

8-20-2006

Arbitrage in closed-end funds: New evidence

Sean M. Flynn

Vassar College, flynn@vassar.edu

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

Citation Information

Flynn, Sean M., "Arbitrage in closed-end funds: New evidence" (2006). *Faculty Research and Reports*. 81.
https://digitalwindow.vassar.edu/faculty_research_reports/81

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

Arbitrage in Closed-end Funds: New Evidence

Sean Masaki Flynn*

August 20, 2006

NOTE: *This paper contains a subtle but crucial data analysis error. As a result, its conclusions are not only incorrect but, in fact, diametrically opposed to what is actually true. The correct analysis has been incorporated into Vassar College Department of Economics Working Paper Number 69, which is entitled, "Noise-trader Risk: Does it Deter Arbitrage, and Is it Priced?" This paper is available as a pdf file at, <http://irving.vassar.edu/VCEWP/VCEWP69.pdf>. As a matter of intellectual honesty, however, I am leaving this paper posted here as a historical document. But, again, its conclusions are in error! The error is described in detail in footnote number 8 of "Noise-trader Risk: Does it Deter Arbitrage, and Is it Priced?"*

ABSTRACT

Arbitrage pressures that could equalize closed-end fund share prices with fund portfolio values appear to be largely absent in an extensive data set. Observed fund behavior violates the static arbitrage bounds of Gemmill & Thomas (2002) and is inconsistent with the dynamic arbitrage bounds of Pontiff (1996). Furthermore, Fama & French (1992) regressions run on arbitrage portfolios designed to profit from closed-end fund mispricings generate excess returns that are either significantly negative or insignificantly different from zero, suggesting that arbitrageurs lack a profit incentive. If arbitrage is absent, observed fund pricing behavior likely reflects changing investor sentiment about fund prospects.

*Department of Economics, Vassar College, 124 Raymond Ave. #424, Poughkeepsie, NY 12604. flynn@vassar.edu *Mea culpa.*

Using a data set that contains nearly every closed-end fund trading in the United States and Canada over the period 1985-2001, this paper finds strong evidence that arbitrage pressures are weak to non-existent in closed-end funds. The 462 closed-end funds exhibit behavior that stands strongly in violation of the static arbitrage bounds said to constrain fund discount and premium levels by Gemmill & Thomas (2002). In addition, discounts and premia also violate the dynamic arbitrage bounds that Pontiff (1996) argues vary with interest rates and the opportunity costs of undertaking arbitrage activities. Most importantly, if you estimate Fama & French (1992) three-factor regressions on the excess returns of portfolios that try to profit from the mispricings of closed-end funds, you find negative risk-adjusted returns in virtually all cases, suggesting that rational traders lack a profit incentive. This would explain why, empirically, the mispricings found in closed-end funds often persist seemingly indefinitely and why they appear to move largely without reference to fundamental factors. Simply put, with rational traders absent, the market is dominated by sentimental noise-traders who move prices unpredictably and without reference to fundamentals, as suggested by Lee, Shleifer & Thaler (1991) and Chopra, Lee, Shleifer & Thaler (1993).

Closed-end funds—mutual funds whose shares trade like stock—have presented a puzzle to financial theorists for decades because they constantly violate the law of one price. In the United States, closed-end funds are required by law to publish the contents of their portfolios weekly. In addition, many funds now present daily updates about their portfolio values on their websites. Consequently, investors have very precise and timely information about fund portfolio values. Given that closed-end funds issue only a fixed number of shares, this should imply that, fund expenses aside, the price of each share should simply be the total portfolio value divided by the number of shares outstanding. Unfortunately, closed-end fund share prices often differ hugely from portfolio values. What is more, these mispricings often last for years at a time, so that there appears to be little pressure for prices to return to par with portfolio values.

Another empirical regularity is that funds tend to trade at discounts rather than premia. Various authors have presented various hypotheses to explain this tendency. Explanations include accrued capital gains liabilities, portfolio illiquidity, poor managerial performance, and management fees. (See Dimson & Marsh-Petersen (1999) for a comprehensive survey.) The first two explanations are unconvincing because on those occasions when closed-end funds announce that they are going to liquidate, share prices jump up toward reported portfolio values. If the discounts had been due to misreported portfolio values, we would not see such jumps. That leaves managerial performance and management fees. But while each of these can explain why funds tend to trade at discounts rather than premia, because each of these variables is rather stable over time, neither can explain why discounts and premia so often change with such stunning rapidity. They are also both incapable of explaining why funds so often trade at substantial premia.

This paper presents evidence that arbitrage activities in closed-end funds are unprofitable. Because closed-end fund price movements appear to be largely random and uncorrelated with fundamentals, this result is consistent with the idea presented by DeLong, Shleifer, Summers & Waldmann (1990) that noise-traders “create their own space.” Under this hypothesis, arbitrageurs are unwilling to set up arbitrage positions that could mitigate violations of the law of one price because they are afraid that irrational noise-traders may cause mispricings to widen further. In their model, the volatility caused by noise traders is an independent source of risk that deters risk-averse arbitrageurs, with the result that the market can be dominated by noise-traders free to move prices about without regard to fundamentals.

This paper divides the evidence against arbitrage into three parts: evidence against static arbitrage bounds, evidence against dynamic arbitrage bounds, and evidence against the profitability of arbitrage portfolios designed to profit from divergences between fund share prices and fund portfolio values.

In a recent paper in this journal, Gemmill & Thomas (2002) use data from a sample of 158 closed-end funds trading in the UK to provide support for their contention that closed-end fund discounts and premia are constrained by static arbitrage bounds. Gemmill & Thomas (2002) argue that discounts

should be no more than thirty percent and premia should be no more than five percent. These constraints imply that the empirical distribution of discounts and premia should be platykurtic and skewed toward discounts rather than premia. Using a subsample of 20 stock funds, they find some support for this empirical prediction.

The prediction is not, however, consistent with the behavior of closed-end funds trading in the United States and Canada. For both bond funds and stock funds, empirical discount distributions not only routinely violate the two Gemmill & Thomas (2002) boundaries, but are also skewed toward premia rather than discounts and are leptokurtic rather than platykurtic.

