Noise-trader Risk: Does it Deter Arbitrage, and Is it Priced?

Sean Masaki Flynn*

September 12, 2005

ABSTRACT

Arbitrage positions that benefit from the reversion of closed-end fund discounts to rational levels show excess returns that increase in magnitude the more funds are mispriced. At the same time, fund trading volumes and bid-ask spreads more than double as funds become increasingly mispriced. These behaviors suggest that non-diversifiable noise-trader risk increases the more funds are mispriced and that market participants are not only aware of this unique risk factor but demand a compensatory rate of return that varies with its magnitude.

JEL codes: G12, G14

Keywords: Closed-end Funds, Noise-trader Risk, Arbitrage

*Department of Economics, Vassar College, 124 Raymond Ave. #424, Poughkeepsie, NY 12604. E-mail: flynn@vassar.edu. Telephone: (845) 437-5209. I would like to thank Osaka University's Institute for Social and Economic Research for support while I completed the final drafts of this paper. I also gratefully acknowledge the diligent and tireless research assistance of Rebecca Forster. All errors are my own.

1. Introduction

As posited by Black (1986), DeLong, Shleifer, Summers, and Waldmann (1990), and Campbell and Kyle (1993), noise-trader risk is the risk faced by rational traders that irrational "noise traders" will cause asset prices to move unpredictably and without reference to information flows. These authors suggest that the presence of noise traders will tend to deter rational arbitrageurs from equalizing asset prices with fundamental values. The problem is that once the random trading activities of noise traders have driven asset prices away from fundamental valuations, rational arbitrageurs will be reluctant to rectify such mispricings because the ongoing, unpredictable trading activity of the noise traders leaves open the possibility that the existing mispricings may widen rather than narrow.

A defining feature of noise-trader risk is that it is an independent risk factor, uncorrelated with either the systematic risk that drives the capital asset pricing model of Sharpe (1964) and Lintner (1965), or the firm size and growth factors that also affect stock returns, as shown by Fama and French (1992). But identifying the independent effect of this risk factor is difficult when studying the returns of typical stock issues because the noise-trader volatility that would deter arbitrage would tend to look much like the noisiness caused by a wide variety of other factors. Even more problematic is the fact that because of reasonable differences in interpreting fundamentals and making forecasts, it is not normally possible to make a convincing case that the typical operating company's stock is at any particular moment mispriced. Since you cannot demonstrate mispricing, it is hard to make an empirical case that noise-trader risk is deterring the arbitrage necessary to rectify mispricings.

Closed-end funds can be used to solve both of these problems. First, because their fundamental values are easily calculated, you can convincingly demonstrate that their market prices often differ from their fundamental values. That of course opens up the possibility that noise-trader risk is deterring the arbitrage activities what would be necessary to equalize their prices with their fundamentals. Second, it is possible to isolate the effects of noise-trader risk because it appears to vary systematically with the degree to which closed-end funds are mispriced.

The reason it is possible to demonstrate that closed-end funds are mispriced is because their fundamental values can be precisely calculated at frequent intervals. Closed-end funds are simply mutual funds that do not redeem their own shares. Instead, their shares trade on large stock exchanges where their prices are determined by supply and demand. At the same time, they voluntarily release on a weekly basis their net asset values (NAVs), which are simply the values of their assets minus their liabilities. As a result, one would expect that an informationally efficient market would set the market value of each fund equal to its NAV, or perhaps to its NAV less the discounted value of expected future management fees.

But such is not the case. Closed-end funds routinely trade at prices that differ significantly from their NAVs or from valuations consistent with discounting out future fund expenses. A natural result of this phenomenon has been a large literature devoted to understanding why funds trade at discounts and premia relative to their portfolio values, as well as whether arbitrage pressures are strong enough to keep closed-end fund share prices at least reasonably linked to portfolio values, management fees, and the standard risk factors (see Dimson and Minio-Kozerski 1999b for an excellent survey.)

But perhaps the strangest thing about the mispricings found in closed-end funds is that they happen despite the absence of any obvious arbitrage barriers and in the presence of very active arbitrage. For instance, large and very liquid funds often trade at large discounts for months or years at a time despite the fact that arbitrage would only involve purchasing fund shares—something that is easily done. At the same time, there is direct evidence of very active arbitrage activities being undertaken against overpriced funds. Indeed, Flynn (2004) shows that there is very intense short selling of NYSE-traded closed-end funds selling at premia and that the intensity of short selling increases the more funds are overpriced relative to their net asset values.

Consequently, while it is clear that arbitrageurs do respond to fund mispricings, it is also clear that the amount of arbitrage in which they engage is not large enough to overcome the large and lingering mispricings routinely observed in closed-end funds. Since there are no obvious barriers to arbitrage—especially for funds trading at discounts, where arbitrage would merely involve buying fund shares—it appears that arbitrageurs must be voluntarily limiting the amount of arbitrage capital that they invest in closed-end funds. This paper makes the case that their self-restraint is due to the deterrent effect of noise-trader risk.

The major evidence in favor of this conclusion comes from examining a cross-section of closed-end fund arbitrage portfolios. The cross section is based on fund discount and premium levels, so that twenty portfolios are created to examine the returns that an arbitrageur would get for arbitrage positions targeting various levels of under- or over-pricing. In addition, each of these mock portfolios goes long the shares of the fund while shorting the fund's underlying portfolio. This is crucial methodological innovation because it serves to isolate the returns that derive solely from the mean reversion of fund share prices towards net asset values. By doing so, we can see if such returns can be accounted for by the standard Fama and French (1992) risk factors. As it turns out, they cannot.

Rather, the portfolios produce a striking pattern of excess returns. Excess returns are near zero for portfolios containing funds trading at prices near rational levels, but grow rapidly in magnitude the more funds are mispriced. For instance, the portfolio that invests only in funds trading at discounts of between zero and five percent generates a statistically insignificant excess return of -0.15 percent per month while the portfolio that invests in funds trading at discounts of between twenty-five and thirty percent generates a highly statistically significant excess return of 2.12% per month.

This pattern suggests very clearly that if you account for only the standard risk factors, arbitrageurs have a strong incentive to engage in arbitrage against mispriced closed-end funds, and that the incentive increases the more funds are mispriced. What then is preventing them from capturing these excess returns? I believe that the answer is noise-trader risk. It deters arbitrageurs from investing the full amount of arbitrage capital that would be necessary to immediately rectify closed-end fund mispricings. But since it is not accounted for by the standard risk factors, there appear to be excess returns when in fact what is happening is that the apparent excess returns serve to compensate rational arbitrageurs for bearing noise-trader risk.

There are several pieces of evidence that support this contention. First, you can show that *fund-specific* noise-trader risk is the likely culprit because the excess returns remain even after adding first-differences of average discount and premium levels as an additional explanatory variable. Since these first differences capture the returns that result from any sort of discount volatility common to closed-end funds as a group, their inclusion makes it much more likely that any remaining volatility is the result of a fund-specific risk factor—and that the observed excess returns are compensation for being exposed to this fund-specific risk factor.

In addition, there is an inverse relationship between R-squared statistics and excess returns across the various arbitrage portfolios that also strongly suggests that fund-specific noise-trader risk (rather than something common to all funds) may be the hidden risk factor by which the excess returns may be explained. For the portfolios consisting of funds trading near rational levels, R-squared statistics are very high while excess returns are near zero. But for portfolios increasingly far from fundamental levels, you see R-squared statistics sharply decrease while excess returns increase. Since the regressions account for both the standard factors as well as movements in discount and premium levels common to all funds, a natural interpretation of this inverse relationship is that the excess returns to arbitrage grow as you move farther away from fundamental discount levels in order to provide compensation for a missing risk factor that also grows as you move farther away from fundamental discount levels.

I conclude the paper with two additional pieces of evidence consistent with fund-specific noise-trader risk increasing as you move farther away from fundamental discount levels. The first is that the average ratio of trading volume to shares outstanding is U-shaped if you plot it against discount and premium levels, with the ratio nearly doubling as you move away from fundamental discount levels to either large discounts or large premia. This pattern jibes with the theoretical models of Varian (1989), Kandel and Pearson (1995), and Harris and Raviv (1993), which link increased divergence of opinion among traders (due, I argue, to increased activity by noise traders) with higher trading volumes.

The second is that the average ratio of the bid-ask spread to the ask price, *(bid – ask)/ask*, is also U-shaped if you plot it against discount and premium levels. Since closed-end funds do not pose any obvious informational asymmetries (because portfolio values are public), it seems very likely that spreads widen as you move away from fundamental discount levels because specialists are reacting to increasingly high price volatility caused by increasingly active noise traders. This behavior on the part of specialists would be consistent with the models of Garbade and Silber (1979) and Ho and Stoll (1981), where increasing price volatility by itself causes risk-averse market makers to widen spreads even in the absence of informational differences between specialists and traders.

