The Closed-End Fund Puzzle: New Evidence From Business Development Companies

The Closed-End Fund Puzzle: New Evidence from Business Development Companies

Christopher Eugene Fenaroli

Presented to Applied Mathematics
in partial fulfillment of the requirements
for a Bachelor of Arts degree with Honors

Harvard College
Cambridge, Massachusetts
March 29, 2019

Abstract

The phenomenon of closed-end funds (CEFs) trading at a discount to their Net Asset Value (NAV) has long been discussed in financial economic literature. While market data from traditional CEFs has been extensively studied, this paper presents novel evidence on the subject from a subset of CEFs known as business development companies (BDCs). This paper shows that, to the extent they are testable, the features of the closed-end fund puzzle are present in the BDC market. Furthermore, the characteristics of BDC discounts generally corroborate the predictions of the behavioral explanation of the closed-end fund puzzle. Qualities observed in the BDC discount data in support of the behavioral hypothesis include: price-driven fluctuations in BDC discounts, empirical independence between observed discounts and commonly cited neoclassical determinants of discounts, and a relationship between proxy noise trader prevalence and sentiment-dominant discount behavior.

Acknowledgements

I am grateful to Professor John Campbell for his constructive guidance, invaluable advice, and aid at all points from deciding on a topic to final edits. I would also like to thank my thesis seminar instructor Dr. Andrés Maggi for his comprehensive feedback and encouragement, as well as Dr. Gregory Bruich and Alex Buzby for their insightful discussions. Finally, this paper greatly benefited from the support of my peers and family. Any views reflected or mistakes made are my own.

Contents

1 Introduction **1**

2 Background **2**

2.1 The Closed-End Fund Puzzle 2

2.2 Proposed Explanations of the CEF Puzzle 3

2.3 Business Development Companies 7

3 Data **10**

3.1 Sources and Construction 11

3.2 Description 12

3.3 Presence of CEF Puzzle Features 15

4 Methodology **16**

4.1 Ordinary Least Squares (OLS) 16

4.2 Fixed Effects (FE) 23

4.3 OLS with Market Capitalization (OLS w/ MC) 28

4.4 OLS with Institutional Ownership (IO) 34

5 Results and Discussion **39**

5.1 Summary of Findings 39

5.2 Discussion 42

5.3 Limitations 44

6 Conclusions **44**

Appendix **46**

References **49**

1 Introduction

A closed-end fund (CEF) is a publicly traded investment company that manages a portfolio of assets and distributes portfolio income to shareholders. CEFs are unique among investment funds in that initial capital is raised through an initial public offering (IPO) on a public stock exchange after which no more shares can be created or redeemed.

The CEF universe today consists of over 500 funds that collectively manage $275 billion (Investment Company Institute 2018). While countless managers have used their unique structure to meet their investment goals, economists and market researchers have been interested in the asset class for a separate reason. A common obstacle in empirically evaluating the efficient market hypothesis is the inability to quantify market efficiency; how can one determine whether the price of a security is reflective of its fundamental value without knowing what it is truly worth? Because closed-end funds present a rare scenario in which both the market price and the fundamental value of a security are directly observable, they have functioned as a sandbox for researchers to test the efficient market hypothesis as well as a number of other market models.

When Pratt (1966) first used CEFs for this purpose and compared CEF prices to that of their underlying portfolios, he obtained the surprising result that CEFs often traded at a discount to their portfolio net asset value (NAV), a finding described by Ross (2002) as an "affront to the most basic principles of neoclassical finance." This behavior, generally referred to together with a number of other anomalous features of CEF premia, has been called the "closed-end fund puzzle."

In this paper I introduce business development companies (BDCs), a form of CEF that makes primarily debt investments in small companies, as a source of new evidence on this long-standing question. While traditional CEF data has been extensively studied, BDCs remain untouched in academia due to the market's opacity and, until recently, small size. The analysis I present first confirms the existence of the features of the close-end fund puzzle in the BDC data and then tests the predictions of different CEF puzzle hypotheses through modeling premium fluctuations.

The rest of this paper is organized as follow. Section 2 presents a history of the closed-end fund puzzle, an overview of proposed competing explanations for its existence, and background on BDCs and their relevance as a new source of data on the CEF puzzle. Section 3 details the construction of the BDC dataset, describes its most salient features, and confirms the presence of features of the CEF puzzle in BDCs. Section 4 details the methodology, justification, and results from each model of BDC premium fluctuations. Section 5 reconciles findings and Section 6 concludes.

2 Background

2.1 The Closed-End Fund Puzzle

The defining feature of the closed-end fund model, and what distinguishes it from open-ended funds, is the inability of CEF managers to issue new shares or redeem existing ones after inception. Their fixed capital structure has led to CEFs presenting fee arrangements, portfolio compositions, and investment strategies distinct from the more popular open-ended fund model. While the history of closed-end funds dates back to 1893, it was after the standardization of the industry brought by the Investment Company Act of 1940 that CEFs began to realize their current popularity.

As discussed in the introduction, closed-end funds have been of tremendous interest to academics as they present a unique market in which both fundamental and market value are directly observable. Pratt (1966) was the first to document CEF discounts within an academic context, noting that CEFs "normally sell at a discount" and claiming there "no mathematical or logical rationale for such a discount characteristic." Lee, Shleifer, and Thaler (1990) deemed this conundrum, together with three other anomalous CEF behaviors, the “closed-end fund puzzle.”

The four features of the closed-end fund puzzle are as follows. The first is the participation of investors in CEF initial public offerings (IPOs). Underwriting costs are typically passed through to investors, meaning that CEF IPOs are usually priced at a premium (Weiss 1989, Peavy 1988). Despite the knowledge that CEFs predominantly trade at a discount, investors still decide to participate in CEF IPOs. The second feature of the puzzle is the long run valuation of CEFs at a mid-single digit discount to their fundamental value (Weiss 1989). The third feature of the puzzle is the wide fluctuations of CEF discounts over time. These fluctuations are mean reverting (Sharp and Sosin 1975, Thompson 1978, Pontiff 1995). The final feature of the puzzle is the tendency for CEF share prices to tend towards their NAV in advance of a fund liquidation event but leave a small discount remaining (Brauer 1984, Brickley and Schallheim 1985).

