

8-30-2005

Sentiment and the interpretation of news about fundamentals

Sean M. Flynn
Vassar College

Follow this and additional works at: http://digitalwindow.vassar.edu/faculty_research_reports

Citation Information

Flynn, Sean M., "Sentiment and the interpretation of news about fundamentals" (2005). *Faculty Research and Reports*. 96.
http://digitalwindow.vassar.edu/faculty_research_reports/96

This Working Paper is brought to you for free and open access by Digital Window @ Vassar. It has been accepted for inclusion in Faculty Research and Reports by an authorized administrator of Digital Window @ Vassar. For more information, please contact library_thesis@vassar.edu.

Sentiment and the Interpretation of News about Fundamentals

Sean Masaki Flynn*

August 30, 2005

ABSTRACT

The reaction of closed-end fund share prices to changes in portfolio values is on average the same whether funds are trading at discounts or premia and whether the changes in portfolio values are positive or negative. If closed-end fund discounts and premia do correctly measure investor sentiment, then these results suggest that investor sentiment does not affect the market's reaction to news about fundamentals. Alternatively, discounts and premia may not in fact measure investor sentiment, or sentiment may play no role at all in closed-end fund pricing. Noise-trader risk and trading costs also fail to explain the observed behavior.

*Department of Economics, Vassar College, 124 Raymond Ave. #424, Poughkeepsie, NY 12604. flynn@vassar.edu I would like to thank Osaka University's Institute for Social and Economic Research for their support while I worked on this paper. I also gratefully acknowledge the sedulous research assistance of Doug Park. All errors are my own.

This paper examines closed-end funds to see whether investor sentiment affects the market's reaction to news about fundamentals. Closed-end funds are the perfect vehicle for this inquiry for three reasons. First, their shares trade on major stock exchanges exactly like the shares of operating companies (Dimson and Marsh 1999). Second, they publish their portfolio values weekly, meaning that there is very little uncertainty about the fundamental valuations of closed-end funds or about how those valuations change over time (Lee, Shleifer, and Thaler 1990). Third, the fact that they typically trade at discounts or premia to their portfolio values gives a natural metric for investor sentiment: funds trading at discounts are funds about which investors are pessimistic, while funds trading at premia are funds about which investors are optimistic (Chopra, Lee, Shleifer, and Thaler 1993).

To examine whether investor sentiment affects the interpretation of news about fundamentals, I examine how closed-end fund share prices react to changes in fund portfolio values when conditioning on discount and premium levels. In addition, I look separately at increases and decreases in portfolio values in order to check whether the state of sentiment colors the market's reaction to positive and negative news. I find that the reaction of closed-end fund share prices to changes in portfolio values is on average the same whether portfolio values increase or decrease and regardless of the discount or premium level at which a fund is trading.

This behavior is confirmed by examining a comprehensive data set containing nearly every US and Canadian closed-end fund trading between 1985 and 2001. Using this data set, I first sort all changes in portfolio values into positive and negative changes. I do this because if sentiment affects the reaction of share prices to changes in portfolio values, sentimental investors might react differently to good and bad news—i.e. to positive and negative changes in portfolio values. I then sort the observations in each group by their associated discount and premium levels. I do this because you might expect optimistic and pessimistic investors to react very differently to changes in portfolio values.

Once these sorts are made, I account for the fact that discounts and premia tend to slowly mean revert towards rational valuations. This tendency means that prices tend to drift towards portfolio values

regardless of how portfolio values evolve over time. Consequently, in order to isolate the response of prices to changes in portfolio values, it is necessary to subtract out from any total price change the expected price change due to mean reversion.

After doing this, the data reveal a simple pattern: Average price reactions to changes in portfolio values are the same regardless of discount or premium level and regardless of whether the changes in portfolio values are positive or negative. Consequently, closed-end funds offer no support to the hypothesis that investor sentiment affects the way markets interpret news about fundamentals.

My second major finding is that prices substantially under-react to changes in portfolio values. This best seen by constructing price-response ratios, which divide the change in a given fund's price over a given period of time by the change in the fund's portfolio value over the same period of time. A price-response ratio of 1.0 means that a fund's price moves dollar-for-dollar with changes in its portfolio value, while ratios less than unity indicate under-reaction.

There is pervasive under-reaction in my data set. In particular, price-response ratios average about 0.70 for weekly changes in portfolio values and 0.69 for monthly changes in portfolio values. This level of under-reaction is of course hard to square with robust arbitrage activities. But since these changes in portfolio values are ex ante unpredictable (because they are caused by the unpredictable changes in the prices of the assets in fund portfolios), you might attribute the low price-response ratios to some sort of delayed reaction to news about changes in portfolio values. However, this hypothesis is hard to square with the fact that you still find substantial under-reaction if you look exclusively at dividend events, which are of course pre-announced reductions in portfolio values. Since the average price-response ratio on ex dividend days is 0.83, there is still substantial under-reaction even when the market can fully anticipate changes in portfolio values.

The fact that price-response ratios do not vary by discount and premium levels is very important because it contradicts the positive correlation between discount and premium levels and price-response

ratios that you would expect to find if discounts and premia were in fact driven by sentiment. To understand why you might expect such a positive correlation, compare the price-response ratios that you would expect to get from two funds, one trading at a premium and another at a discount.

Imagine that the first fund has an initial portfolio value per share of \$10 but a market price per share of \$13, so that it trades at a thirty percent premium. If the portfolio value per share of this fund increases by a dollar to \$11 and if the level of positive sentiment that generated the initial thirty percent premium remains unchanged, then you would expect that the price of the fund's shares would increase by \$1.30 in order to maintain the thirty percent premium. That is, it would have a price-response ratio of 1.30. By contrast, imagine a different fund that has an initial portfolio value per share of \$10, but a market price of only \$7, so that it trades at a thirty percent discount. If the portfolio value of this fund increases by one dollar to \$11 and if the level of sentiment that generated the thirty percent discount remains unchanged, then you would expect the price of the fund to increase by only \$0.70 in order to maintain a thirty percent discount. That is, it would have a price-response ratio of only 0.70.

As you can see, if discount and premium levels were reflective of sentiment, then price-response ratios would be positively correlated with discount and premium levels. Thus, the fact that price-response ratios are actually unrelated to discount and premium levels bears strongly against sentiment having a role in closed-end fund pricing. Indeed, it seems that the only way that a sentiment-based interpretation of closed-end fund discounts and premia could be made consistent with this behavior is if sentiment could somehow affect the market's pricing of levels of portfolio value but not its pricing of changes in portfolio value. Since such a story is rather implausible, the unresponsiveness of price-response ratios to changes in portfolio values should be taken, I feel, as evidence that weighs strongly against discount and premium levels being caused by differences in investor sentiment.

The results found here also provide only ambiguous support for the noise-trader explanation for discounts and premia given by DeLong, Shleifer, Summers, and Waldmann (1990). Under the noise-trader hypothesis, irrational noise traders unpredictably buy and sell closed-end fund shares. Because

their future actions are noisy and unpredictable, they generate a unique source of risk that deters rational arbitrageurs from fully and immediately offsetting the trading activities of the noise traders. As a result, the noise-traders are able to raise and lower closed-end fund share prices relative to portfolio values—thereby generating movements in discount and premium levels.

The fact that price-response ratios are low *is* consistent with arbitrageurs being deterred by noise-trader risk. However, the noise-trader hypothesis runs into trouble with the fact that average price responses do not vary systematically with discount and premium levels. This is a problem for the noise-trader hypothesis because Flynn (2005) gives strong evidence that noise-trader risk varies with discount and premium levels. In particular, noise-trader risk is smallest at moderate discount levels of about six percent, and grows increasingly large as you move to either deeper discounts or towards premia. Consequently, if price-response ratios were affected by the degree of noise-trader risk, you would expect them to also vary with discount and premium levels. The fact that they do not suggests that noise-trader risk is at best an incomplete explanation for the behavior of closed-end fund discounts and premia.

