execution quality of orders, therefore, might have a significant impact on execution quality statistics. Furthermore, we show that the depth improvement rates vary cross-sectionally as a function of volatility and price. We next examine the interaction between price improvement and depth improvement.

4. Interaction between depth improvement and price improvement

To more fully characterize execution quality, we consider price improvement and depth improvement jointly. For orders for fewer shares than the quoted depth, liquidity providers have the option to provide price improvement or execute the order at the quote. For orders with sizes exceeding that quoted, liquidity providers must determine whether to allow the order to exhaust the quote and walk the book or to provide depth/price improvement by executing additional shares at or within the quote. Although we do not provide a formal model of this choice, we do report the frequencies of each occurrence in Table 4.

As reported earlier, the cross-sectional average depth improvement rate is 70.4 percent (14.4 percent + 55.9 percent), i.e., 70.4 percent of orders eligible for depth improvement are improved. Twenty percent (= 14.4/70.4) of these orders also receive price improvement. Of the 30 percent of eligible orders not receiving depth improvement, almost none receive price improvement.[19]

The results in Section 3 show that, on average, one-sixth of NYSE system orders during August 1999 have order sizes exceeding the order-receipt-time size associated with the relevant quoted price. About 70 percent of these

[19] It may seem odd that eligible orders that are not depth improved almost never receive price improvement. However, this finding stems from the fact that most eligible orders receiving price improvement also receive depth improvement. In other words, conditional on an eligible order receiving price improvement, it is almost certain that more than the quoted number of shares execute at the quoted price or better, i.e., that the order receives depth improvement.

Table 4
Interaction between conventional price improvement and depth improvement

We compute cross-sectional means (medians) of the statistics below using New York Stock Exchange system order data (SOD) from August 1999. Tick-sensitive, opening, and Market-On-Close orders are excluded. An order receives price improvement if it executes (at least partially) inside the quote. Depth improvement-eligible orders are those with order sizes exceeding the size of the National Best Bid or Offer computed using the quote companion file to the SOD database. The (gross) price improvement rate is calculated as the percent of orders (shares) receiving better-than-quoted prices. An order receives depth improvement if the order receives more shares than the amount quoted at a price equal to (or better than) quoted price.

	% Of DI-eligible orders receiving depth improvement	% Of DI-eligible orders not receiving depth improvement
% Of DI-eligible orders receiving price improvement	14.4 (13.3)	0.2 (0.0)
% Of DI-eligible orders not receiving price improvement	55.9 (55.6)	29.5 (29.4)

"oversized" orders execute at the quoted (or better) price despite the fact that the specialist need not honor the quoted prices for such orders. Table 4 notes variation in how large orders execute. Some receive both price and depth improvement, some neither, and others receive depth improvement without price improvement. The prevalence of depth improvement and the disparity of the treatment of large orders, suggest that it is important to consider the existence of depth improvement and the interplay between depth and price improvement when measuring execution quality. But how does one simultaneously consider both depth improvement and price improvement in a single measure of execution quality? We address this question in the next section.

5. Adjusted price improvement

One way to evaluate execution quality is to examine price improvement and depth improvement simultaneously as in Table 4. Comparisons across both dimensions may lead to ambiguity, however, especially if liquidity providers treat the two forms of improvement as substitutes. For example, suppose market venues A and B are identical in every way except with respect to their price and depth improvement rates/amounts. Further, suppose venue A provides more price improvement than venue B, but venue B provides more depth improvement. How does one determine which market has better execution quality? A similar problem exists at the order level. Suppose two identical, depth-improvement-eligible orders execute at the same venue.

Further, suppose some of the first order receives price improvement and the remainder exhausts the quote and walks the book, while the second order executes entirely at the quote, i.e., receives depth improvement. Which order receives the better execution? Because examining price improvement and depth improvement separately can produce ambiguous comparisons, we develop an execution-quality metric that considers price and depth improvement simultaneously. Adjusted price improvement (API) provides a single measure of execution quality incorporating quoted depth into the widely used concept of price improvement. Our approach, detailed below, uses the quoted depth to calculate a benchmark price to which we compare the trade price to determine the existence of and amount of price improvement.