Section 2 describes the data and defines discounts and premia. Section 3 examines the cross section of closed-end fund arbitrage returns and how closed-end fund trading volume and bid-ask spreads vary with discount and premium levels. Section 4 concludes.

2. 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. This implies, of course, that the data set suffers from survival bias. However, because this paper is interested in the behavior of returns 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.[1] Of the 462 funds, three 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. It should be noted that this data set is much larger and more comprehensive than any other closed-end fund data set previously examined in the literature.

Below, I concentrate on the Fund Edge data covering 1985-2001 for two reasons. First, some of the time series of older funds are incomplete prior to 1985. Second, there was a huge increase in the number of funds starting in the late 1980s.[2] As a result, the overwhelming majority of the data lies in the post-1985 period anyway.

[1]Whatever *expectation* the market has about funds liquidating or converting to open-end formats should be incorporated into the returns of the remaining funds, so that excluding the funds that did actually liquidate or convert should not bias the results below.

[2]Fewer than 30 funds were listed in the Wall Street Journal in 1985.

Figure 1. Distribution of Weekly Discounts in Fund Edge, 1985-2001

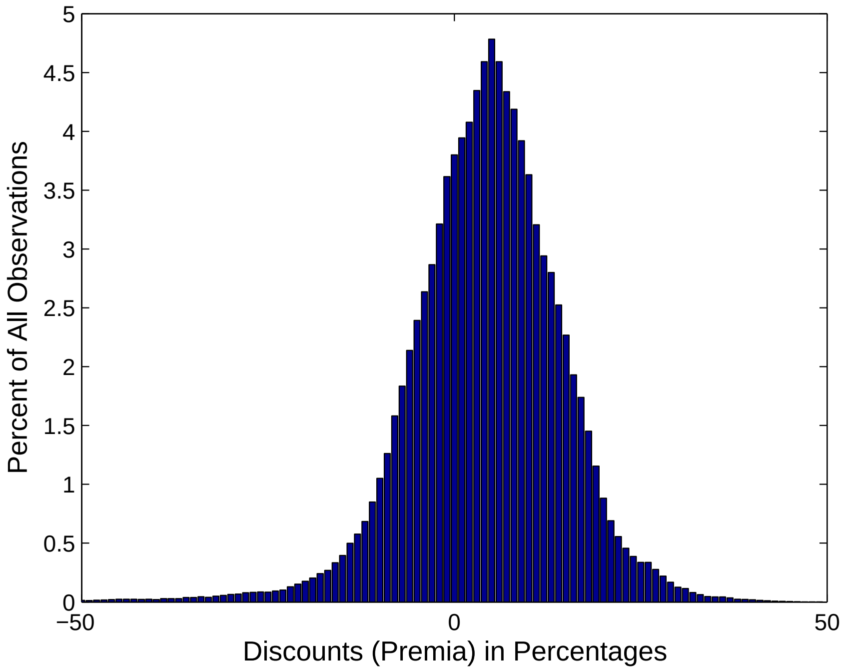

In addition to Fund Edge data, I use CRSP data on bid and ask prices to examine how closed-end fund bid-ask spreads vary with discount and premium levels. CRSP data is also used to examine how closed-end fund trading volume varies with discounts and premia.

In this paper, discounts are defined as positive 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 were the fund to liquidate immediately. Let P_t be the fund's price per share at time t. The discount or premium at which a fund trades at time t is defined as $D_t = N_t/P_t - 1$. Values of $D_t > 0$ are called discounts, while values of $D_t < 0$ are referred to as premia. In this paper, I will multiply D_t by 100 and refer to discounts and premia in percentages.

You can get a good summary of the discount and premia data by examining Figure 1, which plots a relative frequency histogram of the 227,066 weekly D_t observations on the 462 funds found in Fund Edge over the period January 1985 through May 2001. The figure excludes 761 outliers, most of which are in the left (premium) tail since the distribution is skewed towards premia. (The most extreme discount was 66.5% while the most extreme premium was -205.4%.)

The data is bell shaped and centered on a discount of about 6%, consistent with Ross (2002) and Flynn (2002), who argue that funds should rationally trade at discounts of about 7% in order to capitalize out future fund expenses. This means that, on average, fund D_t values are consistent with the markets acting to rationally equate share prices with fundamentals.

What this paper addresses, however, is the spread of the distribution and in particular why D_t values of individual funds often deviate wildly from rational levels. If arbitrageurs can get prices right on average, why can't they get them right more quickly? If they could, deviations from rational prices would be short lived and the distribution of D_t values in Figure 1 wouldn't show such a wide spread.

3. Closed-end Fund Arbitrage Returns

Black (1986) divides market participants into two groups, information traders and noise traders. Information traders make trades based on the latest information about future asset returns. By contrast, noise traders react to noise—signals that contain no new information relevant to assessing future asset returns. For instance, noise traders may falsely extrapolate past trends. They may also be subject to hunches or may incorrectly believe that they have special information as in the model of noise-trader risk applied to closed-end funds by De-Long, Shleifer, Summers, and Waldmann (1990). However, the particular reason that causes them to engage in non-information based trading is not important. The only thing that matters for asset pricing is that their trading behavior be unpredictable.

That's because if their actions are unpredictable, then they will present a unique and independent source of risk to rational information traders, as in Campbell and Kyle (1993). In particular, if the price volatility caused by the noise traders is not correlated with the returns on other assets, it will be unhedgeable. As such, this risk (which I refer to as noise-trader risk) will affect asset prices by tending to discourage rational traders from correcting mispricings: while mispricings offer obvious profit opportunities, they will not be riskless profit opportunities because noise traders may cause mispricings to widen rather than narrow.

Since modern finance theory is based on the idea that all returns in excess of the risk-free rate must be compensation for some form of risk or another, noise-trader risk also has implications for regressions run to test for excess returns. Consider an asset that is affected by noise-trader risk and in which a portion of overall returns is, in fact, a compensation for noise-trader risk. If variables that can capture the effect of noise-trader risk on returns are left out of factor regressions, then you may be misled into believing that there are excess returns when all that is really happening is that you have missed a risk factor.

Now, of course, any regression is apt to turn up excess returns—and it would be foolish to immediately ascribe such excess returns to noise-trader risk. That being said, my strategy in what follows is to cross section the closed-end funds in the data set into portfolios and then look at the excess returns to those portfolios to see if they have any pattern that can reasonably be ascribed to differences in exposure to noise-trader risk. I find that such a pattern emerges if you simply sort funds by their discount and premium levels before running factor regressions.

But before presenting those results on risk-adjusted returns, I will first show you the non-risk-adjusted returns available to arbitrageurs in closed-end funds. Seeing them helps to put the excess returns after risk adjustment into better perspective.

3.1. Non-Risk-Adjusted Arbitrage Returns

Arbitrageurs hoping to benefit from closed-end fund mispricings will set up positions that will benefit if fund prices move toward fundamentals. For funds trading at large discounts, they will go long the shares of the fund and short its underlying, as any reduction in the discount would generate a profit. And for funds trading at large premia, they will go long the underlying and short the fund's shares, as any reduction in premia would generate a profit. By hedging long positions in fund shares with short positions in fund portfolios (or vice versa), arbitrageurs can isolate the returns that come solely from the mean reversion of discounts and premiums, that is from the dissipation of mispricings over time.[3] In what follows, however, I will for simplicity only look at the returns to going long the shares of a fund while shorting its underlying. The returns to the opposite position (shorting the fund while going long the underlying) are of course just the negative of these returns.

The returns are driven by the tendency of discounts and premia to mean revert. It turns out that they revert to the mode of the discount and premium distribution shown in Figure 1 and that they do so following an AR(1) process.[4] This is convenient because, by estimating the AR(1) process, you can get estimates not only for the level to which discounts and premia mean revert, but of the pace of mean reversion as well.

[3]You can see this by rearranging the identity that defines D_t based on NAV and price: $N_t = (1+D_t)P_t$, for all t. Using this definition to construct log returns, you can see that $\log(\frac{P_{t+1}}{P_t}) - \log(\frac{N_{t+1}}{N_t}) = -\log(\frac{1+D_{t+1}}{1+D_t})$. This means that the rate of return to going long the shares of the fund while shorting the underlying is equal to the negative of the rate of change of $1+D_t$.

[4]I tested many processes. AR(1) is best, by far.