Finally, the use of trading costs to explain the pattern of price-response ratios also runs into difficulty given the fact that bid-ask spreads also vary significantly with discount and premium levels while price-response ratios are constant across discount and premium levels.

Section I discusses the data set and formally defines discounts and premia. Section II reviews the literature on sentiment and asset pricing as well as the evidence favor of interpreting discounts and premia as a measure of investor sentiment. Section III reviews closed-end fund behavior, focusing on the behavior of discounts and premia. Section IV explores how closed-end fund share prices react to changes in portfolio values. Section V discusses what this behavior implies about the effect of sentiment on the market's reaction to news as well as whether the observed behavior can be explained by either noise-trader risk or trading costs. Section VI concludes.

I. Data and Definitions

In June of 2001, I purchased a subscription to the Fund Edge data set sold by Weisenberger/Thompson Financial. Fund Edge is used primarily by analysts for its real-time streaming data on fund portfolio values and share prices, which can be utilized to compute the discount or premium at which closed-end funds trade. Fund Edge also contains historical time series of fund prices, net asset values, dividend payments, and other variables.

However, the way the data is sold, a subscriber only receives historical data for the funds currently in existence at the time of subscription. Consequently, my data set only contains historical time series on the 462 closed-end funds trading in the United States and Canada in June of 2001.¹ This implies, of course, that the data set suffers from survival bias. However, because this paper is interested in the behavior of discounts and premia under the normal situation in which a fund is expected to continue operating indefinitely, the survival bias in the data set actually works as a nice filter. Those funds that went through the abnormal process of being liquidated or converted into open-end funds have been eliminated. This is important since fund prices always jump to equal fund portfolio values when liquidations or conversions are announced (Brauer 1984). Such jumps are unrelated to changes in portfolio values and would—if they were included—bias this paper's results. It is consequently convenient that they have been eliminated.

The data set contains daily data on fund share prices but only weekly data on fund portfolio values.² Consequently, the analysis below focuses on price reactions to weekly or longer changes in portfolio values. However, since the data set also contains a full reporting of dividend payments and indicates dividend ex dates, I am able to do some daily analysis—in particular I can see how share prices react to dividend payments on ex dates.

In the analysis below, I use only data from 1985-2001 and ignore all data points before 1985. I do this for three reasons. First, the time series for older funds before that year are in some cases

incomplete. Second, there were very few funds in existence before that time.³ Third, discounts and premia appear to have behaved substantially differently prior to that year. In particular, the average discount or premium was much deeper prior to 1985, with funds trading at an average discount of thirty percent or more through the late 1970's compared with an average discount of only a few percent starting around 1985.⁴ Consequently, it seems prudent in what follows to avoid what may very well be some sort of regime change.

In this paper, discounts will be defined as negative numbers. Let N_t be the net asset value (NAV) per share of a fund at time t . The NAV of a fund is simply its portfolio value less any liabilities the fund may have; it is the value that would be distributed to shareholders if the fund were to liquidate immediately. Let P_t be the fund's price per share at time t . The discount or premium is defined implicitly by the identity $P_t = (1 + D_t)N_t$. Values of $D_t > 0$ are called premia, while values of $D_t < 0$ are referred to as discounts. In this paper, I will multiply D_t by 100 and refer to discounts and premia in percentages.

II. Sentiment, Noise-trader Risk, and Closed-end Funds

Closed-end funds are actively managed investment companies that are similar to mutual funds except that they do not redeem their own shares. Instead, their shares trade on stock exchanges where supply and demand are free to determine their prices. The very interesting result of this institutional difference is that closed-end fund share prices only rarely equal the per-share value of fund NAVs. Rather, funds almost always trade at either discounts or premia relative to their NAVs. Furthermore, since the discounts and premia of individual funds are often quite volatile, it appears that fund share prices move substantially independently of NAVs.

Zweig (1973) was the first to suggest that this decoupling of share prices from NAVs might be caused by investor sentiment. Under this hypothesis, the worse the sentiment about a fund's prospects,

the larger will be the discount at which it trades, while the more positive the sentiment, the larger will be the the premium at which it trades.

This conjecture was fleshed out by Lee, Shleifer, and Thaler (1991), who suggest that discounts and premia are driven by the differential sentiment of small investors. In particular, they note that in the US, closed-end fund shares are mostly owned by small investors, while the assets held by funds in their portfolios are mostly traded by large institutional investors. So, to the extent that the sentiment of small investors varies from that of large investors (who may in fact be above sentiment and trade rationally), closed-end fund share prices will vary independently from the values of their underlying portfolios.

However, this explanation is unsatisfying insofar as it does not explain why rational arbitrageurs do not enter the closed-end fund markets and rectify the discounts and premia caused by differential sentiment. This is particularly true for funds trading at large discounts, where there are no obvious barriers to arbitrage and it would seem easy for arbitrageurs to rectify such underpricings simply by buying up a large number of fund shares. Arbitrage for funds trading at large premia may perhaps be more difficult, but Flynn (2004) shows that there is indeed a large volume of short selling of funds trading at premia, and that the intensity of short selling increases as premia grow larger. Consequently, it is also hard to argue that the existence of large premia is due to explicit arbitrage barriers like institutional short selling constraints.

Instead, it would appear that discounts and premia can exist and persist because something discourages arbitrageurs from undertaking the volume of arbitrage that would be necessary to fully rectify closed-end fund mispricings. The noise-trader model developed by DeLong, Shleifer, Summers, and Waldmann (1990) and applied to closed-end funds by Lee, Shleifer, and Thaler (1991) suggests that arbitrageurs self-limit their arbitrage activities because they are reluctant to bear the price risk generated by unpredictable noise-traders. In this model, the closed-end fund markets are populated by both rational arbitrageurs and irrational noise traders. The noise traders trade at random, independently of information flows. The result is that at any given time, they may trade in a way that would tend to cause

an existing mispricing to widen rather than narrow. Since the noise traders' behavior is completely unpredictable, this possibility constitutes a unique risk that the arbitrageurs cannot diversify away. This causes the risk-averse arbitrageurs self-limit the size of their arbitrage positions. As a result, the volume of arbitrage is insufficient to fully offset the mispricings caused by the noise traders—thereby giving the noise traders the ability to drive funds to discounts and premia inconsistent with fundamentals.

The noise-trader hypothesis is given substantial support by Flynn (2005), who finds extensive evidence in the same data set examined here that fund-specific, non-diversifiable noise-trader risk not only exists but increases the farther a fund's discount or premium diverges from rational levels. Put differently, the more mispriced a fund is, the larger the noise-trader risk that must be borne by arbitrageurs wishing to take out arbitrage positions. This means that the degree to which arbitrage is discouraged increases with the magnitude of mispricings. The more mispriced the fund, the greater the noise-trader risk and, presumably, the greater the deterrent to undertaking arbitrage.

Noise-trader risk could consequently provide the deterrent to arbitrage that would be necessary for sentiment to be able to affect the share prices of closed-end funds. That being said, it is very important to reflect on the relationship between sentiment and noise-trader risk. In particular, it should be noted that the two hypotheses are not the same. The noise-trader hypothesis only requires that noise traders be random. It does not stipulate what causes the randomness. The randomness might be due to unpredictable fluctuations in sentiment, or it might not. And while the sentiment hypothesis requires some sort of arbitrage barrier in order for sentiment to be able to affect closed-end fund share prices but not their underlying portfolio values, it does not specify what the arbitrage barrier actually is. The barrier might be noise-trader risk, or it might be something else.

Keeping the two hypotheses clearly mentally separated before examining the results below is important because it is quite possible to reject one hypothesis without rejecting the other. But before reviewing the results, I will give you some general background on the distribution and evolution of

discounts and premia, as well as the specific aspects of discount and premium behavior that support the sentiment and noise-trader hypotheses.

III. The Behavior of Discounts and Premia

Discounts and premia fluctuate. Over a period of a few years, the D_t value of a given fund may, for instance, move from a discount of twenty percent to a premium of ten percent and then back to a substantial discount. Sudden and drastic movements of fund D_t levels are also seen quite regularly, where there maybe a ten or fifteen percentage point movement taking place in the space of just a few months or even just a few weeks. Extreme discounts and premia of $D_t < -30\%$ or $D_t > 30\%$ are relatively rare, but because there are hundreds of funds currently trading, you can typically see examples of such extreme D_t values every week when updates are released.