5.1. Methodology

The order-weighted (share-weighted) gross unadjusted price improvement rate is the percent of market orders (shares) receiving prices better than the relevant contemporaneous NBBO quoted price (i.e., the bid for sell orders and the offer for buy orders). The net unadjusted price improvement rate is the gross rate less the percent of orders (or shares) executing at prices worse than the quoted price. We define the gross adjusted price improvement rate as the percent of orders (or shares) receiving prices better than the appropriate benchmark price and the net rate as the gross rate less the percent of orders (or shares) receiving prices worse than the benchmark price. For orders with sizes less than or equal to the relevant quoted depth, this benchmark price is simply the quoted price because the order is required to execute entirely at a price no worse than the quoted price. For orders with sizes exceeding the quoted depth, the quoted price is relevant only for the quoted size. In these cases, we redefine the benchmark price as a weighted-average of the quoted price and the price one tick outside the quote (above the offer for buy orders and below the bid for sell orders). The weight on the latter price equals the percent by which the order size exceeds the quoted size.

For example, suppose a buy order for 2,000 shares arrives when the quoted offer is \$20 for up to 400 shares. In this case, only 20 percent of the order is entitled to the offer price, so the offer price is given a weight of 0.20. We assume the remaining 1600 shares are entitled to trade up one tick at \$20.0625. Therefore, the benchmark price equals 0.20 (= 400 shares/2000 shares) times the quoted price of \$20, plus 0.80 (= 1,600 shares/2,000 shares) times \$20.0625, or \$20.05. Suppose 1000 shares execute at \$20 and 1000 shares execute at \$20.0625. The volume-weighted average trade price is \$20.03125. Because the average purchase price (\$20.03125) is less than the benchmark price (\$20.05), the order receives adjusted price improvement. Note that conventional measures of execution quality do not consider this depth-improved order as improved. In fact, because the order executes partially outside the quote, some consider this order as price disimproved, even though 600 shares received depth improvement.[20]

We assume that shares exceeding the quoted size receive a price one tick away from the relevant contemporaneous quoted price when the liquidity supply function may be so steep or the order size so great that the order actually would execute at several prices outside the quote without floor intervention.[21] We make this assumption because the nearly continuous data on the state of the limit order book one would need to more accurately examine actively traded securities are not readily available. Furthermore, because our method assumes infinite depth one tick outside the quoted price, using the limit order book would only make the benchmark price easier to beat (i.e., the benchmark price would be higher for sell orders and lower for buy orders). This, in turn, would increase the adjusted price improvement rate and magnify the differences between adjusted and unadjusted rates that we find using our approach. Therefore, if we find differences between adjusted and unadjusted price improvement rates with our admittedly extreme assumption, then we would find even larger differences if we were to reconstruct the limit order book in computing our benchmark price.[22] We analyze other, less extreme, assumptions regarding the shape of the limit order book to assess the robustness of our findings.

5.2. Results

Table 5 contains the cross-sectional average gross and net unadjusted price improvement (UPI) and adjusted price improvement (API) rates.

[20] In the example above, we show how adjusted price improvement incorporates depth improvement. One may be tempted to conclude that an order receives adjusted price improvement if it receives either depth improvement or conventional price improvement. For most cases, this is true. Suppose, however, that part of a buy order above executes within the quoted spread, part at the quote, and the remainder outside the quote. By conventional measures, the order is price improved as well as disimproved. Using adjusted price improvement, however, we can unambiguously classify all orders.

[21] The limit order book could be estimated using system limit orders, using the methodology of Kavajecz (1999). However, this methodology is extremely data intensive and would severely limit the number of stocks we could analyze. If we were to use the limit order book, our benchmark price would be similar to the CRT measure presented in Irvine et al. (2000) and to the methodology used in Lipson (1999), Corwin and Lipson (2000), and Goldstein and Kavajecz (2000a,b). The implementation of API presented here is similar in principle to Handa et al. (1999), who value depth improvement by assuming shares exceeding the quoted size execute at a price one tick worse than quoted. However, unlike Handa et al., our assumption is made not as a means to estimate the value of depth improvement, but rather as an illustration of the importance of API generally.

[22] The current formulation of API also could serve as a rough benchmark for traders without access to the limit order book. For such traders, our benchmark price represents the best "adjusted" quoted price they can expect to receive given the prevailing quotes.

Table 5
Conventional and adjusted price improvement rates

We compute cross-sectional averages of the statistics below using NYSE system order data from August 1999. Tick-sensitive, opening, and market-on-close orders are excluded. The gross price improvement rate is calculated as the percent of orders (shares) receiving price improvement. The net price improvement rate is equal to the gross price improvement rate less the gross disimprovement rate. A minimum variation market is defined as one where the difference between the offer and bid is equals \$1/16. The gross unadjusted price improvement rate is the percent of orders (shares) receiving price improvement, i.e., buy (sell) orders receiving an execution price below (above) the offer (bid) price. The net unadjusted price improvement rate is the gross rate less the percent of orders (shares) executing outside the quoted price. The adjusted price improvement rate is the percent of orders (shares) executing at a price better than a weighted average of the quoted price and the price \$1/16 worse, where the first weight is the percent of the order that is eligible to execute at the quoted price.