The behavior of American and Canadian closed-end funds is also inconsistent with the regression results that Gemmill & Thomas (2002) obtain for their sample of UK-traded closed-end stock funds. These regressions suggest that arbitrage pressures will, in the long run, tend to drive discounts toward levels consistent with fundamentals such as management fee rates and dividend payout rates. Their regressions also suggest that the level of the discount is inversely related to the difficulty that arbitrageurs have in hedging their closed-end fund positions and that noise-trader risk does not affect discount levels.

Using the same regression methodology used by Gemmill & Thomas (2002) but applying it to US and Canadian data produces very different results: Management fees are *not* significantly related to discount levels for either bond or stock funds; dividend payout rates only affect bond funds; the difficulty that arbitrageurs have in constructing hedge portfolios is *insignificantly* related to discounts; and noise-trader risk *does* significantly increase discount levels.

The Fund Edge data also contradicts the evidence in favor of dynamic arbitrage bounds presented by Pontiff (1996). Pontiff (1996) found lower interest rates to be associated with less extreme deviations of share prices from portfolio values, and took this as evidence that the cost of arbitrage determined the intensity of arbitrage and therefore of the magnitude of the mispricings found among close-end funds.

In particular, as interest rates fall, so do the opportunity costs of collateral invested in short positions. As interest rates fall, arbitrage becomes more intense, constraining the magnitude of mispricings.

Pontiff (1996) tests to see whether the magnitude of the absolute value of discounts and premia varies positively with interest rates. (Discounts are given as positive numbers and premia as negative numbers.) He does this because the transaction costs of arbitrage should increase with interest rates without regard to whether an arbitrageur is shorting a fund trading at a premium while going long its underlying portfolio, or going long a fund trading at a discount while shorting its underlying portfolio. In either case, the higher the interest rates, the less arbitrage will be undertaken. Because it is presumed that discounts and premia will be more extreme without arbitrage pressures, their absolute values should increase when interest rates rise.

Using a sample of 68 closed-end funds trading in the United States over the period 1965-1985, Pontiff finds that interest rates are inversely related to the absolute value of discounts and that the effect is both economically large and highly statistically significant. But if we apply the same regression methodology over the period 1985-2001, the relationship is much smaller in magnitude and not at all statistically significant. The behavior of closed-end fund discounts since 1985 does not, therefore, appear consistent with this measure of dynamic arbitrage pressure.

I next examine the returns to arbitrage in closed-end funds by running Fama French (1992) three-factor regressions on the returns of portfolios designed to profit from closed-end fund mispricings. Using monthly data covering 1985-2001, I place each fund each month into a bin on the basis of its discount or premium level. There are twenty bins, each five percentage points wide, and ranging from a discount of fifty percent to a premium of fifty percent. For example, one bin is for discounts between fifteen and twenty percent and another is for premia between five and ten percent. In any given month, the returns for the month to both share price and portfolio value are calculated for all the funds falling into each given bin. For the discount bins, we then subtract portfolio returns from share price returns, as for funds trading at deep discounts an arbitrageur would go long the shares of the fund and short

the underlying. On the other hand, for funds trading at premia, we subtract share price returns from portfolio returns, as in such cases arbitrageurs would short the fund and go long the underlying. We take the average of the appropriate difference for all funds in each bin each month. Doing this for each month gives us twenty time series of the returns to arbitrage starting at various levels of initial mispricing.

When we run Fama & French (1992) regressions on the excess returns to these twenty portfolios, we find that in only one case (the bin holding premia between twenty-five and thirty percent) is there a significantly positive alpha. In all other cases, the alphas are either statistically indistinguishable from zero or significantly negative. Since only one-third of one percent of all discount/premium observations fall into the bin whose return series generates a positive alpha, it seems fair to conclude that rational traders would not set up arbitrage portfolios in closed-end funds if they hoped to earn positive risk-adjusted excess returns. This conclusion is indeed conservative, as the returns underlying the regressions only take into account bid-ask spreads. Had we also taken account of trading costs and the collateral costs of short positions emphasized by Pontiff (1996), returns would have been even worse.

That closed-end fund arbitrage returns are zero or negative greatly helps to explain closed-end fund pricing behavior. In particular, if rational traders are absent, then the closed-end fund market would be dominated by noise traders who would cause discounts and premia to shift unpredictably. However, the empirical distribution of discounts/premia is concentrated upon a discount of about six percent. This value is consistent with capitalizing out management fees. As such, one is led to the suspicion that while sentimental noise traders may cause discounts and premia to vacillate wildly, they are aware of the correct rational level around which those discounts and premia should vacillate.

This suspicion is strengthened by another crucial difference between Pontiff's data and the data set examined here. While there is a negative but statistically insignificant relationship between interest rates and discount levels over the period 1965-1985, there is a positive and highly statistically signif-

icant relationship between those two variables over the period 1985-2001. This bears crucially upon whether management fees can explain the tendency of closed-end funds to trade at discounts.

Pontiff (1996), and Lee et al. (1991) both rejected the hypothesis that the tendency of funds to trade at discounts was caused by their management fees and trading costs. They reached this conclusion because in their data (they both used the same data) discounts were not sensitive to the level of interest rates, whereas one would expect the capitalized value of future management fees to vary with interest rates. The finding of this paper that there is a very strong positive relationship between interest rates and discount levels suggests that we should reconsider the effects of management fees on discount levels.

Consequently, the final section of this paper presents a simple behavioral model that can explain not only why discounts should be expected to be positively related to interest rates, but also why the empirical distribution of discounts is centered at a discount of about six percent, and why it has the shape that it does. Importantly, the model can explain these facts without appealing to arbitrage.

Section I describes the data set and gives the formulae used to define discounts and premia. Section II tests and rejects the static arbitrage bounds proposed by Gemmill & Thomas (2002). Section III finds no support for the dynamic arbitrage bounds of Pontiff (1996). Section IV uses Gemmill and Thomas' (2002) methodology to test and reject their hypothesis that arbitrage pressures will in the long run tend to move fund prices toward rational levels. Section V demonstrates that the risk-adjusted returns to arbitrage positions in closed-end funds are either significantly negative or insignificantly different from zero for over 99.6% of observations. Section VI presents the behavioral model of discounts. Section VII 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, portfolio 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 the number of funds has exploded in recent years (there were fewer than 30 funds listed in the Wall Street Journal in 1985), and funds only go out of business very infrequently, survivor bias should not be significant. Below, I will only utilize Fund Edge data covering 1985-2001.² Concentrating on more recent years should reduce problems with survival bias.