To see how this works, assume that the discount or premium, D_t, is mean reverting to the level $\bar{D}$. Under an AR(1) process, $D_{t+1} = \bar{D} + \phi(D_t - \bar{D}) + \varepsilon_t$, where ϕ gives the fraction of the month t deviation that remains the next month and ε_t is a Gaussian shock. We can get empirical estimates for ϕ and $\bar{D}$ by re-arranging the algebra on the right-hand side of the equation and running a regression on $D_{t+1} = constant + \phi D_t + \varepsilon_t$. Our estimated constant will be equal to $(1-\phi)\bar{D}$, which will allow us to back out the value of $\bar{D}$ after estimating the equation.

Using monthly discount and premium data for the 462 closed-end funds in the data set, and estimating the equation using pooled least squares on data covering 1985-2001 gives a constant of 0.44 and an estimated value for ϕ of 0.916. The regression has an R^2 statistic of 0.84 and the t-statistics on the constant and ϕ are, respectively, 12.0 and 162.7. Using these estimates, we can back out an estimated value for $\bar{D}$ of 5.2%. That is, this regression methodology indicates that discounts and premia revert to the same level as the mode discount of Figure 1.

In addition, the estimate for ϕ indicates that, on average, 91.6% of the deviation between a current D_t value and the long-run mean-reverting value of 5.2% will remain the following month—meaning that about 8.4% of any such gap will be closed, on average. Consequently, the magnitude of the return available to arbitrageurs in any specific case depends on how far away from the mean-reverting discount level the current discount or premium is—reverting 8.4% of a big gap implies a much bigger return than does reverting 8.4% of a small gap.

To see how this rate of mean reversion affects returns, it is fruitful to plot the monthly returns of both fund shares and fund NAVs against discount and premium levels. I do this in Figure 2 by aggregating all of my monthly data and sorting observations by discount or premium without regard to each observation's date. Specifically, I take each of the 52,188 monthly discount or premium observations running from January 1985 to April 2001 and

Figure 2. Average one-month share and NAV returns when sorting each observation by its discount or premium level the previous month.

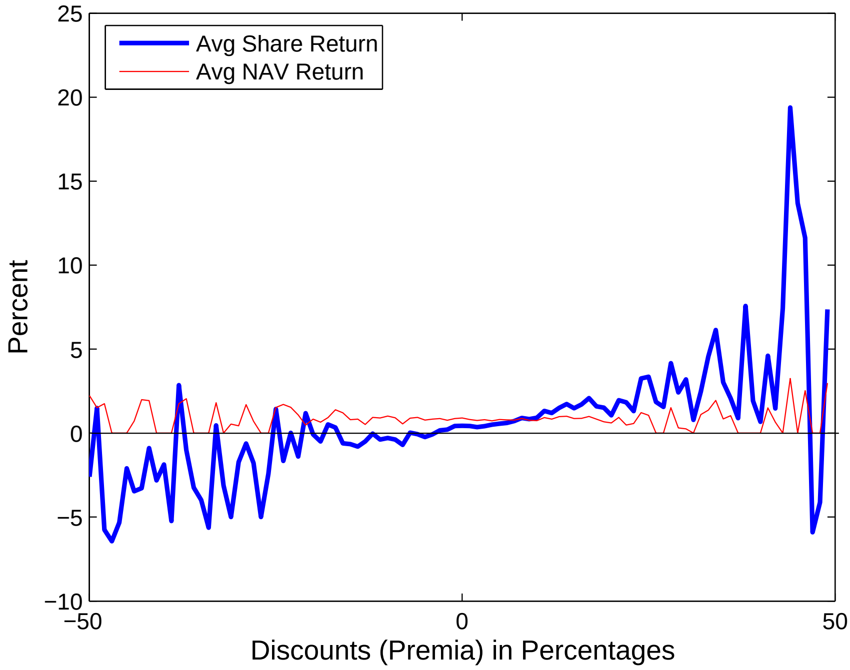

place them into one-percent wide bins running from a premium of -50% to a discount of 50%. For all of the observations in a given bin, I then separately calculate the return over the next month to holding a long position in the associated fund's shares as well as the return over the next month to holding a long position in the fund's underlying NAV, being sure to properly account for dividend payments. I call the former share returns and the later NAV returns and plot their respective bin averages in Figure 2.

As you can see, NAV returns are a basically a horizontal line at about one percent per month, meaning that they are unrelated to discount and premium levels.[5] On the other hand,

[5]This is consistent with Malkiel (1977), Pontiff (1995), and Dimson and Minio-Kozerski (1999) who find that discount and premium levels are unrelated to future NAV performance. However, Fund Edge contains mostly bond funds (since they are much more numerous than stock funds) and this must be noted because Chay and Trzcinka (1999) found that discount and premium levels were significantly related to future NAV performance for stock funds, but not for bond funds.

share returns are positively related to D_t levels. This is because of mean reversion. For instance, as fund prices move back towards the center of the distribution in Figure 1, share returns will exceed NAV returns for funds trading at $D_t > \bar{D}$ because of capital gains caused by the mean reversion of prices upward toward NAVs. Contrariwise, capital losses will accrue to funds trading at $D_t < \bar{D}$ as mean reversion causes prices to fall downward toward NAVs.

The difference between share returns and NAV returns is important because it serves to attract arbitrageurs to closed-end funds (and, under this paper's working hypothesis, to compensate them for noise-trader risk.) The magnitude of that attraction is best seen by looking at the how the difference between share and NAV returns is distributed across D_t levels. I do this by again placing each of the 52,188 monthly discount and premium observations into one-percent wide bins running from a premium of -50% to a discount of 50%. But this time, instead of averaging the share and NAV returns separately, I take their difference for all observations in each bin and then calculate each bin's average and standard deviation of those differences. These are plotted separately in Figure 3.

The average difference between share price and NAV returns by bin (the thicker line in Figure 3) is of course equal to the vertical difference between the share return and NAV return lines in Figure 2. It is upward sloping and linear in the center of the figure where there are a substantial number of observations in each bin, reinforcing the fact that arbitrage returns clearly depend on discount and premium levels. Its high volatility at both ends of the graph is due to there being very few observations in the outlying bins. For instance, the big spikes on the right side of the figure are for bins containing fewer than six and sometimes just one observation. The mode bin, by contrast, contains 2,308 observations.

Very importantly, the standard deviation increases the further you move away from the mean-reverting discount level. This means that while the returns to arbitrage positions increase the further you move away from the mean reverting level, so does the amount of risk associated

Figure 3. Average one-month share less NAV returns and standard deviations when sorting each observation by its discount or premium level the previous month.

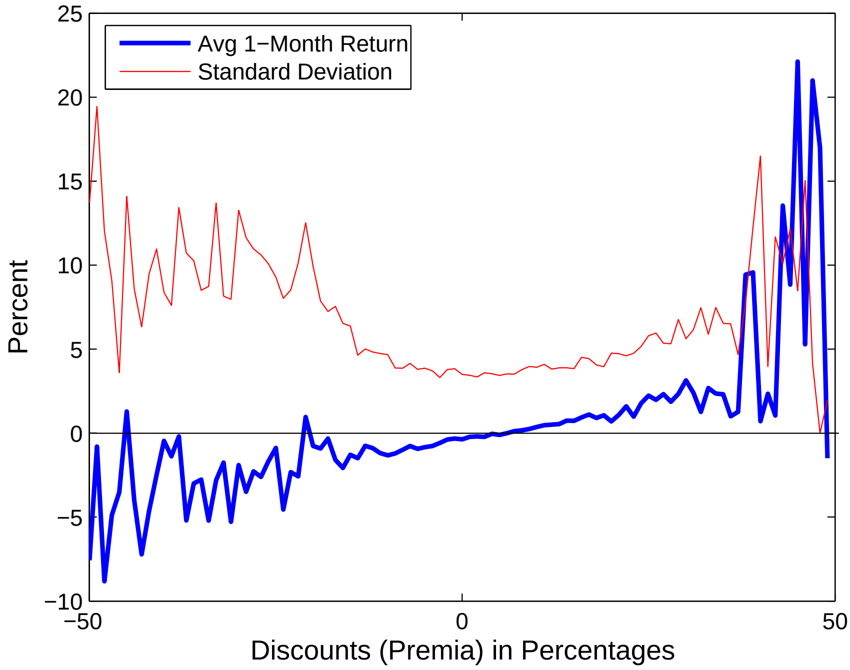

with such positions. This is, of course, consistent with the intuitions of modern finance theory, and the question that will be explored in the next section is whether the risk associated with arbitrage in closed-end funds can be explained by standard risk factors.