	Mean (%)	Median (%)
Panel A: All orders		
Gross price improvement rates		
Unadjusted price improvement		
Order-weighted	33.4	33.3
Share-weighted	23.2	22.7
Adjusted price improvement		
Order-weighted	41.5	41.6
Share-weighted	39.4	39.2
Net price improvement rates		
Unadjusted price improvement		
Order-weighted	25.8	25.1
Share-weighted	3.3	3.9
Adjusted price improvement		
Order-weighted	35.5	34.6
Share-weighted	24.5	24.8
Panel B: Price improvement rates in minimum variation markets		
Gross price improvement rates		
Unadjusted price improvement		
Order-weighted	6.3	5.6
Share-weighted	4.9	4.1
Adjusted Price Improvement		
Order-weighted	14.4	13.9
Share-weighted	20.0	19.6
Net price improvement rates		
Unadjusted price improvement		
Order-weighted	−3.3	−1.9
Share-weighted	−15.1	−12.6
Adjusted price improvement		
Order-weighted	7.0	6.8
Share-weighted	5.6	7.0

Table 5 (*continued*)

	Mean (%)	Median (%)
Panel C: Price improvement rates in greater than Minimum Variation Markets		
Gross price improvement rates		
Unadjusted price improvement		
Order-weighted	47.2	47.4
Share-weighted	33.3	33.5
Adjusted price improvement		
Order-weighted	55.1	55.4
Share-weighted	49.6	50.2
Net price improvement rates		
Unadjusted price improvement		
Order-weighted	40.4	41.0
Share-weighted	14.1	15.4
Adjusted price improvement		
Order-weighted	49.8	50.0
Share-weighted	35.2	37.0

Panel A analyzes all orders. The mean gross and net order-weighted UPI rates are 33.4 and 25.8 percent, respectively[23] UPI rates ignore the fact that some orders exceed the quoted depth and execute outside the quote simply because they exhaust the quote. Using our suggested adjustment, we find mean gross and net API rates of 41.5 and 35.5, respectively. The larger gross API rate (relative to the gross UPI rate) occurs because API considers depth-improved orders as price improved and UPI does not. Note also that the average unadjusted *disimprovement* rate is about 7.6 (i.e., the difference between the gross and net UPI rates is 7.6). The difference in the gross and net API rates, however, is only six percent. This disparity between UPI and API disimprovement rates is because API also correctly classifies orders erroneously viewed as disimproved with conventional measures. That the disimprovement rate fell by only 1.6 percentage points while the net rate increased by 9.6, suggests that the bulk of the difference between API and UPI comes from depth-improved orders classified as "neutral" executions under UPI being classified as improved using API.

[23] Although our sample period is identical to that of Bacidore et al. (1999), the averages for conventional price improvement reported here differ from their rates because most of the analysis in Bacidore et al. (1999) focuses on orders not exceeding the quoted depth and because our averages are cross-sectional averages, i.e., averages of stock-by-stock price improvement rates. With respect to the latter, we choose to use equally-weighted, cross-sectional averages because we wish to document the importance of considering depth in assessing execution quality on a stock-by-stock basis. The focus of Bacidore et al. (1999), however, is to assess the overall execution quality on the NYSE, and, as such, volume-weighted averages are more appropriate in the context of their study.

In terms of share-weighted statistics, the disparities are more dramatic. The mean gross UPI rate is 23.2 percent, but the net UPI rate is only 3.3 percent. The disparity between order-weighted and share-weighted UPI rates exists because small orders tend to receive price improvement, leading to lower gross price improvement numbers relative to order-weighted statistics, and large orders tend to execute outside the quote, inflating the disimprovement numbers. When we use API, the gross price improvement rate is 39.4 percent and the net rate is 24.5. Both rates are considerably higher than the corresponding UPI rates. With share-weighted statistics, the 21.2 percentage-point increase in the net price improvement rate results from 16.2 percent more shares classified as improved and 5 percent fewer shares classified as disimproved. Our results suggest that the choice of metric has an enormous impact on our assessments of execution quality.