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 & Thomas (2002). The choice of definition will depend on whether we are comparing Fund Edge results with those from Pontiff (1996) or Gemmill & 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 generally multiply D_t by 100 and refer to discounts and premia in percentages.

¹With the exception of one fund, all the funds traded on major exchanges, either the NYSE, AMEX, NASDAQ, or Toronto Exchanges. The one exception is the NAIC Growth Fund, which trades on the Midwest Exchange.

²For all closed-end stock funds trading in the USA over the period 1960-1999, only an average of 3.64% of those in business on January 1st of a given year went out of business (through liquidation or conversion to an open-end fund) by December 31st of the same year. This implies a half-life for a fund of around 20 years, given that you cannot reject a Bernouilli death process for funds. See footnote 8.

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 which guarantee to redeem shares at par with portfolio value, closed-end funds do not engage in redemptions of any sort. Because of this, 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 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 contents of their portfolios at least once per week, and some funds now even use their websites 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.

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 can be explained rationally or is the result of irrational noise-traders driving share prices away from fundamental values. (See, for example, Zweig (1973), DeLong et al. (1990), Chen, Kan & Miller (1993), Chopra et al. (1993), and Pontiff (1997). Dimson & Marsh (1999) provide a comprehensive overview of the debate.)

The model constructed by Gemmill & Thomas (2002) (hereafter GT) can be viewed as one of many attempts to explain the observed distribution of discounts 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 in any way attempt to explain the causes of those fluctuations. Rather, they simply 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 of discounts would also be log normal and centered on zero. GT parameterize the log normal distribution of discounts 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 for discounts 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 for discounts 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.

We will discuss both of these arbitrage barriers below. But first, let us compare the actual distribution of discounts with the one 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.

The problem with GT’s predicted distribution is that it looks little like the actual distribution found in the Fund Edge data set. 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 looks nothing 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. Indeed, nearly 16% of the observations lie outside the bounds assumed by GT. These facts suggest that the GT bounds do not hold and that discounts do 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 particularly 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, GT make no attempt to 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 GT assertion that the "discount can often reach 30% percent before the upper bound is reached" appears conveniently specific.

As for the supposed lower arbitrage bound at a premium of -5%, GT justify it 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 & Thomas (1995) and Lee et al. (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 GT's 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 & Thomas (1995) and Lee et al. (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 et al. (1991) and of about 11.5% in Levis & Thomas (1995). Neither paper gives any suggestion about an arbitrage barrier for individual funds at a premium of -5%.

Indeed, substantial evidence exists 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 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 the 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 at least 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 massive violation of GT's proposed -5% premium barrier without any sign of IPO's of similar funds.

While the boundaries used by GT are not well justified, the values they 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 their log normal model. Rather, it likely stems from using the wrong model, the incorrect boundary assumptions, or both.

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 GT 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. The distribution is leptokurtic and skewed toward premia rather than discounts. For the 225,306 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%. (Most 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%.) 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 in Fund Edge which were 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.³

This section has given ample evidence that GT's static arbitrage boundaries are routinely violated in the Fund Edge data. 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 funds discounts, it would have to do so in a different fashion.

³The 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.

III. Testing The Dynamic Arbitrage Bounds Hypothesis

As we turn to consider the dynamic arbitrage bounds hypothesis, 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. 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 that price funds by capitalizing out future management fees (such as Ross 2002 and the model presented below in Section VI) 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 et al. (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 et al. (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, using 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 percent.

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 1985-2001.⁴

Table II shows that in the more recent period, interest rates had a much different effect on both average absolute discounts and average discounts. 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.

The reader 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 the relationship has the wrong sign, as one might expect that if interest rates rose, the

⁴I chose this time period because it continues where the Pontiff (1996) data set leaves off and because there is twelve months' overlap. This overlap allowed me to verify that the Pontiff and Fund Edge average discount series blend smoothly into each other.

present value of future management fees would decrease, thereby leading to smaller rather than larger discounts.

In Section VI, I demonstrate that one can build a very simple behavioral model of closed-end fund pricing that takes account of management fees but which also features discounts increasing as interest rates rise. Moreover, this model is capable of explaining the shape of the discount distribution of Figure 2 without resorting to arbitrage pressures. This is important because we have now rejected both the static and dynamic arbitrage bounds hypotheses and will, in Section V, present evidence that rational investors cannot earn positive risk-adjusted returns by attempting to set up arbitrage portfolios in closed-end funds. But before moving on, the next section performs and rejects one further test of the presence of arbitrage in closed-end funds.

IV. Do Arbitrage Pressures Drive Discounts toward Fundamentals?

The previous two sections have presented evidence that closed-end fund discounts violate both static and dynamic arbitrage bounds. These rejections may be evidence against arbitrage, or may just be due to low power. Consequently, it is wise to check for arbitrage 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 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 mimicing 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.⁵ 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 & 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 & 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. 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

⁵The 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 mimicing 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 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 different 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 & Thomas (2002).

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 et al. (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 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 et al. (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 over the period 1985-2001.

The results of columns (2) and (3) also cast doubt on the robustness of GT's conclusion that arbitrage pressures will 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 GT 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).

It is the case that none of the GT regression results generalize with any significance to both bond and stock funds traded in the US and Canada. This casts significant doubt on GT's 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 Negative Returns to Arbitrage in Closed-end Funds

The previous section showed that arbitrage does not appear to push closed-end fund share prices toward discount levels consistent with fundamental factors like management fees. But why does arbitrage fail? As it turns out, arbitrage is unattractive because the risk-adjusted excess returns to arbitrage are either

negative or insignificantly different from zero, even without taking into account broker's fees or the opportunity costs associated with maintaining short positions.

This can best be seen by setting up portfolios which would benefit from the movement of fund prices toward fundamental levels. For funds trading at large discounts, one 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, one would go long the underlying and short the fund's shares, as any reduction in premia would generate a profit.

But before examining the returns to such portfolios, we should examine the rate at which discounts and premia mean revert, as it is this rate that determines the profitability of such positions. As it turns out, profits are low because discounts and premia revert only very slowly, on average, to their mean level.