2.2 Proposed Explanations of the CEF Puzzle

Over the past 40 years over 100 articles have been published on closed-end funds, the vast majority in relation to explaining portions of the puzzle described in the previous section. Most of these proposed explanations for the closed-end fund puzzle fall into two schools of thought, both detailed in this section.[1]

Neoclassical Model

The first, commonly referred to as the neoclassical theory of closed-end fund pricing (Ross 2002), states that the significant discounts observed in closed-end fund markets are attributable to "offsetting factors" such as fees, managerial skill, asset liquidity, and tax liabilities (Alexander and Peterson 2016). These factors, either individually or in combination, can push a fund towards a premium or a discount to NAV. As such, this model implies that a fund trading at a discount is still trading at its fundamental value, only its intrinsic value has deviated from NAV due to offsetting factors. The most commonly discussed factors are:

[1] The survey done by Cherkes (2012) summarizing the literature on the CEF puzzle was invaluable in navigating this body of research and much of the content in the following section (2.2) was derived from his paper.

Managerial Fees: Investors in closed-end funds pay small annual fees to the fund managers. Neoclassical theory dictates that, when discounted in perpetuity, managerial fees diminish future cash flows such that the fundamental value of a closed-end fund is less than that of its net asset value. When Malkiel (1977) published the first evidence of the closed-end fund puzzle, he discussed manager fees as a possible explanation but ultimately rejected it as the magnitude of these fees (usually less than 1%), when discounted, was not enough to explain the observed closed-end fund discounts (around 15%). Ross (2002) reexamined Malkiel's findings on managerial fees and found that, using more modern discounting methods, managerial fees could come to explain the entirety of the discount. Ross's evaluation was purely theoretical and, while it held up in a small number of case studies (such as Tri-Continental Corp.), it fails to explain why open-ended funds with similar fees trade at net asset value (Berk and Stanton 2007).

Managerial Skill: As detailed by Berk and Stanton (2007), the funds of managers with outstanding security selection ability should be priced at a premium, as the expected earnings of a well-managed fund is higher than that of a poorly-managed one. Therefore, the intrinsic value of a well-managed fund is higher than a poorly-managed one despite not being reflected in its net asset value. The Berk and Stanton (2007) model attributes discounts specifically to the "tradeoff between managerial ability and fees," and is successful in explaining some of the features of the closed-end fund puzzle, however model requirements with respect to return predictability contradict the findings of Pontiff (1995).

Asset Liquidity: Asset liquidity within a neoclassical model can cause both premiums and discounts in closed-end funds. Discounts arise from the barriers to arbitrage that arise when assets are infrequently traded (Gemmill and Thomas 2002). Premiums in closed-end funds holding illiquid assets can arise from clientele effects. Closed-end funds provide a vehicle for individual investors to hold assets they would otherwise be unable to (Cherkes 2003, Malkiel and Xiu 2005), and the cost of owning illiquid assets is reduced as shares of closed-end funds can be traded without the underlying assets changing hands (Cherkes, Sagi, and Stanton 2009).

Taxes: Because closed-end funds pass capital gains taxes through to investors, large amounts of unrealized gains on assets can hurt closed-end fund prices because of the taxes to be passed onto shareholders upon the assets' sale (Malkiel 1977, 1995).

Behavioral Model

The behavioral theory of closed-end fund pricing attempts to address some of the shortcomings of the neoclassical school of thought. The key distinguishing factor of the behavioral theory of closed-end fund pricing is that the price of a closed-end fund does not always align with its fundamental value. Instead, the presence of irrational investor behavior can cause closed-end funds to trade away from their fundamental value.

De Long, Shleifer, Summers, and Waldmann (1990) proposed a market model in which market participants are either rational investors or noise traders. Rational investors trade based off of reasonable expectations about future asset cash flows and look for and take advantage of arbitrage opportunities. Furthermore, De Long et al. assume rational investors to have short-term horizons and are exposed to short-term fluctuations in asset prices, as is suggested by real-world limitations such as margin or liquidity needs.

In contrast, noise traders are those market participants that form irrational expectations based off of sentiment. Summarizing this model, and drawing off of Black (1986), De Long et al. explain noise traders mistake public or inconsequential information as an investing edge and trade off of these irrelevant signals, overvaluing securities in some periods and undervaluing them in others. De Long et al. make the assumption that the sentiment upon which noise traders act is stochastic and unpredictable. Because the short-horizon rational investors are exposed to short-term price fluctuations, the stochastic influence of noise traders presents an additional risk.

This added risk from noise traders, according to the behavioral theory, is what causes closed-end funds to trade mostly at a discount. The increased risk of closed-end funds as a result of noise traders would require a higher rate of expected return, pushing their prices down. De Long et al. propose the pricing-in of this additional sentiment risk is a driver of persistent discounts, with fluctuations above and below the equilibrium discount driven by stochastic changes in the sentiment of noise traders. The behavioral explanation for the premia seen in a minority of funds is that a large concentration of irrational investors with a large enough positive sentiment swing can dominate the equilibrium discount and result in a price above NAV.

Lee, Shleifer and Thaler (1990), in addition to highlighting some of the above relevant features of the De Long et al. model, find that an additional element, differential clienteles, is needed in order to apply it to closed-end funds. Because changes in noise trader market sentiment for closed-end fund must move with that of other securities, changes in sentiment will be reflected in both the closed-end fund market price as well as that of its underlying securities. If the NAV and the share price of closed-end funds are both affected equally, sentiment would ultimately have no effect on discount. However, if noise traders were more active in buying and selling closed-end funds than in the funds' underlying assets, changes in sentiment would have a greater impact in the closed-end fund market than that of the underlying. Therefore, it is the different clienteles that dictate how changes in sentiment translate into closed-end fund discount fluctuations.