[Insert Figure 1 Approximately Here]

Figure 1 plots a relative frequency histogram of the 227,066 weekly D_t observations of the 462 funds found in Fund Edge over the period January 1985 through May 2001. It was constructed by placing each observation into a one-percentage point wide bin, where the bins range from a discount of -50% to a premium of 50%. Because of the limited number of bins, the figure excludes 761 outliers, most of which lie in the right (premium) end of the distribution. The largest discount in the data set is -66.5%, while the largest premium is 205.4%.

The two most important things to notice about the bell-shaped distribution are where it is centered and how it is distributed around that central value. The mode is located at the -6% to -5% discount bin. This is important since the models of both Ross (2002) and Flynn (2002) predict that fully rational investors would price funds at about a -7% discount. This prediction is made under two assumptions. First, managers are not gifted and their stock picking will result in returns that are no better than what

you could get from an index fund invested in similar assets. Second, investors rationally capitalize out future expected management fees. Doing so using the average fee rate results in a discount of about -7%.

The location of the mode is thus a victory for arbitrage and rational pricing. But what is harder to explain is the spread of the distribution around the mode. There are, of course, both behavioral and rational explanations set forth in the literature. On the one hand, Lee, Shleifer, and Thaler (1990) argue that discounts and premia that differ from the rational level are due to changing sentiment about whether or not fund managers will beat or trail the market in the future. On the other hand, rational considerations like outstanding tax liabilities, agency problems, and management fees may go some way to explaining the cross-sectional variation in fund D_t levels. (See Dimson and Minio-Kozerski 1999 for an extensive survey of the debate over whether discounts and premia can be explained rationally.)

Such rational factors, however, have a very difficult time explaining three prominent features of closed-end fund behavior. First, none of these factors changes with enough rapidity to explain the rapid changes that are often seen in the D_t levels of individual funds. Second, they have a very difficult time explaining why funds would ever trade at premia (or, for that matter, at any D_t level in excess of the rational discount level that capitalizes out expected future management fees.) Third, since these explanatory factors are fund-specific, they have a very hard time explaining the high level of correlation found among both levels and changes of fund D_t values, as was first pointed out by Lee, Shleifer, and Thaler (1991).

By contrast, the sentiment hypothesis can handle each of these three considerations rather neatly. First, rapidly changing sentiment about a fund's own outlook can explain rapid changes in a fund's D_t . Second, positive sentiment about the ability of a fund's manager to beat the market would cause optimistic investors to bid a fund's share price up to a premium that would reflect their optimism. Third, the high correlation of D_t levels across funds is easily explained if you posit that sentimental investors have a general level of sentiment that affects their feelings about all funds, in addition to whatever

feelings they have about the manager of any specific fund. To the extent that general sentiment changes, D_t levels across funds will be correlated.

The ability of the sentiment hypothesis to explain these three aspects of discount and premium behavior weighs in favor of interpreting fund D_t levels as sentiment metrics. But the sentiment explanation requires some sort of arbitrage barrier to make sense. DeLong, Shleifer, Summers, and Waldmann (1990) and Lee, Shleifer, and Thaler (1991) argue for segmented markets such that small investors dominate the markets for closed-end fund shares while large institutions dominate the markets for the stocks and bonds held by funds in their underlying portfolios. The segmentation that keeps the two groups apart is supposed to be something along the lines of closed-end funds being too small or too illiquid for large investors to bother with.

The market segmentation theory, however, is weak on two counts. The first is that you still see the same behavior of discounts and premia in extremely large funds with billions of dollars in assets and lots of liquidity.⁵ The second problem is that the small vs. large investor story only applies to the USA. In the UK, you still get the same behavior of discounts and premia despite the fact that both the IPOs and secondary markets of UK-traded closed-end funds are dominated by large institutional investors who also dominate the markets for the assets held by UK-traded closed-end funds in their portfolios (Brown 1998). Consequently, it does not appear that segmented markets provide the arbitrage barrier necessary for sentiment to be able to affect closed-end fund share prices.

That being said, other aspects of discount and premium behavior strongly suggest that there must be some effective barrier to arbitrage in closed-end funds. The strongest evidence for this is the fact that fund D_t levels show an average rate of mean reversion that is too slow to be consistent with strong and effective arbitrage pressures acting to drive fund D_t values towards levels consistent with fundamentals.

The slowness of mean reversion can be estimated econometrically by noting that mean reversion appears to follow an AR(1) process. That is, the evolution of D_t appears to be well described by the

process $D_t = \bar{D} + \phi(D_t - \bar{D}) + \varepsilon_t$, where \bar{D} is the level to which discounts and premia mean revert, ϕ is the rate of decay, and ε_t represents noisy shocks (perhaps caused by noise trading.) Flynn (2005) finds that if you estimate this equation using monthly data on the 462 funds, you find that $\bar{D} = -5.20\%$ and $\phi = 0.916$. This means that discount and premium levels mean revert to the mode discount bin in Figure 1, but that they revert so slowly that, on average, 91.6% of any gap between the current D_t level and $\bar{D} = -5.20\%$ remains after one month. A similar regression run on weekly data shows that on average 97.2% of any gap between the current D_t level and the mean reverting discount level remains after one week.

Flynn 2005 makes a strong case that this slow rate of mean reversion is due to noise-trader risk. But it should be noted that he finds a U-shaped pattern for noise-trader risk, with the bottom of the U located at a discount of about -6% (consistent with the mode in Figure 1.) As funds move to either deeper discounts or towards premia, noise-trader risk increases substantially. This has an important implication for price-response ratios. If price-response ratios reflect the intensity of arbitrage, and if arbitrage varies with the intensity of noise-trader risk, then price-response ratios should vary in a U-shaped fashion with discount and premium levels since noise-trader risk varies in a U-shaped fashion with discount and premium levels. As I will show you, however, price-response ratios show no such pattern, a fact that weakens the noise-trader explanation for discount and premium levels.

IV. The Reaction of Closed-end Fund Share Prices to Changes in Net Asset Values

This section shows that once you correct for discount and premium mean reversion, closed-end fund share prices react the same to changes in NAV regardless of whether those changes are positive or negative and regardless of the discount or premium level at which funds are trading when NAV changes.

To perform this analysis, I construct price-response ratios by dividing a fund's price change over a given period of time, ΔP , by its change in NAV over the same period of time, ΔNAV , to form the ratio $\frac{\Delta P}{\Delta NAV}$. Please note that *ceteris paribus*, you would expect such price-response ratios to be positive because you would expect fund prices to move in the same direction as fund NAVs. For instance, if $\Delta NAV > 0$, then you would expect the price to rise as well, so that the ratio would have positive terms in both the numerator and the denominator. Similarly, if $\Delta NAV < 0$, you would expect the price to fall, so that the numerator and the denominator would both be negative.

Before examining how price-response ratios vary over discount and premium levels and by whether or not they are responding to positive or negative changes in NAVs, it is instructive to look at Table I, which gives average price-response ratios for various subsets of the data *before* disaggregating the observations in those subsets by discount and premium levels and by whether or not they occur in response to positive or negative changes in NAVs.

[Insert Table I Approximately Here]

The first column of Table I contains row numbers to make referencing the table easier. The second column contains the three lengths of time over which price reactions are observed—daily, weekly, and monthly. The third column indicates whether the price-response ratios used in a given row were adjusted for the tendency of discounts and premia to mean revert. The fourth column gives any other restrictions. For instance, the row may include only observations where dividends were paid, or where the changes in NAV to which prices are reacting are greater in magnitude than a given value, such as twenty cents. The fifth column gives the number of observations meeting all the requirements stipulated by columns two, three, and four. And the sixth and seventh columns give, respectively, the mean and standard deviation of the price-response ratios for each row's observations. Please note that in all cases, only observations where the relevant daily, weekly, or monthly change in NAV was non-zero are counted in the table. Cases where NAV did not change were eliminated as the point of this paper is to

investigate how prices respond to changes in NAV. In addition, the data used in Table I only includes data points for which $-30\% < D_t < 30\%$ in order to exactly match the data plotted in Figures 2 through 6. This eliminates the more extreme data points but does not affect the averages given in Table I because there are so few observations in the tails of the distribution.