The slowness of mean reversion can be quantified by running an AR(1) regression of current values of D_t on those one year prior. Assume that D_t is mean reverting to the level \bar{D} . Then it should follow $D_{t+1} = \bar{D} + \phi(D_t - \bar{D}) + \varepsilon_t$, where ϕ gives the fraction of the year t deviation that remains the next year, and ε_t is a Gaussian shock. We can get empirical estimates for ϕ and \bar{D} by running the regression, $D_{t+1} = \text{constant} + \phi D_t + \varepsilon_t$. 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/premium data for the 462 closed-end funds in our sample, and estimating the equation using pooled least squares on data covering 1985-2001 gives a constant of 2.19 and an estimated value for ϕ of 0.64. The regression has an r-squared statistic of 0.41 and the t-statistics on the constant and ϕ are, respectively, 53.3 and 184.2. Using these estimates, we can back out an estimated value for \bar{D} of 6.08%. That is, this regression methodology indicates that discounts/premia revert to the same level as the mode discount of Figure 2. But the regression also indicates that discounts/premia are so slow getting back to 6% that, on average, 0.64 of any deviation from the mean remains after one year. This rate of mean reversion is inconsistent with there being significant arbitrage activities that quickly rectify deviations from fundamental levels.

The slow rate also implies, as we will see, that the risk-adjusted returns to arbitrage portfolios that attempt to profit from mean reversion are usually negative.

I examined the profitability of arbitrage on monthly data ranging from January 1985 through May 2001. I did this by creating twenty portfolios set up to generate profits from the mean reversion of discounts and premia. Each of the portfolios corresponds to a five percentage-point wide discount or premium bin. The first bin corresponds to premia falling between -50% and -45%, and 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 fourteen 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 are denoted, respectively, r^{spot} and r^{NAV} . Excess returns over the following month to the appropriate arbitrage portfolio were then calculated. Because the long-run mean reverting value of discounts is about 6%, for those bins containing discounts larger than 5%, one would want to go long the fund and short the underlying in order to set up a hedged position that would profit purely from mean reversion. The excess return to this position over the month is, $r^{spot} - r^{NAV}$. For those bins containing discounts or premia less than 5%, one would want to short the fund and go long its underlying. Consequently, for all bins containing discounts less than 5%, $r^{NAV} - r^{spot}$ was calculated for each fund.⁶

Please note that the round-trip returns used in these calculations take into account bid-ask spreads, thereby giving a more realistic measure of the attractiveness of potential arbitrage profits. For r^{spot} , the returns were calculated using bid and ask prices obtained from the CRSP daily stock file. For r^{NAV} , NAV spreads had to be imputed, and they were assumed to be 20 cents per share. This value was chosen because it was the mean bid-ask spread on fund shares when funds traded near the long-run discount

⁶Although the mean reverting discount level of 6% falls into the 5% to 10% bin, for convenience we assume that arbitrageurs would want to go long the fund and short the underlying for all discounts falling into this bin.

level of 6%—my assumption being that the effective spreads if one were to buy and sell the stocks and bonds in the underlying portfolios would be about the same as the spreads found on fund shares. While this justification may seem unconvincing, please note that a spread of 20 cents is also consistent with typical average stock and bond spreads. For instance, Bessimbinder (2003) finds spreads of about 17.4 cents for all NYSE stocks and 31.4 cents for all NASDAQ stocks. And for bonds, Chakravarty & Sarkar (1999) find that the average spreads on US municipal, corporate, and federal government bonds are, respectively, 22 cents, 21 cents, and 11 cents. Since most closed-end funds invest in either stocks or municipal bonds, my assumption of a 20 cent spread on the NAV seems reasonable; if investors were to attempt to directly replicate a fund's underlying portfolio they would face spreads of about 20 cents for all the stocks or bonds that they would have to buy or short.

Taking spreads into account is extremely important given the slowness of mean reversion. Above it was noted that, on average, 0.64 of the gap between D_t and the mean reverting value of 6% remains after one year. A similar analysis shows that 0.94 of any such gap remains after one month. Because this translates into only small movements in fund share prices, the potential gain is overwhelmed by bid-ask spreads. For instance, imagine a fund that has a constant NAV of \$20. If it was initially trading at a discount of 25%, it would have a spot price of \$15. Because the fund is trading at a discount 19 percentage points above the mean-reverting discount level of 6%, and because, on average, 0.94 of that gap will remain after a month, we would expect that the fund would be trading at a 23.7% discount the next month. That would imply that its share price would only rise from \$15 to \$15.26 over the course of the month.

A simple example shows that this 26 cent increase in price would be overwhelmed by bid-ask spreads. Assume that the quoted share prices are ask prices and that you face 20 cent bid-ask spreads. Then, at the beginning of the month, you would have to buy at the ask price of \$15 per share but at the end of the month you would only be able to sell at the bid price of \$15.06. In the mean time, you short the NAV at the beginning of the month, selling at the bid price of \$19.80 but at the end of the month

have to cover at the ask price of \$20. So you lose \$0.20 on the short, and only gain \$0.06 on the long. That makes for a loss of \$0.14.

The average of the arbitrage portfolio returns of all the funds in each bin in a given month was then taken, and that average constitutes the bin's one-month return. Proceeding in this fashion for each of the twenty bins over the 196 months yields twenty time series that give the returns to attempting to profit from the mean reversion of D_t .

Table IV presents the results of running the Fama-French (1992) three factor model on the excess returns of each of the twenty bin portfolios. Yields on one-year US Treasuries were used to proxy for the risk-free interest rate when calculating excess returns. The regression results show that the twenty arbitrage portfolios fail to produce significant positive risk-adjusted returns. The constant α is significantly negative for three of the twenty bins, negative but insignificant for another twelve of the twenty bins, insignificantly positive for six of the bins and only significantly positive for premia falling into the -25% to -30% bin. Because less than one-third of one percent of the monthly discount and premium observations in the sample fell into this bin, we can say with confidence that for the vast majority of cases, the rate of regression toward the long run mean of 6% is too slow to generate attractive profits to arbitrageurs after adjusting for the standard risk factors.⁷

Please note, too, that the returns to the twenty arbitrage portfolios only account for bid-ask spreads. Real-world investors hoping to benefit from the mean reversion of fund prices would also have to deal with trading costs, the difficulty of shorting fund shares, problems with finding suitable hedge vehicles for fund portfolios, and the holding costs associated with maintaining short positions. Therefore, arbitrage activities in closed-end funds would likely be even less attractive than suggested by the results of Table III. Trading costs alone would likely render any high-frequency trading strategy unprofitable given the slow pace of mean reversion. In exchange for these poor returns, arbitrageurs also have to

⁷Another way to see that the positive returns in the -25% to -30% bin are anomalous is to look at the second-to-last column in Table IV. That column gives the mean of the excess returns of each bin's time series. The average excess return for the -25% to -30% is extremely large compared to that of other bins.

contend with noise-trader risk, the risk that discounts will widen rather than narrow—a risk that is not accounted for by the three Fama & French (1992) risk factors.