Subsequent studies, as summarized by Gasbarro, Johnson, and Zumwalt (2003), have brought up evidence that supports the behavioral model. Hanley, Lee, and Seguin (1995) found that closed-end funds IPOs are marketed to poorly informed investors. Targeting of poorly informed customers by closed-end funds could be a potential driver of differential clientele, as such advertising is not necessarily happening in the markets for the underlying assets. Furthermore, it offers an explanation for one part of the closed-end fund puzzle - the existence of closed-end fund IPOs. Despite closed-end fund initial offerings mostly being priced at a premium (due to underwriting costs being passed through), and all evidence indicating funds will settle at a discount, investors still participate, as pointed out be Lee et al. IPOs marketed to poorly informed investors with irrational expectations would help explain this behavior. Further evidence in support of the behavioral model is presented in Brown (1999), mainly the finding that investor sentiment has a strong correlation with closed-end fund price volatility.

According to the survey by Cherkes (2012), a consensus has yet to be reached for this long-standing puzzle. For more background on the closed-end fund puzzle and proposed explanations, reference Cherkes (2012).

2.3 Business Development Companies

Traditional closed-end funds have been the subject of extensive economic research, however business development companies, a specific subset of closed-end fund, have yet to receive any level of investigation with respect to the CEF puzzle. The contribution of this paper is in the application of the BDC market as new evidence for the closed-end fund puzzle, and the following subsection presents the justification for doing so.

What are Business Development Companies?

Business development companies (BDCs) are closed-end funds whose assets consist mostly of illiquid growth-oriented loans to small and medium sized businesses. These loans are typically rated as sub-investment grade due to the uncertainty associated with lending to growing small businesses. Shares of many, but not all, BDCs are publicly traded, and portfolio proceeds in the form of dividends and capital gains are distributed to fund investors. Due to their regulatory status as a regulated investment company, BDCs are not required to pay income taxes but must adhere to a number of limitations regarding leverage, asset diversity, and income distribution as well as provide operational assistance to portfolio companies.

As summarized in Beltratti and Bock (2018), the emergence of BDCs began in the 1980s, at a time when bank balance sheets were recovering from excess lending of the 1970s. Small and medium sized business had difficulty accessing credit as the constrained banks were no longer able to adequately service middle-market demand. Alternative sources of credit, such as mutual funds or traditional closed-end funds, were constrained as well; mutual funds were capped at 15% illiquid investments and traditional closed-end funds were unwilling to take on the high costs of servicing and maintaining loan portfolios given CEF leverage regulations. It was at this point Congress passed an amendment to the Investment Company Act of 1940 (ICA) which laid out the framework for the creation of BDCs to help bring public money into smaller businesses (Beltratti and Bock 2018; Boehm and Friedberg 2008). While remaining a relatively niche industry for the first 30 years of its existence, BDCs have experienced rapid growth in the decade since the global financial crisis. The drivers of this growth and its connection to the salience of studying BDCs is discussed later in this section. Today, publicly traded BDCs cumulatively manage around $60 billion in capital, or around 7% of the so-called "leveraged loan" market within which they operate (SP LCD and Wells Fargo Securities, LLC in Beltratti and Bock 2018).

Finally, given the similar role BDCs play in my analysis when compared to how CEFs feature in existing CEF puzzle literature, it is important to note the most significant distinctions between the two fund structures. First, while both CEFs and BDCs are considered investment companies under the ICA, and both pay no income tax as a result of this distinction, BDCs are generally given more flexibility in leverage use and management structures in exchange for more stringent SEC reporting requirements (Boehm and Friedberg 2008). Furthermore, CEFs tend to invest in markets (such as public equities and bonds) more liquid and less opaque than the middle-market loan sector in which BDCs operate.

Why Explore BDCs?

The justification behind my investigation of the BDC market with respect to the closed-end fund puzzle has two primary components.

First, the most clear-cut reason for exploring BDC data is that BDCs generally have had very little coverage within academia and none with respect to the closed-end fund puzzle.

This is despite BDCs having the key features required for studying the puzzle including an observable NAVs and market price. It requires mention, however, that there are legitimate reasons for the BDC sector's lack of coverage. First, the BDC data is small relative to CEFs; my cohort consists of quarterly data for 56 BDCs over 4 years while there are over 500 actively traded CEFs, with daily pricing data dating back decades. Furthermore, whether or not NAV is an accurate representation of fundamental value is a more serious question in the BDC data than in the CEF market, as CEF NAV is more valuable due to up-to-date pricing from the liquidity of underlying. BDC NAV values are self-reported and require estimating the fair market value for loans which rarely trade hands.

Despite these limitations, relevant findings can still be drawn from the data. The BDC cohort I describe in the next section consists of 961 premia observations, which while much smaller than the CEF data typically used is still enough for significant conclusions. Moreover, any findings would be made more interesting by possible comparisons across markets. Because there are important differences between CEFs and BDCs, conclusions drawn in the BDC market can elicit further information by comparing to similar findings in the CEF market. As a specific example, consider the ICA requirement that BDCs provide significant operational assistance to portfolio companies. This confers responsibility to fund managers beyond just portfolio selection, and as a result might overweight manager ability as a driver of premium in comparison to CEFs. Comparisons across markets like this would require intimate knowledge of the CEF data and as a result is outside the scope of this paper, but nonetheless highlight the contribution of exploring the CEF puzzle within distinct fund structures like BDCs.[2]

Beyond providing new and relevant evidence on the closed-end fund puzzle, BDCs are worth understanding more than ever due to recent developments in the loan markets in which they operate. The dominant trend in middle-market loans since the global financial

[2]Another BDC - CEF distinction with which to compare results across the two markets is liquidity of underlying, as Cherkes, Sagi, and Stanton's (2009) theory predicts the relative illiquidity of BDC portfolios when compared to equity and bond CEFs ought to yield a higher premium as a result of bridging a more significant liquidity gap.

crisis has been the retreat of regulation-constrained banks as middle-market lenders and their replacement by non-bank sources of credit (Beltratti and Bock 2018). The agents replacing the regulation-constrained banks include collateralized loan obligations, loan mutual funds, and BDCs.