Because price improvement may be affected by the extent to which a stock trades in minimum variation markets, we compute the price improvement statistics separately for orders arriving in minimum variation markets (i.e., the quoted spread equals $0.0625) and those arriving when the spread exceeds the minimum variation.[24] The results are presented in Panel B of Table 5. Here the differences between UPI and API rates are more pronounced than in the overall sample. In minimum variation markets, the mean order-weighted (share-weighted) gross UPI rate is 6.3 percent (4.9 percent) and the net rate is −3.3 percent (−15.1 percent). The gross API rate is 14.4 percent (20.0percent), while the net rate is 7 percent (5.6 percent). This former result is striking because using the conventional measure suggests that, on average, investors are more likely to receive poor-quality executions than high-quality ones. This result, however, is because orders for more than the quoted depth are inappropriately benchmarked to the quoted price.

For those cases where the spread exceeds the minimum variation (Panel C), the mean order-weighted (share-weighted) gross UPI rate is 47.2 percent (33.3 percent) and the net rate is 40.4 percent (14.1 percent). The gross API rate in this case is 55.1 (49.6), while the net API rate is 49.8 percent (35.2 percent). Again, the difference between the two measures is substantial in each case, especially when share-weighted

[24] Specifically, if a market buy (sell) order arrives in a minimum variation market, it cannot receive price improvement by buying (selling) at the bid (offer) unless all other previously placed limit orders at the same price or better are filled. In other words, strict time priority assures that market orders cannot "step ahead" of limit orders posted at the same price (or better). Consequently, it is relatively less likely that a market order receives price improvement in minimum variation markets.

statistics are used. This provides additional evidence that the choice of execution-quality measure may have a significant influence on estimates of execution quality.

Our results depend on our assumptions regarding the liquidity supply function. We conduct two sensitivity analyses: one to examine the sensitivity of our results to the assumption of infinite liquidity one tick away from the inside quote and another to examine the sensitivity regarding the assumed NBBO depth. The sensitivity of our conclusions to assuming infinite depth one tick from the NBBO is examined in two stages. Firstly, we assume that one-half of the excess size is absorbed by liquidity supplied one tick outside of and the other half by liquidity supplied two ticks outside of the NBBO. Secondly, we assume one-third of the excess size finds liquidity at each of the first three ticks outside the NBBO. UPI rates reported on Panel A in Table 5 are unaffected by this alternative approach. Overall gross order-weighted API rates are just under (over) two percentage points higher than the 41.5 percent rate reported in Panel A of Table 5 for the two-tick (three-tick) assumption. The share-weighted gross UPI is 44.4 (45.6) percent using two (three) ticks outside the NBBO compared to 39.4 percent assuming infinite depth at one tick. Net API rates are more sensitive to this assumption. Assuming two (three) ticks are needed to absorb the excess shares increases the order-weighted net API by 2.3 (4.2) percentage points over the 35.5 percent reported in Panel A of Table 5. Finally, for the share-weighted statistics, the two-tick (three-tick) net API is 30.8 (36.5) percent compared to 24.5 percent with the base case assumption.

The second sensitivity analysis focuses on the fact that one could argue that the correct size to associate with the inside quote is not the NBBO size, but rather the cumulative size available from all markets at the best price. Because our data provide the bid/offer price and size of only the best off-NYSE market (if two non-NYSE markets are tied for the best price, only the one with the most quoted size is displayed), such aggregation is not possible. To determine how sensitive our analysis is to this potential limitation, we assume that all non-NYSE markets are posting quotes identical to that of the best non-NYSE market. If the non-NYSE quoted price is the NBBO, we multiply that quoted size by six under the assumption that all five regional exchanges and Nasdaq are posting the same quote. This gives us the largest possible quoted size available at the quoted price. Furthermore, when the NYSE also is at the inside quote, the NYSE depth is added to this potentially inflated off-NYSE quote. While the mean percent of depth-improvement-eligible orders (shares) falls to 11.9 percent (17.7 percent) with this definition of quoted size, the rate at which these orders (shares) receive depth improvement is little changed. As a result, the difference between overall UPI and API rates is reduced, but the mean difference between gross (net) order-weighted and share-weighted API and UPI rates is still approximately 5.8 percent (6.9 percent) and 12 percent (15 percent). All of these differences are significant at the 0.01 significance level.[25]