VI. A Very Simple Behavioral Model of Discounts

This paper has given evidence that closed-end funds in the US and Canada routinely violate both dynamic and static arbitrage bounds. Furthermore, they appear to violate these bounds because the returns to arbitrage are so poor that arbitrageurs do not enter the closed-end fund market. The resulting lack of arbitrage pressure, in turn, helps to explain why the long-run average discounts of funds do not appear to be related to fundamental factors like fund expense ratios and replication risks, as with arbitrage absent, closed-end fund pricing is free to be driven by the sentiments of noise traders.

Yet, it remains the case that the histogram of fund discounts is centered at a level consistent with simply discounting out expected future management fees. The mode discount of 6% seen in Figure 2 is what we would expect if we priced a fund by simply capitalizing out fees. The model presented in this section takes account of this.

Implicit in the conjecture that the share price should just be NAV less capitalized future management fees is the assumption that managers cannot beat the market. However, it is clear from the financial markets that many investors fail to share this assumption. Rather, they often grow optimistic or pessimistic about potential managerial performance. The model presented below incorporates investor beliefs about the ability of managers to beat the market, and can therefore generate discounts and premia that deviate from the value consistent with simply capitalizing out management fees. When investors think that a manager will produce inferior returns, they price the fund at a discount deeper than that consistent with simply capitalizing out management fees. And when investors think a manager will produce superior returns, they price the fund at a lesser discount or even at a premium.

The model therefore explains the shape of Figure 2 as arising from the distribution of investor sentiment about fund managers. The distribution is free to reflect these changing sentiments because of the lack of arbitrage. Also, because arbitrage is not causing the mean reversion of discounts and premia to the long-run value, such mean reversion must simply be a reflection of the natural decay process of investor optimism and pessimism.

The model has the further benefit of being able to explain the positive relationship between interest rates and discounts reported in Table II. In a nutshell, if investors think that they can earn more outside the fund than the fund manager can generate through his portfolio picking, discounts widen. This is because, for any given level of confidence in the manager, rising interest rates are a tempting alternative that cause fund prices to fall and discounts to widen as investors reallocate out of funds.

The model prices a fund by equalizing the expected future cash flows that sentimental investors *believe* are available outside the fund with those that they *believe* they would get by owning shares of the fund. The model assumes that when estimating the cash flow that fund shares are expected to generate, investors capitalize out expected future management fees. To do this, the model must take account of the expected lifetime of the fund until it either liquidates or converts to an open-end format—either of which would force the fund's share price to equal NAV. An examination of all closed-end stock funds trading on the NYSE from 1960 to 1999 reveals that one cannot reject the hypothesis that closed-end funds die off in a Bernoulli fashion, with an annual death probability of $\gamma = .0364$.⁸ That is, if a fund is in business on January 1st, it has a 3.64% chance of going out of business by the end of the year and a $1 - \gamma = 96.36\%$ chance of continuing in business into the next year.⁹

⁸ Let X_t be the number of funds alive at the start of year t and O_t be the number of those that die during year t . Assuming that fund deaths are Bernoulli, with death rate γ , the expected number of deaths in year t is γX_t . A Pearson's Chi-squared test statistic can therefore be constructed as $D^2 = \sum_{t=1960}^{1999} \frac{(O_t - \gamma X_t)^2}{\gamma X_t}$. D^2 is distributed approximately χ^2 with $1999 - 1960 - 1 = 39$ degrees of freedom. Our estimated D^2 is 30.52 which is significantly less than the 90% critical value of 51.81. We fail, therefore, to reject the hypothesis that fund deaths follow a Bernoulli process.

⁹ Annual fund death rates were regressed against macro variables, fund returns, fund discount levels and other variables that might plausibly affect the decision to liquidate a closed-end fund or convert it to an open-end fund. All of the variables were found to be uncorrelated with fund death rates.

As above, N_t stands for the fund's NAV at time t , and P_t for the fund's price at time t . If a fund goes out of business or converts to an open-end format at some particular time T , then $P_T = N_T$. Either event will be referred to below as a liquidation. In any period $t < T$ before liquidation takes place, the closed-end fund may choose to pay a dividend, d_t . The current price of a fund should simply be the present value of expected dividend payments over the period that the fund is expected to remain in business plus the present value of the liquidation payment that will eventually be made when the fund goes out of business. Given the Bernoulli death process, the current price of a closed-end fund is consequently the sum of two terms¹⁰:

$$\begin{aligned}
P_t = & \sum_{i=1}^{\infty} \gamma(1-\gamma)^{i-1} \left(\frac{1}{R^{out}} \right)^i E_t[N_{t+i}] \\
& + \sum_{i=1}^{\infty} (1-\gamma)^i \left(\frac{1}{R^{out}} \right)^i E_t[d_{t+i}].
\end{aligned} \tag{1}$$

The first term in equation (1) is the expected present value of liquidation payments, where discounting is done relative to R^{out} , the expected period-on-period, total rate of return that fund investors believe is available outside of the fund in alternative investments. The second term in equation (1) gives the expected present value of dividend payments, also discounted by R^{out} .

Equation (1) is incomplete because it fails to take into account the fact that a dividend payment made in period t lowers the NAV of the firm in later periods. We must also take account of the fact that management fees paid at time t also reduce the NAV of the fund in later periods, and that our sentimental investors care about the rate of return that the fund manager is expected to generate.

Define R to be the period-on-period total rate of return that the fund manager is expected to earn. Then, if a manager begins with N_t dollars at the beginning of period t , investors will expect that his stock

¹⁰We must also impose a transversality condition to rule out speculative bubbles. Formally, we require that $\lim_{n \rightarrow \infty} (1-\gamma)^n \left(\frac{1}{R^{out}} \right)^i E_t[P_{t+n} + d_{t+n}] = 0$.