But before addressing that issue, it is interesting to plot out reward-to-risk ratios against discount and premium levels. I do this in Figure 4 by dividing each bin's average share less NAV return by the bin's standard deviation of those returns—that is, I divide the values given by the thick line in Figure 3 by those given by the thin line in Figure 3.

The result in Figure 4 is very striking. It shows that for all the bins in the center of the discount and premium distribution, there is a clear linear relationship between discount and premium levels and the reward-to-risk ratio. This is true all the way from a -10% premium

Figure 4. Ratio of average one-month share less NAV returns to the standard deviation of those returns (i.e., the reward-to-risk ratio.)

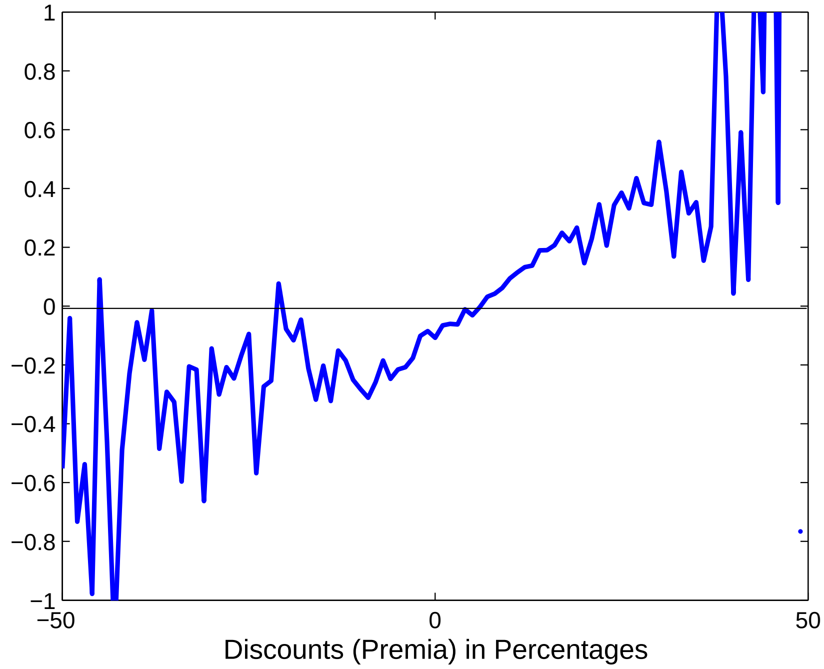

to a 30% discount and it suggests that closed-end fund markets have some systematic method of pricing the risk associated with holding closed-end fund arbitrage positions. In addition, it suggests quite clearly that there are discount and premium levels for which arbitrage many not be worth undertaking. In particular, if the volatility inherent in the reward-to-risk ratios plotted in Figure 4 cannot be accounted for by the standard risk factors and is indeed a unique form of risk, then the reward-to-risk ratios near the center of the discount and premium distribution are too low to warrant arbitrage activity. Only as you move into the tails of the distribution will reward-to-risk ratios be high enough to justify investing substantial capital into arbitrage activities. Consequently, Figure 4 tells a very clear story about how mispricings can persist in closed-end funds given the deterrent effects of noise-trader risk.

3.2. The Cross Section of Risk-Adjusted Arbitrage Returns

In this section, I examine the risk-adjusted profitability of arbitrage in closed-end funds and the effects of noise-trader risk on arbitrage returns by cross sectioning closed-end fund arbitrage returns by discount and premium levels. I do this by constructing twenty portfolios set up to isolate the profits that can be achieved by arbitraging against the mean reversion of discounts and premia.These portfolios are constructed using monthly data ranging from July 1985 through May 2001 and each portfolio corresponds to a five percentage-point wide discount or premium bin.[6] The first bin corresponds to premia falling between -50% and -45%. The twentieth bin corresponds to discounts falling between 45% and 50%. For each month, funds were placed into the bins based upon their D_t levels that month. For instance, all of the funds with discounts between 10% and 15% in a given month were placed into the thirteenth bin.

For each fund in a given bin, the returns over the following month to its spot share price and to its NAV were calculated. These returns take account of dividend payments and are denoted r^{share} and r^{NAV}, respectively.[7] Excess returns over the following month to the arbitrage portfolio that goes long fund shares and short the underlying, $r^{share} - r^{NAV}$, were then averaged with those of all other funds in the bin to get that bin's monthly portfolio return. By doing this for all funds in each bin each month, I generated twenty time series that form a cross section of the returns available to arbitragers of US and Canadian funds over the period 1985 to 2001.

Please note that this methodology of going long the fund while shorting the NAV in order to isolate the returns to arbitrage in closed-end funds is an innovation. Previous authors be-

[6]January through June 1985 can be included without affecting the results. I left them out because I eliminated the first six months of data on all funds (including, as overkill, those in existence in January 1985) in order to avoid any of the predictable mean reversion that happens after funds have their IPOs. Weiss (1989) and Peavy (1990) both find that closed-end funds begin trading immediately after their IPOs at premia of about eight to ten percent in order to pay investment bankers. But these typically turn to discounts within a few months after an IPO.

[7]The NAV returns also properly account for the reductions in NAV caused by fund expenses and management fees. Thus, you do not have to worry about NAV returns being overstated as they would be if such costs were not taken into account.

ginning with Thompson (1978) and including Anderson (1986) and Pontiff (1995) run CAMP and Fama-French three-factor regressions on long positions in closed-end funds. These authors find excess returns but because the returns to long positions are affected by both changes in fund NAVs as well as changes in fund D_t levels relative to NAVs, it is not possible to isolate the returns available to arbitrageurs attempting to benefit from the reversion of mispriced funds towards rational valuations. By going long fund shares while going short their NAVs, the portfolios examined here make such an isolation possible. This is important because isolation allows you to properly examine whether the standard risk factors can explain the returns to arbitraging mispriced funds.

The twenty time series that isolate arbitrage returns are used as the independent variables in the Fama-French (1992) regressions shown in Table 1. Each time series was regressed on a constant, α, and the three Fama and French (1992) factors: $r^m - r^f$ is the market excess return, SMB is the return to small capitalization stocks less the return to large capitalization stocks, and HML is the return to value stocks less the return to growth stocks. All variables are given in percents.

Fama and French (1992) famously argue that their three factors can explain returns to long positions in stocks. But as you can see from the cross-section of closed-end fund arbitrage returns in Table 1, these three factors do a poor job of explaining the returns available to arbitrageurs in closed-end funds. This is obvious by looking at the third column of the table, which gives the excess returns to each portfolio. With the exception of the zero to five percent discount bin and a couple of the most extreme bins, the excess returns are all robustly significant, indicating that the standard risk factors are not accounting for all of the risk facing closed-end fund arbitrage positions. This, of course, suggests that there may be a missing risk factor.

Table 1
Fama French Regressions on Closed-end Fund Arbitrage Portfolios Defined by Discount Levels

This table reports the results of Fama and French (1992) three-factor regressions performed on 20 arbitrage portfolios whose returns depend entirely on changes in discount and premium levels. The regressions are run on monthly data covering July 1985 to May 2001. The 20 portfolios are defined by discount and premium levels, so that each month each fund is placed into one of 20 five-percentage point wide bins, the bins ranging from a premium of -50% to a discount of 50%. For each fund in a given bin, I take the difference between share returns and NAV returns, $r^{share} - r^{NAV}$, since this difference isolates the return that derives solely from changes in discount and premium levels. I then average all the differences in each bin each month. Doing so provides twenty time series that together generate a cross section of the returns to arbitrage in closed-end funds. These returns are then regressed in the normal way on a constant, α, and the three Fama and French (1992) factors: $r^m - r^f$ is the market excess return, *SMB* is the return to small capitalization stocks less the return to large capitalization stocks, and *HML* is the return to value stocks less the return to growth stocks. Standard errors are calculated using the method of Newey and West (1987) and t-statistics are given in parentheses. These are OLS regressions. All variables are defined in percents.