Often referred to as "shadow banks" due to the lack of oversight and regulation, these agents have grown rapidly in size over the course of the past decade. For example, in 2009, BDCs cumulatively held only around $11 billion in loans but by 2016 BDC cumulative AUM had increased to around $60 billion (SP LCD and Wells Fargo Securities, LLC in Beltratti and Bock 2018). This growth, in all forms of non-bank lending, has been accompanied by a considerable deterioration in lending standards and protections (Haunss 2018), and regulators have taken notice; the US Federal Reserve, the International Monetary Fund, and the Bank of International Settlements have all issued blunt warnings over the last few months (Warnings over Leveraged Loans 2019).

BDCs are still small in comparison to the other forms of non-bank lending that have sprung up since the global financial crisis, however they have potential to serve a big role as an additional window from which regulators can monitor the opaque smaller end of the non-investment grade loan market. Publicly-traded BDC shares offer a real-time indirect market valuation of portfolios of some of the most opaque and illiquid leveraged loans. Should fluctuations these market valuations reflect a reasonable expectation of changes in the value of respective BDC portfolios, BDC prices could serve as an additional perspective and an early warning indicator with which to oversee the lower end of the leveraged loan market.

3 Data

This section is divided into three parts. The first details the sources and construction of the data used in the analysis, the second describes key features of the data, and the third provides a preliminary discussion of the extent to which the features of the closed-end fund puzzle are present in BDC markets.

3.1 Sources and Construction

I obtain data for this thesis from two main sources. BDC net asset value, expense ratio, percent institutional ownership, average loan size, and NAV filing release date were all acquired from the Closed-End Fund Advisors (CEFA) database. CEFA is Registered Investment Advisory firm whose primary business is asset management for individual and institutional clients[3]. This data was obtained through CEFA's research and data service offerings. Net asset value is defined as the net assets of the fund divided by shares outstanding. Expense ratio is defined as the total operating expenses for a fund expressed as a percentage of average net assets. Percent institutional ownership is the portion of outstanding common shares owned by institutional investors. Average loan size is the total fair market value of a BDCs portfolio, divided by the number of assets, expressed in $ million. NAV filing release date is the earliest date the quarterly NAV data was made public, typically the SEC quarterly filing release date.

The CEFA data was obtained in the form of weekly cross-sectional Excel reports. The format of these Excel reports changed a number of times over the course of the sample, where the earliest configuration lacked certain fields including NAV filing release date. NAV filing release date is critical to preventing look-ahead bias, so in order to account for this I used the SEC's Electronic Data Gathering, Analysis, and Retrieval system (EDGAR) to manually populate the missing data. For each BDC, the three quarters with missing NAV release date items were populated with the release date of the relevant SEC quarterly filing.

BDC share price and shares outstanding data used in this paper is from The Center for Research in Security Prices (CRSP) database on the Wharton Research Data Services

[3]Copyright 2019 Closed-End Fund Advisors, Inc. All rights reserved. For institutional use only. The information contained herein: (1) is proprietary to Closed-End Fund Advisors and/or its content providers; and (2) is not warranted to be accurate, complete, or timely. Neither Closed- End Fund Advisors, Inc. nor its content providers are responsible for any damages or losses arising from any use of this information. Past performance is no guarantee of future results.

portal. The CRSP US Stock database contains end-of-day prices for equities on all major US exchanges.

Aggregating the data involved choosing a specific date for every quarter on which to calculate premium. Because the NAV release date for a given quarter was usually at month after the quarter's end, calculating premium at any point within a quarter or immediately after would involve using NAV data that would not be known at the time. To avoid this pitfall and source of bias, I calculated premium on the next business day following the release of a BDC's NAV. For example, the Q4 2017 NAV of PennantPark Investment Corp (NASDAQ: PNNT) of $9.1 per share was released on February 7, 2018. Using the February 8, 2018 share price of 6.77, I determine that PNNT had a Q4 2017 discount of $6.77/9.1 - 1$ or -25.6%. Rarely, NAV will be released through a press release in the weeks prior to the BDCs SEC filing. In these few cases, the date on which premium is calculated will not be the business day directly after NAV is made public. While slightly inconsistent, I do not expect this to pose any major issues.

The next step in aggregation was adjusting for name changes and other firm actions. In the sample there are eight instances of funds changing their names and tickers while continuing to operate the same loan portfolio. For the purposes of this paper, there is no use in maintaining these nominal distinctions and I therefore consolidate all data for a fund with multiple names over the sample under its most recent ticker. There are also a number of fund liquidations (2) and mergers (2) resulting in funds missing data after certain dates, as well as nine funds founded within the sample period missing data before their founding dates. While not requiring any specific manipulation, as a result of these BDCs the fund panels are unbalanced.

3.2 Description

After the above adjustments, the data consisted 20 quarters of data across 56 BDCs, totalling 961 independent observations. Table 1 provides, for each field, (1) the number of observations, (2) the equally weighted mean of within-fund averages, (3) the fund average market cap weighted mean of within-fund averages, (4) the standard deviation of the within-quarter equally weighted average across funds, (5) the standard deviation of fund averages, (6) the mean standard deviation within funds, and the (7) minimum and (8) maximum fund averages.

