

Arbitrage, Noise-trader Risk, and the Cross Section of Closed-end Fund Returns

Sean Masaki Flynn*

This version: April 18, 2005

ABSTRACT

I find that despite active arbitrage activity, the discounts of individual closed-end funds are not driven to be consistent with their respective fundamentals. In addition, arbitrage portfolios created by sorting funds by discount level show excess returns not only for the three Fama and French (1992) risk factors but when a measure of average discount movements across all funds is included as well. Because the inclusion of this later variable soaks up volatility common to all funds, the observed inverse relationship between the magnitude of excess returns in the cross section and the ability of these variables to explain overall volatility leads me to suspect that fund-specific risk factors exist which, were they measurable, would justify what otherwise appear to be excess returns. I propose that fund-specific noise-trader risk of the type described by Black (1986) may be the missing risk factor.

*Department of Economics, Vassar College, 124 Raymond Ave. #424, Poughkeepsie, NY 12604. flynn@vassar.edu I would like to thank Steve Ross for inspiration, and Osaka University's Institute for Social and Economic Studies for research 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.

Closed-end mutual funds have been closely studied because they offer the chance to examine asset pricing in a situation in which all market participants have common knowledge about fundamental valuations. By law, closed-end funds trading in the USA must disclose their portfolio values at least once per week. As a result, one would expect that an informationally efficient market would set the market value of each fund equal to its publicly announced portfolio value, or perhaps to its portfolio value 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 the values of their underlying portfolios. A natural result of this phenomenon has been a large literature devoted to understanding why funds trade at discounts and premia, 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 1999 for an excellent survey.)

In this paper, I explore arbitrage in closed-end funds using the Fund Edge data set sold by Weisenber/Thompson Financial. It is largest data set ever used in the academic literature on closed-end funds and contains nearly every closed-end bond and stock fund trading in the USA and Canada over the relevant time period, 1985-2001. From it, I draw three major conclusions.

The first is that the funds in this data set routinely violate both the static arbitrage bounds proposed by Gemmill and Thomas (2002) as well as the dynamic arbitrage bounds proposed by Pontiff (1996). In a data set of 168 UK-traded closed-end stock funds running from 1991 to 1997, Gemmill and Thomas (2002) found support for their static arbitrage bounds hypothesis, under which discounts should not exceed thirty percent while premia should not exceed five percent. These bounds are, however, routinely—and often extremely—violated by the US and Canadian funds in the Fund Edge data set.

Using a data set of 53 US-traded stock and bond funds running over the period 1965-1985, Pontiff (1996) found evidence in favor of dynamic arbitrage bounds that widen and narrow depending on the interest rate and how it affects the opportunity cost of capital invested in closed-end fund arbitrage positions. This relationship between interest rates and discounts is completely absent in the Fund Edge data, which begins in 1985 (where Pontiff's data ends) and runs through 2001. In particular, if you apply Pontiff's econometric methodology to the Fund Edge data, interest rates (and thus short selling costs) appear to have no effect on the magnitude of discounts and premiums during the more recent period.

The second finding of this paper is that the Fund Edge data does not support Gemmill's and Thomas' (2002) conclusion that, at least over the long run, arbitrage pressures tend to push individual funds to trade at prices consistent with their own particular fundamentals. If you apply their methodology to the US and Canadian funds in the Fund Edge data set, you get insignificant or contradictory results. In particular, fund expenses (which should affect the present value of future fund disbursements) and a measure of the difficulty of replicating a given fund's portfolio (something very important when considering arbitraging against a mispriced fund) are both insignificant. In fact, the only variable that is consistently correlated with long-run average discount and premium levels is a measure of each fund's exposure to noise-trader risk. The higher the exposure to noise-trader risk, the deeper the discount.

This fact bears notice because it hints that noise-trader risk—the risk that irrational “noise traders” will cause equity prices to move unpredictably and without relation to news about fundamentals—may be the reason that we do not see funds driven, even over the long run, to have share prices consistent with their fundamentals. If the volatility induced by noise trading is significant, closed-end fund share prices may deviate substantially from rational levels even in the presence of substantial arbitrage activities: the noise-trader induced volatility in such cases is simply too great for arbitrageurs to fully overcome.

That there *are* substantial arbitrage activities is confirmed by Flynn (2004), who finds that short ratios of NYSE-traded closed-end funds increase from zero to very high levels as you move from looking at funds trading at discounts to those trading at successively larger premia. Flynn finds that arbitrageurs pay close attention to fund discount and premium levels and intensely short sell funds trading at large premia. This fact is important because it means that the large deviations of fund prices from fundamentals take place despite intense arbitrage activities.

This is a fact calling for an explanation, and the one I propose is price volatility caused by noise traders. More particularly, I suggest that noise-trader induced volatility is strong enough to limit the size of arbitrage positions to a level where they cannot immediately and forever keep fund share prices equal to fundamentals. As suggested by DeLong, Shleifer, Summers, and Waldmann (1990), “noise traders create their own space.”

My third finding is that unpredictable noise-trader induced volatility would also offer an independent source of risk that would serve to rationalize what otherwise appear to be excess returns available to anyone who wishes to arbitrage mispriced funds. These show up strongly if you run regressions on model portfolios created by sorting closed-end funds each month cross sectionally by their respective discount and premium levels. These portfolios go long fund shares and short fund portfolios, thereby isolating the returns generated when discount levels change. The conventional Fama and French (1992) risk factors cannot eliminate these excess returns.

Neither can they be eliminated by introducing a variable that captures changes in average discount levels across funds. Since this variable will tend to capture any sort of volatility that affects closed-end funds as a group, its inclusion makes it much more likely that any remaining volatility is due solely to fund specific noise-trader volatility.

After including this variable, the pattern of *R*-squared statistics and excess returns across the various portfolios strongly suggests that fund-specific noise-trader risk (rather than something common to all

funds) may be the hidden risk factor for which the otherwise excessive returns to arbitrage in closed-end funds may be explained. For the portfolios that reflect funds trading near rational levels, R -squared statistics are very high while excess returns are near zero. But for portfolios farther and farther away from fundamental levels, you see R -squared statistics dropping precipitously while excess returns grow larger and larger.

Since the regressions account for movements in discount and premium levels that affect 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 risk factor that also grows as you move farther away from fundamental discount levels. I conclude by suggesting that this risk factor is fund-specific noise-trader risk of the type envisioned by Black (1986) and provide additional evidence for it by showing that both trading volume and bid-ask spreads also increase significantly as you move farther and farther away from fundamental discount levels.

Section I describes the Fund Edge data set and defines discounts and premia. Section II tests and rejects the static arbitrage bounds proposed by Gemmill and Thomas (2002). Section III finds no support for the dynamic arbitrage bounds of Pontiff (1996). Section IV uses the methodology of Gemmill and Thomas (2002) to test and reject their hypothesis that arbitrage pressures are strong enough to push fund prices to be consistent with fundamental factors like replication risk and dividend yield. Section V examines the cross section of closed-end fund arbitrage returns. 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.

Below, I will only utilize Fund Edge data covering 1985-2001. I do this for two reasons. First, the bulk of the data lies in that time period because there was a huge increase in the number of funds starting in the late 1980s.² Second, this time period allows for a useful comparison with the results of Pontiff (1996) and other prominent papers such as Lee, Shleifer, and Thaler (1990) that share the use of a data set of US closed-end funds covering the period 1965-1985.

In this paper, discounts will be 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 . We can define the discount or premium at which a fund trades at time t to be either $D_t = \log(N_t/P_t)$, consistent with the convention of Pontiff (1996), or $D_t = N_t/P_t - 1$, consistent with the convention of Gemmill and Thomas (2002). The choice of definition will depend on whether we are comparing Fund Edge results with those from Pontiff (1996) or Gemmill and Thomas (2002). 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.

II. Testing The Static Arbitrage Bounds Hypothesis

Closed-end funds are mutual funds whose shares trade like common stock on major stock exchanges. Unlike the much more numerous open-end mutual funds that guarantee to redeem shares at par with portfolio value, closed-end funds do not engage in redemptions. As a result, the only way for a current shareholder to cash out her shares is by selling them on the stock exchange at whatever price the market will carry.

It is typically the case, however, that a fund's price per share does not equal its portfolio value per share. This is surprising because, management fees aside, purchasing the shares of a closed-end fund entitles the owner to the same stream of future payments as she could earn mimicking the fund's underlying portfolio. One would expect arbitrage pressures to either set the price of a fund's shares equal to the portfolio value per share or to the portfolio value per share less the capitalized value of expected future management fees.

This is especially true given that there do not appear to be any substantial barriers to arbitrage among closed-end funds. For instance, many of the equity funds have betas near 1 and are consequently easy to hedge. Transparency is also not an issue because closed-end funds in the United States are required by law to publicly disclose the value of their portfolios at least once per week, and some funds now even use their Web site to update investors in real time about any changes to their portfolios. The shares also trade on major exchanges and most of them are quite liquid. And, as I mention in the introduction, Flynn (2004) does in fact find very active short selling of NYSE funds trading at premia.³

Despite the apparent ease with which arbitrage might be conducted in this market, closed-end funds trade at large and lingering discounts and premia relative to their portfolio values. This has spawned a substantial literature debating whether the distribution of discounts and premia can be explained rationally or is the result of irrational noise-traders driving share prices away from fundamental values. See, for example, Boudreaux (1973), Zweig (1973), Thompson (1978), DeLong, Shleifer, Summers,

and Waldmann (1990), Chen, Kan, and Miller (1993), Chopra, Lee, Shleifer, and Thaler (1993), and Pontiff (1997). Dimson and Minio-Kozerski (1999b) provide a comprehensive overview of the debate.