Bin Lower Bound	Bin Upper Bound	α	$r^m - r^f$	SMB	HML	R-sq	Average Dependant Variable	Std. Dev. Dependant Variable	Obs.
-50	-45	-3.36	0.69	1.11	0.77	0.25	-3.59	8.26	17
		(-1.88)	(1.79)	(1.05)	(0.69)				
-45	-40	-5.58	0.61	-0.15	-0.11	0.11	-4.95	8.09	42
		(-3.75)	(1.47)	(-0.26)	(-0.23)				
-40	-35	-2.57	-0.15	0.29	0.34	0.03	-2.59	8.92	45
		(-1.92)	(-0.55)	(0.42)	(0.71)				
-35	-30	-4.19	-0.08	0.18	0.28	0.02	-4.06	8.08	65
		(-4.19)	(-0.45)	(0.64)	(1.21)				
-30	-25	-3.41	0.25	0.81	0.60	0.14	-2.99	7.90	77
		(-3.68)	(1.27)	(2.13)	(1.67)				
-25	-20	-1.88	0.44	0.20	0.17	0.05	-1.51	8.01	100
		(-2.39)	(2.34)	(0.94)	(0.72)				
-20	-15	-1.29	0.43	0.19	0.36	0.11	-0.93	5.67	138
		(-3.69)	(2.86)	(1.79)	(2.58)				
-15	-10	-1.19	0.13	0.01	0.14	0.06	-1.10	2.53	170
		(-6.13)	(3.30)	(0.16)	(2.48)				
-10	-5	-0.92	0.13	0.07	0.23	0.18	-0.82	1.88	178
		(-6.62)	(4.13)	(1.39)	(3.80)				
-5	0	-0.43	0.07	0.10	0.18	0.13	-0.39	1.60	179
		(-4.13)	(2.50)	(2.18)	(3.83)				
0	5	-0.15	0.10	0.08	0.17	0.16	-0.08	1.51	179
		(-1.54)	(3.51)	(1.65)	(3.70)				
5	10	0.27	0.06	0.11	0.14	0.08	0.32	1.70	179
		(2.35)	(1.92)	(2.34)	(3.18)				
10	15	0.65	0.08	0.13	0.15	0.09	0.72	1.98	175
		(4.40)	(2.41)	(3.13)	(3.86)				
15	20	0.84	0.13	0.15	0.13	0.08	0.93	2.65	164
		(3.96)	(2.17)	(2.58)	(2.73)				
20	25	1.83	0.20	0.04	0.14	0.06	1.98	3.54	147
		(6.23)	(1.73)	(0.42)	(1.98)				
25	30	2.12	0.49	0.10	0.21	0.11	2.60	5.92	99
		(4.07)	(3.02)	(0.71)	(2.52)				
30	35	2.83	0.57	-0.10	0.10	0.11	3.47	7.18	56
		(3.05)	(3.02)	(-0.47)	(0.94)				
35	40	1.32	0.87	0.34	0.38	0.42	1.56	4.72	23
		(1.61)	(4.95)	(1.11)	(1.60)				
40	45	8.23	0.91	-0.98	0.04	0.43	6.28	6.88	9
		(2.55)	(2.11)	(-1.57)	(0.62)				
45	50	Insufficient	Observations						

What is more, the excess returns vary radically from one end of the cross section to the other. They are large, negative, and very significant for the extreme premium bins at the top of the column, fall towards zero as you move near the mean-reverting discount level of six percent, and then grow into large, positive, and highly significant values as you continue moving down the table into bins containing larger and larger discounts. This pattern strongly suggests that the portfolios are exposed to varying levels of the missing risk factor. In particular, they are consistent with larger exposure for portfolios farther away from the mean-reverting discount level.

Belief in this possibility is reinforced by looking at the numbers in the ninth column, which give the standard deviations over time of each portfolio's dependent variable, $r^{share} - r^{NAV}$. These standard deviations are smallest for the portfolios near the mean-reverting discount level, and grow much larger as you move to either deep discounts or large premia. Since the dependent variables are returns to arbitrage portfolios, this is clear evidence that the level of risk associated with attempting to profit from the mean reversion of discounts and premia increases the farther you move away from the mean-reverting discount level.[8] difference between share and NAV returns The Fama-French factors fail to account for this risk pattern across portfolios.

Finally, the pattern found in the R-squared statistics as you look across the twenty portfolios also suggests that the Fama-French factors explain very little of the time variation within any given portfolio. With the exception of a few portfolios that have very few observations because they are in the tails of the discount and premium distribution, R-squared statistics are

[8] In an earlier paper entitled,"Equilibrium Arbitrage in the Presence of Noise-trader Risk: Evidence from Closed-end Funds" I made the error of subtracting off the risk-free interest rate from the dependent variable in the regressions reported in Table 1. That is, I used $r^{share} - r^{NAV} - r^f$ as the dependent variable, where r^r is the risk free interest rate. Subtracting the risk-free rate incorrectly introduced a bias in the estimated alphas, making it incorrectly appear that no excess returns were available on closed-end fund arbitrage positions for most of the bins. Subtracting off the risk-free rate, though, is incorrect since these long/short arbitrage positions are self-financing. The correct alphas, reported in Table 1, are in fact significantly different from zero for nearly all bins.

modest. The highest R-squared statistics of 0.18 and 0.16 happen for funds near the long-run mean-reverting discount level. As you move away from this level in either direction, the R-squared statistics fall (with the exception already noted of the extreme bins with very few observations). This pattern suggests that the modest explanatory power possessed by the Fama-French factors is concentrated near the center of the discount and premium distribution. As you move into the tails of the distribution, the three Fama-French factors do an increasingly poor job of explaining the time variation of the returns within each of the various portfolios.

3.2.1. Accounting for Volatility that Affects all Funds Simultaneously

Lee, Shleifer, and Thaler (1991) find that discount and premium levels across US closed-end funds from 1965 to 1985 are highly correlated. This is also true for the Fund Edge data from 1985 to 2001 that is used here. Consequently, it is important to see if the cross-sectional pattern of excess returns found in Table 1 is robust to accounting for changes in fund discount and premium levels that affect all funds simultaneously.

Table 2 regresses a constant and the monthly change in the capital-weighted average discount level across all funds on the twenty arbitrage portfolio return series used in Table 1. More specifically, the capital-weighted average discount, W_t, was calculated for all funds in existence each month by combining Fund Edge data on D_t values and NAVs with CRSP data on shares outstanding. The independent variable in each regression in Table 2 is simply this variable's first difference, $\Delta W_t = W_t - W_{t-1}$. This lag structure matches the timing used to sort funds into portfolios. The return on each arbitrage portfolio during month t is the result of sorting funds into the twenty portfolios based on their discounts at the end of month $t-1$. In a similar way, $\Delta W_t = W_t - W_{t-1}$ gives the change in the average discount level from the end of month $t-1$ to the end of month t. Hence, by running the regression $r^{share} - r^{NAV} = constant + \Delta W_t + \varepsilon_t$ for each portfolio, you can see how the common factor that causes discounts and premiums to be correlated affects the returns to arbitrage positions.

Table 2
Regressions of Changes in the Monthly Capital Weighted Average Discount or Premium Level on the Returns to Closed-end Fund Arbitrage Portfolios Defined by Discount Levels

This table reports the results of regressing a constant and the change in the capital-weighted average discount level across all funds each month on the twenty closed-end fund arbitrage portfolios used in Table 1. That is, if W_t is the capital-weighted average discount across all funds in a given month t, then the independent variable in this regression is, $\Delta W_t = W_t - W_{t-1}$. This timing convention for ΔW_t is consistent with that used for the dependent variable since the return on the arbitrage portfolios during month t, $r_t^{share} - r_t^{NAV}$, is the result of sorting the funds into the twenty portfolios based on their discounts at the end of month $t-1$. Hence, for each portfolio, the regression is $r_t^{share} - r_t^{NAV} = constant + \Delta W_t + \varepsilon_t$, where ε_t is the error term. Standard errors are calculated using the method of Newey and West (1987) and t-statistics are given in parentheses. These are OLS regressions. All variables are defined in percents.