Table 1: Summary of BDC Data

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
x:	Obs	$\bar{\bar{x}}$	$\bar{\bar{x}}_{MCW}$	$\sigma(\bar{x}_t)$	$\sigma(\bar{x}_i)$	$\overline{\sigma_i(x_{i,t})}$	$\min(\bar{x}_i)$	$\max(\bar{x}_i)$
Premium	961	-0.114	-0.046	0.059	0.220	0.119	-0.537	0.573
Market Cap $M	961	641		69.0	1020	201	20.9	5,760
Expense Ratio	851	9.335	8.954	0.472	3.600	2.188	4.257	19.860
Inst. Ownership %	791	28.701	36.544	1.249	14.275	5.318	9.394	70.758
Avg. Loan Size $M	454	8.933	17.204	1.246	7.602	1.367	0.184	32.959

The mean BDC premium across all observations is -10.5%, however this does not take into account the unbalanced nature of the fund panels; premia from funds with more observations have larger influence on the global mean. Normalizing for this effect by giving each fund equal weight yields a mean fund average premium of $\bar{\bar{x}} = -11.4\%$ while using market cap weights yields a mean premium of -4.6% as shown in Column 3 of Table 1. Figure 1 shows the distribution of average fund premia is approximately normal. The mean within-fund standard deviation is 0.119, indicating BDCs experience significant premium fluctuation across time. Despite this, BDCs experience even more variation between funds (0.220) than within as BDCs exhibit a wide range of average premia. Certain funds, such as Main Street Capital Corp. (NYSE: MAIN) averaged large premia over the sample period (Mean: 57.3%, Std: 14.6%), while others, such as Firsthand Technology Value (NASDAQ: SVVC), averaged large discounts (Mean: -53.7%, Std: 13.69%)[4].

The global average BDC market cap is $649 million. Market cap varies very widely

[4]For more detail on premia of specific BDCs, see Appendix Table 13

Figure 1: Distribution of Fund Avg Premium and Market Cap

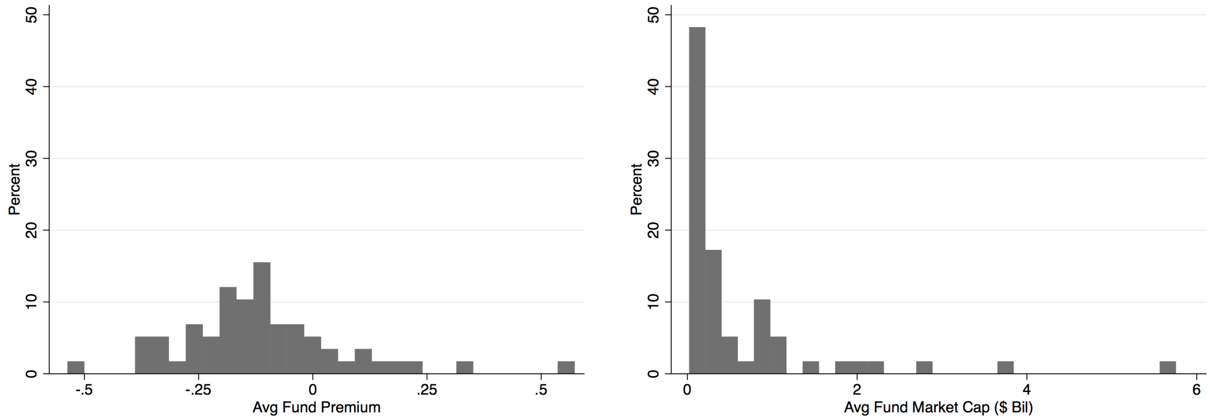

($\sigma(\bar{x}_i) = 1020$) between BDCs, with the smallest, Rand Capital Corp., valued on average at only $20 million and the largest, Ares Capital Corp., at almost $6 billion. The distribution is heavily right-skewed (See Figure 1). Furthermore, fund average log market cap seems to be correlated to average institutional ownership ($R = 0.487$), average mean log loan size ($R = 0.769$), and average log premium ($R = 0.538$). The source of this relationship and its implications are discussed in the next section.

Of the remaining data fields, some relevant observations include the following. The global average expense ratio is 9.46% of respective fund net assets. While the average BDC expense ratio has stayed stable over time ($\sigma(\bar{x}_t) = 0.0472$), it can vary to a greater degree between funds and within funds over time. Fund average levels of institutional ownership in the BDC sample range from as high as 70.8% to as low as 9.3%, with an average around 36.5%. This is very low in comparison to typical equity markets; the S&P 500 has an institutional ownership percentage of around 80% (Mcgrath 2017). The global mean average portfolio loan size is $9.1 million. Intra-fund variation is in average portfolio loan size is low ($\overline{\sigma_i(x_{i,t})} = 1.367$), meaning funds stick to investing in a certain facility size and they typically do not deviate from their preferred market. However, the large between-fund variation ($\sigma(\bar{x}_i) = 7.602$) means BDCs differ widely in the size of loans they invest in on average. Finally, it is important to note, for all data fields, that both Column 5 and Column 3 of Table 1 are greater than Column 6, indicating there is both less variation within funds and less variation in the equal weighted mean than variation between funds. The relatively high between-fund heterogeneity is significant as it implies that BDCs can generally be described as having distinct characteristics.

3.3 Presence of CEF Puzzle Features

The four tenents of the closed-end fund puzzle as detailed by Lee, Shleifer, and Thaler (1990) are: CEFs IPOs being priced at a premium, CEFs trading at a mid single-digit discount in the long run, CEF discounts fluctuating in a mean reverting manner over time, and CEF prices trading up towards NAV but not reaching NAV in advance of liquidation.

The first and the last observations are difficult to corroborate in the BDC data due to the limitations of the sample, both in terms of its quarterly frequency and constrained temporal scope. However, cursory examination using SEC EDGAR of the Form N-2 Registration Statement of the BDCs founded within the sample indicate that sales loads and underwriting expenses were passed through to initial investors through the IPO price, meaning the funds were initially purchased at a premium. Most of these funds now trade at a discount, providing indefinite evidence that the first behavior described by Lee et al. is corroborated by the BDC data.