Another approach to quantifying adjusted price improvement is to compute the difference between the effective spread and the benchmark spread. The effective spread (see, e.g., Huang and Stoll, 1996) is defined as twice the distance between the transaction price and the midpoint of the contemporaneous quoted spread. For example, suppose the order-receipt-time quoted prices are \$20.00 bid and \$20.125 offered. The effective spread for a buy (sell) order is the trade price less \$20.0625 times 2 (-2). The benchmark spread varies depending on whether you consider the order's size relative to the size associated with the relevant quoted price. For orders with sizes not exceeding the relevant quoted size, the benchmark spread (both adjusted and unadjusted) is simply the quoted spread. For large orders, the adjusted benchmark spread is twice the distance between the spread's midpoint and the benchmark price, which is the volume-weighted average of the relevant quoted price and a price \$0.0625 (one tick) worse for the trade initiator. Consider an example. Suppose the quoted prices are \$20.00 bid for 500 shares and \$20.125 offered for 700 shares when a market buy order for 1000 shares arrives. The adjusted benchmark price is $\$20.14375 (= \$20.125 \times (700/1000) + \$20.1875 \times (300/1000))$. If the entire order executes at the quoted price, then the effective spread is $\$0.125 (= (\$20.125 - \$20.0625) \times 2)$. Without considering the size of the order, we compare the effective spread of \$0.125 to the quoted spread of \$0.125 and conclude that the investor received the expected price. Considering the order's size, however, we compare the benchmark spread of $\$0.1625 (= (\$20.14375 - \$20.0625) \times 2)$ to the effective spread and conclude that the execution is a favorable one.

The results from examining effective spreads are presented in Table 6, where effective and quoted spreads are standardized by the quote-midpoint to control for variation in spreads due to differences in stock prices.

Note that we reach different conclusions regarding market quality with share-weighted numbers using the adjusted benchmark than using the unadjusted quoted spread. The share-weighted effective spread is 85.1 basis points. Comparing that to the quoted spread of 79.1 basis points, one concludes that the average execution price is outside of the benchmark spread. Conversely, comparing the effective spread to the adjusted quoted spread of 95.2 basis points, one concludes that, on average, investors receive favorable executions. Order-weighted mean conventional price improvement is 22.3 basis

[25] Bessembinder (2001) finds that non-NYSE quoted sizes typically are less than NYSE quoted sizes, even when the non-NYSE quote is the NBBO. This suggests that the difference between the UPI and API measures may be less when the NYSE is part of the NBBO. Focusing on the cases in which the NYSE is on the relevant side of the NBBO suggests that there are small increases (at most 2.2 percentage points) in the rates reported in Panel A of Table 5. Of more relevance, the *differences* between the UPI and API rates are virtually unchanged.

Table 6
Conventional and adjusted price improvement as percent of price

We compute cross-sectional averages of the statistics below using NYSE system order data (SOD) from August 1999. Tick-sensitive, opening, and market-on-close orders are excluded. The conventional price improvement rate is calculated as difference between the effective and quoted spread relative to the quote mid-point price (in basis points). The adjusted price improvement rate is calculated as difference between the effective and adjusted quoted spread relative to the quote mid-point price (in basis points). The adjusted quoted spread is defined as twice the difference between the volume-weighted execution price and the volume-weighted quote price.

	Mean	Median
Quoted spread		
Order-weighted	79.6	59.5
Share-weighted	79.3	59.5
Adjusted quoted spread		
Order-weighted	85.7	63.7
Share-weighted	95.2	70.6
Effective spread		
Order-weighted	57.4	42.5
Share-weighted	85.1	60.9
Unadjusted price improvement		
Order-weighted	22.3	15.4
Share-weighted	−5.8	−0.1
Adjusted price improvement		
Order-weighted	28.3	19.8
Share-weighted	10.0	8.9

points, while the mean adjusted price improvement is 28.3 basis points. With respect to share-weighted numbers, the differences are more pronounced, as one may expect given our earlier findings. The mean conventional price improvement rate is -5.8 basis points, while the adjusted price improvement rate is 10.0 basis points.[26] This implies that using adjusted price improvement generates a mean increase in estimated price improvement of 6.0 basis points per order and a mean increase of 15.9 basis points per share. The differences between conventional and adjusted price improvement are significantly different from zero, indicating that considering the relative size of the order is important in assessing execution quality.

[26] Because the basis point results consider the *size* of the improvement, the sign of the net price improvement in basis points need not be the same as that of the net price improvement rate, as is the case for our share-weighted numbers. Our results suggest that while a given share is more likely to receive improvement than disimprovement, the amount of improvement an improved share receives is less than the amount of disimprovement the disimproved shares receive.

