WHITE PAPER #12

Calculating NAVs with Bid Prices; Implications of a New Process

Robert Haddad, Interactive Data Pricing and Reference Data, Inc.
February 22, 2007

Introduction

The Canadian Institute of Chartered Accountants (CICA) has adopted new standards (Section 3855) for recognizing and measuring the valuation of financial instruments effective for financial years beginning on or after October 1, 2006.[1] The Canadian Securities Administrators have recently granted an exemption which permits all Canadian mutual funds to continue to calculate net asset value (NAV) for purposes other than financial statements in accordance with Canadian GAAP without giving effect to Section 3855. This was done in part to allow additional time for dialogue between the fund industry and the relevant local securities regulators. Under Section 3855 of the CICA Handbook, with respect to active markets, "the appropriate quoted market price for an asset held or a liability to be issued is usually the current bid price and, for an asset to be acquired or a liability held, the asking price."[2] For means of financial reporting in Canada, this means valuing long positions based on the bid price and valuing short positions based on the ask price. This paper considers some of the implications of Section 3855. In particular, we examine how this change might affect Canadian mutual fund NAVs.

Our research examines the implications of this change with a focus on the bid side of the equation. First, we review bid-close relationships, and analyze the dynamics of bid-close spreads on three hypothetical portfolios. Next, we examine the potential for a new arbitrage opportunity based purely on bid-close spread dynamics. Here, we create a model to forecast bid-close spread changes, and measure the strength of this model by simulating the use of our predictions in an arbitrage strategy. We then discuss the connection between our observations and fair valuation of international equity securities, if based on local bid prices instead of local closing prices.

Bid-Close Spread Variability

Most exchanges establish official bid and ask prices at or near the close of trading or auction period based on the actions of market participants. However, methodologies for determining end-of-day bid and ask quotes can vary among exchanges. For instance, bid-ask quotes for equities trading on the SETS market segment of the London Stock Exchange (LSE) are actually *derived* from the official closing price. In this situation, each security has a fixed spread attached to it, contingent upon price level, which determines bid and ask prices equidistant from the official closing price.

In the following charts, we illustrate these subtleties by comparing two hypothetical portfolios of differing "official bid" methodologies. Chart 1 shows the daily bid-close spreads (log differences) of a market capitalization weighted portfolio of 100 securities listed on the Tokyo Stock Exchange, called Japan100. Chart 2 shows the daily bid-close spreads of a market capitalization weighted portfolio of 100 securities listed on the London Stock Exchange, called UK100. Japan100 shows a fairly normal distribution around the mean over time, suggesting that the market randomly establishes bid prices each day. On the other hand, UK100 shows a clear pattern. The structure of controls in bid-ask spreads employed by the LSE yields a dramatically different picture, with very little variation from the mean bid-close spread.

---

[1] CICA Handbook, Section 3855 *Financial Instruments – Recognition and Measurement*

[2] *Id*. at A45.