Stronger statements can be made with respect to the second and third claims. Similar to CEFs, BDCs tend to trade at a mid single-digit discount in the long run ($\bar{\bar{x}}_{MCW} = -4.6\%$), supporting the presence of the persistent discount feature of the CEF puzzle in the BDC data. Similarly, the BDC data corrborates the presence of significant fluctuations in premium within funds, as the mean premium standard deviation is relatively high ($\overline{\sigma_i(x_{i,t})} = 11.9\%$). Whether or not these fluctuations are reversions to a stable mean or follow some sort of other trend is evaluated in the next section.

4 Methodology

The previous section showed that BDCs trade at a persistent discount and that this discount presents large fluctuations within funds over time. In this section, I present the methodology behind determining whether these fluctuations are mean reverting and whether the mean reversion is drive primarily by changes in price or by changes in NAV.

As discussed in earlier sections, neoclassical finance theory dictates that BDCs, after accounting for rational price distortions from features such as management fees and unrealized tax payments, should not trade at a discount or premium and instead should be priced in line with their fundamental value. Should a BDC trade at a discount or premium to NAV, one can expect the discount or premium to disappear as market forces bring market value (price) in line with fundamental value (NAV). The point of interest lies in whether this theoretical reversion is observed in the BDC data, and whether the reversion happens mainly due to movements in price or movements in NAV.

I use four different regression models to attempt to answer the above questions. For each model, I briefly motivate the choice, discuss the assumptions of the model in the context of BDC premia, and evaluate the performance of the model over the sample period. Results are presented for each model individually and interpreted as they relate to motivating changes from one model to the next.

4.1 Ordinary Least Squares (OLS)

The first model is an ordinary least squares regression of $\Delta \log \text{premium}_{i,t+1}$ on $\log \text{premium}$. This model makes no distinctions between funds and pools together all observations for analysis (with the exception of calculating clustered standard errors). The OLS model is represented by

$$\Delta \log \text{premium}_{i,t+1} = \alpha + \beta \log \text{premium}_{i,t} + \epsilon_{i,t+1}, \tag{1}$$

where $\Delta \log \text{premium}_{i,t+1}$ is the change in log premium of fund i from quarter t to quarter $t+1$, and log premium is defined by $\log \text{premium}_{i,t} \equiv \log \text{price}_{i,t} - \log \text{NAV}_{i,t}$. The coefficient β represents the relationship between log premium at time t and the change in log premium from quarter t to $t+1$. Should mean reversion behavior be present in the data, β will be negative and significant. Such a result would imply that one can reject the hypothesis that $\log \text{premium}_{i,t}$ has no relationship with $\Delta \log \text{premium}_{i,t+1}$. The error term is represented by $\epsilon_{i,t+1}$, and under the assumptions of OLS is assumed to be normally distributed and have a conditional mean $E(\epsilon_{i,t+1} | \log \text{premium}) = 0$. The model also contains a constant α, which signifies the amount by which one would expect log premium to change from one quarter to the next should the current log premium be 0 (meaning the BDC trades at exactly NAV). This is a rare scenario, as BDCs will rarely trade at an exactly 0 premium for any significant amount of time; the constant's more relevant interpretation is in the calculation of the model's equilibrium premium. The model's equilibrium premium occurs when a BDC has reached the premium level at which the premium is not expected to change. The equilibrium premium for this model, $\log \overline{\text{premium}}$, can be found be setting $\Delta \log \text{premium}_{i,t+1} = 0$ and solving for $\log \text{premium}_t$:

$$\log \overline{\text{premium}_{i,t}} = -\frac{\alpha}{\beta}. \tag{2}$$

An important note for this model is that α and β are assumed to be constant over time and the same for all BDCs, and $\epsilon_{i,t+1}$ is assumed to be a normally distributed random variable with a mean of 0. This means that the model assumes BDCs trade at some persistent equilibrium premium and that this equilibrium premium is expected to be consistent across funds. This is an important assumption of this preliminary OLS model, and whether or not this is an accurate representation of BDC premium behavior is discussed later.

The results of this regression can be found in the first column of Table 2. Change in log premium has a significant inverse relationship with $\log \text{premium}_t$, meaning that a fund with a premium above $\overline{\text{premium}}$ will typically see their premium decrease and a fund discounted below $\overline{\text{premium}}$ will typically see its premium rise. Furthermore, using the values of α and β from Table 1, it can be seen that $\log \overline{\text{premium}} = -\frac{-0.0180}{-0.0693} = -0.259$. This means that the equilibrium discount to which all BDCs are assumed to revert is $e^{-0.259} - 1 = -0.22$.

Table 2: Regression Table (OLS)

	(1) $\Delta \log \text{premium}_{i,t+1}$	(2) $\Delta \log \text{price}_{i,t+1}$	(3) $\Delta \log \text{NAV}_{i,t+1}$
$\log \text{premium}_{i,t}$	-0.0693***	-0.0395**	0.0298*
	(0.0150)	(0.0138)	(0.0124)
Constant	-0.0180***	-0.0248***	-0.00679***
	(0.00390)	(0.00471)	(0.00183)
Observations	903	903	903

Note: Standard errors clustered at fund level in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Now that, according to the OLS model and the BDC data, mean reversion in BDC discounts has been determined to exist, the next step is to determine the amount that each component of premium, price and NAV, contribute to the mean reversion. If a fund experiences, for example, an increase in the size of its premium, this means that price has risen relative to NAV. What one does not immediately know is how price and NAV have changed independently of one another; if, for example, premium rises, this could indicate that price has increased and NAV has decreased, that both increased but price increased more, or that both decreased and price decreased less. This distinction is important because, in the real world, it is the independent movements of price and NAV that are the most salient.

If the findings are that premium mean reversion is reliably driven by changes in price, then there are two important conclusions. First, it would serve as evidence of the behavioral model of closed-end fund pricing, as it is the price fluctuating around a relatively stable fundamental value, indicating investors are ignorant of the true value of the fund. Second, it would mean that knowing both the magnitude of the premium of a fund as well as to its equilibrium $\overline{\text{premium}}$ would enable someone to confidently future predict changes in the price of a BDC.