Table I makes it clear that closed-end fund share prices consistently under-react to changes in NAVs. Indeed, a quick glance down column six shows that average price-response ratios are less than unity in all cases. They are 0.70 or less for weekly data no matter what subsets of the weekly data are examined in rows (1) through (6), while the largest average price response of 0.84 happens in row (8), which gives the average price-response ratio on dividend ex dates to dividend payments exceeding twenty cents. That these average values truly reflect price-response ratios being less than unity is strongly supported by t-tests. While it is true that the standard deviations in the final column are relatively large, due to the large number of observations it is the case that t-tests decisively reject the hypothesis that price-response ratios are equal to unity. The p-values are less than 0.0001 for most rows, and the hypothesis is rejected at the one-percent significance level for all rows.

Comparisons of individual rows also make it clear that the under-reaction is robust to correcting for mean reversion, robust to whether or not dividends were paid, and robust to the magnitude of the changes in NAVs. For instance, since rows (1) and (2) vary only by whether or not price-response ratios are corrected for mean reversion, it is clear by comparing the average price-response ratio of 0.67 in row (1) to the value of 0.70 for row (2) that correcting for mean reversion makes no substantial difference. In the same way, a comparison of the average price-response ratio of 0.69 in row (4) with the 0.70 value in row (5) makes it clear that price-response ratios in weeks where dividends were paid are much the same as those in weeks where dividends were not paid. Finally, a comparison of the weekly average price-response ratios in rows (2) and (3)—which differ only because row (3) restricts itself to changes in NAVs exceeding twenty cents—shows that price-response ratios remain substantially below unity even if you restrict yourself to larger changes in NAVs. The same fact holds true when comparing the

daily average price-response ratios to dividend payments in rows (7) and (8) as well as monthly average price-response ratios in rows (9) and (10).

The rows that restrict attention to only larger changes in NAVs are especially important because price-response ratios for small changes in NAVs are less convincing than those for large changes in NAVs. For instance, if a fund's NAV changes by one cent during a week where its market price falls by 23 cents, it will have a price-response ratio of -23. But because the change in NAV was so small, it is most likely that the price change was due to the market reacting to other factors. By contrast, if there is a large change in NAV, the price change that accompanies it is much more easily ascribed to being a response to the change in NAV. It is thus reassuring that average price-response ratios are still substantially lower than unity when looking only at observations where there were substantial changes in NAVs.

The elimination of observations with small changes in NAVs is also reassuring because it greatly reduces standard deviations, as can be seen by comparing the values in the final column for rows (2) and (3), for rows (5) and (6), for rows (7) and (8), and for rows (9) and (10). For each pair, the elimination of observations that result from small changes in NAVs causes a huge reduction in standard deviation. The decline is due to the fact that large price-response ratios are almost inevitably the result of cases where prices move a lot when NAVs only change by a little. It is satisfying that when these observations are eliminated, the resulting distributions are substantially tighter while still being centered on average values that are well below unity and which in most cases are nearly the same as when all observations are included.

Under-reaction is thus a robust finding. But does price reaction vary with discount and premium levels so that sentiment might be shown to affect how markets react to fundamentals? The next two subsections address this question.

A. Unadjusted Price Reactions to Changes in Net Asset Values

Figure 2 takes the 174,364 weekly observations reported in row (1) of Table I and sorts the associated price-response ratios by both D_t levels and by whether or not there was an increase or a decrease in NAV. These observations constitute all instances where a fund had successive weeks where both NAV and price were reported in the data set, where the weekly change in NAV was non-zero, and where the fund's D_t falls into the central $-30 < D_t < 30$ range displayed in Figure 2.

[Insert Figure 2 Approximately Here]

For each of the 174,364 observations, I first sort the associated price-response ratio by whether its denominator (ΔNAV) is positive or negative. I then use each observation's D_t at the beginning of the week to place it into a one-percent wide bin ranging from a -30% discount to a 30% premium.⁶ Finally, for each bin, I take separate averages for all of the $\Delta NAV > 0$ cases and all of the $\Delta NAV < 0$ cases that fall into that bin.⁷ Figure 2 shows the results of plotting those two sets of averages.

What is of course striking is that the price-response ratio not only varies by discount or premium level, but that it also varies depending on whether ΔNAV is positive or negative. For observations where the portfolio value falls ($\Delta NAV < 0$), the average response ratio steadily rises from less than zero at deep discounts to a value of 1.0 at a small premium before continuing to rise to values near 2.0 at a 20% premium. (After that, much higher values are sometimes observed, but because there are so few observations in these bins, it is best to ignore them.) For observations where portfolio values rise ($\Delta NAV > 0$), the price-response ratio makes a slow decline, starting at values above 1.0 for deep discounts before falling down to zero at a small premium, and then falling even farther to values around -1.0 at a 20% premium. (Again, it is best to ignore values after that point due to there being only a small number of observations in the high premium bins.)

The differences between the two lines are also highly statistically significant. If you run ANOVA F-tests for each of the 61 bins in Figure 2 to see whether the means of the $\Delta NAV > 0$ and $\Delta NAV < 0$ observations within each bin are different, you find that the means are different at a 10% confidence level for 49 of the bins, and that the means are different at the 1% confidence level for 35 of the bins. What is more, those numbers include the bins in the middle of the figure where the two lines converge and you would not expect to be able to reject equality. For the large majority of the bins more than a few percentage points away from the crossing point, ANOVA tests strongly reject equality.

As I will show you in the next subsection, adjusting for the mean reversion of discounts and premia completely obliterates the results of Figure 2. But before doing that, it is important, I think, to consider how a person who only sees Figure 2 might incorrectly interpret it as making a strong case for sentiment modulating the market's reactions to news about changes in fundamentals.

Consider first the left side of Figure 2. There, you find very strong price responses to positive changes in NAV, but very weak price responses to negative changes in NAV. It is as if sentimental investors notice and react to good information but ignore and don't react to bad information. Since this is the discount side of the graph, it suggests that the pessimistic investors who have driven the funds to discounts show a large response to positive information that contradicts their pessimism, while showing only a very small response to information that confirms their pessimism.

The right side of Figure 2 is also consistent with investors only reacting to information that contradicts their expectations. There, we see average price-response ratios above unity for bad news, while seeing price-response ratios of -1.0 and below for good news. Since this is the positive end of the figure, with funds presumably trading at premia because of positive investor sentiment, it again seems that when changes in portfolio values confirm expectations, there is a small price response, but that when changes in portfolio value contradict expectations, there is a substantial price response.

As you can see, a coherent behavioral story can be built around the crossing lines of Figure 2. The next section shows, however, that the X-shape plotted out by the two lines is driven entirely by the mean reversion of discounts and premia. Once that tendency is accounted for, the X collapses and along with it any behavioral interpretations that you might be tempted to ascribe to Figure 2. Consequently, I feel that Figure 2 should serve as a caveat to anyone examining how prices react to information in other assets whose prices may also be affected by mean reversion toward rational levels.

B. Price Reactions to Changes in Net Asset Values After Correcting for Mean Reversion

Figure 3 is the same as Figure 2 except that it corrects price-response ratios for expected mean reversion. As you can see, once this correction is made, the price responses for both increases and decreases in portfolio values are virtually identical for nearly all of the discount and premium bins. Indeed, ANOVA F-tests run separately for each of the bins find that, at the 10% significance level, you cannot reject equality of means for 52 of the 61 bins.