To show how the difference between API and UPI varies in the cross-section, we estimate a regression equation similar to that in Section 3. The dependent variable is the difference between the order-weighted net API and UPI rates. To account for restrictions on the ability to receive conventional price improvement in minimum variation markets, we include a variable to account for the fraction of orders arriving in minimum variation markets. Specifically, we estimate the following regression equation:

$$Dif_i = \beta_0 + \beta_1 \,\mathrm{Ln}(Volume_i) + \beta_2 \,\mathrm{Ln}(MktCap_i) + \beta_3 \,Volatility_i + \beta_4 \,InvPrice_i + \beta_5 \,MinVar_i, \qquad (2)$$

where *Dif* is the API rate minus the UPI rate and *MinVar* is the percent of orders arriving in minimum variation markets. In interpreting the coefficient estimates, it is important to note that the API rate generally exceeds the UPI rate because API incorporates depth improvement and UPI incorrectly classifies some large orders as disimproved. Therefore, a significant negative (positive) coefficient implies that, for larger values of the independent variable, the UPI rate become closer to (further from) the API rate.[27] The results are presented in Table 7.

We find that the difference between the API and UPI rates does not vary systematically with respect to market capitalization and inverse price level, ceteris paribus. However, the remaining independent variables' coefficient estimates differ significantly from zero. We find that the coefficient estimates on *MinVar* and trading volume are negative, while the volatility coefficient estimate is positive. We re-estimate (2) using the percentage difference in the API and UPI rates as the dependent variable (instead of the raw difference in API and UPI rates) to incorporate the magnitude of the improvement in addition to the frequency of improvement. We find qualitatively similar results, although the slope coefficients on inverse price and market capitalization are significantly positive. The significance of these coefficients with the alternative definition of improvement suggests that these variables are important determinants of the size of the improvement. In the case of inverse price level, this is intuitive because it proxies for the relative tick, i.e., the amount by which a trader must improve the quote in order to step ahead of standing limit orders. Nevertheless, this suggests that the difference between API and UPI is most pronounced in the least-actively traded, lowest-priced, and most volatile stocks as well as those that rarely trade in minimum variation markets.

[27] One may expect these results to mirror those in Section 3 because the key difference between API and UPI is depth improvement. However, depth improvement is calculated as a percent of eligible orders, while API is calculated as a percent of *all* orders. Therefore, if the percent of eligible orders varies as a function of the conditioning variables, we may find different results than those in Section 3.

Table 7
Cross-sectional analysis of the bias in unadjusted price improvement rates

We estimate cross-sectional regressions using the difference between the adjusted and unadjusted price improvement rates as the dependent variable. In Panel A, the independent variable is the difference in adjusted and unadjusted price improvement rates, while Panel B contains the difference in improvement measured relative to the stock price (in basis points). The independent variables are the natural logarithm of average daily New York Stock Exchange volume, the natural logarithm of market capitalization, the average inverse quoted mid-point price, the average percentage difference between the intraday high and low price (a proxy for intraday volatility), and the percent of orders arriving when the width of the NBBO spread equals the minimum variation. The adjusted and unadjusted price improvement rates (i.e., API and UPI) are calculated using proprietary NYSE system order data from August 1999. Tick sensitive, opening, and market-on-close orders are excluded. Data on the conditioning variables come from the NYSE Master file and the Consolidated Trade Summary file. *P*-values based on *t*-tests are reported in parentheses.

	Independent variables						
Dependent variable	Intercept	Ln (avearage daily trading volume)	Ln (market cap)	Intraday volatility	Inverse price	% in Min. variation markets	Adj. R^2
A. Difference in price improvement rates							
Difference between API and UPI rates (in %)	23.56 (0.000)	−1.169 (0.000)	−0.037 (0.748)	0.869 (0.000)	2.433 (0.379)	−5.389 (0.000)	0.337
Probability that slopes coefficient jointly equal zero (*F*-test)	$P=0.000$						
B. Difference in price improvement measured as a percentage of stock price							
Difference between API and UPI in basis points	2.035 (0.241)	−0.779 (0.000)	0.394 (0.001)	1.137 (0.000)	93.25 (0.000)	−9.39 (0.000)	0.624
Probability that slope coefficients jointly equal zero (*F*-test)	$P=0.000$						

6. Conclusion

Although quoted size represents the specialist's maximum trading obligation, traders frequently find that they can execute orders for more than the quoted number of shares at the quoted (or better) price. About 11 percent of the NYSE system market orders in August 1999 have sizes exceeding the relevant quoted size but trade at prices no worse than the quoted prices. Floor brokers and specialists supply non-displayed liquidity. The existence of non-displayed liquidity suggests many research topics. In this paper, we argue that non-displayed liquidity must be considered to properly assess execution quality. We posit that "depth improvement", the ability to trade more than the quoted number of shares at the quoted price or better, is a form of price improvement and is an important consideration when measuring execution quality. If the coming decimalization of U.S. equity markets leads to a significant reduction in the minimum price variation, then a greater proportion of orders will have sizes exceeding quoted depth. As shown here, conventional execution quality metrics may provide misleading assessments of market quality, suggesting that we must refine our measures to correctly capture depth improvement.