The model constructed by Gemmill and Thomas (2002) (hereafter GT) can be viewed as one of many attempts to explain the observed distribution of discounts and premia and why funds trade on average at discounts rather than at par with their underlying portfolio values. GT “begin by assuming that the discount is subject to fluctuations.” They do not attempt to explain the causes of those fluctuations. Rather, they assume that the price at which a fund trades and the value of its portfolio will each follow a log normal distribution with mean zero. Given their definition of discounts as $D_t = \log(N_t/P_t)$, this implies that the distribution would also be log normal and centered on zero. GT parameterize this log normal distribution for discounts and premia by assuming that the underlying price and portfolio distributions each have a volatility of 25% and a correlation of 0.9. This implies that the log normal distribution would have a standard deviation of 11.2% as well as a mean of zero.

The authors next propose that the full log normal distribution will not be seen in its entirety in real-world data because it will be censored at both ends. The deepest discounts will be censored because of an arbitrage pressure having to do with the possibility that funds trading at deep discounts will be either liquidated or converted to open-end funds by angry shareholders. Similarly, most of the premia will be censored because of an arbitrage pressure having to do with the creation of new funds through IPOs, the new funds putting downward pressure on the share prices of existing funds, thereby preventing their premia from becoming very large.

I will discuss both of these arbitrage barriers below. But first, compare the actual distribution found in the Fund Edge data set with the distribution predicted by the GT bounds. GT set their suggested upper bound at a discount of 30% and their suggested lower bound at premium of -5%. Applying these bounds to the assumed log normal distribution centered on zero and having a standard deviation of 11.2% produces Figure 1.

GT's predicted distribution does not fit the Fund Edge data. If we plot a relative frequency histogram of the 227,066 weekly discount observations on the 462 funds found in Fund Edge over the period January 1985 through May 2001, we get Figure 2. Figure 2 does not look like a log normal distribution of mean zero being censored. In particular, there is no sign that the distribution suddenly comes up against an arbitrage barrier on either side. The tails are gently tapering and nearly 16% of the observations lie outside the bounds assumed by GT. These facts suggest that the GT bounds do not hold and that the distribution found in Fund Edge does not follow the truncated log normal distribution assumed by GT.

However, it must be noted that GT's choice of upper and lower bounds along with their assumption that the discount distribution is mean zero with a standard deviation of 11.2% do imply that the censored distribution should have a mean discount of 5.87%, which is in fact very close to the empirical mode of about 6% found in Figure 2. However, the justifications given for the bounds as well as the assumption regarding the standard deviation of the log normal distribution are not strongly persuasive.

For instance, GT give two reasons why discounts should not exceed 30%. The first is that extremely deep discounts invite arbitrage activities that will tend to reduce discounts. Arbitrageurs will go long fund shares and short fund portfolios (or sufficiently similar hedge portfolios) in order to profit from the price divergence. GT contend that the price pressures that result from such arbitrage activities should keep discounts narrower than 30%. The second reason given by GT that discounts should never exceed 30% is that funds which trade at deep discounts are more likely to be liquidated or converted to open-end funds after a boardroom revolt by shareholders angry at the extremely deep discounts. However, the authors do not quantify the value at which either of these two pressures would become effectual. Both might be expected to increase in intensity gradually as discounts deepen, but the level at which they would become so strong as to deter even deeper discounts is not obvious. Thus the statement that the "discount can often reach 30% percent before the upper bound is reached" appears overly specific.

As for the lower arbitrage bound at a premium of -5%, it is justified by the notion that if a fund were to trade at a large premium it would attract competition in the form of IPOs of similar closed-end funds. While it is true that Levis and Thomas (1995) and Lee, Shleifer, and Thaler (1991) find that there are more fund IPOs when the average discount across all funds decreases (i.e. when the average discount moves from, say, 15% to only 8%), there is no justification in the literature for the assumption that potential IPO activities imply a boundary for individual funds of “somewhere around” -5%.

This is especially true given that the evidence presented by Levis and Thomas (1995) and Lee, Shleifer, and Thaler (1991) has to do with IPOs increasing when average values of discounts across all funds decrease, not with whether or not the discount of any particular firm reaches the -5% premium boundary. Furthermore, the average values which appear to trigger increased IPOs are actually *discounts* of about 5% in the sample of Lee, Shleifer, and Thaler (1991) and of about 11.5% in Levis and Thomas (1995). Neither paper gives any suggestion about an arbitrage barrier for individual funds at a premium of -5%.

Indeed, there is substantial evidence that no such barrier exists. For instance, if competition for investor dollars meant that seasoned funds trading at large premia would engender IPOs of similar funds, then there should have been a plethora of IPO's in funds investing in Taiwan. That is because the original Taiwan Fund began trading in December 1986 at a -205% premia and stayed at more than a -50% premium for most of the next 18 months. That -205% premium was one of the largest seen among closed-end funds in the USA since the run up to the Crash of '29. However, the only other competitive fund to get started, the ROC Taiwan Fund, was not started until May of 1989—by which time the original Taiwan Fund was trading at par. Even worse for any theory that new funds should provide competition for existing funds is the subsequent behavior of the two funds. Six months after the new ROC Taiwan Fund began trading, the seasoned Taiwan Fund jumped to a -104% premium while the new ROC Taiwan Fund fell to an 8% discount. And for most of the next 16 months, the seasoned

Taiwan Fund traded at D_t values at least 20 percentage points more negative than those of the newer ROC Taiwan Fund.

The original Korea Fund also went to a very large premia of just over -150% in 1986 and continued to trade at more than a -30% premia for the next three years. Yet there were no new IPO's of funds investing in Korea until 1992 and 1993, when, respectively, the Korean Investment Fund and the Korean Equity Fund got started. Here, again, we see a large violation of the suggested -5% premium barrier without any sign of IPO's of similar funds.

While the suggested arbitrage boundaries are not consistent with the Fund Edge data, the parameter values that GT assume when generating their log normal distribution for discounts are consistent with the values found in Fund Edge. They assume "a 25 percent annual volatility for both net asset value and price and a correlation [between them] of 0.90." The actual figures for the 462 funds in Fund Edge using monthly data over the period 1985-2001 are 22.9% and 29.8%, respectively, for the yearly volatilities of NAVs and share prices, and 0.895 for the correlation. Consequently, the poor fit that Figure 2 has with the predictions of Figure 1 cannot be due to incorrectly parameterizing the log normal distribution. Rather, it likely stems from the inapplicability to the Fund Edge data of the log normal model and the given boundary assumptions.

The premium boundary assumption is particularly suspect. Of the 284 bond funds that, as of June 2001, had been in business at least 60 months, 87% of them had surpassed a -5% premium at least once. Even more extreme premia were quite common. 44% had exceeded a -10% premium at least once, 22% had exceeded a -15% premium at least once, and 10% had exceeded a -20% premium at least once. Of the 114 stock funds that had, as of June 2001, been in business at least 60 months, 93% had broken the -5% premium barrier at least once. 72% had exceeded a -10% premium at least once, 53% had exceeded a -15% premium at least once, and 38% had exceeded a -20% premium at least once. If a premium barrier exists, it is much less rigid than suggested by GT, or is at a much deeper premium.

The discount and premium barriers also imply that the discount distribution should be platykurtic (because the tails should be chopped off) and skewed toward discounts (because more of the left tail than the right tail would be chopped off). In fact, just the opposite is true in the Fund Edge data. The distribution is leptokurtic and skewed toward premia rather than discounts. For the 227,066 weekly discount and premia observations in Fund Edge between 1985 and 2001, the skewness and kurtosis for the entire sample are, respectively, -2.00 and 20.96. Skewness toward premia and leptokurtosis remain even if we exclude the 761 most extreme observations. These are the ones that are not included in Figure 2 because they were in the extreme tails of the distribution, with either premia of less than -50% or discounts greater than 50%.⁴ Excluding the 761 outliers, the skewness and kurtosis decline dramatically to, respectively, -0.42 and 5.07. But these values still contradict the predictions of skewness toward discounts and platykurtosis implied by the GT arbitrage barriers.

The same pattern emerges if we look not only at the aggregate discount distribution, but at the individual discount distributions of each of the 284 bond funds and 114 stock funds which had been in business for at least 60 months. For each individual fund, we can calculate the mean discount and the standard deviation of discounts as well as the skewness and kurtosis of the fund's discounts. After we have done this for all funds, we can average across, respectively, the bond funds and the stock funds. For the bond funds, we find that the average mean discount is 3.35%, with an average standard deviation of 5.59%, an average skewness of -0.13, and an average kurtosis of 3.21. For the stock funds, we find that the average mean discount is 7.53%, with an average standard deviation of 10.68%, an average skewness of -0.70, and an average kurtosis of 4.23. From these statistics, we see that both bond and stock funds are skewed toward premia rather than discounts, and that both are leptokurtic, with stock funds being highly leptokurtic.⁵

This section has given ample evidence that the static arbitrage boundaries suggested by GT in their paper on UK funds are routinely violated by the US and Canadian funds found in the Fund Edge data set. Moreover, the gentle taper of both tails of the discount distribution of Figure 2 appears

inconsistent with discounts or premia suddenly coming up against static arbitrage barriers of any sort at any point. Consequently, if arbitrage constrains closed-end fund discounts and premia, it must do so in a different fashion.

III. Testing The Dynamic Arbitrage Bounds Hypothesis

As we turn to consider the dynamic arbitrage bounds hypothesis argued for by Pontiff (1996), we should first note that it has the potential to easily account for the shape of the empirical discount distribution of Figure 2. The gentle taper of the distribution (rather than its suddenly stopping near static arbitrage bounds) would be a natural consequence of such boundaries waxing and waning over time. And the fairly symmetric shape (again compared to the predictions of the static bounds hypothesis) would result if the dynamic bounds were about evenly spaced above and below the mean as they expanded and contracted. That leaves only the centering of the empirical distribution on a discount of about 6% to be explained.