Chart 1: Japan100 Daily Bid-Close Spreads
August 2005 to July 2006

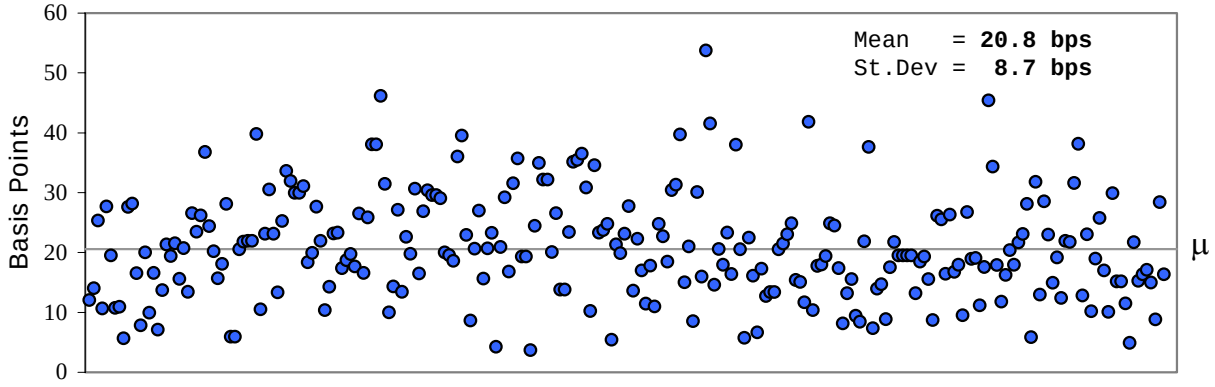

August 1, 2005 ---------- *July 31, 2006*

Chart 2: UK100 Daily Bid-Close Spreads
August 2005 to July 2006

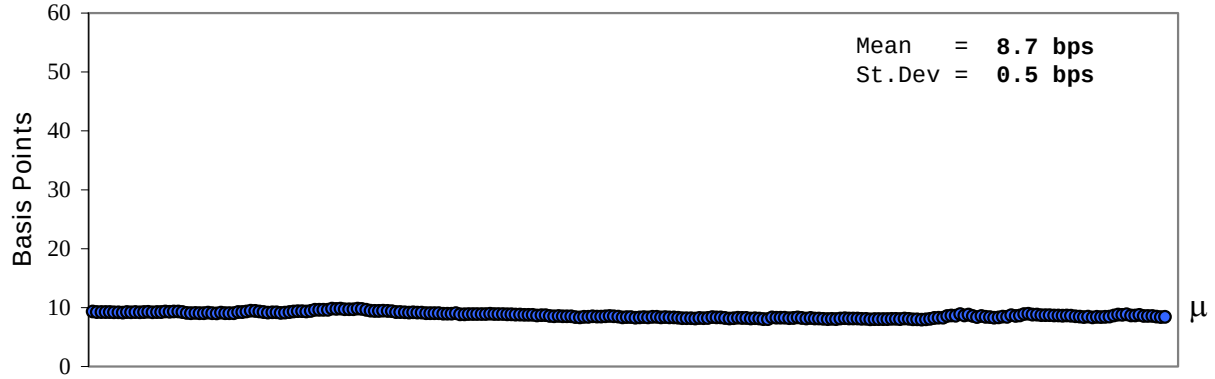

Additionally, we examined the daily bid-close spreads for a hypothetical market capitalization weighted portfolio created from 750 global securities, called International750, compiled from 18 international stock markets and exchanges, excluding securities in Japan100 and UK100, and find that the results are approximately in the middle of the Japan100 and UK100 results. In light of the variability in Japan100, this suggests that diversification seems to have lowered the variability of bid-close spreads in International750. Table 1 below provides our results for all three hypothetical portfolios.

Table 1.

Portfolio	*Mean of Bid-Close Spreads*	*Standard Deviation of Bid-Close Spreads*
Japan100	20.8	8.7
UK100	8.7	0.5
International750	12.9	4.0

A "New" Arbitrage Opportunity

A new arbitrage opportunity may be created if Canadian mutual funds value equity portfolios with bid prices. This opportunity results from the bid-close spread, with the hypothesis that a wide or narrow spread will revert towards the mean bid-close spread the next day. For instance, if the average bid-close spread for a particular portfolio of securities is 10 basis points (bps), and there happens to be a bid-close spread of 50 bps on a given day, the strategy is to buy into the fund, with the expectation that the bid-close spread will tighten towards its mean of 10 bps the next day. In this example, if the bid-close spread moves to the mean of 10 bps the next day, holding all other variables constant, the arbitrageur would realize a return of 40 bps simply on bid-close spreads tightening. Buying into the fund on a day when there exists an exceptionally wide bid-close spread can be considered comparable to receiving a modest discount on the purchase. Conversely, the strategy would be to redeem shares on days when there exists an exceptionally narrow bid-close spread.