[Insert Figure 3 Approximately Here]

I adjust for mean reversion in the following way. You will recall that mean reversion is modeled as an AR(1) process in which $D_{t+1} = \bar{D} + \phi(D_t - \bar{D}) + \varepsilon$, where \bar{D} is the value to which discounts and premia mean revert, ϕ gives the rate of decay, and ε is a random shock of mean zero. This equation is very handy because once you have run a regression and estimated \bar{D} and ϕ , you can use an observed value of D_t along with the fact that ε is mean zero to make a prediction about what the expected discount or premium at period $t + 1$ will be.

Let me denote this prediction as \hat{D}_{t+1} to distinguish it from D_{t+1} , which is the actual value of the discount or premium at period $t + 1$. \hat{D}_{t+1} is simply the expected value of the AR(1) process with the ε dropping out because it is mean zero. That is, $\hat{D}_{t+1} = \bar{D} + \phi(D_t - \bar{D})$. If you run a panel regression on

the weekly Fund Edge data over the period 1985-2001, you find that $\bar{D} = -4.85$ percent and $\phi = 0.972$. Consequently, given any value of D_t , we can predict that the expected value of D_{t+1} in percents will be $\hat{D}_{t+1} = -4.85 + 0.972(D_t + 4.85)$.

Next, it is important to realize that the AR(1) mean reversion operates completely independently of fund NAVs. Indeed, you can calculate it without having to know either N_t or N_{t+1} . That is very convenient because once period $t + 1$ arrives and you know the actual value of N_{t+1} , you can combine it with \hat{D}_{t+1} to get the expected price consistent with the expected rate of mean reversion. Let me call this expected price \hat{P}_{t+1} . To derive \hat{P}_{t+1} , start with the period $t + 1$ version of the identity that implicitly defines the discount or premium given that period's price and NAV: $P_{t+1} = (1 + D_{t+1})N_{t+1}$. Given N_{t+1} and the expected discount or premium, \hat{D}_{t+1} , this formula implies that $\hat{P}_{t+1} = (1 + \hat{D}_{t+1})N_{t+1}$. Consequently, you can determine the price that would result just from the expected rate of mean reversion, independent of any effect that the change in NAV from N_t to N_{t+1} might have on the total change in price from P_t to P_{t+1} .

That, in turn, allows you to divide the total price response accompanying a change in NAV into two parts. When NAV changes by $\Delta NAV = N_{t+1} - N_t$, the total price change, $\Delta P = P_{t+1} - P_t$, can be divided into an expected part due to mean reversion, $\Delta P^{exp} = \hat{P}_{t+1} - P_t$, and an unexpected part, $\Delta P^{unexp} = P_{t+1} - \hat{P}_{t+1}$, that can be interpreted as the market's response to the change in NAV. Consequently, since the expected pace of mean reversion proceeds the same no matter what ΔNAV is, the proper way to measure the price response to a change in NAV is to subtract out the change due to mean reversion and concentrate only on the unexpected change in price. That is precisely what I do in Figure 3, where for each of the 61 bins I plot the average unexpected price-response ratios net of mean reversion, $\Delta P^{unexp} / \Delta NAV$, separately for cases where $\Delta NAV > 0$ and $\Delta NAV < 0$.

As you can see, once expected mean reversion is accounted for, the price-response ratios are about the same whether or not there was an increase or a decrease in NAV and without regard to the D_t level. Indeed, ANOVA tests run separately on each of the 61 bins show that you cannot reject equality at the

10% significance level for 53 of the 61 bins. Consequently, it appears that investors react symmetrically to changes in fundamentals whether the changes themselves are positive or negative and whether the existing state of investor sentiment (to the extent that it can be measured by discounts and premia) is positive or negative.

The dramatic effect of correcting for mean reversion in Figure 3 points out that the X shape traced out by the two price-response lines in Figure 2 has nothing at all to do with sentiment affecting the interpretation of news about fundamentals. Rather, it is an artifact generated entirely by mean reversion.

You can see how the artifact is generated by looking at the left side of Figure 2. There, funds are trading at large discounts, meaning that mean reversion will tend to cause $D_{t+1} > D_t$. Stated slightly differently, prices will tend to rise relative to NAVs. To see how this causes price-response ratios to be different for increases and decreases in NAVs, suppose that at time t , there is a fund trading at a thirty percent discount, with $N_t = 10$ dollars and $P_t = 7$ dollars. Then imagine that over the next week, NAV stays the same so that $N_{t+1} = N_t = 10$ dollars. Given that the AR(1) formula predicts that $\hat{D}_{t+1} = -4.85 + 0.972(D_t + 4.85)$, you should expect that since on average $\varepsilon = 0$, it will be the case that on average $D_{t+1} = \hat{D}_{t+1} = -29.3$ percent. With NAV still equal to \$10, this implies that the expected price due to mean reversion will be \$7.07. So the effect of mean reversion will be for the price of the fund to increase by seven cents, even if NAV stays the same.

But suppose that NAV doesn't stay the same. Suppose, first, that it rises to \$10.50. Given that the average rate of mean reversion is independent of NAV changes, you would still expect a discount of -29.3 percent. Applying that to an NAV of \$10.50 gives a predicted price at time $t + 1$ of \$7.42. So in this case where NAV rises, the total expected price response will be 42 cents, implying a price-response ratio of $42/50 = 0.84$.

Contrast that with the price-response ratio that results if NAV falls by 50 cents down to \$9.50. Applying the expected discount of -29.3 percent to that NAV gives an expected price of \$6.72. This implies a fall in price of 28 cents, leading to a price response ratio of $-28/-50 = 0.56$.

As you can see, the fact that mean reversion is taking place means that there will be a bias in the price-response ratios. At the left end of Figure 2, the bias will cause the ratios of responses to cases where NAV falls to be lower than those where NAV rises. At the other end of the figure, the bias will be reversed, which is why the lines cross and you get an X shape.

A good way to make the bias even more concrete is to compare what the price response ratios would be if there were no mean reversion. For the case where NAV increases by 50 cents, imagine that there is no mean reversion and that $D_{t+1} = D_t = -30$ percent. If so, the price would rise only to \$7.35, implying a price-response ratio of $35/50 = 0.70$. Similarly, if the discount stays at -30 percent while NAV falls down to \$9.50, the price would fall to \$6.65, implying a price-response ratio of $-35/-50 = 0.70$. In the absence of mean reversion, the price-response ratios are equal.

C. Robustness Checks

The previous section shows that once you account for mean reversion, discount and premium levels appear to have no effect on the average values of price-response ratios. Because this may indicate that investors sentiment has no effect on the market's interpretation of news about fundamentals, it is important to know if the results shown above are truly robust. In particular, are the results the same if you eliminate the observations featuring only small changes in NAV, are they robust to price-response ratios examined over time periods longer than a week, and do they generalize to fully predictable movements in NAV, namely dividend payments? I examine these questions in the following subsections.

C.1. Eliminating Observations with Only Small NAV Changes

To see that price-response ratios remain unrelated to D_t levels even after eliminating all observations featuring small NAV changes, please look at Figure 4, which uses a smaller vertical scale than Figures 2 and 3 in order to give you a close-up view. It is constructed by eliminating most of the 192,198 observations used to make Figure 3. What remains are the 38,195 observations for which weekly NAV changes exceed 20 cents in magnitude.

[Insert Figure 4 Approximately Here]

These observations are the same as those used in row (3) of Table I, where the average price-response ratio over these observations is 0.64. It is thus no surprise that the lines giving the average price-response ratios for positive and negative changes in portfolio values in Figure 4 vary around that number. More importantly, as in Figures 2 and 3, the averages show no relationship with D_t levels and are statistically indistinguishable from each other given that, for 51 of the 61 bins, ANOVA tests cannot reject equality of means at the 10% significance level. Consequently, the results I present above are not driven by the price-response ratios associated with small NAV changes.

This is important because for many observations, weekly NAV changes are small, often only a cent or two. For such small changes in NAV, it is hard to argue convincingly that the accompanying price changes are the result of the associated NAV changes since other factors do affect prices and may tend to dominate the overall change in price when NAV changes by only a small amount. By contrast, if you restrict attention to larger NAV movements, a presumably substantial fraction of any associated price movement will be attributable to the change in NAV. What is reassuring about Figure 4 is that in such cases where NAV moves a lot and gives prices something substantial to react to, we find the same pattern as when we examine all observations.