Management fees are the most obvious candidate. Thinking in terms of a discounting model, if a fund's manager is not able to beat a buy-and-hold strategy but charges fees for managing the fund's portfolio, the present value of future fund disbursements will necessarily be less than the current portfolio value by the present value of future management fees. This suggests that a rational investor would only be willing to buy shares of the fund if it traded at a discount to its current portfolio. As it turns out, models like those of Ross (2002) and Flynn (2002) that price funds by capitalizing out future management fees predict that funds should on average trade at discounts of around 7%, which is reasonably close to the mode discount of 6% seen in the data. Fee-based models are, however, rejected by GT because "discounts are not sensitive to the level of interest rates."

For evidence in support of this contention, GT cite Lee, Shleifer, and Thaler (1991) and Pontiff (1996), who utilize the same data set covering 68 closed-end stock and bond funds trading in the USA

during the period 1965-1985. Lee, Shleifer, and Thaler (1991) utilize a subset of only 20 stock funds to construct a time series of average discounts that they find to be uncorrelated with unanticipated changes in the term structure of interest rates. They take this as evidence “counter to the agency cost argument which predicts that when long rates fall the present value of future management fees rise, so discounts should increase.”

Pontiff (1996) also uses only a subset of the data, as he chooses to use just the 52 stock and bond funds for which there were at least six months of data. He argues that interest rates affect the profitability of investment activities that would force fund share prices toward fund NAVs. Under this hypothesis, higher interest rates imply higher opportunity costs when undertaking an arbitrage, and since this is true for both arbitrages against discounts as well as against premia, he expects to find that higher interest rates will lead to a wider spread of observed discounts. That is, there will be both larger discounts as well as larger premia when higher interest rates reduce the profitability arbitrage. Or, utilizing our notation, the cross-sectional average of the absolute value of D_t observations should increase when interest rates rise. Pontiff’s arbitrage bounds are, therefore, time-varying unlike the static arbitrage bounds suggested by GT.

Pontiff tests for time-varying arbitrage bounds by regressing short-run Treasury yields on a monthly time series that averages the absolute values of D_t across all funds trading in each given month over the period January 1965 through December 1985. His results are reported in Table III of Pontiff (1996), which is reproduced here as Table I. The regressions are performed in both levels and first differences in columns (1) and (2). Slope coefficients on interest rates have the predicted positive sign and are strongly statistically significant. Their magnitude is such that each 1-percentage point increase in short-term interest rates increases the average absolute value of D_t by one-half of one percentage point.

Pontiff also reports in columns (3) and (4) that regressions in both levels and first differences of short-term interest rates on the levels and first differences of the average value of D_t itself (rather than the average of the absolute value of D_t) indicate a positive but statistically insignificant relationship

between interest rates and discount levels. That is, higher interest rates cause D_t to insignificantly decrease in value (e.g., from a discount of 5% to a discount of 3.5%). Since his costly-arbitrage model only makes predictions about the response of the absolute value of D_t to interest rates, and not about the response of D_t itself, this insignificant result is consistent with Pontiff's theory. But it is also what is cited by GT as evidence to reject fees-based discount models.

The Fund Edge data, however, demonstrate a strikingly different relationship between interest rates and discount levels. Table II presents the results of running the exact same specifications as given in Pontiff's Table III, but this time using Fund Edge data. As with Pontiff (1996), I construct discounts as the log ratio of portfolio value to market price, use secondary-market yields on 4-week Treasuries for interest rates, include only funds that had been in business at least six months, and use the same AR(4) model for error autocorrelation. The only difference is the time period. Pontiff's data set covered 1965-1985, while I use the section of Fund Edge covering 1985-2001. While the more recent period shows less interest rate volatility than the former, there were still large movements in rates during the more recent period so that if a robust relationship does exist between interest rates and arbitrage bounds, it should be readily observable.

Table II shows that in the more recent period, interest rates had a much different effect on both average absolute discounts and average discounts than they did during the earlier period. Columns (1) and (2) of Table II show that the positive relationship found by Pontiff between interest rates and average absolute discounts is no longer statistically significant. This suggests that discounts may not be constrained by interest-sensitive arbitrage bounds. In the previous section, I gave evidence that the stationary arbitrage bounds suggested by GT did not bind. The regression results of Table II suggest that the time-varying arbitrage bounds suggested by Pontiff also fail to bind, or that they bind so weakly that their effect cannot be clearly determined empirically.

Even more interesting are the results of columns (3) and (4) of Table II, which show an economically large and statistically significant positive relationship between interest rates and discount levels.

Both the levels specification and the first differences specification indicate that a 1 percent increase in short-term interest rates leads to an increase in the average discount level of about 2.2 percent.

You may wonder about the implications of this positive relationship between interest rates and discount levels, especially in terms of a fees-based explanation for discounts. In particular, it may seem that the relationship has the wrong sign, as one might expect that if interest rates rose, the present value of future management fees would decrease, thereby leading to smaller rather than larger discounts. However, such a relationship is consistent with the behavioral model of discounts presented by Flynn (2002).

In that model, investors price a fund by how much they think managers will be able to beat alternative rates of return. If interest rates are a measure of the alternative rate of return, then the model predicts that increases in short-term interest rates will be associated with increases in discount levels. Holding investor expectations about managerial performance constant, an increase in interest rates decreases the amount by which investors think that managers will beat the market (or increases the amount by which they feel managers will trail the market.) This causes discounts to increase, consistent with the behavior found in columns (3) and (4).

The model also predicts that if investors are neutral about a manager's prospects, then his fund will trade at a discount equal to the present value of his expected future management fees. Thus, the model suggests that the placement of the center of the distribution of Figure 2 and its overall shape are due to sentiment varying around the neutral level consistent with discounting out expected future management fees. Extreme pessimism causes large discounts. Extreme optimism causes large premia. And ambivalence—the more normal state of sentiment—leaves you near a discount equal to the present value of expected future management fees.

IV. Do Arbitrage Pressures Drive Discounts Toward Fundamentals?

The previous two sections have presented evidence that discounts and premia violate both static and dynamic arbitrage bounds. These violations may be evidence against strong and effective arbitrage—or they may just be due to the statistical tests having low power. Consequently, it is wise to measure the effectiveness of arbitrage pressures in aligning fund share prices with fundamentals in yet another way. This section applies GT's regression methodology to test whether or not long-run average discounts are consistent with fundamentals. If so, this could be taken as evidence that rational arbitrageurs effectively constrain discount levels to be consistent with fundamentals.

GT provide evidence in favor of this hypothesis by examining 158 UK-traded closed-end stock funds over the period 1991 to 1997. For each fund, they calculate the fund's average discount over that period and then regress it on a constant, the average expense ratio of the fund over those years, a measure of the fund's noise-trader risk over those years, the log of the fund's age at the end of those years, a measure of the difficulty of mimicking the fund's portfolio over those years, the fund's average dividend yield over those years, and the log of the fund's average market capitalization over those years.⁶ When this regression is run on their UK data, they find that six of the seven variables are significant at the one-percent level, and that the seventh, the log of market capitalization, is significant at the five-percent level. This regression result is presented in the first column of Table IV of Gemmill and Thomas (2002) and is reproduced here as the first column of Table III. The regression was estimated using weighted least squares, with the volatility of each fund's discount series over the period 1991-1997 used as the weighting variable. This procedure gives the regression a very robust weighted R^2 of 0.52, and an un-weighted R^2 of 0.34.

Besides being highly statistically significant, each of the variables has the expected sign, except for the measure of noise-trader risk. The positive coefficient on average expense ratios makes sense because higher fund expenses should be capitalized out by rational investors, thereby increasing discount

levels. The positive coefficient on the log of fund age is consistent with the work of Weiss (1989), and Levis and Thomas (1995), who find evidence that funds begin trading at their IPOs at ten percent premia which subsequently evaporate over the next six months as the funds mature and move toward discounts.⁷ The positive coefficient on replication risk is consistent with the idea that the harder a fund is to replicate, the more averse arbitrageurs will be to betting against deep discounts by going long the fund and short an imperfect hedge portfolio (that is, the lower is the replication risk, the more strongly GT expect their discount arbitrage bound to hold). The negative coefficient on the dividend yield is consistent with the idea that higher dividends should cause lower discounts because higher dividends reduce the amount of fund capital that will in the future be taken away from shareholders in the form of management fees. The negative coefficient on the log of market capitalization is expected by GT because they believe that larger funds enjoy a liquidity premium because they can be traded rapidly and with low bid-ask spreads, and because there may be economies of scale in fund management such that larger funds have lower expense ratios. Finally, GT interpret the negative coefficient on the noise-trader risk variable as evidence against the Lee, Shleifer, and Thaler (1991) hypothesis that noise-trader risk is a priced risk because if it were a priced risk, then one would expect higher levels of noise-trader risk to be associated with larger rather than smaller discounts.

Columns (2) and (3) of Table III give, respectively, the results of running the same specification on the closed-end stock funds and closed-end bond funds found in Fund Edge. Because Fund Edge does not provide back data on fund fees and expenses, these were gathered by hand from fund annual reports. Because complete data could not be obtained for all funds, and because the averages were taken over the ten year period running from January 1991 to December 2000, only 69 stock funds could be included in the stock fund regression reported in column (2), and only 123 bond funds could be included in the bond fund regression reported in column (3). (Only funds in business over the entire period and for which there was complete data were used.) The results of these regressions are very different from those of column (1), which GT got from their sample of UK-traded stock funds.