(or bond) picking will increase the NAV of the fund to RN_t . Not all of that amount will be retained by the fund, however. At some point, the manager must be paid. Closed-end fund management contracts typically specify that managers will receive a fraction $f < 1$ of fund NAV each year. Therefore, the NAV of the fund after the manager is paid will be $(1 - f)RN_t$. If a fund liquidated, this amount of NAV would be returned to shareholders. If, however, the fund continues in business into the next year, dividends must be paid out of $(1 - f)RN_t$.¹¹ I model dividends by assuming that a constant fraction α of NAV is paid out each year to investors in the form of dividends. Therefore, if a fund continues in business, its NAV at the start of the next period, after paying a dividend of $\alpha(1 - f)RN_t$, will be $N_{t+1} = (1 - \alpha)(1 - f)RN_t$. If we substitute into equation (1) this model of management fees, dividend payments, and expected managerial return, we obtain a convergent geometric series which, when simplified, gives the period t price of the fund as a simple function of period t NAV:

$$P_t = \frac{\frac{R}{R^{out}}(1 - f)\gamma}{1 - \frac{R}{R^{out}}(1 - f)(1 - \gamma)(1 - \alpha)}N_t + \frac{\frac{R}{R^{out}}(1 - f)(1 - \gamma)\alpha}{1 - \frac{R}{R^{out}}(1 - f)(1 - \gamma)(1 - \alpha)}N_t. \quad (2)$$

As with Equation (1), the first term is the expected present value of liquidation payments, while the second term is the expected present value of dividend payments. However, equation (2) incorporates both fundamental and non-fundamental factors into the price of a closed-end fund. The fundamental factors are management fees, dividend payout rates, the current value of the underlying portfolio, and the fund's Bernoulli death rate. The non-fundamental factors are summarized by $\frac{R}{R^{out}}$, which captures investor expectations about the manager's ability to beat the market.

A nice feature of this pricing model is that it can capture rational expectations as a special case of the sentiment ratio. Malkiel (1995) provides evidence that mutual fund managers cannot systematically

¹¹ Under US securities law, closed-end funds can avoid paying taxes on the capital gains and dividend payments generated by their underlying portfolios by passing on at least 90% of said gains and dividend payments to the shareholders of the fund. These disbursements are made as quarterly dividend payments to fund shareholders. A consequence of this tax law is that fund NAVs are usually more or less constant over time as most funds simply pass through capital gains and dividend earnings from their portfolios as they are accrued.

beat the market over the long run, while Zheng (1999) offers evidence that mutual fund investors are unable to predict which managers will beat the market even in the short run. Consequently, an investor having rational expectations should assume that $\frac{R}{R_{out}} = 1$. That is, the rational investor assumes that the manager of a given fund will on average do neither better nor worse than the market—both because managers cannot systematically beat the market and because the rational investor would not believe that he had the power to predict which managers would in the future beat the market by sheer chance. With rational expectations imposed, our pricing model is reduced to only fundamental factors. Setting $\frac{R}{R_{out}} = 1$ in equation (2) and combining terms gives the rational price of the fund as,

$$P_t = \frac{[\gamma + (1 - \gamma)\alpha](1 - f)}{1 - (1 - f)(1 - \gamma)(1 - \alpha)} N_t. \quad (3)$$

The three parameters f , γ , and α affect the rational price of the fund in equation (3) by modulating the future stream of management fee payments. The higher are fee rates, f , the lower is the price of the fund, as higher rates imply that more of the fund's capital will flow over time to managers rather than investors. A higher value of the death rate, γ , implies a higher price for the fund because the sooner the fund goes out of business, the fewer times management fees will be paid out, thereby leaving more capital to shareholders. And the higher the value of the dividend rate, α , the higher the current price, because any capital paid out to shareholders in the form of dividends is capital out of which management will not be able to take fees during later periods.

If investors are optimistic about the prospects of the fund, then $\frac{R}{R_{out}} > 1$ and the price of the fund exceeds the rational price given in equation (3). On the other hand, if investors are pessimistic, then $\frac{R}{R_{out}} < 1$ and the price of the fund is less than the rational price. In this way, the shape of Figure 2 can be explained in terms of sentiment fluctuating around the rational level $\frac{R}{R_{out}} = 1$. This can be seen by substituting empirical values into equation (3) to determine the rational discount level. The average

management fee rate of the 462 closed-end funds trading in the USA in 2001 was $f = .0081$, while the average annual dividend payout rate was $\alpha = .0690$. As noted above, the empirically estimated Bernoulli death rate for closed-end funds is $\gamma = .0364$. If we substitute these values, we obtain a model-predicted discount of 7.2%. This prediction is quite close to the mode discount of 6% found in Figure 2, so that the model's prediction for the center of the distribution is quite close to the actual center.

The observed spread around the center is of course caused by fluctuating sentiment, and the model would interpret the skewness of the distribution toward premia as arising from the fact that irrational, noise trading investors are biased toward believing that managers will beat the market. That is, the premium tail is longer than the discount tail because investors are more likely to be extremely optimistic than extremely pessimistic.

This model also gives the prediction that higher rates of interest will lead to larger discounts, consistent with the regression results of Table II. To see this, simply fix the expected fund return, R , at some value in equation (2) while increasing the outside rate of return, R^{out} , that investors use to discount the fund's future stream of payments. This will cause the current price of the fund, P_t , to fall and hence generate a larger discount for any given portfolio value, N_t . If increases in the risk-free rate are positively correlated with the discount rate R^{out} that fund investors use to discount the fund's future payout stream, then this model gives a clean explanation for the positive association between discounts and risk-free interest rates.

It should be noted that the fact that the 6% mode discount of Figure 2 is less than the rational expectations discount of 7.2% likely confirms the model's fundamental intuition. The most simple alternative outside rate of return that investors could obtain would be generated by simply buying and holding the manager's current portfolio. In a world in which the mean sentiment toward managers is slightly positive because it is believed that active management can beat buy and hold, investors would

expect $R > R^{out}$. They would therefore price funds such that discounts would typically be less than the rational expectations discount that assumes $R = R^{out}$. This is exactly what we find in Figure 2.

Note, however, that the empirical distribution of discounts and premia also suggests that while investors may typically believe that managers will beat buy and hold, they are not usually expected to do so to such an extent as to make up for their management fees. The fact that only 31% of D_t observations are of premia suggests that only 31% of the time do investors believe that managers will produce returns that are enough in excess of buy and hold to make up for the losses caused by management fees .