We argue that considering price and depth improvement separately may not allow meaningful execution quality comparisons because rankings based on price improvement may be inconsistent with depth-improvement rankings. To address this problem, we develop the concept of adjusted price improvement, which incorporates both conventional price improvement and depth improvement into a single metric. Our approach compares the volume-weighted average execution price to a benchmark price conditioned on the order's size relative to the quoted depth. Conventional price improvement uses the quoted price as a benchmark price regardless of the order's relative size. We show that adjusted price improvement rates can differ significantly from conventional rates, especially for share-weighted numbers, a finding that highlights the importance of using a measure that incorporates the size of the order relative to the quoted depth. Decimalization may render the joint use of conventional price improvement and depth improvement even less effective in characterizing execution quality because it captures additional liquidity only at the quoted price. If additional depth is provided outside the posted quote, conventional price improvement will not capture it because conventional price improvement ignores depth. Likewise, depth improvement considers only additional depth *at the quoted price or better*. Thus, capturing the additional liquidity supplied at prices beyond the quoted price using API will become more important in the future.

Finally, we document that the differences between traditional market quality estimates and estimates considering quoted size vary systematically in the cross-section as a function of trading volume, market capitalization, and price.

Specifically, we find that conventional metrics tend to overestimate the trading costs of the less-actively traded, low-priced, volatile stocks as well as those that rarely trade in minimum price variation markets.

Our findings document differences across stocks trading on the NYSE. We must take care with certain types of across-market comparisons of execution quality. For auction markets like the NYSE, the trading "crowd" provides the majority of non-displayed liquidity. Markets that are dominated by dealers (e.g., the regional stock exchanges and Nasdaq) cannot rely on the crowd to provide non-displayed liquidity. As such, using conventional price improvement statistics will bias comparisons in favor of dealer markets. On the other hand, some dealers provide additional liquidity in excess of their quoted depth, often in the form of size guarantees that are independent of quoted size. If this willingness to provide additional liquidity offsets the lack of floor-based liquidity, then conventional statistics bias in favor of the dealers' floor-based competitors. Nevertheless, to the extent that API corrects the deficiencies in conventional price improvement, API provides more meaningful cross-market comparisons of execution quality. Furthermore, no measure based on price improvement can provide an absolute across-market measure of execution quality. For example, when comparing stocks trading in two markets without a common benchmark price, it is possible that one market may have more price improvement than the other market but not be viewed as providing better executions because of differences in the quoted spreads.

References

Ahn, H., Cao, C., Choe, H., 1996. Tick, size spread, and volume. Journal of Financial Intermediation 5, 2–22.

Ahn, H., Cao, C., Choe, H., 1998. Decimalization and competition among stock markets: evidence from the toronto stock exchange cross-listed securities. Journal of Financial Markets 1, 51–87.

Angel, J., 1997. Who gets price improvement on the NYSE? Working paper, Georgetown University.

Bacidore, J., 1997. The impact of decimalization on market quality. Journal of Financial Intermediation 6, 92–120.

Bacidore, J., Ross, K., Sofianos, G. 1999. Quantifying best execution at the New York Stock Exchange: market orders. New York Stock Exchange Working Paper #99-05.

Bacidore, J., G. Sofianos, 2000. NYSE Specialist Trading in Non-US stocks. Working paper #00–04, New York Stock Exchange, Inc. Exchange Information memo #99-38.

Battalio, R., Greene, J., Jennings, R. 1998. Order flow distribution, bid-ask spreads, and liquidity costs: Merrill Lynch's decision to cease routinely routing orders to regional stock exchanges. Journal of Financial Intermediation 7, 338–358.

Battalio, R., B. Hatch, R. Jennings, 1999. All else equal? A multidimensional analysis of retail, market order execution quality. Working paper, Indiana University.

Beneviste, L., Marcus, A., Wilhelm, J., 1992. What's special about the specialist? Journal of Financial Economics 32, 61–86.

Bessembinder, H., Kaufman, H. 1997. A cross-exchange comparison of execution costs and information flow for NYSE-listed stocks. Journal of Financial Economics 46, 293–319.

Bessembinder, H., 2001. Price-time priority, order routing, and trade execution cost on NYSE-listed stocks. Unpublished paper, Emory University.