To test this arbitrage opportunity, we created a regression model to forecast the *change* in bid-close spreads of Japan100, and measured the strength of this model through an out-of-sample application. Since our hypothesis is that a wide or narrow bid-close spread will revert towards the mean bid-close spread the next day, our independent variable is the bid-close spread (on day_t) less the mean bid-close spread from the prior 120 days (mean from day_{t-121} to day_{t-1}), which we refer to as "excess spread". The excess spread serves to predict the change in bid-close spread from day_t to day_{t+1}, our dependent variable in the model. Specifically, the hypothesis is that if the excess spread is greater than zero today, we expect that the bid-close spread will generally move downwards tomorrow in the direction of the mean bid-close spread, and vice versa, thus implying a negative correlation between the two variables.

We utilized a time series of the 120 days prior to August 1, 2005 for each regression variable to calculate an initial beta estimate. Accordingly, we rolled the time series forward to calculate all subsequent daily estimations of beta. We then multiplied each beta estimate to its corresponding excess spread to forecast a bid-close spread change for each day over the period of August 2005 to July 2006. For instance, we multiplied the initial beta estimate to the excess spread on August 1, 2005 to create a forecast of the change in bid-close spread from August 1, 2005 to August 2, 2005. The model's average beta estimate over the period was -0.85 (0.09 standard error), which was statistically significant at the 1% level throughout the period.[3] Intuitively, this indicates that a positive excess spread is generally followed by a tightening of bid-close spread the next day, while a negative excess spread is generally followed by a widening of bid-close spread the next day. Next, we examined the strength of this model by comparing our predicted bid-close spread change to the actual bid-close spread change. We simulated an arbitrage strategy by incorporating our predicted bid-close spread change as the trading signal.[4] The strategy is to buy into the fund when bid-close spreads are predicted to tighten beyond the trigger the next day, and redeem shares when bid-close spreads are predicted to widen beyond the trigger the next day. The simulation was executed at triggers of 0%, 0.05%, and 0.10%, and Table 2 displays the results of this strategy with three performance tests: (1) directional accuracy of predicted movements, (2) risk-adjusted excess return, and (3) correlation between predicted and actual bid-close spread changes.

[3] This model was also tested on each of the 100 securities individually, all of which were statistically significant at the 1% level.

[4] We test in the construct of "triggers" for simplicity. A more sophisticated algorithmic trading signal, based on exceeding one standard error of the prediction, yields moderately stronger results than the 0.10% trigger. The approximate standard error of the prediction on August 1, 2005 was 0.09%.

Table 2.

Trigger	*No. of Days Exceeding Trigger*	*Percentage of Days when Predicted Change was in the Correct Direction*	*Risk-Adjusted Excess Return*	*Correlation Between Predicted and Actual Spread Changes*
0.00%	247	71.7%	7.3%	0.64
0.05%	125	80.8%	6.1%	0.73
0.10%	48	83.3%	2.8%	0.80

To illustrate our model's predictive strength graphically, we isolate all days when our forecasted bid-close spread change (trigger) exceeds +/- 0.10%, and show the actual bid-close spread changes that transpired on those days. Chart 3 shows a clear trend, demonstrating that for 40 out of the 48 days (83.3% of the time) the model accurately predicted the direction of the next day's change in bid-close spreads. Hypothetically, holding all other variables constant, an arbitrageur could have earned a risk-adjusted excess return of 2.8% on these days by following this strategy.[5]

Chart 3: Japan100 Bid-Close Spread Changes Following Trading Signal of +/- 0.10%
August 2005 - July 2006