Whereas all six independent variables were significant for GT's UK data, only two variables each are significant for stock and bond funds in the Fund Edge data. For the stock funds, the noise risk variable and the log of fund age are significant at the five percent level, while for the bond funds the noise risk variable and the dividend yield are significant at the one percent level.

Notably, the only variable that is significant in all three regressions is the noise risk variable. However, while it is negative for GT's UK data, it is positive for the US and Canadian data. Consequently, GT's decision to "reject very clearly the view of Lee, Shleifer, and Thaler (1991) that noise-trader risk is a priced factor which causes a discount" is not supported by the US and Canadian data. Rather, noise-trader risk is the only tested factor that is priced into *both* bond and stock fund discounts in the US and Canadian data.

The results of columns (2) and (3) also cast doubt on the robustness of GT's conclusion that arbitrage pressures are strong enough to consistently drive discounts toward fundamental values. Not a single variable has the same sign and significance for both columns (2) and (3) as it does for column (1). While arbitrage pressures may be present and properly functioning in the sample of UK stock funds examined by GT, those pressures appear not to generalize to the Fund Edge data, and also do not seem to generalize to bond funds as opposed to stock funds. The expense ratio, replication risk, and market capitalization are all insignificant in columns (2) and (3). Worse yet for the generalizability of the UK results, the signs of the coefficients are different in columns (2) and (3). As for the other two independent variables, fund age is of the right sign for both US bond and stock funds in columns (2) and (3), but is only significant for stock funds. And whereas the coefficient on dividend yield is significant and of the expected negative sign in column (3) for bond funds, it is of the wrong sign and insignificant for stock funds in column (2).

To summarize, none of the GT regression results for UK stock funds generalize with any significance to both bond and stock funds traded in the US and Canada. This casts significant doubt on the contention that arbitrage pressures cause discounts to move toward levels consistent with fundamentals

like fund expense ratios and trading costs. The insignificance of the replication risk variable for both stock and bond funds in the US and Canada also casts doubt on whether the ease of replicating a fund's portfolio in any way affects arbitrage activities in closed-end funds. And the positive and statistically significant coefficients on noise-trader risk directly contradict GT's contention that noise-trader risk is not a priced factor in closed-end funds.

V. The Cross Section of Closed-end Fund Arbitrage Returns

The previous section showed that, across funds, arbitrage does not appear to push closed-end fund share prices toward discount levels consistent with fundamental factors. In particular, there is no evidence that funds with higher average expense ratios trade at larger long-run average discounts. How does one reconcile this with the nice bell shape of Figure 2? Because its mode is at a moderate discount consistent with capitalizing out fund expenses, it seems to suggest that arbitrage pressures *do* push discounts towards fundamentals, at least in the aggregate. How can this apparent contradiction be resolved?

This section presents evidence that noise-traders may be responsible. In particular, imagine a situation in which rational arbitrageurs interact with irrational noise traders in the closed-end fund markets. Any time discounts or premiums vary significantly from fundamentals, arbitrageurs take positions. Their actions will tend to keep discounts and premiums bounded. But the ongoing effect of noise-traders—who buy and sell at random without considering fundamentals—will tend to scramble fund discounts and premiums.

Unless the arbitrageurs are able to immediately and always overcome the mispricings caused by noise traders, the distribution of discounts and premiums across funds will tend to have little to do with the distribution of fundamentals across funds. Consequently, while Figure 2 is centered more-or-less correctly on the rational discount level that is consistent with capitalizing out fund expenses,

there is simultaneously no statistically significant relationship between individual fund discounts and individual fund fundamentals. Arbitrageurs have enough power to keep things more or less centered on fundamentals in the aggregate, but not enough power to ensure that funds with worse fundamentals will trade at deeper discounts than funds with better fundamentals.

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 randomly mis-interpret publicly available information as in the model of noise-trader risk applied to closed-end funds by DeLong, 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. 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 regressions run on this asset's returns, 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.

There is, however, a metaphysical problem with looking for a measure of noise-trader risk to use as an independent variable in regressions. Since noise-trader risk is, by definition, uncorrelated with anything else, you won't be able to find anything with which it is correlated to use as an independent variable in a regression. As such, the best that one can hope to do even in theory is to account for all other sources of risk and then attribute any excess return that remains as being the result of failing to include a measure of noise-trader risk in the regression.

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. 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.⁸

This strategy is of value because you should only find a pattern if noise-trader risk is, in fact, a priced risk factor. If it were not a priced factor, then it would not matter how you grouped funds into cross sections since if cross sections are unrelated to noise-trader risk, no pattern will emerge. On the other hand, if noise-trader risk *is* a priced factor, then those portfolios having more exposure to noise-trader risk will have larger excess returns. The trick is to cross-section the funds in such a way that you find a pattern of excess returns that can be plausibly associated with differences in exposure to noise-trader risk. My finding below is that such a pattern emerges if you simply sort funds by their discount and premium levels.

Below, I create twenty portfolios of closed-end fund returns based on discount and premium levels. For each month in the data set, I sort each fund into one of these twenty portfolios and keep track of the portfolio's return over the following month. I then run regressions on the returns from these portfolios over 1985-2001, using the standard Fama and French (1992) risk factors as well as a factor that will tend to account for any volatility that is common to all funds. That means that any noise-trader risk that remains will be fund-specific (and thus portfolio specific).

As I will show you, the excess returns from this cross section are consistent with the hypothesis that fund-specific noise-trader risk is a priced risk factor. Portfolios drawn from the tails of the discount and premium distribution of Figure 2 show much more volatile returns. They also show excess returns. In fact, the farther a portfolio lies from the center of the distribution, the larger are both its volatility of returns and its excess returns. Since the regressions account for other risk factors, I interpret this to mean that the higher volatility is caused by noise-traders, and that the noise-trader risk that they generate is a priced risk factor which shows up in the cross section as higher excess returns.

But before presenting the cross-sectional regressions, it is good to first look at 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.

A. Non-Risk-Adjusted Arbitrage Returns

Arbitrageurs hoping to benefit from closed-end fund mispricings must set up positions that will benefit from the movement of fund prices toward fundamental levels. For funds trading at large discounts, they would go long the shares of the fund and short its underlying, as any reduction in the discount would generate a profit. For funds trading at large premia, they would 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.⁹ 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 2 and that they do so

following an AR(1) process.¹⁰ 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.

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 $D_{t+1} = \text{constant} + \phi D_t + \varepsilon_t$. 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 5.20%. That is, this regression methodology indicates that discounts and premia revert to the same level as the mode discount of Figure 2.

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 5.20% will remain the following month—meaning that about 8.4% of any such gap will be closed, on average.¹¹ 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. It is thus fruitful to plot actual one-month returns to going long fund shares while shorting fund NAVs against discount and premium levels.

I do this in Figure 3 by aggregating all of my monthly data, sorting observations by discount or premium without regard to each observation's date. Specifically, I took each of the 52,188 monthly

discount or premium observations running from January 1985 to April 2001 and placed 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 calculated 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.¹² I call the former share returns and the later the NAV returns and plot their respective bin averages in Figure 3.

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.¹³ On the other hand, 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 2, 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 4.

The average difference between share price and NAV returns by bin (the thicker line in Figure 4) is of course equal to the vertical difference between the share return and NAV return lines in Figure 3. 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 line 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 with such positions. This is, of course, consistent with the intuitions of modern finance theory. 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, or whether it is a form of noise-trader risk that is uncorrelated with other factors.

Before addressing that issue, however, it is interesting to plot out reward-to-risk ratios against discount and premium levels. I do this in Figure 5 by dividing each bin's average share less NAV return by the bin's standard deviation of those returns—put differently, I divided the values given by the thick line in Figure 4 by those given by the thin line in Figure 4.

The result in Figure 5 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 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.

Even more striking is how the ratio behaves in the tails of the distribution. If you were to use your imagination to extend the line in the middle of the graph right and left, you would find that nearly all the values of the ratio in the right side of the graph lie below the imaginary line you've drawn, while nearly all those in the left side of the graph lie above the imaginary line. This means that the reward-to-risk ratios of arbitrage positions undertaken for the most extreme discounts and premiums are worse than

those in the middle of the distribution.¹⁴ This fact should be kept in mind as you look at risk-adjusted excess returns in the next section. While they grow increasingly large in magnitude as you move to more extreme discounts and premia, they don't grow excessively compared to the risk involved.

What the next section will show is that this risk is not accounted for by the standard risk factors. In fact, disaggregating the data into time series and running regressions with time varying risk factors cannot explain away the association between share less NAV returns and D_t levels that you get by aggregating all of the data as in Figure 4. Even after accounting for time varying risk factors, this relationship remains entirely robust: regardless of the time period, the best predictor of the returns to arbitrage in closed-end funds appears to be the distance between the current discount or premium and the mean-reverting discount level, \bar{D} . And regardless of the time period, there appears to be a systematic relationship between the returns to arbitrage positions that bear noise-trader risk and the magnitude of that risk.

B. 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 returns by discount or premium level. 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.¹⁵ 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.¹⁶ 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 discount 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, respectively, r^{share} and r^{NAV} . 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 arbitrageurs of US and Canadian funds over the period 1985 to 2001.

These time series are used as the independent variables in the Fama-French (1992) regressions shown in Table IV. 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 argued 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 IV, these three factors do a poor job of explaining the returns available to arbitrageurs in closed-end funds.