C.2. Price Responses to Dividend Changes

A different robustness check can be done by examining the price-response ratios that accompany dividend events. These changes in NAVs are especially interesting because they are pre-announced. The markets know well in advance of the ex dividend dates by exactly how much NAV will change. Consequently, any systematic under-reaction cannot be blamed on information only being slowly impounded into prices.

[Insert Figure 5 Approximately Here]

Figure 5 exploits the fact that Fund Edge contains daily data to isolate and sort the same-day price responses to dividend payments by discount and premium levels. To construct this figure, I begin with the 37,988 dividend payments made by the 462 funds over the period 1985-2001. I record the size of each dividend and then look to see if I have the two pieces of price data needed to calculate the price change that accompanies the dividend—the closing price on the trading day before the ex date and the closing price on the ex date. There are 35,783 observations for which this data is also available.

For each of these observations, I next check to see if there is an NAV and a same-day price reported during the previous week.⁸ There are 35,252 observations for which this information is also available and with which I can construct D_t values with which to sort the observations by discount or premium level. Of these observations, 35,113 fall into the 61 one-percent wide bins ranging from a -30% discount to a 30% premium that you see in Figure 5; these are the same observations as in row (7) of Table I. After sorting, I construct a price-response ratio for each observation by dividing the change in closing prices from the day before the ex dividend date to the ex dividend date by the negative of the dividend payment made on the ex dividend date.⁹ I then, finally, calculate the bin averages for those ratios. They are plotted in Figure 5.

As you can see, average price-response ratios remain unrelated to discount and premium levels even when prices are responding to changes in NAV caused by dividend payments. In fact, from a -20% discount to a 10% premium, the ratio is remarkably stable, varying within a tight range around a value of 0.83. Only in the more extreme D_t bins where there are very few observations in each bin do you see any substantial deviation from this behavior.

Please note that the price-response ratios that I report in Figure 5 are not adjusted for mean reversion. Therefore, the flatness of the line is completely inherent. I feel that this makes for especially strong confirmation that price-response ratios do not vary with discount and premium levels. No matter what the D_t level is before the dividend is paid, prices fall by about 0.83 of the dividend.¹⁰

Also note that the average price-response ratio of 0.83 found here for closed-end funds when they pay dividends is consistent with the price-response ratios of operating companies when they pay dividends. Elton and Gruber (1970), Michaely (1991), and Eades, Hess, and Kim (1994) all find operating company price-response ratios to be less than 1.0, and the value of 0.83 found here falls right in the middle of the range found by Graham, Michaely, and Roberts (2003) for NYSE issues over the period 1996-2001. Within that literature, price-response ratios less than 1.0 are referred to as ex date premiums. Consequently, the contribution of this paper to that literature is to demonstrate that ex date premiums do not appear to vary by whether or not a company is over- or under-valued at the time it pays a dividend. And to the extent that over- or under-valuation is reflective of investor sentiment, sentiment appears to have no effect on ex date premiums.

Finally, by comparing rows (7) and (8) of Table I you can also see that the average price-response ratio resulting from dividend payments is not affected by the size of dividend payments. Indeed, the average price-response ratio of 0.84 in row (8) in response to dividends larger than 20 cents is nearly identical to that of 0.83 in row (7) where price-response ratios to dividends of all sizes are included in the average.¹¹ And, while I have left it out in order to save space, it is the case that if you sort the

price-response ratios that result from dividends exceeding 20 cents, you again get a figure like Figure 5 where it is clear that price-response ratios do not vary with D_t levels.

C.3. Longer Time Horizons

Figure 6 shows that the results obtained for weekly data in Figure 3 hold up when examining longer time periods. Figure 6 gives the mean-reversion corrected price-response ratios that result from monthly changes in NAVs. It was constructed the same way as Figure 3, except that estimated values of an AR(1) process run on monthly (rather than weekly) data were used to correct for mean reversion.

[Insert Figure 6 Approximately Here]

If you do not correct for mean reversion in the monthly data, you get an X-shaped figure similar to that found for weekly data in Figure 2. But by once again correcting for mean reversion, the X shape is eliminated and the the two price-response lines in Figure 6 are flat with respect to D_t and show no systematic differences between price responses to increases and decreases in NAVs.¹² It is consequently the case that the patterns found for weekly data extend to monthly data as well. Price-reaction ratios remain low, and because they do not vary with discounts and premia, there is again no evidence that sentiment affects how closed-end fund share prices react to changes in fundamentals. This suggests that whatever it is that retards arbitrage in the short-run also appears to retard it in the longer run and that this factor does not appear to dissipate over time. This can also be seen by comparing the average price-response ratio for monthly data of 0.69 in row (9) of Table I with that of 0.70 for weekly data in row (2) of Table I.

V. Discussion with Respect to Sentiment, Noise-trader Risk, and Trading

Costs

I have shown that once you correct for the average rate at which discounts and premia mean revert, price-response ratios do not vary by discount or premium level and also do not vary by whether or not they are in reaction to positive or negative changes in NAV. These facts bear strongly on the plausibility of both the sentiment and noise-trader risk hypotheses. Additional evidence about the behavior of fund bid-ask spreads also casts doubt on the ability of trading costs to explain the behavior of price-response ratios.

A. Constant price-response ratios and investor sentiment

As discussed above, if investor sentiment determines closed-end fund discounts and premia, then you would expect to find a positive relationship between D_t and the price-response ratio. The basic idea is that if a given level of sentiment has driven a fund to trade at, for instance, only 70 cents per dollar of NAV, then any change in the fund's NAV should cause the fund's share price to change at a similar ratio. Since this ratio will obviously increase as D_t increases, there should be a positive correlation between price-response ratios and D_t levels. More broadly speaking, sentiment should affect the market's pricing of changes in NAV in the same way that it is alleged to determine the market's pricing of levels of NAV. The evidence, however, clearly contradicts this prediction. Price-response ratios bear no relationship to D_t levels. Consequently, to the extent that sentiment can be measured using D_t levels, this paper offers no support for the hypothesis that sentiment affects how closed-end fund share prices react to changes in fundamentals.

Since I am in effect testing the joint hypothesis that sentiment affects closed-end funds and that it can be measured using D_t levels, several possibilities are consistent with the fact that price-response

ratios do not vary with D_t levels. First, sentiment may play no role at all in closed-end funds. That would explain why the variable that attempts to measure sentiment, D_t , has no relationship with price-response ratios. Second, sentiment may play a role, but is not properly measured using D_t levels. That would also explain why D_t levels do not affect price-response ratios. Third, sentiment may be very strange in that it can affect the market's pricing of levels of NAV without affecting how the market reacts to changes in NAV. That would allow sentiment to affect D_t levels but not affect price-response ratios. Since this final option seems rather implausible, either of the first two possibilities seems more likely.

B. Constant price-response ratios and noise-trader risk

Since investor sentiment appears to be incompatible with the pattern of price-response ratios found in this paper, can they be explained by alternative hypotheses? At first glance, noise-trader risk looks like a good candidate. First, average price-response ratios are always less than one. This would be consistent with noise-trader risk deterring arbitrage so that prices do not respond robustly to changes in NAVs. Second, price response ratios are similar for both increases and decreases in NAVs. This would make sense since noise-trader risk would presumably deter arbitrage symmetrically for both increases and decreases in NAVs. Third, price-response ratios don't vary with D_t levels. This would be consistent with Flynn's (2005) contention that closed-end fund arbitrageurs face noise-trader risk at all D_t levels.

However, Flynn (2005) also finds that the amount of noise-trader risk varies substantially with D_t levels. This seems quite inconsistent with the results presented here because if noise-trader risk deters arbitrage and if the intensity of noise-trader risk varies with D_t levels, then presumably price-response ratios would also vary with D_t levels since arbitrageurs would be deterred to a greater extent when noise-trader risk was more intense. The fact that price-response ratios do not vary with D_t levels

consequently suggests that noise-trader risk is at best a partial explanation for the pattern of price-response ratios found in closed-end funds.