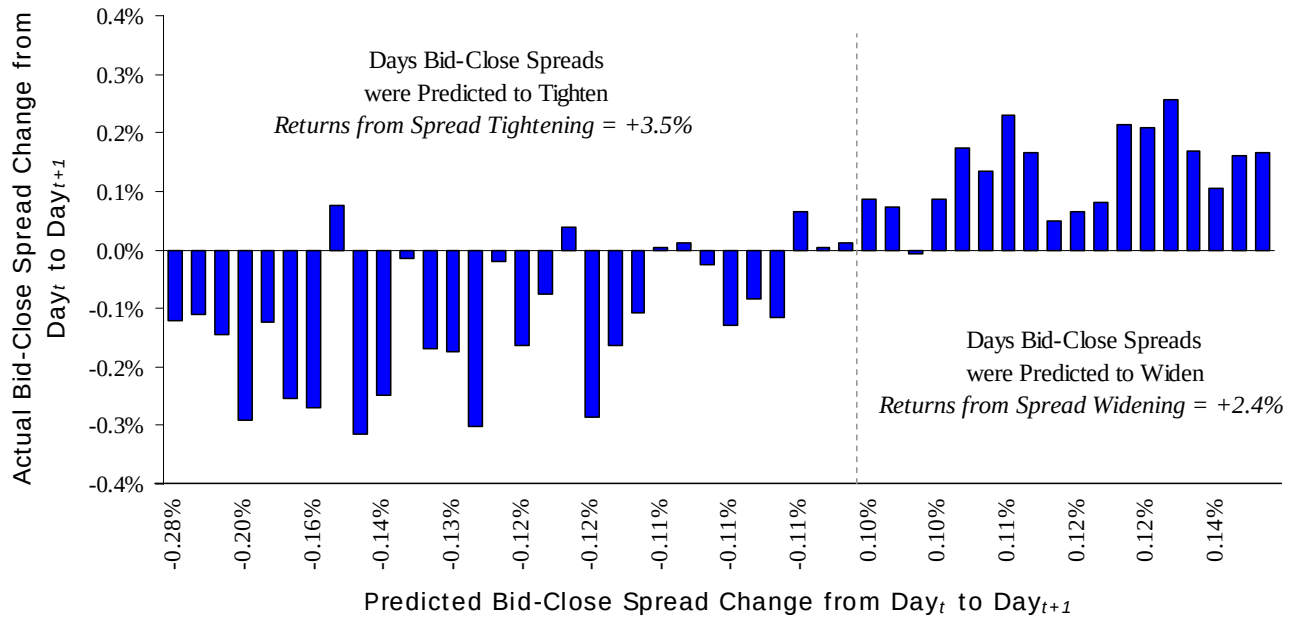

[5] We also replicated our model to test the predictability of bid-close changes in International750 and UK100. The model for International750 resulted in similar statistical significance as Japan100, although arbitrage returns were smaller, which is likely attributable to the lower variance in bid-close spreads. Results were statistically insignificant for UK100.

Fair Value Estimation Using Bid Prices

Next, we examine the effects of using bid prices in fair valuation models for international equity securities, and discuss the performance testing implications if these fair valuation models are applied to bid prices for NAV calculation. Today, funds use models, including Interactive Data Pricing and Reference Data's Fair Value Information Service, to assist them in estimating a price at the time of NAV calculation for international securities trading in markets that have closed. These models are currently based on the closing prices in the relevant markets. Our research has demonstrated that such fair value methodologies can reduce the predictability of NAV movements and the potential opportunities for arbitrage profits.[6]

In contrast, our studies with Japan100 and International750 indicate that bid-close spreads may introduce significant levels of predictability and variability into NAV movements. If fair value adjustments are applied to bid prices, the predictable features of bid-close spreads may potentially create new opportunities for arbitrage. For instance, if on day 1 there is an aggregate positive fair value adjustment to a fund's NAV, yet the aggregate bid-close spread of the portfolio happens to be exceptionally wide, the fair value adjustment on the NAV may seemingly under-correct to the NAV on day 2, since the aggregate bid-close spread of the fund is expected to tighten on day 2 on average. Hypothetically, an arbitrageur could make the bet that the fair value adjustment on day 1 will not go far enough, since the expected tightening of bid-close spreads will synthetically raise the NAV on day 2 even higher.