This is most obvious by looking at the third column of the table, which gives the excess returns to each portfolio. These 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 there is a risk factor that the Fama-French factors are failing to capture. This feeling is strengthened by comparing each portfolio's risk-adjusted excess return, α , with

the unadjusted average value of its dependent variable (share less NAV returns) in column eight. For instance, look at the sixth row of the table, which gives the regression results for the portfolio that bins together funds with premia between -25% and -20% each month. The average value of the arbitrage returns for that portfolio across the 100 months for which there was data was -1.51 percent per month. After regressing these returns on the three Fama-French factors, the excess returns were -1.88 percent per month. This means that the Fama-French factors are not accounting for all of the risks facing this portfolio over time (since if they were, the excess return would be statistically indistinguishable from zero.) A similar conclusion follows by making the same comparison for each portfolio. The Fama-French factors do not do a good job explaining the time variation of closed-end fund arbitrage returns.

But the pattern as you move down the third column (the excess returns column) also suggests that the risk that is not being captured varies systematically with fund discount and premiums. In particular, it seems to be larger the farther away you move from the mean-reverting discount level. This would be consistent with a situation in which noise-trader risk increased as you moved 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 deeper discounts or to 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. The three Fama-French risk factors fail to capture how this risk varies *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* a 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 modest. The highest R -squared statistics of 0.18 and 0.16 happen for funds near the long-run mean-reverting discount level, and 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 whatever explanatory power the Fama-French factors have 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.

B.1. Accounting for Volatility that Affects all Funds Simultaneously

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

Table V regresses a constant and the monthly change in the capital-weighted average discount level across all funds on the twenty arbitrage portfolio returns examined earlier in Table IV. 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 V is simply this variable's first difference, $\Delta W_t = W_t - W_{t-1}$. This lag structure makes sense since it 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 $r^{share} - r^{NAV} = constant + \Delta W_t + \varepsilon_t$ (where ε_t is the error term) for each portfolio,

you can see how the common factor that causes discounts and premiums to be correlated affects the returns to arbitrage positions.

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 IV. 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 the same portfolio.

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. This means that while changes in discount levels have great explanatory power for funds trading near the center of the distribution in Figure 2, 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 small investor sentiment argument 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.¹⁷ In addition, they argue that these movements in small investor sentiment may be a priced risk factor for assets in general.

But Elton, Gruber, and Busse (1998) and Doukas and Milonas (2004) have given strong evidence against small investor sentiment (as measured by 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 V 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. This would explain why the *R*-squared statistics in Table V fall so quickly as you move away from the mean-reverting discount level.

Table VI 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.¹⁸What remains, therefore, is fund (and, therefore, portfolio) specific.

Since modern finance theory rests on the idea that excess returns must be a compensation for risk, the cross-sectional pattern of excess returns found in Table VI 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 into the tails of the distribution in Figure 2. As noise-trader risk increases, so must compensation. That is why excess returns increase in Table VI for portfolios further and further away from the long-run mean-reverting discount level of about six percent.

The only question that remains is whether it is proper to ascribe the increasing volatility of arbitrage portfolio returns as you move away from the center of the distribution in Figure 2 to increasing levels of noise-trader risk. The next subsection gives more evidence in favor of this attribution by examining how trading volume and bid-ask spreads vary with discount and premium levels.

B.2. Other Evidence for Fund-Specific Noise-trader Risk

The cross-sectional pattern of excess returns found in Tables IV and VI 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 if you plot bid-ask spreads against discount and premium levels.

In Figure 6, I plot the ratio of average monthly trading volume to shares outstanding against discount and premium levels. To construct the figure, I aggregated monthly Fund Edge discount and premium observations for all funds available from January 1985 to May 2001 and then grouped them into one-percent wide discount and premium bins. For each fund in a given bin, I used CRSP data to

calculate the fund's ratio of trading volume that month to total shares outstanding that month. I then took the average for each bin and plotted them in Figure 6.

As you can see, the average ratio of trading volume to shares outstanding has a very pronounced U-shape. The U is centered on the mean-reverting discount level and shows substantial increases 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. The ratio more than doubles as you move from the center of the distribution in Figure 2 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. Obviously, as funds move further away from fundamentals, 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 tend to generate divergence of opinion and help to generate the U-shaped volume pattern observed in Figure 6.

It is also important to note that you could only with difficulty attribute the U-shaped volume pattern found in closed-end funds to any sort of informational asymmetry among traders, as the nature of closed-end funds makes them extremely transparent to all market participants. The driving factor appears to be the distance from the mean-reverting discount level. Thus, models such as those 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 7.¹⁹ I constructed this figure by placing the monthly discount and premium observations into bins and then using CRSP ask and bid prices to construct the ratio of spread to ask price for each observation—that is, $(\text{ask} - \text{bid})/\text{ask}$ for each observation.²⁰ I then averaged all the ratios in each bin and plotted them in Figure 7.

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, $(\text{ask}-\text{bid})/\text{ask}$, 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, as in as in the models of Kyle (1985), Easley and O'Hara (1992), and Glosten and Milgrom (1985). These models all feature market makers who adjust their bid and ask prices in response to order flow. Since the market makers maintain their own inventories, they are worried about being taken advantage of by informed investors. As a result, the greater the potential informational asymmetry, the larger market makers will make their spreads.

This explanation doesn't hold up well for closed-end funds where market makers are just as well informed about fundamentals—fund NAVs—as any other market participants (see Neal and Wheatley (1995).) So something else must be causing market makers to widen their spreads as fund prices move further and further 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).²¹

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 that lies at the center of the discount and premium distribution of Figure 2. I believe that this increasing risk explains the cross-sectional pattern of excess returns observed in Tables IV and VI.

VI. 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 existence of such a risk factor would explain the cross-sectional distribution of excess returns found for arbitrage portfolios when using only the Fama-French factors to adjust for risk. The standard risk factors cannot account for these excess returns, which increase in magnitude the farther you move away from the center of the discount and premium distribution. Given that noise-trader risk also appears to increase as you move away from the center of the discount and premium distribution, this paper argues that the observed excess returns serve to compensate arbitrageurs for bearing noise-trader risk.

The paper begins by using Fund Edge data on US and Canadian funds to test and reject two hypotheses about how arbitrage activities may constrain closed-end fund discounts and premia. First,

I reject both the upper and lower static arbitrage bounds proposed by Gemmill and Thomas (2002). Contrary to their model, the data is not asymmetric due to censoring, does not have a mode at par, does not show skewness toward discounts, and does not show low kurtosis due to the tails of the discount distribution being censored. Rather, the discount distribution shows no signs of censoring, has a mode discount at approximately the value consistent with capitalizing out future management fees, is skewed toward premia rather than discounts, and shows excess kurtosis rather than leptokurtosis.

Second, closed-end fund discounts do not appear to be constrained by interest-sensitive arbitrage bounds as suggested by Pontiff (1996). In particular, average absolute discounts are not significantly positively correlated with short-term interest rates as they would be if arbitrage activities slackened when rising interest rates increased the opportunity costs faced by risk-averse arbitrageurs.

In addition, contrary to Gemmill and Thomas (2002), long-run average discount and premium levels across funds are not cross-sectionally correlated with fundamental factors like expense ratios and dividend yields. Rather, the only variable that is significant for both bond and stock funds is a measure of noise-trader risk. The higher the noise-trader risk, the larger the discount, consistent with the formal model of DeLong, Shleifer, Summers, and Waldmann (1990).

The last result suggest that noise-trader risk not only exists, but that it directly impacts discount and premium levels. What is less clear is whether and how this potential risk factor affects arbitrage in closed-end funds and whether or not it is a priced risk factor.

The final section of this paper explores these possibilities by examining a cross-section of closed-end fund arbitrage returns. Twenty portfolios were 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 are in fact compensation for increasing levels of noise-trader risk as you move toward the tails of the distribution. If this risk could be 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 in the regressions run using the standard risk factors (which fail to capture noise-trader risk.)

As additional evidence in favor of noise-trader risk increasing as you move away from the center of the discount and premium distribution, the paper closes with 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 volume and spreads 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-shape for volume is consistent with the idea that noise-trader risk increases as you move away from the mean-reverting discount level. But the U-shape for spreads is even more indicative of increasing noise-trader risk because spreads are set by risk-averse market makers. Given that closed-end fund markets offer no obvious informational asymmetries, the best explanation for the widening of spreads is that they are the rational, 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 can explain the cross sectional return pattern. 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 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

- Black, Fisher, 1986, Noise, *Journal of Finance* 41, 529–543.
- Boudreaux, K. J., 1973, Discounts and Premiums on Closed-end Mutual Funds: A study of valuation, *Journal of Finance* 28, 515–522.
- 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.
- Chen, Nai-Fu, Raymond Kan, and Merton H. Miller, 1993, Are Discounts on Closed-end Funds a Sentiment Index?, *Journal of Finance* 48, 795–800.
- 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.
- D’Avolio, Gene, 2002, The Market For Borrowing Stock, *Journal of Financial Economics* 66, 271–306.
- 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, 1999a, The Closed-end Fund Discount and Performance Persistence, *Working Paper, London Business School*.
- Dimson, Elroy, and Carolina Minio-Kozerski, 1999b, Closed-end Funds: A survey, *Financial Markets, Institutions, and Instruments* 8.
- 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.
- Gemmill, Gordon, and Dylan C. Thomas, 2002, Noise-Trading, Costly Arbitrage, and Asset Prices: Evidence from Closed End Funds, *Journal of Finance* 57, 2571–2594.
- 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, 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.
- Levis, Mario, and Dylan C. Thomas, 1995, Investment Trust IPOs: Issuing Behavior and Price Performance. Evidence from the London Stock Exchange., *Journal of Banking and Finance* 19, 1437–1458.
- 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.
- Pontiff, Jeffrey, 1996, Costly Arbitrage: Evidence from Closed-end Funds, *Quarterly Journal of Economics* 111.
- Pontiff, Jeffrey, 1997, Excess Volatility and Closed-end Funds, *American Economic Review* 87, 155–169.
- Ross, Stephen A., 2002, Neoclassical Finance, Alternative Finance, and the Closed End Fund Puzzle, *European Financial Management* 8, 129–137.
- 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.
- Zweig, Martin E., 1973, An Investor Expectations Stock Price Predictive Model Using Closed-end Fund Premiums, *The Journal of Finance* 28, 67–78.