Bin Lower Bound	Bin Upper Bound	Constant	ΔW_t	R-sq	Average Dependant Dep. Variable	Std. Dev. Dependant Variable	Obs.
-50	-45	-2.68	1.99	0.12	-3.93	8.41	16
		(-1.62)	(1.29)				
-45	-40	-4.49	1.51	0.04	-4.95	8.09	42
		(-3.66)	(0.97)				
-40	-35	-2.73	-0.09	0.00	-2.74	8.97	44
		(-1.85)	(-0.09)				
-35	-30	-4.03	0.32	0.00	-4.06	8.08	65
		(-4.10)	(0.51)				
-30	-25	-2.89	1.92	0.12	-3.18	7.78	76
		(-2.70)	(1.84)				
-25	-20	-1.32	1.36	0.05	-1.51	8.01	100
		(-1.57)	(2.53)				
-20	-15	-0.88	0.33	0.01	-0.93	5.67	138
		(-2.20)	(0.54)				
-15	-10	-1.02	0.95	0.28	-1.10	2.53	170
		(-5.29)	(5.09)				
-10	-5	-0.77	0.79	0.34	-0.82	1.89	177
		(-5.15)	(8.30)				
-5	0	-0.33	0.89	0.59	-0.39	1.60	178
		(-4.12)	(10.88)				
0	5	-0.02	0.83	0.57	-0.08	1.52	178
		(-0.32)	(10.60)				
5	10	0.36	0.92	0.56	0.30	1.70	178
		(3.91)	(11.29)				
10	15	0.77	0.80	0.32	0.74	1.97	174
		(5.45)	(6.75)				
15	20	0.95	0.93	0.23	0.93	2.65	164
		(5.07)	(6.92)				
20	25	1.94	0.59	0.05	1.98	3.54	147
		(7.48)	(2.23)				
25	30	2.55	0.98	0.05	2.60	5.92	99
		(4.51)	(1.66)				
30	35	3.20	1.07	0.05	3.47	7.18	56
		(3.48)	(1.96)				
35	40	1.51	0.35	0.02	1.56	4.72	23
		(1.54)	(0.69)				
40	45	6.10	0.17	0.00	6.28	6.88	9
		(1.76)	(0.14)				
45	50	Insufficient	Observations				

The results of running these regressions are very striking. First of all, look at the *R*-squared statistics in the fifth column. They are much higher than those generated by the Fama-French factors in Table 1. In particular, the *R*-squared statistic for the zero percent to five percent discount bin (row 11 of the table) is an extremely high 0.59, much higher than the 0.16 achieved when regressing the Fama-French factors on that portfolio's return series.

But the pattern of *R*-squared statistics in cross-section is even more interesting. As you move away from the mean-reverting discount level in either direction, the explanatory power of changes in the cross-sectional average discount level falls to zero very quickly and stays near zero except for the first two bins which have very few observations. This means that while changes in discount levels have great explanatory power for funds trading near the center of the distribution in Figure 1, they have no explanatory power for funds trading in the tails of the distribution. This suggests that the returns to funds in the tails of the distribution are subject to fund-specific risk factors (i.e. fund-specific noise-trader risk) that are independent of whatever common risk factors cause closed-end fund discounts and premia to move in unison.

This is an important point because I would like to distinguish clearly between fund-specific risk factors and the argument about small investor sentiment made by Lee, Shleifer, and Thaler (1991). Under their hypothesis, average discount and premium levels across funds are correlated because they are a measure of small-investor sentiment. They suggest that larger discounts happen when sentiment turns negative and people are willing to pay less for closed-end fund shares. Similarly, they argue that smaller discounts (and premiums) reflect positive sentiment which bids up fund share prices relative to NAVs.[9] In addition, they argue that these movements in small investor sentiment may be a priced risk factor for assets in general.

[9] This hypothesis requires the assumption of segmented markets. The authors assume that small (irrational) investors dominate the closed-end fund markets while large (rational) investors dominate the markets for the securities in which closed-end fund invest. As a result, when small investor sentiment changes, it changes fund share prices without affecting fund NAVs. Consequently, discounts and premia vary with small investor sentiment.

But Elton, Gruber, and Busse (1998) and Doukas and Milonas (2004) have given strong evidence against small investor sentiment (as measured by average discounts and premia) being a priced risk factor in asset markets. Consequently, I want to make clear that the volatility found for the tails of the discount and premium distribution is independent of whatever it is that causes discounts and premia across funds to be correlated (be it small investor sentiment or something else.) Rejection of that common factor as a priced risk factor is a separate matter from whether or not the volatility found in the tails of the discount and premium distribution is itself priced.

Indeed, the results of Table 2 suggest that whatever it is that causes closed-end fund discounts and premia to be correlated, there appears to be something else affecting the movement of discounts and premia in the tails of the distribution. I suggest that this "something else" is fund-specific noise-trader risk. As you move farther away from the center of the discount and premium distribution, fund specific noise-trader risk increases, and as it does, funds in the tails of the distribution are less and less likely to move with any overall change in sentiment (or whatever it is that causes discount and premium levels near the center of the distribution to be correlated.) This would explain why the R-squared statistics in Table 2 fall so quickly as you move away from the center of the distribution.

Table 3 reports the results of using both the three Fama-French risk factors as well as changes in the capital-weighted average discount level as independent variables for the twenty arbitrage portfolios. As you can see, the cross-sectional pattern of excess returns is unchanged. This is important because by adding ΔW_t as an explanatory variable, we have accounted for movements in discounts and premia that are common to all funds.[10] What remains, therefore, must be portfolio specific.

Since modern finance theory rests on the idea that returns in excess of the risk-free rate must be a compensation for risk, the cross-sectional pattern of excess returns found in Table 3 can only be explained by appealing to a missing risk factor. Since ΔW_t will tend to account for any noise-trader risk that is common to all funds, I conjecture that the missing risk factor must be *fund-specific* noise-trader risk—the possibility that uninformed traders will drive a fund's share price further away from fundamentals than it already is.

Under this hypothesis, the observed excess returns are the compensation necessary to get rational investors to invest in a specific fund and bear the fund-specific noise-trader risk that its share price will move even farther away from fundamentals than it already has. The cross-sectional pattern of excess returns is then explained by noise-trader risk increasing as you move away from the center of the distribution in Figure 1. As noise-trader risk increases, so must compensation. That is why excess returns increase in Table 3 for portfolios increasingly far from the mean-reverting discount level.

But is it really proper to ascribe the increasing volatility of arbitrage portfolio returns as you move away from the center of the distribution in Figure 1 to increasing levels of noise-trader risk? The next subsection gives additional evidence in favor of this attribution by examining how trading volume and bid-ask spreads vary with discount and premium levels.

3.2.2. Spreads, Volume, and Noise-trader Risk

[10] I also experimented with using first differences of the monthly standard deviation of discounts and premia as well as first differences of the average absolute discount or premium to see if there was any tendency for the distribution to suddenly come together or expand—as if the rate of mean reversion sometimes sped up or slowed down. These and other variables had no explanatory power whatsoever which is why I only report results on ΔW_t.

Table 3
Regressions on Closed-end Fund Arbitrage Portfolios Defined by Discount Levels Including the Fama French Factors as well as a Factor to Account for Volatility Shared by All Funds

This table reports the results of expanding the Fama-French regression results of Table 1 by adding changes in the capital-weighted average discount across all funds as an additional explanatory variable. The notes to Table 1 explain the three Fama-French factors while the notes to Table 2 explain the capital-weighted average discount series, W_t. Standard errors are calculated using the method of Newey and West (1987) and t-statistics are given in parentheses. These are OLS regressions. All variables are defined in percents.

Bin Lower Bound	Bin Upper Bound	α	$r^m - r^f$	SMB	HML	ΔW_t	R-sq	Average Dependent Variable	Std. Dev. Dependent Variable	Obs.
-50	-45	-3.53	0.68	0.74	0.75	0.53	0.26	-3.93	8.41	16
		(-1.21)	(1.27)	(0.57)	(0.66)	(0.24)				
-45	-40	-5.19	0.58	-0.20	-0.05	1.08	0.13	-4.95	8.09	42
		(-3.73)	(1.41)	(-0.35)	(-0.1)	(0.72)				
-40	-35	-2.76	-0.13	0.26	0.38	-0.16	0.03	-2.74	8.97	44
		(-2.10)	(-0.48)	(0.31)	(0.78)	(-0.15)				
-35	-30	-4.17	-0.09	0.18	0.26	0.09	0.02	-4.06	8.08	65
		(-3.92)	(-0.37)	(0.60)	(0.70)	(0.10)				
-30	-25	-3.22	0.11	0.65	0.35	1.39	0.19	-3.18	7.78	76
		(-2.94)	(0.57)	(1.79)	(1.01)	(1.53)				
-25	-20	-1.63	0.34	0.11	0.02	1.18	0.08	-1.51	8.01	100
		(-2.02)	(1.71)	(0.49)	(0.09)	(1.94)				
-20	-15	-1.31	0.44	0.20	0.37	-0.11	0.11	-0.93	5.67	138
		(-4.65)	(2.21)	(1.73)	(1.94)	(-0.14)				
-15	-10	-1.07	0.06	-0.06	0.01	0.93	0.29	-1.10	2.53	170
		(-5.73)	(1.82)	(-1.06)	(0.14)	(5.06)				
-10	-5	-0.85	0.08	0.01	0.14	0.67	0.40	-0.82	1.89	177
		(-6.27)	(2.69)	(0.28)	(2.71)	(6.74)				
-5	0	-0.34	0.01	0.02	0.06	0.84	0.61	-0.39	1.60	178
		(-4.51)	(0.52)	(1.02)	(.00)	(10.58)				
0	5	-0.07	0.05	0.01	0.06	0.77	0.60	-0.08	1.52	178
		(-0.88)	(1.73)	(0.57)	(2.31)	(9.58)				
5	10	0.36	-0.01	0.03	0.01	0.91	0.57	0.30	1.70	178
		(3.91)	(-0.16)	(1.31)	(0.65)	(10.26)				
10	15	0.75	0.02	0.08	0.04	0.75	0.34	0.74	1.97	174
		(5.12)	(0.82)	(2.51)	(1.53)	(5.35)				
15	20	0.91	0.05	0.07	0.01	0.88	0.25	0.93	2.65	164
		(4.73)	(1.29)	(1.32)	(0.18)	(6.77)				
20	25	1.83	0.16	0.02	0.08	0.45	0.09	1.98	3.54	147
		(6.54)	(1.53)	0.01	(1.17)	(2.03)				
25	30	2.12	0.44	0.01	0.09	0.86	0.15	2.60	5.92	99
		(4.04)	(2.41)	(0.04)	(0.90)	(1.55)				
30	35	2.64	0.47	-0.23	-0.05	1.08	0.14	3.47	7.18	56
		(2.82)	(2.74)	(-1.05)	(-0.38)	(1.78)				
35	40	1.30	0.86	0.32	0.37	0.17	0.43	1.56	4.72	23
		(1.57)	(4.54)	(1.03)	(1.48)	(0.38)				
40	45	8.17	0.90	-0.99	0.03	0.08	0.43	6.28	6.88	9
		(2.00)	(1.57)	(-1.23)	(0.12)	(0.03)				
45	50	Insufficient	Observations							