Additionally, the bid-close spread phenomenon should be taken into account when assessing the effectiveness of a fair value methodology through performance testing. To date, fair valuation performance tests have focused primarily on comparisons to the next day's opening trade price. For instance, a common test is to measure how much closer a fair value adjusted closing price is to the next day's opening price than the actual closing price. If we modify this test to measure the performance of a fair value adjusted bid price, the next day's opening trade may not be the most appropriate benchmark for comparison. Since it is a trade price, the next day's opening price will implicitly carry a bid-close spread of zero. Given this, a portion of the bid to next day open movement will generally include an inherent positive bias. For example, consider an international equity which has a closing trade price of $10.05 and a closing bid price of $10.00. Assuming nothing else influences the equity's opening price the next day, and the first trade is at $10.05, the close to next day open movement will be 0% and the bid to next day open movement will be 0.50%. In this case, the positive movement of 0.50% from the bid to next day open can be deceptive, effectively creating the perception of a changing market when no actual changes have taken place.[7] This additional noise could produce misleading fair value performance testing results. Even if a bid price at the open was available to use as a basis for testing, the variability of these quotations could still cause inconsistency in test results.

---

[6] See e.g. Peter Ciampi and Robert Haddad, *White Paper #10, Using NAV to Measure the Effectiveness of a Fair Value Methodology* (2005).

[7] Intuitively, if a model incorporating bid prices was designed for the sole purpose of creating values that are closer to the next day open than the actual bid price, this positive drift would be taken into account in the form of a positive intercept. Utilizing this positive drift to estimate a fair value is not viable in practice, since hypothetically, a positive adjustment would be applied to a security even while all relevant information pertaining to that security is suggesting that it is already fair valued.

Summary

We examined some of the potential effects of calculating NAVs using bid prices. First, we showed that the variability in bid-close spreads can differ between securities, with explanatory factors ranging from liquidity of the stock to the use of an "official bid" methodology by an exchange. Next, we demonstrated with our hypothetical Japan100 portfolio that changes in bid-close spreads may be predictable. Specifically, that wide or narrow spreads tend to revert towards their mean. We tested this hypothesis with a model that forecasted changes in bid-close spreads, and measured the accuracy of our predictions in an out-of-sample application over a one-year period. The results were statistically significant at the 1% level. This suggests that an arbitrage opportunity may exist if NAVs are calculated using bid prices. Future studies are necessary to determine if there is any substantial opportunity when considering the variability of other elements, such as the movements of a portfolio's underlying securities. Acknowledging the characteristics of bid-close spreads is necessary when testing the performance of a fair value procedure.

Addendum to White Paper 12

Robert Haddad, Interactive Data Pricing and Reference Data, Inc.
February 22, 2007

Abstract

This addendum continues our research of bid-close spread dynamics and the potential impact on recognizing and measuring the valuation of financial instruments under Section 3855. We expand on our original ideas by showing how a portfolio's bid-close spread variability, also referred to herein as volatility, relates to the market capitalization of the equities in the portfolio. Our hypothesis is that between small and large capitalization equity securities, large-cap securities should demonstrate higher liquidity, associating with narrower bid-ask spreads and therefore tighter bid-close spreads. Conversely, small-cap equities should be generally less liquid, with higher variability in bid-close spreads and therefore higher potential returns from arbitrage.