If the data show that the premium mean reversion is reliably driven by changes in NAV, then it would imply that BDC premia contain information on the value of the underlying loans not yet made public. This would support a neoclassical model of closed-end fund pricing, where investors understand the fundamental value and changes in price are a reaction to unpublished changes in NAV rather than irrational swings. Such a conclusion would also mean that anyone who knows the premium of a fund and its equilibrium $\overline{\text{premium}}$ would be able to reliably predict changes in the NAV of a BDC. This would be valuable to regulators attempting to monitor the middle-market loan space.

Just as premium mean reversion was modeled using OLS, the issue of whether price or NAV drive changes in premium can also be viewed through an OLS regression. It is a similar model to previously, however now using $\Delta \log \text{price}_{i,t+1}$ and $\Delta \log \text{NAV}_{i,t+1}$ as the endogenous variables instead of $\Delta \log \text{premium}_{i,t+1}$. The new OLS models are

$$\Delta \log \text{price}_{i,t+1} = \alpha + \beta \log \text{premium}_{i,t} + \epsilon_{i,t+1} \tag{3}$$

and

$$\Delta \log \text{NAV}_{i,t+1} = \alpha + \beta \log \text{premium}_{i,t} + \epsilon_{i,t+1}, \tag{4}$$

where α represents the expected change in log price or log NAV given a log premium of 0, the error term $\epsilon_{i,t+1}$ is still assumed to be a normally distributed random variable with a mean of 0, and β signifies the relationship between log premium at quarter t and the change in log price or log NAV from quarter t to quarter $t+1$. Because $\Delta \log \text{premium}_{i,t+1}$ can be linearly separated into the difference of $\Delta \log \text{price}_{i,t+1}$ and $\Delta \log \text{NAV}_{i,t+1}$, the following must hold:

$$\beta_{\text{premium}} = \beta_{\text{price}} - \beta_{\text{NAV}} \tag{5}$$

$$\alpha_{\text{premium}} = \alpha_{\text{price}} - \alpha_{\text{NAV}} \tag{6}$$

The results of the regressions can be found in Table 2. The results suggest, that within the OLS model, the discount mean reversion behavior is driven to a significant extent by both changes in price and NAV. The changes in price are somewhat steeper than the changes in NAV, and as a result changes in price compose a slightly larger share of the discount mean reversion than do changes in NAV.

This indicates some support for both the neoclassical theory and the behavioral theory of closed-end fund pricing. Furthermore, it suggests that, knowledge of both the premium of a BDC and its $\overline{\text{premium}}$ would allow someone to reliably predict some changes in both price and NAV. Such a conclusion is useful to both investors and regulators alike, however its implementation is held back by the impossibility of knowing the equilibrium premium for a fund at a given time. In this model, the equilibrium $\overline{\text{premium}}$ is calculated using data from across the entire sample period (Q3 2014 - Q2 2018). It would be impossible to calculate this equilibrium $\overline{\text{premium}}$ at any point during the sample given the impossibility of have the future data required. To circumvent this look-ahead bias, as well as to simulate use of this model in a more realistic prediction scenario, I partition the data into a training sample (in-sample) and a testing sample (out-of-sample), and see what predictive power remains estimates of the equilibrium $\overline{\text{premium}}$ can be made using only the data known at a given time.

The training sample consists of the first 10 quarters of data (Q3 2013 through Q4 2015), while the test sample is the other 10 quarters of data (Q1 2016 through Q2 2018). To evaluate the predictive power of these models, I estimate parameters for the OLS model in the in-sample period before attempting to predict out-of-sample changes in log price and log NAV using those same parameters that were estimated using the in-sample data. The predicted values that this yields, represented by $\Delta \widehat{\log \text{price}}_{i,t+1}$ and $\Delta \widehat{\log \text{NAV}}_{i,t+1}$, are then compared to the observed values over the out-of-sample period by regressing

$$\Delta \log \text{price}_{i,t+1} = \alpha + \beta \Delta \widehat{\log \text{price}}_{i,t+1} + \epsilon_{i,t+1} \tag{7}$$

and

$$\Delta \log \text{NAV}_{i,t+1} = \alpha + \beta \Delta \widehat{\log \text{NAV}}_{i,t+1} + \epsilon_{i,t+1}. \tag{8}$$

If the models trained in the in-sample period were perfect estimators of changes in log price and log NAV in the out-of-sample period, then one would expect both these regressions to have an intercept of 0, a β of 1, and a nonexistent error term. The results can be found in Table 3.

Table 3: Out-of-Sample Predicted vs. Observed Regression (OLS)

	(1) $\Delta \log \text{price}_{i,t+1}$	(2) $\Delta \log \text{NAV}_{i,t+1}$
$\Delta \widehat{\log \text{price}}_{i,t+1}$	1.134	
	(0.530)	
$\Delta \widehat{\log \text{NAV}}_{i,t+1}$		0.561
		(0.296)
Constant	0.0478	0.00610
	(0.0232)	(0.00435)
Observations	428	428
RMSE	0.1278	0.0395

Note: Standard errors clustered at the fund level in parentheses

Clearly, the predictive power of the OLS estimator leaves much to be desired when used out-of-sample. Although the β coefficients are both positive, significant, and somewhat close to 1, and the constant is near 0, Figure 2 shows that for any given predicted change in log price or log NAV, the actual value can vary widely from what the model predicted.

While these results match expectations in terms of mean reversion, the predictive power of the auxiliary regressions on the components of premium was small when taken out-of-sample. The OLS model from which these results were derived is likely a bad fit for the data for two reasons.