VII. Conclusion

Closed-end fund pricing behavior in the USA and Canada over the period 1985-2001 is inconsistent with there being strong or consistent arbitrage pressure. Closed-end fund discounts routinely violate the static arbitrage bounds proposed by Gemmill & Thomas (2002) and are inconsistent with the dynamic arbitrage bounds proposed by Pontiff (1996). That arbitrage pressures are weak is also suggested by the fact that long-run average discounts for both stock and bond funds are insignificantly correlated with fundamental factors like expense ratios and replication risk. Indeed, the only variable that is significant for both bond and stock funds is their exposure to noise-trader risk.

The lack of arbitrage pressure appears to derive from a lack of profit incentive. Hedged positions that go long the fund and short its underlying if the fund is trading at a discount, or go short the fund and long its underlying if it is trading at a premium, generate negative excess risk-adjusted returns when using the three Fama-French (1992) factors to capture risk. This finding is consistent with the noise-trader model of closed-end fund discounts proposed by DeLong et al. (1990). In that model, risk-averse arbitrageurs do not fully offset pricing errors because mispricings may widen rather than narrow due to the unpredictable buying and selling behavior of irrational noise traders. Because instantaneous arbitrage is not possible between the shares of a closed-end fund and its underlying portfolio, the

presence of noise-traders—and the independent source of risk that they create—is enough to discourage rational investors. Because they are not willing to stand ready to invest whatever amount of capital would be necessary to fully offset the noise traders, prices can deviate from rational levels.

This paper also presents evidence that the cross-sectional average discount on closed-end funds increases 2.2 percentage points for every single percentage point increase in the 1-year Treasury rate. This behavior is shown to be consistent with a simple behavioral model of closed-end fund discounts. In that model, sentimental noise traders price a closed-end fund by comparing the rate of return that the fund's manager is expected to generate with the returns that investors expect alternative assets to generate. As the Treasury rate increases, closed-end funds become relatively less attractive, with the result that their prices fall and their discounts widen.

The model is also useful because it can explain the shape of the empirical discount distribution without recourse to arbitrage. The empirical discount distribution is centered on a discount of 6%, and tapers off gently in both directions, with extreme premia being more common than extreme discounts. The rate of mean reversion is extremely slow, with 0.64 of a deviation from the long-run discount level of 6% remaining after one year, on average. This rate of mean reversion is too slow to be consistent with the presence of strong arbitrage pressures, but is consistent with the idea that closed-end fund discounts wax and wane in response to evolving investor sentiment.

Viewed this way, the slow average rate of mean reversion is simply a reflection of the slow average rate of sentiment reversion. When sentiment is neutral and investors believe that managers will do no better and no worse than the market, they price the fund by simply discounting out future management fees. When sentiment is negative and investors are pessimistic about a manager's prospects, the fund's share price falls so that the fund trades at a discount that is deeper than the discount consistent with simply discounting out future management fees. And when sentiment is positive and investors are optimistic about a manager's prospects, the fund's share price rises above this level so that the fund trades at a lesser discount or even at a premium.

The model predicts that the empirical discount distribution should be centered on a discount of 7.2%, given average management fee rates, the empirical half life of funds, and the average dividend payout rate. This value is quite close to the empirical mode discount of 6%, suggesting that the model fits the data rather well. Moreover, the fact that the empirical mode discount is less than the discount generated by simply capitalizing out future management fees is consistent with the model's fundamental intuition: If investor perceptions about managerial ability are what drive discount levels, then we should expect such a result because real-world investors are in fact biased toward believing that managers can beat the market.

However, the correct point of this paper may not be that the absence of arbitrage means that prices will deviate wildly from fundamentals. Rather, the correct point may be that when people mostly agree on fundamentals, we do not need arbitrage to keep prices from deviating too far from fundamentals. Indeed, despite the fact that closed-end funds appear to be very little constrained by arbitrage, fund prices for the most part do not deviate that far from rational levels. In particular, just over 50% of the discount and premium observations in Fund Edge over the period 1985-2001 are within five percentage points of the long-run mean-reverting discount level of 6%, and just over 75% are within ten percentage points of the mean-reverting level. Since the mean reverting level is very nearly the level that would rationally price a fund by capitalizing out future management fees, it would seem that, for the most part, the absence of arbitrage doesn't lead to wild mispricings in closed-end funds.

By contrast, in other markets where arbitrage is constrained, we see much larger violations of the law of one price. For instance, in the options markets, Ofek, Richardson & Whitelaw (2002) find that synthetic long positions constructed from options often cost much less than actual long positions in underlying stocks when arbitrage is impeded by short-sales constraints. In fact, mispricings that are an order of magnitude larger than any ever seen in closed-end funds are relatively common in the options markets.

The key difference may be that whereas closed-end funds have agreed upon fundamental values (pegged to well known portfolio values), operating company stocks do not. Therefore, a much greater spread of investor valuations is likely to exist. When short-sales constraints segment the stock and options markets by constraining arbitrage, the prices that each of those markets set on the the same stream of payments can therefore be quite different. As a result, the more interesting question when we see a large departure from the law of one price may not be to ask why arbitrage is absent, but instead to ask why people have such differing valuations that the absence of arbitrage leads to such large violations of the law of one price .