Data and Portfolio Creation

Our test sample includes 2000 equities; 1000 listed on the New York Stock Exchange (NYSE1000) and 1000 listed on the Toronto Stock Exchange (TSX1000). These equities were randomly selected from a larger population containing all active equities on each exchange from October 2004 to September 2006. NYSE1000 and TSX1000 were subdivided by market capitalization as of September 30, 2005 into 10 deciles of 100 equities. This resulted in 20 hypothetical portfolios, 10 from NYSE1000 and 10 from TSX1000. Next, equity bid-close spreads were equally weighted and aggregated to create a bid-close spread for each portfolio for each trading day. Thus, each decile represents a portfolio time series of bid-close spreads as of market close, categorized by market capitalization of the underlying equities.[8]

Forecasting Bid-Close Spread Changes

Next, we used the regression model described in White Paper 12 to predict bid-close spread changes for the decile portfolios. Similar to our results with Japan100, the model was statistically significant at the 1% level for each of the 20 portfolios, with the average t-statistic across the portfolios ranging from -8.3 to -10.4 over the period. We then measured the strength of our predictions out-of-sample by simulating the hypothetical arbitrage strategy described in White Paper 12 (buying shares on days when the bid-close spread is predicted to tighten the next day, and redeeming shares on days when the bid-close spread is predicted to widen the next day). Tables 3 and 4 provide summary results by decile for NYSE1000 and TSX1000, respectively, with maximum trading frequency (0% trigger) from October 2005 to September 2006.

[8] We believe that using bid-trade spreads at a time when an arbitrageur could process mutual fund transactions (e.g. 3:55 p.m. ET) for this analysis would have been ideal, but that data is not readily available. Using bid and close prices as of market close assumes that bid-trade spreads do not change significantly from 3:55 p.m. ET to 4:00 p.m. ET. See e.g. Thomas H. McInish and Robert A. Wood, *An analysis of Intraday Patterns in Bid/Ask Spreads for NYSE Stocks (1992),* which analyzes mean bid-ask spreads for each minute of the trading day. Their analysis shows that spreads at 4:00 p.m. are slightly wider on average then spreads at 3:55 pm. In general, McInish and Wood find the daily course of mean intraday spreads to follow a reverse J-shaped pattern. These findings are corroborated by Kee H. Chung and Xin Zhao, *Intraday Variation in the Bid-Ask Spread: Evidence after the Market Reform (2002).*

Table 3. NYSE1000

Decile	*Average Market Capitalization (USD)*	*Standard Deviation of Spread Changes*	*Risk-Adjusted Excess Return*	*Percentage of Days when Predicted Change was in the Correct Direction*	*Correlation Between Predicted and Actual Spread Changes*
1	$35.1 B	0.02%	1.6%	74.4%	0.62
2	$5.1 B	0.03%	1.7%	72.0%	0.62
3	$2.7 B	0.03%	2.1%	72.4%	0.65
4	$1.6 B	0.06%	3.3%	74.4%	0.69
5	$976 M	0.06%	3.2%	67.6%	0.67
6	$557 M	0.08%	4.2%	71.2%	0.67
7	$316 M	0.09%	5.2%	69.2%	0.60
8	$178 M	0.08%	5.1%	72.0%	0.67
9	$95 M	0.11%	6.4%	75.6%	0.66
10	$31 M	0.10%	5.8%	66.0%	0.64

Table 4. TSX1000

Decile	*Average Market Capitalization (USD)*	*Standard Deviation of Spread Changes*	*Risk-Adjusted Excess Return*	*Percentage of Days when Predicted Change was in the Correct Direction*	*Correlation Between Predicted and Actual Spread Changes*
1	$8.7 B	0.09%	4.8%	71.6%	0.60
2	$892 M	0.13%	7.4%	70.8%	0.59
3	$371 M	0.11%	6.1%	68.4%	0.59
4	$216 M	0.12%	7.1%	70.0%	0.63
5	$141 M	0.16%	9.1%	68.8%	0.57
6	$96 M	0.15%	9.0%	72.8%	0.59
7	$61 M	0.21%	11.7%	69.6%	0.58
8	$37 M	0.35%	19.8%	68.0%	0.60
9	$21 M	0.37%	22.2%	68.0%	0.65
10	$8 M	0.61%	36.1%	67.6%	0.61