C. Constant price-response ratios and trading costs

Market frictions offer another potential explanation for the flatness of price-response ratios with respect to D_t levels. If such frictions deterred arbitrage but did not vary with D_t levels, then they would provide a good explanation not only for the fact that price-response ratios are on average less than one, but also for the fact that price-response ratios are not related to D_t levels. Unfortunately, Flynn (2005) provides evidence about bid-ask spreads and trading volume that makes it hard to attribute the flatness of price response ratios to market frictions.

Using the same data set analyzed here, Flynn finds that bid-ask spreads and trading volume are U-shaped when plotted against D_t levels. They have a nadir near the long-run mean-reverting discount level of -6% and grow steadily higher as you move away from that level in either direction. Indeed, both trading volume and spreads about double by the time you reach either discounts of -30% or premia of 30%.

These increases in spreads and volume make it very difficult to attribute the observed pattern of price-response ratios to market frictions. For instance, the behavior of spreads makes it hard to argue simultaneously that market frictions can explain why price-response ratios are low and why price-response ratios do not vary with discount and premium levels. The problem is that if price-response ratios are low because of costly arbitrage, then you would expect price-response ratios to vary with bid-ask spreads since wider spreads reduce the profitability of arbitrage activities. But while bid-ask spreads are U-shaped with respect to D_t levels, price-response ratios are flat with respect to D_t levels.

This fact not only casts doubt on the ability of bid-ask spreads in particular to retard arbitrage, but also makes it hard to argue that market frictions in general are responsible for the price-response

patterns seen in the data. While it is true that other market frictions like broker's fees are presumably independent of D_t levels and might therefore provided a way for market frictions to retard arbitrage but not vary with D_t levels, it is not clear that their ability to deter arbitrage should be presumed when the ability of a similar market friction to deter arbitrage is not apparent in the data.

Further doubts arise when you consider volume. That's because volume increases along with spreads. If arbitrage were being deterred by higher spreads, then you would tend to expect less volume *ceteris paribus*. But since you actually find more volume, it is hard to argue that arbitrage is being deterred by higher spreads. Consequently, the volume data also provides no support for low price-response ratios being the result of trading costs or market frictions limiting the intensity of arbitrage.

VI. Conclusion

Does sentiment affect the market's response to news about fundamentals? Judging from the reaction of closed-end fund share prices to changes in fund NAVs, the answer is no. To the extent that sentiment can be measured using closed-end fund discounts and premia, there is no evidence that sentiment affects either the magnitude of price reactions to changes in NAVs or the way that prices react to increases rather than decreases in NAVs.

Rather, prices tend to under-react to changes in NAVs by the same amount regardless of the discount or premium at which a fund is trading, while increases and decreases in NAVs lead to under-reactions that are statistically indistinguishable. These findings are robust to examining various time horizons and hold true for both changes in NAVs caused by fund dividend payments as well as changes caused by fluctuations in net asset values. Consequently, to the extent that investor sentiment is correctly measured by discounts and premia, this behavior is inconsistent with sentiment playing any role in how markets react to changes in fundamental valuations. Simply put, if discounts and premia do measure

sentiment, then sentiment has no effect on how markets react to fundamentals because such reactions are totally unrelated to discount and premium levels.

In addition, the observed behavior also makes it very difficult to maintain the belief that discounts and premia reflect investor sentiment. The problem is that if sentiment drives discounts and premia by determining the ratios of fund prices to *levels* of NAVs, then you would expect sentiment to also cause similar ratios to be applied to *changes* in NAVs. Indeed, if sentiment were driving discounts and premia, you would expect price reactions to changes in NAVs to vary positively with discount and premium levels. But since price reactions in fact show no relationship to discount and premium levels, this paper also finds no support for the hypothesis that closed-end fund discounts and premia are determined by investor sentiment.

The behavior of price reactions to changes in net asset values also fails to offer strong support to the noise-trader risk explanation for closed-end fund discounts and premia. The problem is that while noise-trader risk can explain the under-reaction of fund prices to changes in NAVs as being the result of arbitrageurs not being willing to fully adjust share prices to changes in fundamentals because they are discouraged by noise-trader risk, you would expect the magnitude of under-reaction to be related to discount and premium levels. That's because Flynn (2005) provides extensive evidence that noise-trader risk in closed-end funds varies strongly with discount and premium levels. Consequently, if noise-trader risk is causing under-reactions by deterring arbitrageurs, then the magnitude of those under-reactions should vary with discount and premium levels. The fact that they do not weighs against the noise-trader risk hypothesis.

The behavior of bid-ask spreads in closed-end funds also reduces one's confidence that trading frictions might be able to explain the pattern of under-reactions. The problem is that while trading costs like the bid-ask spread would presumably reduce the profitability and intensity of arbitrage and therefore lead to price under-reactions of the type seen in the data, such frictions would also have to be constant across discount and premium levels in order to explain the fact that price reactions to

changes in fundamentals are constant across discount and premium levels. Instead, bid-ask spreads increase dramatically the more funds move to either large discounts or large premia. As a result, if bid-asks spreads were affecting the intensity of arbitrage, you would expect price reactions to vary with discount and premium levels. Since this is not true, it is hard to argue that the constant pattern of price under-reactions in closed-end funds is due to trading frictions.

The major result of this paper is that sentiment as measured by discounts and premia appears to have no effect on how closed-end fund prices react to changes in fundamental values. In addition, neither noise-trader risk nor bid-ask spreads appear to be able to explain the nature of those reactions. As a result, two lines of further research seem potentially fruitful. The first would be to see if alternative measures of investor sentiment have any effect on how stock prices react to changes in fundamental values. And the second would be to explain price reactions in closed-end funds. In particular, why do prices under-react to changes in fundamental values, and why are those under-reactions unrelated to discount and premium levels. Discoveries in either line of research would tell us interesting things about information impounding and its potential limitations.

References

- Brauer, Gregory A., 1984, Open-ending Closed-end Funds, *Journal of Financial Economics* 13, 491–507.
- Brown, Christopher, 1998, Discounts: Theory & Evidence, Investment Trusts, *Cazenove & Co.*
- Chopra, Navin, Charles M. Lee, Andrei Shleifer, and Richard Thaler, 1993, Yes, Discounts on Closed-end Funds are a Sentiment Index, *Journal of Finance* 48, 801–808.
- DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, Noise Trader Risk in Financial Markets, *Journal of Political Economy* 98, 703–738.
- Dimson, Elroy, and Carolina Minio-Kozerski, 1999, Closed-end Funds: A survey, *Financial Markets, Institutions, and Instruments* 8.
- Eades, Ken, Pat Hess, and Han E. Kim, 1994, Time Series Variation in Dividend Pricing, *Journal of Finance* 49, 1617–1638.
- Elton, Edwin, and Martin Gruber, 1970, Marginal Stockholder Tax Rates and the Clientele Effect, *Review of Economics and Statistics* 52, 68–74.
- Flynn, Sean Masaki, 2002, A Model of the Discounts on Closed-end Mutual Funds, the Quantification of Investor Sentiment, and the Inability of Arbitrage to Force Closed-end Fund Share Prices to Par, Ph.D. thesis University of California, Berkeley.
- Flynn, Sean Masaki, 2004, Short-selling Behavior When Fundamentals are Known: Evidence from NYSE-traded Closed-end Funds, *Vassar College Department of Economics Working Papers, No. 63.*
- Flynn, Sean Masaki, 2005, Arbitrage, Noise-trader Risk, and the Cross-section of Closed-end Fund Returns, *Vassar College Department of Economics Working Papers, No. 69.*

- Graham, John R., Roni Michaely, and Michael R. Roberts, 2003, Do Price Discreteness and Transactions Costs Affect Stock Returns? Comparing Ex-Dividend Pricing before and after Decimalization, *Journal of Finance* 58, 2611–2636.
- Lee, Charles M., Andrei Shleifer, and Richard Thaler, 1990, Closed End Mutual Funds, *The Journal of Economic Perspectives* 4, 153–164.
- Lee, Charles M., Andrei Shleifer, and Richard Thaler, 1991, Investor Sentiment and the Closed-end Fund Puzzle, *The Journal of Finance* 46, 75–109.
- Michaely, Roni, 1991, Ex Dividend Day Stock Price Behavior: The Case of the 1986 Tax Reform Act, *Journal of Finance* 46, 845–860.
- Ross, Stephen A., 2002, Neoclassical Finance, Alternative Finance, and the Closed End Fund Puzzle, *European Financial Management* 8, 129–137.
- Zweig, Martin E., 1973, An Investor Expectations Stock Price Predictive Model Using Closed-end Fund Premiums, *The Journal of Finance* 28, 67–78.