VII. 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.

²In 1985, there were fewer than 30 funds listed in the Wall Street Journal.

³The shares sold short are apparently borrowed from individual investors' accounts because D'Avolio (2002) finds that the large institutional share clearing houses (why normally do most share lending to short sellers) don't typically have any closed-end fund shares available. This is because the clearing-houses get their stocks of lendable shares from large institutional holders like pension funds, which don't typically invest in closed-end funds in the USA.

⁴Most of these 761 outliers lie in the left tail. The most extreme discount was 66.5% while the most extreme premium was the Taiwan Fund's -205.4%.

⁵The averages across funds are virtually identical if they are taken across all funds in the sample rather than just those just those in business for at least 60 months.

⁶The measure of the exposure of a fund to noise-trader risk is how highly correlated the fund's own discount is with the average discount series. For bond funds, these are betas when regressing the fund's own discount against the average discount on bond funds, and for stock funds, these are the betas when regressing the fund's own discount against the average discount on stock funds. The variable that captures the difficulty of mimicking a fund's portfolio is the log of the residual standard error obtained by regressing a fund's monthly NAV returns on those of 11 different open-end funds, following the method of Pontiff (1996). These open-end funds were chosen because they have no loads and because they have different portfolio compositions which would give them different risk exposures. The 11 funds that I used were the T. Rowe Price International Bond Fund, Vanguard Short-term Corporate Bond Fund, Vanguard High-yield Corporate Bond Fund, Vanguard GNMA Bond Fund, Vanguard Intermediate Term Tax Exempt Bond Fund, Vanguard Index 500 Fund, Vanguard International Growth Fund, Vanguard International Value Fund, Vanguard Equity Income Fund, Vanguard Windsor II Fund, and the T. Rowe Price New Horizon Fund. The regressions results of Table III are only slightly differ-

ent if you instead use log residual standard errors that are obtained by regressing fund NAV returns on large market indices, as was done by Gemmill and Thomas (2002).

⁷The initial premiums at which funds trade are necessary to raise the money needed to pay the investment bankers who float fund IPOs.

⁸This is similar to the cross-sectioning strategy used by Fama and French (1992) to show that firm size and value orientation (price-to-book ratio) are priced risk factors.

⁹You can see this formally by recalling the definition that $P_t = (1 + D_t)N_t$, for all t . Using this definition to construct log returns, you can see that $\log\left(\frac{P_{t+1}}{P_t}\right) - \log\left(\frac{N_{t+1}}{N_t}\right) = \log\left(\frac{1+D_{t+1}}{1+D_t}\right)$. This means that the rate of return to going long the shares of the fund while shorting the underlying is equal to the rate of change of $1 + D_t$.

¹⁰I tried several other mean reversion processes, but the AR(1) works best.

¹¹The mean reversion of closed-end fund discounts and premia, and whether this mean reversion could generate excess returns, was first studied by Zweig (1973) and Thompson (1978).

¹²When a fund enters the data set (IPO for the vast majority, pre-existence for those in the data set in January 1985), I create for it two separate mock positions to track returns. For the first, I take a \$100 position in the fund's underlying by purchasing pseudo-shares of the underlying at the initial NAV. For the other, I invest \$100 by buying shares of the fund at the spot price on the same day as the initial NAV. I then follow the values of each of these positions over time. I did the following to account for dividends. I re-invested dividends using the price and NAV on first day in the data set (after a given dividend's payment date) on which both share price and NAV were available. This was typically the Friday immediately following the payment date. The dividend value was then used to buy more fund shares at that day's share price (in order to keep track of returns to owning fund shares) and to buy more "shares" of the underlying portfolio at the given NAV value (in order to keep track of returns to mimicking the fund's underlying portfolio.)

¹³This is consistent with Malkiel (1977), Pontiff (1995), and Dimson and Marsh (1999a) 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.

¹⁴Remember that for funds trading at $D_t < \bar{D}$, an arbitrageur would want to go short the fund and long the underlying portfolio, which is the opposite of what is graphed in the figures since, for simplicity, I graphed share returns less NAV returns for all D_t levels. So when interpreting Figure 5, you should remember that the returns would be positive on the left-side of the graph for anyone actually engaged in arbitrage and taking the reverse position for those D_t levels.

¹⁵This approach differs from that used by (?), (?) and others whereby they run regression on the excess returns to holding fund shares, rather than on the difference between the returns to fund shares and their portfolios. While that method can clearly demonstrate that discount/premium levels generate excess returns, it cannot isolate the return due solely to mean reversion. Subtracting NAV returns from share returns does this.

¹⁶January through June of 1985 are not used in this section for the following reason. Weiss (1989) and Peavy (1990) both find that closed-end funds begin trading immediately after their IPOs at premia of about eight to 10 percent. These premia are necessary to pay investment bankers for floating the IPOs. They also find that such premia typically dissipate within six months of the IPO, with funds typically moving to discounts of around six to eight percent. Since this movement in D_t levels is predictable, I do not want it to bias my results. So I have thrown out the first six months of data on each fund that enters the data set by IPO after January 1985. And for good measure, I also threw out the first six months of all funds in existence in January 1985 as well. That's why the time series begins in July 1985. But please note that the results presented below are not, in fact, sensitive to the inclusion of the fund months that I threw out.

¹⁷This 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.

¹⁸I 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 .

¹⁹Pontiff (1995) also found a U-shaped pattern of spreads the 1965-1985 data set of US funds first used by Lee, Shleifer, and Thaler (1991). See the rightmost column of Pontiff's Table 3.

²⁰Since I have the dates for each fund's monthly D_t observation (the last Friday of each month unless there's a holiday), I was able to use the CRSP daily file to find the fund's bid and ask price on the same date.

²¹These results also tend to 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.

Figure 1. Censored Log Normal Distribution of Discounts Assumed by GT

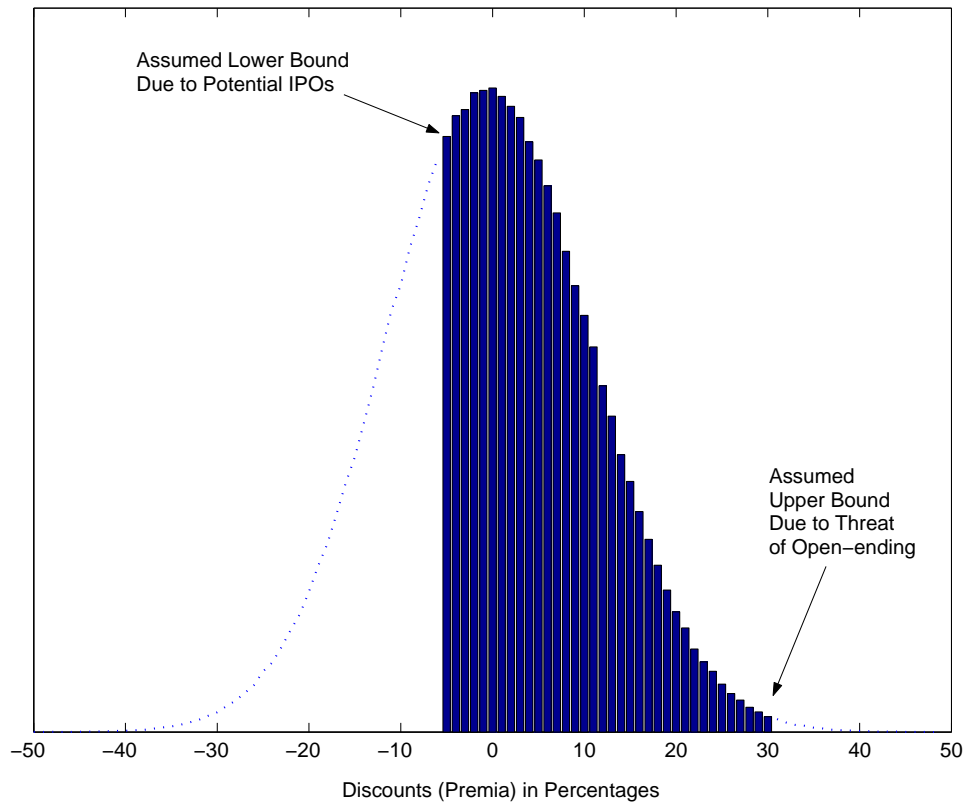


Table I
Time Series Interest Rate Regressions, Reproduction of Pontiff (1996) Table III

Time series estimation with mean cross-sectional absolute discount and mean discount as dependent variables, and Treasury-bill yield, changes in Treasury-bill yields, and lagged dependent variables as independent variables (*t*-statistics are in parentheses).