The results for directional accuracy and correlation between predicted and actual bid-close spread changes indicate that the model's predictive strength is reasonably consistent across the NYSE1000 and TSX1000 deciles. However, a noticeable trend is that the volatility of bid-close spread changes, as indicated by the standard deviation of spread changes, for both NYSE1000 and TSX1000 gradually increases as we move from the large-cap deciles to the small-cap deciles. The same pattern exists for risk-adjusted excess returns; where the magnitude of potential arbitrage returns is directly related to the variability of bid-close spread changes. These results suggest a natural risk and return tradeoff. That is, even though the returns in decile 1 are of lower magnitude than the returns in decile 10, the bid-close spread changes in decile 1 contain substantially less volatility and thus less risk of a significant loss.

To illustrate, we plotted our results side-by-side, with the NYSE1000 deciles on the left presented in gray (charts 4.a, 4.b, and 4.c) and the TSX1000 deciles on the right presented in blue (charts 5.a, 5.b, and 5.c). Charts 4.a and 5.a show the risk-adjusted excess returns, charts 4.b and 5.b display the volatility of bid-close spread changes, and charts 6.a and 6.b provide the correlation between the predicted and actual spread changes.

NYSE1000 deciles

Chart 4.a: Risk-Adjusted Excess Returns

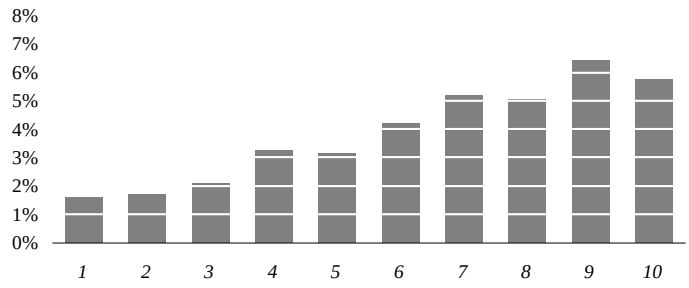

Chart 4.b: Volatility of Bid-Close Spread Changes (σ)

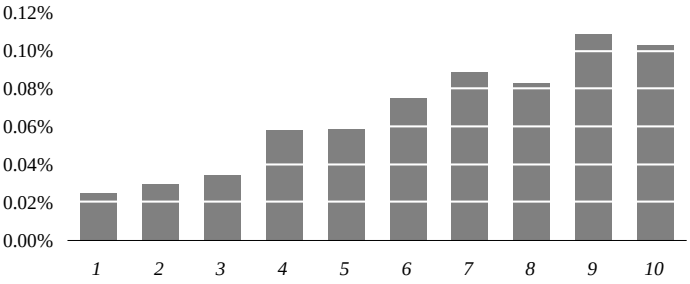

Chart 4.c: Correlation of Predicted and Actual Spread Changes

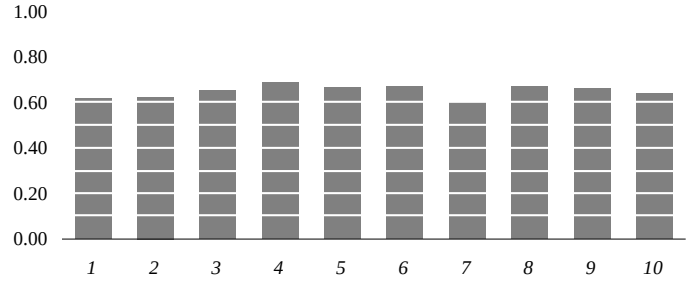

TSX1000 deciles

Chart 5.a: Risk-Adjusted Excess Returns

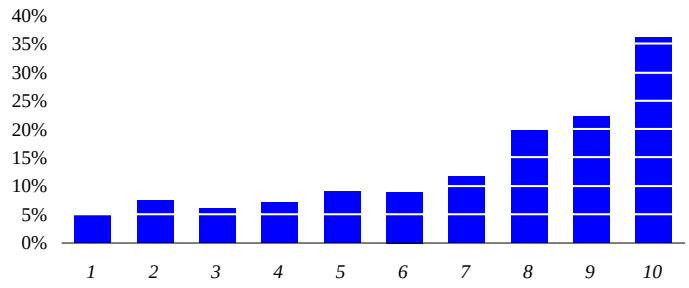

Chart 5.b: Volatility of Bid-Close Spread Changes (σ)