The cross-sectional pattern of excess returns found in Tables 1 and 3 suggests that a risk-factor is not being properly accounted for. In this section, I give further evidence in favor of the hypothesis that the risk factor in question is fund-specific noise-trader risk.

There are two pieces of evidence. The first is the fact that if you plot closed-end fund trading volume against discount and premium levels, you get a U-shaped curve centered on the mean-reverting discount level of six percent. The second is that you also get a U-shaped curve centered on the mean-reverting discount level when you plot bid-ask spreads against discount and premium levels.

In Figure 5, I plot the ratio of average monthly trading volume to shares outstanding against discount and premium levels. To construct the figure, I group monthly Fund Edge discount and premium observations for all funds available from January 1985 to May 2001 into one-percent wide discount and premium bins. For each fund in a given bin, I use CRSP data to calculate the fund's ratio of trading volume that month to total shares outstanding that month. I then take the average for each bin and plot them in Figure 5.

As you can see, the average ratio of trading volume to shares outstanding has a very pronounced U-shape. The U has its lowest point near the mean-reverting discount level and increases substantially as you move away from the mean-reverting level in either direction. For instance, whereas the average ratio of monthly trading volume to shares outstanding is 3.9% for the five-to-six percent discount bin, it is 8.8% for the 29% to 30% discount bin, and 9.5% for the -29% to -30% premium bin. Consequently, the ratio more than doubles as you move from the center of the distribution in Figure 1 to either of the tails.

The observed increase is consistent with several theoretical models such as Varian (1989), Kandel and Pearson (1995), and Harris and Raviv (1993) that link increased divergence of opinion among traders with higher trading volumes. As funds move further away from fun-

Figure 5. Average ratio of monthly trading volume to shares outstanding (expressed as a percent) when sorting by discount or premium level, 1985 to 2001.

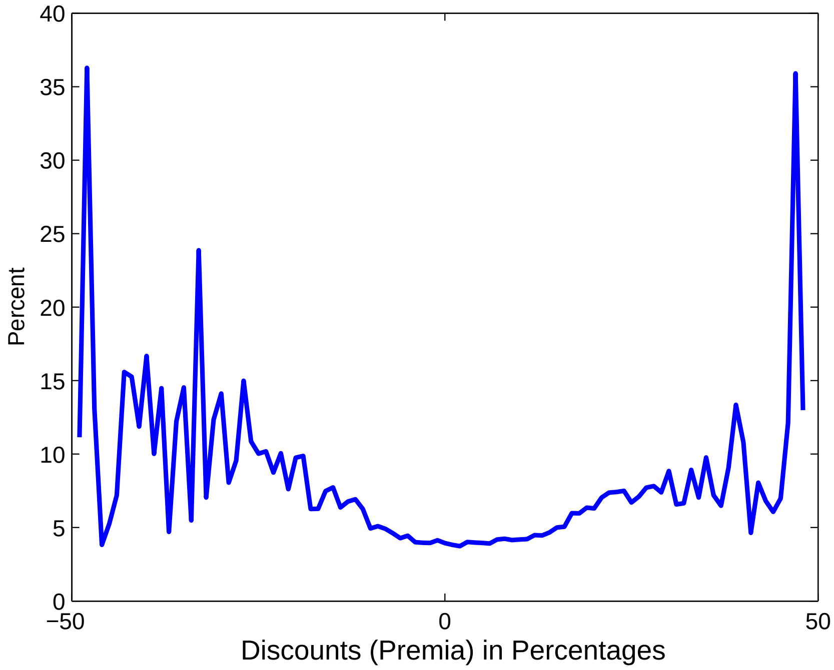

damentals, it is likely that more rational traders will enter closed-end fund markets in order to attempt to profit from the mispricings. Their opinions will differ greatly from those of irrational noise traders who believe that the current mispricings are valid or that even greater mispricings are necessary. In addition, there may also be increasingly large differences of opinion among the noise-traders themselves as fund prices move farther away from fundamental valuations. This, too, would increase the divergence of opinion and help to generate the U-shaped volume pattern observed in Figure 5.

It is also important to note that you could only with difficulty attribute the U-shaped volume pattern to informational asymmetries among traders, since all market participants have access to the same public information about fund portfolio values. Thus, models such as those

Figure 6. Average ratio of bid-ask spread to ask price (expressed as a percent) when sorting by discount or premium level, 1985 to 2001.

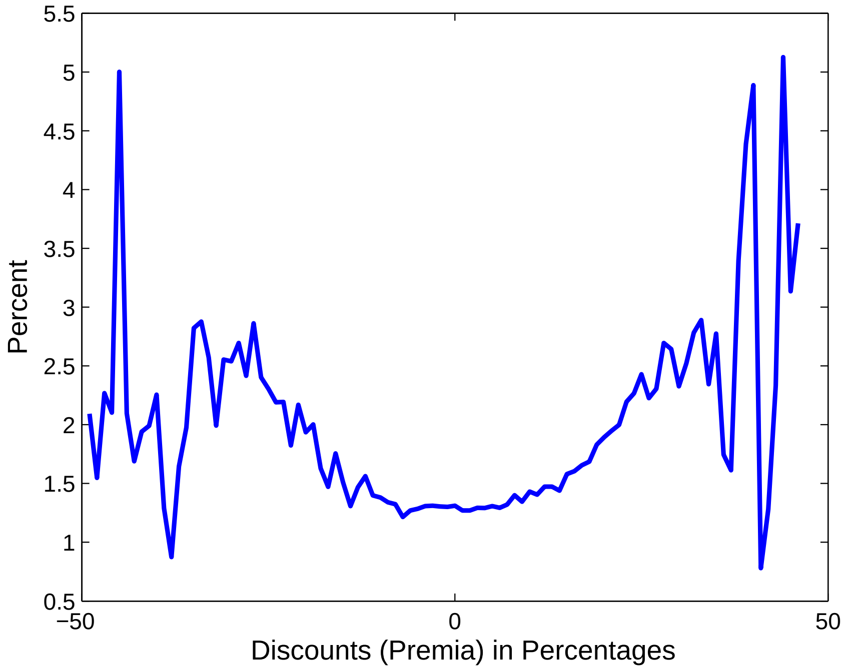

by Kyle (1985) and Easley and O'Hara (1992) in which trading volume increases when informational asymmetries between informed and uninformed traders increase do not appear to be applicable here.

For the same reason, it is also difficult to attribute to informational asymmetries the U-shaped pattern of bid-ask spreads that is evident in Figure 6. To build this figure, I place the monthly discount and premium observations into bins and then use daily CRSP ask and bid prices matched to the date of each monthly observation to construct the ratio of spread to ask price for each observation—that is, $(ask - bid)/ask$ for each observation. I then average all the ratios in each bin and plot them in Figure 6.

As you can see, there is a large increase in spreads as you move from the center of the figure to the tails. At the mean reverting discount level, the average spread to ask ratio, $(ask - bid)/ask$, is 1.3%. For the 29% to 30% discount bin, it is 2.3%. For the -29% to -30% premium bin it is 2.7%. These figures are, respectively, 77 percent larger and 108 percent larger than the ratio at the mean-reverting discount level.