References

- Bessimbinder, H. (2003), 'Issues in assessing trade execution costs', *Journal of Financial Markets* **6**(3), 233–257.
- Chakravarty, S. & Sarkar, A. (1999), 'Liquidity in u.s. fixed income markets: A comparison of the bid-ask spread in corporate, government and municipal bond markets', *Staff Reports, Federal Reserve Bank of New York* **73**.
- Chen, N.-F., Kan, R. & Miller, M. H. (1993), 'Are discounts on closed-end funds are a sentiment index?', *Journal of Finance* **48**(2), 795–800.
- Chopra, N., Lee, C. M., Shleifer, A. & Thaler, R. (1993), 'Yes, discounts on closed-end funds are a sentiment index', *Journal of Finance* **48**(2), 801–808.
- DeLong, J. B., Shleifer, A., Summers, L. H. & Waldmann, R. J. (1990), 'Noise trader risk in financial markets', *Journal of Political Economy* **98**(4), 703–738.
- Dimson, E. & Marsh, P. (1999), 'Closed-end funds: A survey', *Financial Markets, Institutions, and Instruments* **8**(2).
- Fama, E. & French, K. (1992), 'The cross-section of expected stock returns', *Journal of Finance* **47**(2), 427–65.
- Gemmill, G. & Thomas, D. C. (2002), 'Noise-trading, costly arbitrage, and asset prices: Evidence from closed end funds', *Journal of Finance* **57**(6), 2571–2594.
- Lee, C. M., Shleifer, A. & Thaler, R. (1991), 'Investor sentiment and the closed-end fund puzzle', *The Journal of Finance* **46**(1), 75–109.
- Levis, M. & Thomas, D. C. (1995), 'Investment trust ipos: Issuing behavior and price performance. evidence from the london stock exchange.', *Journal of Banking and Finance* **19**, 1437–1458.
- Malkiel, B. (1995), 'Returns from investing in equity mutual funds 1971 to 1991', *The Journal of Finance* **50**(2), 549–572.
- Newey, W. & West, K. (1987), 'A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix', *Econometrica* **55**, 703–708.
- Ofek, E., Richardson, M. & Whitelaw, R. F. (2002), 'Limited arbitrage and short sales restrictions: Evidence from the options markets', *NBER Working Papers* **#9423**.
- Pontiff, J. (1996), 'Costly arbitrage: Evidence from closed-end funds', *Quarterly Journal of Economics* **111**(4).
- Pontiff, J. (1997), 'Excess volatility and closed-end funds', *American Economic Review* **87**(1), 155–169.
- Weiss, K. (1989), 'The post-offering price performance of closed-end funds', *Financial Management* **18**(3), 57–67.
- Zheng, L. (1999), 'Is money smart? A survey of mutual fund investor's fund selection ability', *The Journal of Finance* **54**(3), 901–933.
- Zweig, M. E. (1973), 'An investor expectations stock price predictive model using closed-end fund premiums', *The Journal of Finance* **28**, 67–78.

This page blank.

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).

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	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 - 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 NAV 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 French (1992) regressions performed on 20 hedge portfolios whose returns are completely driven by the mean reversion of discounts and premia to their long-run value of a 6% discount. The regressions were run on monthly data covering January 1985 to May 2001. The 20 portfolios are defined by initial discount or premium levels, so that each month each fund is placed into one of 20 five-percentage point wide bins, the bins ranging from a discount of 50% to a premium of -50%. For each fund in a given bin, the returns to the appropriate hedge position over the next month were computed and then averaged with those of the other funds in the bin. The appropriate hedge position for funds trading at discounts deeper than 6% is to go long the fund's shares and short its underlying, while that for funds trading at lesser discounts or premia is to do the opposite. Proceeding in this way, we obtain for each of the 20 portfolios a returns time series from which 1-year Treasury yields were subtracted to obtain excess returns. These excess returns were then regressed in the normal way on a constant, α , and the three Fama & French (1992) factors: $R_m - R_f$ is the market excess return, SMB is the return to small capitalization stocks less the return to big 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 & West (1987) and t-statistics are given in parentheses. The final two columns give the mean and standard deviation of the dependent variable.

Bin Lower Bound	Bin Upper Bound	α	$R_m - R_f$	SMB	HML	R-sq	DW	Obs.	Mean Dependant Variable	SD Dependant Variable
-50	-45	-1.43 (-0.79)	-0.60 (-3.38)	-2.43 (-3.01)	-1.85 (-3.19)	0.39	1.45	33	0.01	11.66
-45	-40	1.14 (0.65)	-0.08 (-0.31)	0.03 (0.07)	-0.55 (-0.77)	0.02	0.96	54	0.79	9.88
-40	-35	-1.65 (-1.27)	0.54 (1.62)	-0.26 (-0.54)	0.16 (0.42)	0.09	1.61	65	-0.90	8.77
-35	-30	-1.09 (-0.82)	0.26 (1.39)	-0.07 (-0.27)	0.05 (0.19)	0.01	2.23	81	-1.07	11.88
-30	-25	2.59 (2.68)	-0.57 (-2.91)	0.01 (0.02)	-0.55 (-1.89)	0.13	1.46	90	2.26	8.45
-25	-20	0.79 (0.77)	-0.30 (-1.40)	-0.25 (-1.65)	-0.02 (-0.14)	0.02	0.58	119	0.57	10.27
-20	-15	0.44 (1.07)	-0.49 (-3.56)	-0.02 (-0.25)	-0.52 (-4.84)	0.13	1.00	154	-0.04	6.07
-15	-10	-0.04 (-0.14)	-0.04 (-0.55)	-0.01 (-0.09)	-0.17 (-2.37)	0.04	1.36	185	-0.09	3.02
-10	-5	-0.03 (-0.18)	-0.15 (-3.15)	-0.07 (-1.19)	-0.20 (-3.36)	0.17	1.75	196	-0.16	1.76
-5	0	-0.36 (-3.23)	-0.11 (-3.72)	-0.09 (-2.07)	-0.17 (-3.37)	0.13	2.05	196	-0.46	1.65
0	5	-0.61 (-4.73)	-0.14 (-3.33)	-0.10 (-2.29)	-0.18 (-4.07)	0.21	1.78	196	-0.73	1.50
5	10	-0.29 (-2.30)	0.14 (3.96)	0.17 (3.63)	0.20 (5.32)	0.20	2.15	196	-0.18	1.76
10	15	-0.07 (-0.49)	0.14 (3.68)	0.10 (2.30)	0.20 (4.59)	0.11	2.10	189	0.05	2.17
15	20	-0.28 (-1.35)	0.08 (1.34)	0.16 (2.65)	0.12 (2.55)	0.04	1.70	182	-0.22	2.96
20	25	-0.41 (-1.42)	0.19 (2.24)	0.08 (0.97)	0.18 (2.17)	0.05	1.24	170	-0.26	3.93
25	30	0.14 (0.29)	0.42 (3.12)	0.32 (2.64)	0.39 (4.24)	0.13	1.50	121	0.33	5.97
30	35	-1.00 (-1.07)	0.36 (1.83)	-0.02 (-0.08)	0.18 (1.23)	0.04	1.56	85	-0.42	7.52
35	40	-0.05 (-0.04)	0.74 (4.15)	0.71 (2.22)	0.64 (2.73)	0.30	0.34	33	-0.34	8.17
40	45	4.61 (1.28)	-0.08 (-0.10)	-0.12 (-0.15)	0.15 (0.49)	0.01	0.24	20	4.60	13.27
45	50	1.23 (0.26)	0.23 (0.24)	-1.33 (-0.68)	4.07 (2.00)	0.29	1.43	11	1.53	11.14

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




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