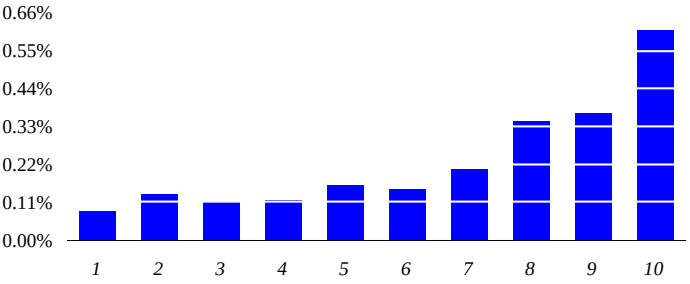

Chart 5.c: Correlation of Predicted and Actual Spread Changes

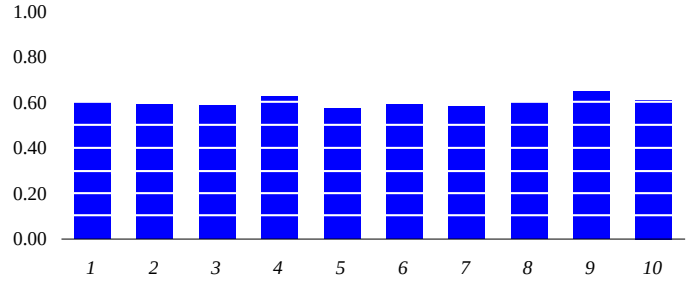

A clear pattern emerges from the images, as excess returns and volatility gradually become larger as we move out towards the smaller capitalization deciles.[9] The correlation figures indicate that the model sufficiently and consistently predicts bid-close spread changes, irrespective of market capitalization. Hypothetically, if bid prices were used to calculate these portfolio NAVs, holding all other variables constant, an arbitrageur could have earned statistically significant excess returns with this trading strategy.

[9] In charts 4.c and 5.c, the correlations are drawn to the same scale for NYSE1000 and TSX1000; while in charts 4.a, 4.b, 5.a, and 5.b, excess returns and volatility are not. Although the same pattern exists, the excess returns and volatility are dramatically higher for the TSX1000 deciles. These differences are somewhat proportional to the average market capitalization of equities within each decile.

Summary

In this addendum to White Paper 12, we explored the variability and predictability of bid-close spreads for equities listed in the U.S. and Canada. Our initial investigation methods were expanded to examine the general relationship between equity market capitalization and the variability and predictability of bid-close spreads. After categorizing equities by market capitalization, we constructed decile portfolios from each of our samples, NYSE1000 and TSX1000, and analyzed each portfolio's bid-close spread dynamics. We created a model to predict bid-close spread changes, and imposed a hypothetical trading strategy to measure the scope of these returns across decile portfolios. In particular, we demonstrated that bid-close spread changes may be predictable, suggesting that an arbitrage opportunity may exist if NAVs are calculated using bid prices. Moreover, we determined that while the model's predictive strength is fairly consistent across deciles, the increased bid-close spread volatility in the small-cap decile portfolios results in higher potential risk-adjusted excess returns.

References

Chung, Kee H. and Zhao, Xin, 2002, "Intraday Variation in the Bid-Ask Spread: Evidence after the Market Reform," *State University of New York (SUNY) at Buffalo, Buffalo, NY*. (on file with author)

McInish, Thomas H. and Wood, Robert A., 1992, "An Analysis of Intraday Patterns in Bid/Ask Spreads for NYSE Stocks," *Journal of Finance*, 47, 753-764.