Blume, M., Goldstein, M. 1992. Differences in execution prices among the NYSE, the regionals, and the NASD. Working Paper 4-92, University of Pennsylvania.

Chakravarty, S., Harris, S., Wood, R., 2001. Decimal trading and market impact. Unpublished paper, Purdue University.

Corwin, S., Lipson, M., 2000. Order flow and liquidity around trading halts. Journal of Finance 55, 1771–1801.

Easley, D., Keifer, N., O'Hara, M., 1996. Cream-skimming or profit-sharing: the curious role of payment for order flow. Journal of Finance 51, 811–834.

Easley, D., O'Hara, M., 1987. Price, trade size, and information in securities markets. Journal of Financial Economics 19, 69–90.

Easley, D., O'Hara, M., Saar, G., 2000. How stock splits affect trading: a microstructure approach. Journal of Financial and Quantitative Analysis 36, 25–52.

Goldstein, M., Kavajecz, K., 2000a. Eighths, sixteenths, and market depth: changes in tick size and liquidity provision on the NYSE. Journal of Financial Economics 56, 125–149.

Goldstein, M., Kavajecz, K. 2000b. Liquidity provision during circuit breakers and extreme market movements. New York Stock Exchange Working Paper #2000-02.

Handa, P., Schwartz, R., Tiwari, A., 1999. Price improvement and price discovery on a primary market. Journal of Portfolio Management 1999, 55–64.

Harris, L., 1994. Minimum price variations, discrete bid-ask spreads, and quotation sizes. Review of Financial Studies 7, 149–178.

Harris, L., 1997. Does a large minimum price variation encourage order exposure? Working paper, University of Southern California.

Ho, T., Stoll, H., 1981. Optimal dealer pricing under transaction and return uncertainty. Journal of Financial Economics 9, 47–73.

Huang, R., Stoll, H., 1996. Dealer vs. auction markets: a paired comparison of execution costs on Nasdaq and the NYSE. Journal of Financial Economics 41, 313–357.

Irvine, P., Benston, G., Kandel, E. 2000. Liquidity beyond the inside spread: measuring and using information on the limit order book. Working paper, Emory University.

Kavajecz, K., 1999. A specialist's quoted depth and the limit order book. Journal of Finance 56, 747–771.

Lee, C., 1993. Market integration and price execution for NYSE-listed securities. Journal of Finance 48, 1009–1038.

Lee, C., Mucklow, B., Ready, M., 1993. Spreads, depths, and the impact of earnings information: an intraday analysis. Review of Financial Studies 6, 733–746.

Lightfoot, L., Martin, P., Peterson, M., Sirri, E. 1998. Order preferencing and market quality on United States equity exchanges. Working paper, United States Securities and Exchange Commission.

Lipson, M., 1999. Stock splits, liquidity, and limit orders. New York Stock Exchange Working Paper #99-04.

Macey, J., O'Hara, M., 1997. The law and economics of best execution. Journal of Financial Intermediation 6, 188–223.

Madhavan, A., Sofianos, G., 1998. An empirical analysis of NYSE specialist trading. Journal of Financial Economics 48, 189–210.

New York Stock Exchange Information Memo #99-38, July 31, 1999.

Petersen, M., Fialkowski, D., 1994. Posted versus effective spreads. Journal of Financial Economics 35, 269–292.

Porter, D., Weaver, D., 1997. Tick size and market quality. Financial Management 26 (4), 5–26.
Ready, M., 1999. The specialist's discretion: Stopped orders and price improvement. Review of Financial Studies 12, 1075–1112.
Ronen, T., Weaver, D., 1998. The effect of tick size on volatility, trader behavior, and market quality. Working paper, Rutgers University.
Ross, K., Shapiro, J., Smith, K., 1996. Price improvement of SuperDot market orders on the NYSE. New York Stock Exchange Working Paper #96-02.
Schultz, P., 2000. Stock splits, tick size, and sponsorship. Journal of Finance 55, 429–450.
Sofianos, G., Werner, I. 1997. The trades of NYSE floor brokers. New York Stock Exchange Working Paper #97-04.
Stoll, H., 1978. The supply of dealer services in securities markets. Journal of Finance 33, 1133–1151.
United States Securities and Exchange Commission (SEC), 1997. Report on the practice of preferencing.
Wall Street Journal, 2000. NYSE Seeks to Start Decimal Stock Trade Beginning on August 28, July 26, p. C6.
White, H., 1980. A heteroskedasticity-consistent covariance matrix estimator and a direct test for heteroskedasticity. Econometrica 48, 817–838.