	(1)	(2)	(3)	(4)
Dependent variable	Level of avg $ D_t $	Change in avg $ D_t $	Level of avg D_t	Change in avg D_t
Number of Monthly observations	246	245	246	245
Intercept	14.74 (7.64)	1.02 (1.97)	10.41 (1.87)	-0.02 (-0.20)
Level of 1-yr Treas	0.46 (3.01)		-0.22 (-1.21)	
Change in 1-yr Treas		0.49 (3.14)		-0.21 (-1.11)
Autoregressive Error parameters				
First Order	0.78 (12.03)	-0.20 (-3.08)	0.79 (12.17)	-0.20 (-3.14)
Second Order	0.03 (0.35)	-0.17 (-2.58)	0.06 (0.68)	-0.16 (-2.37)
Third Order	0.14 (1.69)	-0.07 (-1.01)	0.12 (1.53)	-0.04 (-0.60)
Fourth Order	-0.02 (-0.35)	-0.08 (-1.23)	0.08 (0.12)	-0.03 (-0.47)
R^2	0.84	0.11	0.93	0.05
Durbin-Watson	1.96	2.01	2.00	2.00

Level of average absolute discount is the monthly average of the absolute value of the log ratio of fund price to fund net asset value, expressed in percentage terms. Level of average discount is the monthly average log ratio of fund net asset value to fund price, expressed in percentage terms. Level of one-month T-bill yield is the one-month T-bill yield, expressed as an annualized percent. A fourth-order moving average process is used to model autocorrelated residual behavior. The model can be expressed as $y_t = x_t' \beta + v_t$, where y are the dependent values, x_t is a vector of independent values, and β is the vector of slope coefficients. $v_t = e_t + \delta_1 v_{t-1} + \delta_2 v_{t-2} + \delta_3 v_{t-3} + \delta_4 v_{t-4}$, where δ_i is the i th-order autoregressive parameter and e_t is assumed to be independently distributed.

Table II
Time Series Interest Rate Regressions on Fund Edge Data 1985-2000

Time series estimation with average cross-sectional absolute discount, $|D_t|$, and mean discount, D_t , as dependent variables, and one-month Treasury bill yields and changes in yields as dependent variables, using the same AR(4) model for autocorrelated residual behavior as used by Pontiff (1996). (*t*-statistics are in parentheses.)

Dependent variable	(1) Level of avg $ D_t $	(2) Change in avg $ D_t $	(3) Level of avg D_t	(4) Change in avg D_t
Number of Monthly observations	182	181	182	181
Intercept	8.38 (4.07)	0.03 (0.63)	-6.43 (-1.58)	0.05 (0.63)
Level of 1-yr Treas	0.32 (1.10)		2.20 (4.53)	
Change in 1-yr Treas		0.26 (0.92)		2.29 (4.90)
Autoregressive Error parameters				
First Order	0.71 (9.33)	-0.28 (-3.66)	0.80 (10.77)	-0.20 (-2.61)
Second Order	0.15 (1.64)	-0.12 (-1.56)	0.06 (0.67)	-0.12 (-1.60)
Third Order	0.13 (1.38)	0.01 (0.15)	-0.08 (-0.81)	-0.20 (-2.64)
Fourth Order	-0.04 (-0.52)	-0.03 (-0.41)	0.17 (2.24)	-0.07 (-0.94)
R^2	0.87	0.08	0.82	0.14
Durbin-Watson	1.98	1.98	2.02	2.00

The regressions reported above were set up to be directly comparable to Table III of Pontiff (1996). In particular, I utilize the same specification for modelling autocorrelated residual behavior. Also consistent with Pontiff, I use the secondary market yields on one-month T-bills, construct discounts as the log of the ratio of fund NAV divided by price, and when constructing the average log discount for each month, I include only funds that had been in business a least six months. Because the Federal Reserve terminated the construction of the time series giving one-month T-bill yields in June 2000, these regressions span the period January 1985 through June 2000.

Table III
Comparison of GT and Fund Edge Results from Cross-Sectional Regressions to Explain Which Factors Explain the Discount

The table reports cross-sectional regressions for closed-end funds. Data are averaged for each fund over the relevant sample period. The first column reproduces column (1) of Table IV of Gemmill and Thomas (2002), which reports the results of a cross sectional regression run on UK-traded closed-end stock funds over the period 1991-1997. The second and third columns report the results of running the same specification on, respectively, bond and stock funds traded in the USA over the period 1991-2000. The discount is measured as (net asset value less share price)/(net asset value). The expense ratio is annual expenses divided by net asset value. The individual fund noise beta is the individual fund sensitivity to the average discount of the funds in the sample; the replication risk is the residual error from a regression of net asset value returns on a wide variety of no-load mutual funds. Numbers in parentheses are t-values. The symbol * denotes significance at the five percent level and ** denotes significance at the 1 percent level.

Dependent Variables	GT's UK Stock Fund Avg. Discs (1)	Fund Edge Stock Fund Avg. Discs (2)	Fund Edge Bond Fund Avg. Discs (3)
Independent Variables			
Constant	+0.049 (0.81)	+0.172 (0.48)	-0.187 (-1.67)
Expense Ratio	+2.992** (2.98)	+2.627 (1.01)	-0.085 (-1.31)
Noise Risk Beta	-0.029** (5.03)	+0.020* (2.58)	+0.031** (3.91)
Log of Age	+0.040** (8.77)	+0.064* (2.06)	+0.018 (1.21)
Replication Risk	+0.087** (4.16)	-0.020 (-0.77)	+0.0129 (0.92)
Dividend Yield	-0.0073** (-3.60)	+0.0059 (1.61)	-0.0071** (-2.52)
Log of Size	-0.011* (-2.11)	-0.019 (-1.05)	+0.009 (1.81)
R-sq (weighted)	0.52	0.06	0.42
R-sq (unweighted)	0.34	-0.29	0.29
Weighting Variable	Volatility of Discount	Volatility of Discount	Volatility of Discount
Number of funds	168	69	123

Table IV

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

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

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

This table reports the results of expanding the Fama-French regression results of Table IV by adding changes in the capital-weighted average discount across all funds as an additional explanatory variable. The notes to Table IV explain the three Fama-French factors while the notes to Table V explain the capital-weighted average discount series, W_t . Standard errors were 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 (-1.21)	0.68 (1.27)	0.74 (0.57)	0.75 (0.66)	0.53 (0.24)	0.26	-3.93	8.41	16
-45	-40	-5.19 (-3.73)	0.58 (1.41)	-0.20 (-0.35)	-0.05 (-0.1)	1.08 (0.72)	0.13	-4.95	8.09	42
-40	-35	-2.76 (-2.10)	-0.13 (-0.48)	0.26 (0.31)	0.38 (0.78)	-0.16 (-0.15)	0.03	-2.74	8.97	44
-35	-30	-4.17 (-3.92)	-0.09 (-0.37)	0.18 (0.60)	0.26 (0.70)	0.09 (0.10)	0.02	-4.06	8.08	65
-30	-25	-3.22 (-2.94)	0.11 (0.57)	0.65 (1.79)	0.35 (1.01)	1.39 (1.53)	0.19	-3.18	7.78	76
-25	-20	-1.63 (-2.02)	0.34 (1.71)	0.11 (0.49)	0.02 (0.09)	1.18 (1.94)	0.08	-1.51	8.01	100
-20	-15	-1.31 (-4.65)	0.44 (2.21)	0.20 (1.73)	0.37 (1.94)	-0.11 (-0.14)	0.11	-0.93	5.67	138
-15	-10	-1.07 (-5.73)	0.06 (1.82)	-0.06 (-1.06)	0.01 (0.14)	0.93 (5.06)	0.29	-1.10	2.53	170
-10	-5	-0.85 (-6.27)	0.08 (2.69)	0.01 (0.28)	0.14 (2.71)	0.67 (6.74)	0.40	-0.82	1.89	177
-5	0	-0.34 (-4.51)	0.01 (0.52)	0.02 (1.02)	0.06 (.00)	0.84 (10.58)	0.61	-0.39	1.60	178
0	5	-0.07 (-0.88)	0.05 (1.73)	0.01 (0.57)	0.06 (2.31)	0.77 (9.58)	0.60	-0.08	1.52	178
5	10	0.36 (3.91)	-0.01 (-0.16)	0.03 (1.31)	0.01 (0.65)	0.91 (10.26)	0.57	0.30	1.70	178
10	15	0.75 (5.12)	0.02 (0.82)	0.08 (2.51)	0.04 (1.53)	0.75 (5.35)	0.34	0.74	1.97	174
15	20	0.91 (4.73)	0.05 (1.29)	0.07 (1.32)	0.01 (0.18)	0.88 (6.77)	0.25	0.93	2.65	164
20	25	1.83 (6.54)	0.16 (1.53)	0.02 (0.01)	0.08 (1.17)	0.45 (2.03)	0.09	1.98	3.54	147
25	30	2.12 (4.04)	0.44 (2.41)	0.01 (0.04)	0.09 (0.90)	0.86 (1.55)	0.15	2.60	5.92	99
30	35	2.64 (2.82)	0.47 (2.74)	-0.23 (-1.05)	-0.05 (-0.38)	1.08 (1.78)	0.14	3.47	7.18	56
35	40	1.3 (1.57)	0.86 (4.54)	0.32 (1.03)	0.37 (1.48)	0.17 (0.38)	0.43	1.56	4.72	23
40	45	8.17 (2.00)	0.90 (1.57)	-0.99 (-1.23)	0.03 (0.12)	0.08 (0.03)	0.43	6.28	6.88	9
45	50	Insufficient	Observations							