Such large differences are hard to reconcile with theories that ascribe the size of spreads to informational asymmetries between market makers and other traders, as in as in the models of Kyle (1985), Easley and O'Hara (1992), and Glosten and Milgrom (1985). As pointed out by Neal and Wheatley (1995), the informational asymmetry story is simply not very convincing for closed-end funds given the fact that market makers are as well informed about fundamentals—fund NAVs—as any other market participants.

So something else must be causing market makers to widen their spreads as fund prices move away from fundamental valuations. Fund-specific noise-trader risk is a reasonable explanation since increasing price volatility by itself will cause risk-averse market makers to widen spreads—even in the absence of informational differences—as in the models of Garbade and Silber (1979) and Ho and Stoll (1981). My results also support Van Ness, Van Ness, and Warr (2001), who conclude—after testing five theoretical models that incorporate asymmetric information as an explanation for the size of spreads—that volatility itself (rather than asymmetric information) is the major factor influencing the size of spreads.

To summarize, the U-shaped patterns of both volume and bid-ask spreads imply that there are increasing levels of volatility and risk in closed-end fund markets the further you move away from the mean-reverting discount level. Of particular interest is the widening of spreads. Since these are determined by risk-averse human beings—the NYSE specialists who make markets for the vast majority of the funds in the Fund Edge data set—they are the strongest evidence that there is a form of risk that increases in magnitude the farther you move away from the mean-reverting discount level. I believe that the cross-sectional pattern of excess returns observed in Tables 1 and 3 are the result of the market pricing this risk so that those who choose to bear it receive due compensation.

4. Conclusion

This paper examines the effects of arbitrage and the returns to arbitrage in closed-end funds. It concludes that the behavior of discounts and premia, arbitrage returns, fund trading volume, and fund bid-ask spreads all argue in favor of a risk factor that increases in magnitude the farther you move away from the center of the discount and premium distribution. This paper argues that this risk factor is a form of noise-trader risk—the risk that fund mispricings may widen rather than narrow due to the unpredictable trading activity of noise-traders.

The strongest evidence supporting this contention comes from an examination of the cross section of closed-end fund arbitrage returns. Twenty portfolios are created to examine how the returns to arbitrage vary by discount or premium level. When the returns to these twenty arbitrage portfolios are regressed on the standard Fama-French factors, they produce a striking pattern of excess returns. The excess returns are near zero for the portfolio closest to the center of the discount and premium distribution, and then grow rapidly in magnitude as you move away from the center of the distribution in either direction.

Since this pattern matches the growth in the volatility of arbitrage returns as you move away from the center of the distribution, it suggests that the excess returns serve to compensate investors for noise-trader risk that also increases as you move away from the center of the distribution. If this risk were properly accounted for, the excess returns would disappear. But because we have no variable capable of capturing the differing levels of noise-trader risk to which the various portfolios are exposed, we get the observed pattern of excess returns.

Additional evidence consistent with noise-trader risk increasing as you move away from the center of the discount and premium distribution is provided by graphs showing, respectively, how trading volume and bid-ask spreads vary with discount and premium levels. Both are U-shaped. Their nadirs occur at the same discount level as the center of the discount and premium distribution and both increase rapidly as you move away from the center of the distribution.

Since higher volume is consistent with a larger diversity of trader opinions, the U-shaped pattern for volume is consistent with the idea that noise-trader risk increases as you move away from the mean-reverting discount level. The U-shaped pattern for spreads is even more indicative of increasing noise-trader risk because spreads are set by risk-averse market makers. Given that the markets for shares of closed-end funds offer no obvious informational asymmetries (since fund portfolio values are public), the best explanation for the widening of spreads is that they are the self-protecting response of risk-averse market makers to increasing levels of noise-trader induced volatility.

In conclusion, something is causing the cross-sectional pattern of excess returns found in this paper. Further research should examine whether it is in fact due to noise-trader risk, or whether some other more orthodox risk factor is responsible. If noise-trader risk is to blame, then many new research avenues will be opened up to see whether the phenomena noted in this paper are restricted to closed-end funds, or whether they affect other assets as well. Of particular interest may be whether the observed rate of mean reversion for discounts and premia is an equilibrium: Does mean reversion proceed just fast enough to generate a rate of return just large enough to just compensate risk-averse arbitrageurs for the level of noise-trader risk that they must bear when engaging in arbitrage? Or do they earn excess returns even after accounting for noise-trader risk?

References

Anderson, Seth, 1986, Closed-end Funds versus Market Efficiency, *Journal of Portfolio Management* 13, 63–67.

Black, Fisher, 1986, Noise, *Journal of Finance* 41, 529–543.

Campbell, John, and Albert Kyle, 1993, Smart Money, Noise Trading, and Stock Price Behavior, *Review of Economic Studies* 60, 1–34.

Chay, J.B., and Charles A. Trzcinka, 1999, Managerial Performance and the Cross-sectional Pricing of Closed-end Funds, *Journal of Financial Economics* 52, 379–408.

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, The Closed-end Fund Discount and Performance Persistance, *Working Paper, London Business School*.

Doukas, John, and Nikolaos Milonas, 2004, Investor Sentiment and the Closed-end Fund Puzzle: Out-of-sample Evidence, *European Financial Management*.

Easley, David, and Maureen O'Hara, 1992, Time and the Process of Security Price Adjustment, *Journal of Finance* 47, 577–605.

Elton, Edwin, Martin Gruber, and Jeffrey Busse, 1998, Do Investors Care About Sentiment?, *Journal of Business* 71, 477–500.

Fama, Eugene, and Kenneth French, 1992, The Cross-Section of Expected Stock Returns, *Journal of Finance* 47, 427–65.

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*.

Garbade, Kenneth, and William Silber, 1979, Structural Organization of Secondary Markets: Clearing Frequency, Dealer Activity, and Liquidity Risk, *Journal of Finance* 34, 577–593.

Glosten, Lawrence, and Paul Milgrom, 1985, Bid-ask Spreads and Transactions Prices in a Specialist Market, *Journal of Financial Economics* 14, 71–100.

Harris, Milton, and Artur Raviv, 1993, Differences of Opinion Make a Horse Race, *Review of Financial Studies* 6, 473–506.

Ho, Thomas, and Hans Stoll, 1981, Optimal Dealer Pricing Under Transactions and Return Uncertainty, *Journal of Financial Economics* 9, 47–73.

Kandel, Eugene, and Neil D. Pearson, 1995, Differential Interpretation of Public Signals and Trade in Speculative Markets, *Journal of Political Economy* 103, 831–872.

Kyle, Albert, 1985, Continuous Auctions and Insider Trading, *Econometrica* 53, 1315–1335.

Lee, Charles M., Andrei Shleifer, and Richard Thaler, 1991, Investor Sentiment and the Closed-end Fund Puzzle, *The Journal of Finance* 46, 75–109.

Lintner, John, 1965, The Valuation of Risk Assets and the Valuation of Risky Investments in Stock Portfolios and Capital Budgets, *Review of Economics and Statistics* 47, 13–37.

Malkiel, Burton, 1977, The Valuation of Closed-end Investment Company Shares, *The Journal of Finance* 32, 847–859.

Neal, robert, and Simon Wheatley, 1995, How Reliable Are Adverse Selection Models of the Bid-Ask Spread, *Kansas City Federal Reserve Bank Working Paper Series*.

Newey, Whitney, and Kenneth West, 1987, A Simple Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703–708.

Peavy, John W., 1990, Returns on Initial Public Offerings of Closed-end Funds, *Review of Financial Studies* 3, 695–708.

Pontiff, Jeffrey, 1995, Closed-end Fund Premia and Returns: Implications for Financial Market Equilibrium, *Journal of Financial Economics* 37, 341–370.

Ross, Stephen A., 2002, Neoclassical Finance, Alternative Finance, and the Closed End Fund Puzzle, *European Financial Management* 8, 129–137.

Sharpe, William F., 1964, Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk, *Journal of Finance* 19, 425–442.

Thompson, Rex, 1978, The Information Content of Discounts and Premiums on Closed-end Fund Shares, *Journal of Financial Economics* 6, 151–186.

Van Ness, B, R Van Ness, and R. Warr, 2001, How Well Do Adverse Selection Components Measure Adverse Selection?, *Financial Management* 30, 77–98.

Varian, Hal, 1989, Differences of Opinion in Financial Markets, in Courtenay C Stone, eds.: *Financial Risk: Theory, Evidence, and Implications* (Kluwer Academic Publishers,).

Weiss, Kathleen, 1989, The Post-offering Price Performance of Closed-end Funds, *Financial Management* 18, 57–67.