Notes

¹Three hundred eighty nine were listed on the NYSE, sixty one on the AMEX, seven on the NASDAQ, four on the Toronto Exchange, and one—the NAIC Growth Fund—on the Mid-west Exchange.

²458 of the 462 funds are American and subject to the Investment Companies Act of 1940, which requires them in Section 30(e) to report their net asset values to shareholders semi-annually. However, funds have developed the tradition of reporting their net asset values weekly, and nearly all funds follow this tradition. The net asset values are published every Monday in the Wall Street Journal and in other financial newspapers.

³Back issues of the Wall Street Journal list fewer than 30 funds as trading in 1985.

⁴I know this by comparison of the Fund Edge data after 1985 with the Wall Street Journal data from 1965-1985 that was used in Lee, Shleifer, and Thaler (1990) and which Charles Lee graciously offered me. A proper analysis of why average discounts decreased so much between 1980 and 1985 is beyond the scope of this paper but two possible candidates are improving sentiment and the drastic reduction in marginal taxes rates during the first Reagan administration.

⁵For instance, Tri-Continental Corporation has two billion dollars under management and an average daily trading volume in excess of 100,000 shares. But it typically trades at large discounts.

⁶There are 2,579 observations with D_t values outside this range (in addition to the 174,364 within this range.) I ignore them because when they are placed into more extreme bins, there are too few

observations in each of the extreme bins for their averages to be meaningful. Consequently, for Figure 2 and the subsequent analysis, I concentrate on the center of the distribution.

⁷There were 79,575 cases where $\Delta NAV < 0$ and 94,889 cases where $\Delta NAV > 0$. Increases outnumber decreases because NAVs tend to drift up for several months until a dividend is paid. Thus, many positive observations are countered by a single negative (dividend) observation.

⁸NAVs are typically calculated using closing prices on Fridays and then reported on Mondays. So if a dividend was paid any time during the next week, the previous Friday's NAV and price are used to construct the discount used to sort the observation. For funds using closing prices on other days, the matching day's price is used.

⁹Since dividend payments are reductions in NAVs, I divide by the negative of each dividend payment to make these ratios consistent in sign with the other price-response ratios reported in this paper.

¹⁰To the naked eye, there may appear to be a slight upward trend in Figure 5, but if you regress the price-response ratios of all observations on their associated D_t levels, the slope is insignificantly different from zero. Furthermore, if you do make a correction for mean reversion using plausible estimates of the one-day rate of mean reversion, the apparent upward slope as seen by the naked eye flattens out.

¹¹Please note that a 20 cent dividend is a substantial dividend in this context given that the median dividend for bond funds is seven cents while the median dividend for stock funds is 11 cents.

¹²At the 10-percent significance level, ANOVA tests cannot reject the hypothesis that the price-response ratios from increases and decreases in NAV are the same for 51 of the 61 bins.

Table I
Averages and standard deviations of price-response ratios for various subgroups of the data

This table presents averages and standard deviations of price-response ratios for various subgroups of the Fund Edge data. The subgroups are constructed by conditioning on several filters including time horizon, whether the data was adjusted for the average rate of mean reversion, whether or not dividends were paid, and the magnitude of NAV movements.

Row No.	Time Horizon	Adjusted for Mean Reversion?	Other Restrictions?	Number	Avg.	SD
(1)	Weekly	No	None	192,198	0.67	7.85
(2)	Weekly	Yes	None	192,198	0.70	7.79
(3)	Weekly	Yes	$ \Delta NAV > 0.20$	38,195	0.64	1.25
(4)	Weekly	Yes	Dividends	33,069	0.69	7.12
(5)	Weekly	Yes	No dividends	159,129	0.70	7.92
(6)	Weekly	Yes	No dividends; $ \Delta NAV > 0.20$	33,339	0.63	1.28
(7)	Daily	No	Dividends	35,113	0.83	1.89
(8)	Daily	No	Dividends; $ \Delta NAV > 0.20$	2,511	0.84	0.74
(9)	Monthly	Yes	None	50,266	0.69	8.92
(10)	Monthly	Yes	$ \Delta NAV > 0.50$	9,031	0.80	0.90

Figure 1. Relative frequency histogram of 227,066 weekly closed-end fund discounts and premia, 1985-2001.




Figure 2. Ratio of the change in weekly share price to the change in weekly NAV, $\Delta P/\Delta NAV$, plotted separately for cases where $\Delta NAV > 0$ and $\Delta NAV < 0$.




Figure 3. Ratio of the unexpected change in weekly share price (net of the expected change due to mean reversion) to the change in weekly NAV, $\Delta P^{unexp} / \Delta NAV$, plotted separately across various discount and premium levels for cases where $\Delta NAV > 0$ and $\Delta NAV < 0$.




Figure 4. Close-up view of the ratio of the unexpected change in weekly share price (net of the expected change due to mean reversion) to the change in weekly NAV, $\Delta P^{unexp} / \Delta NAV$, plotted separately across various discount and premium levels for cases where $\Delta NAV > 0.20$ and $\Delta NAV < -0.20$




Figure 5. The ratio of average changes in share prices on dividend ex dates divided by the negative of their associated dividend values, $\Delta P/(-Dividend)$, plotted across discount and premium levels.

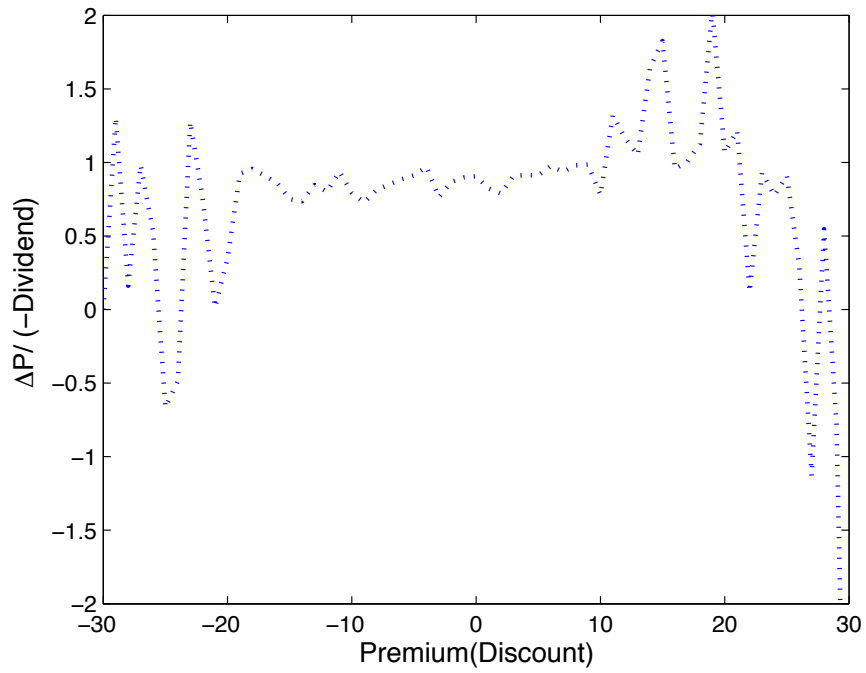


Figure 6. The ratio of the unexpected change in monthly share price (net of the expected change due to mean reversion) to the change in monthly NAV, $\Delta P^{unexp} / \Delta NAV$, plotted separately across various discount and premium levels for cases where $\Delta NAV > 0$ and $\Delta NAV < 0$