Figure 2. Distribution of Weekly Discounts in Fund Edge, 1985-2001

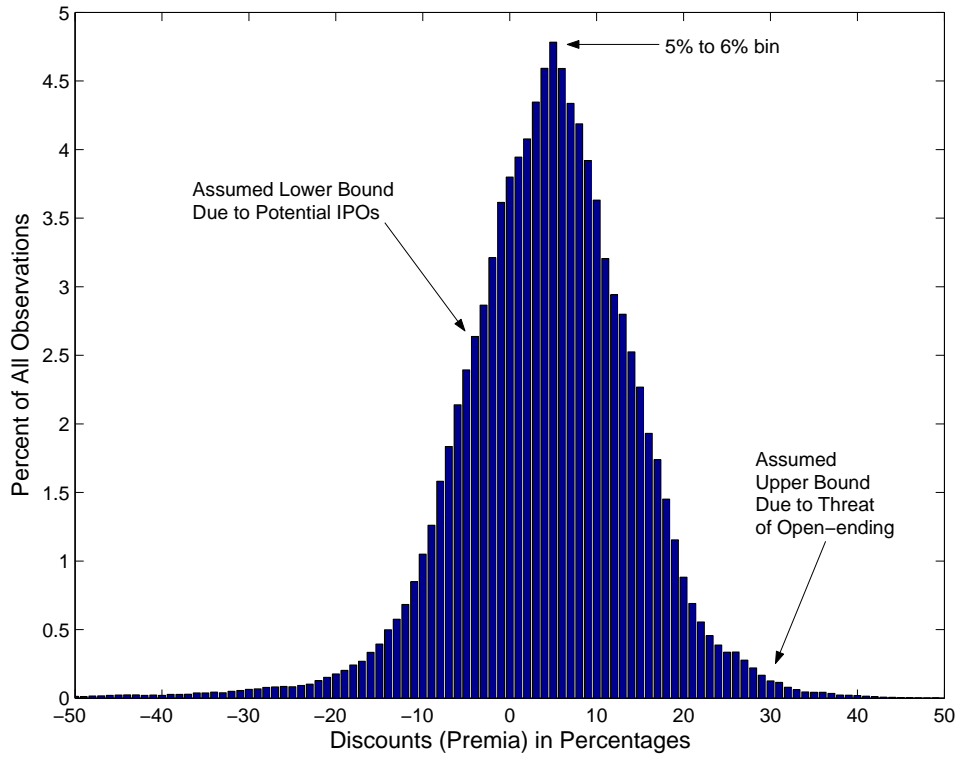


Figure 3. Average one-month share and NAV returns when sorting each observation by its discount or premium level the previous month.

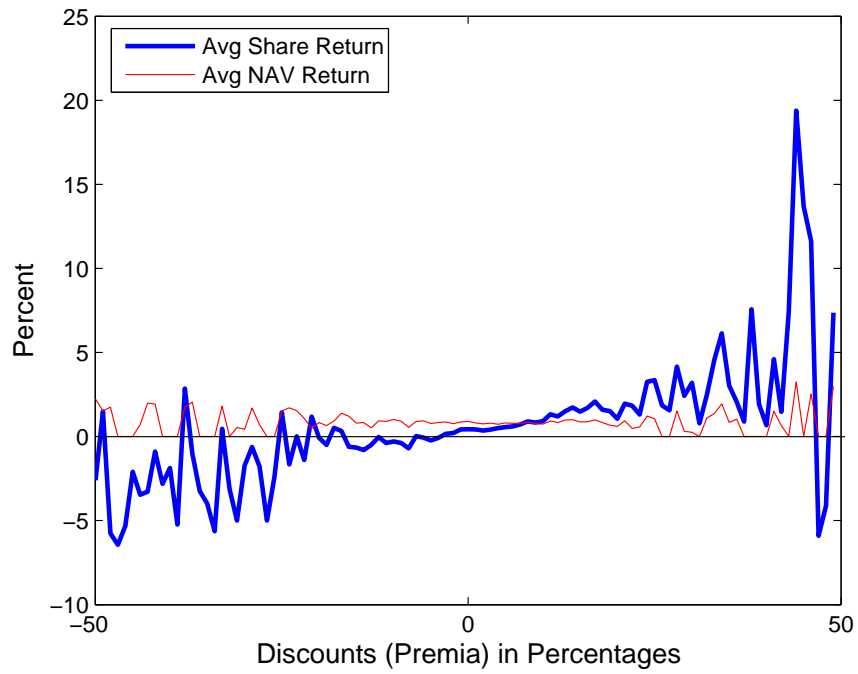


Figure 4. Average one-month share less NAV returns and standard deviations when sorting each observation by its discount or premium level the previous month.

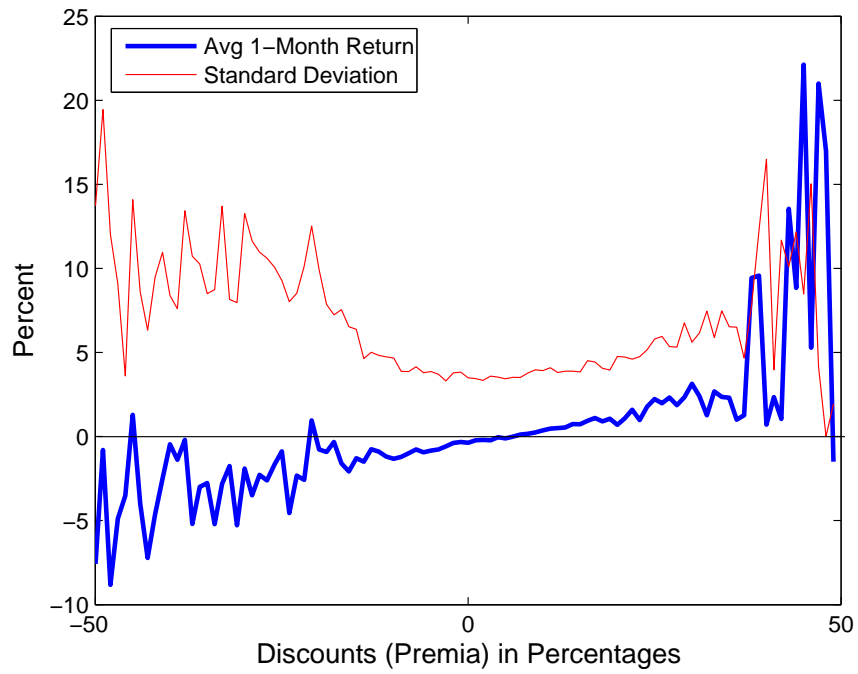


Figure 5. Ratio of the average one-month share less NAV returns to the standard deviation of those returns (i.e., the reward-to-risk ratio.)

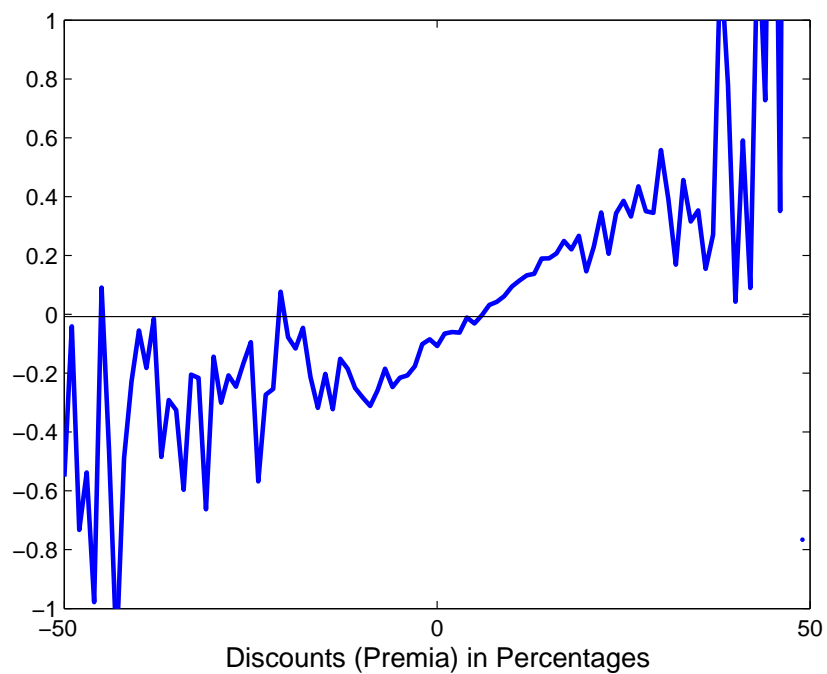


Figure 6. Average ratio of monthly trading volume to shares outstanding when sorting by discount or premium level, 1985 to 2001.

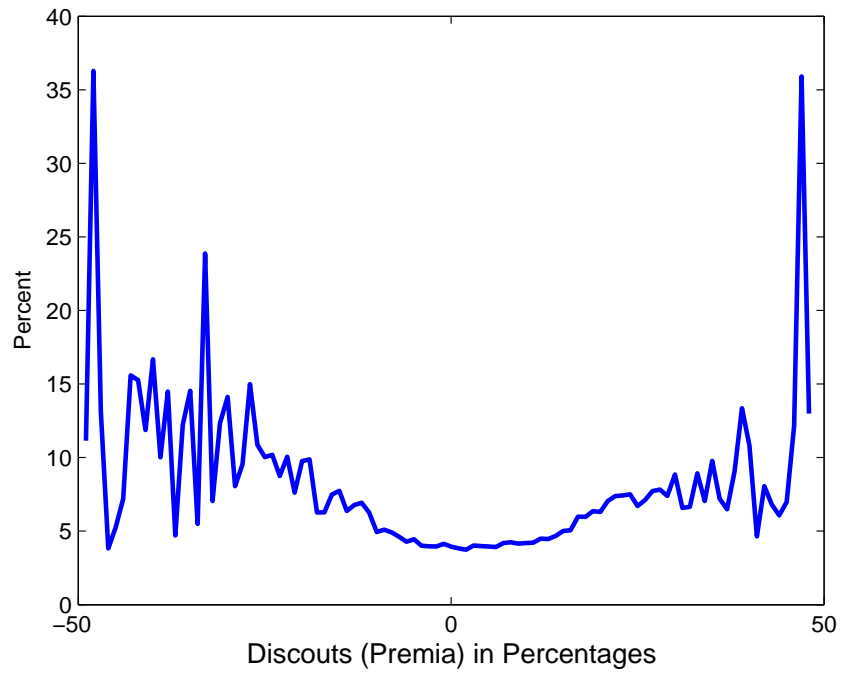


Figure 7. Average ratio of bid-ask spread to ask price when sorting by discount or premium level, 1985 to 2001.