First, the assumption that α and β remain constant across funds seems inconsistent with both the data and theory. Figure 3 shows the average discount across the sample period for every BDC. It is clear that, while many funds do trade at a persistent moderate discount,

Figure 2: OOS Predicted Values vs. Observed Values (OLS)

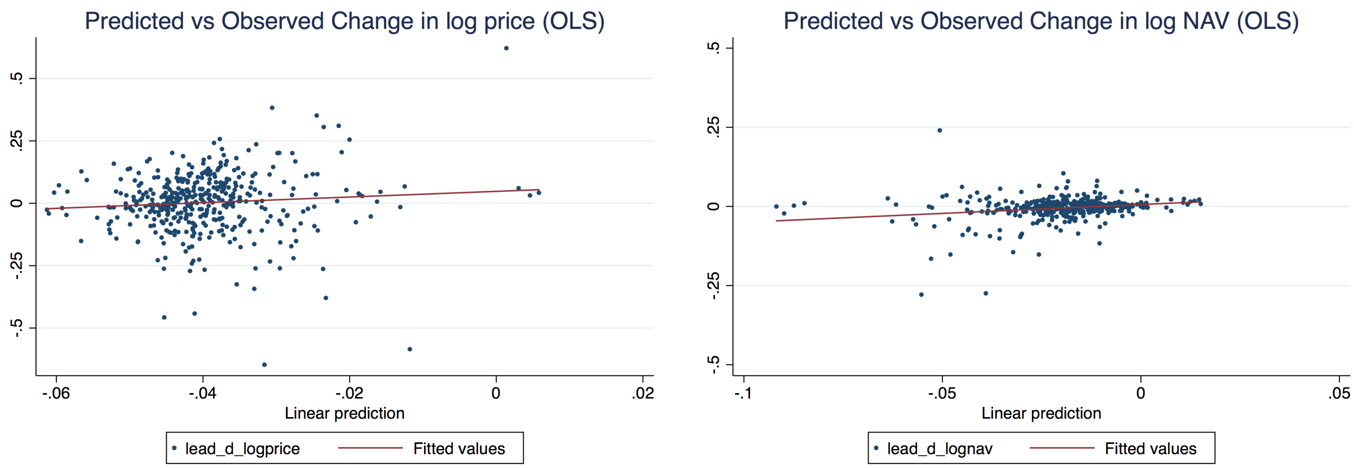

Note: x-axis scale distorted for interpretability

there is a large amount of variation in average discounts, and many trade at a premium. This seems to run counter to the OLS model's assumption of a market-wide equilibrium discount. A possible way around this would be to include a fixed effects term which represents the time-invariant fund idiosyncrasies that might cause equilibrium $\overline{\text{premium}}$ to be different for each fund.

Figure 3: Distribution of Historical BDC Premia

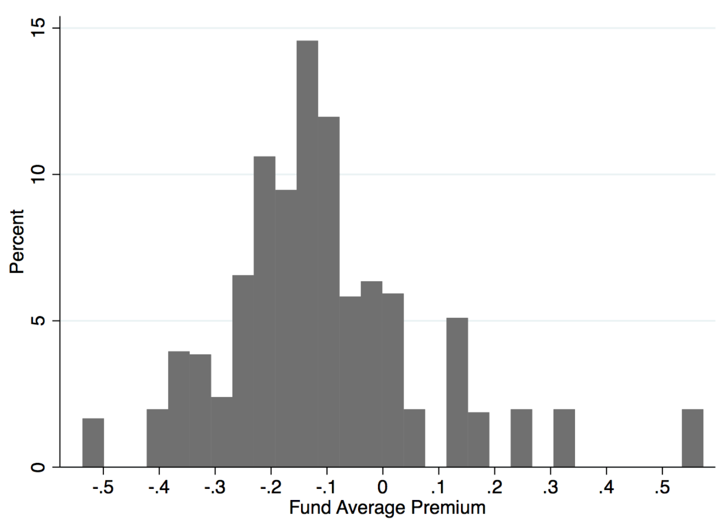

Second, neoclassical theory acknowledges that certain fund characteristics, such as expense ratio or portfolio asset liquidity, can influence the equilibrium discount. BDCs expense ratio, market cap, and similar features vary significantly between funds, meaning that the theoretical equilibrium $\overline{\text{premium}}$ ought to similarly vary between funds. Inclusion of these variables would be supported by both theory and data. As shown in Figure 4, cursory evaluation of market cap data in relation to premium seems to indicate a positive correlation.

Figure 4: Average log market cap by fund vs. average log premium by fund

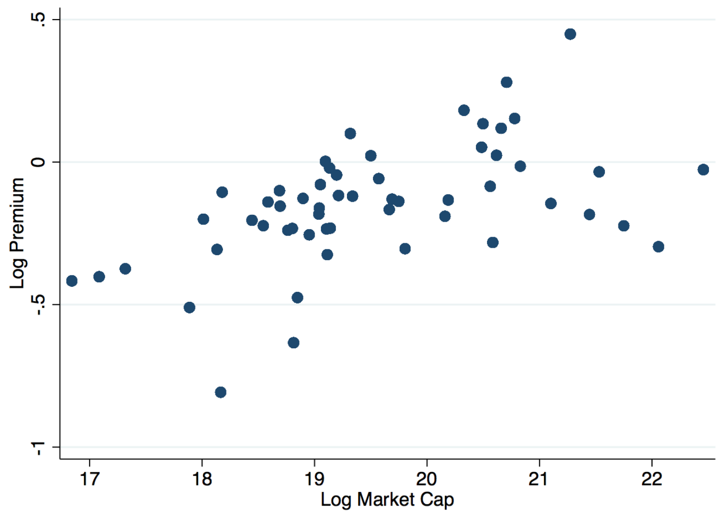

The OLS model has strength in its simplicity. The relatively small dataset means overfitting is a risk, and a parsimonious model is important to consider. That being said, the above shows two clear directions of improvement: a fixed effects model and the addition of fund characteristics to the original OLS model.

4.2 Fixed Effects (FE)

As mentioned in the OLS subsection, a major weakness of the OLS model in the context of modeling BDC premium behavior is that it doesn't allow for variation in equilibrium $\overline{\text{premium}}$ between funds. This contradicts theory, due to the number of time-invariant fund characteristics such as management fee or brand name that could influence discount. It also contradicts reasonable interpretation of the data, given the large amount of variation in premium between BDCs compared to the variation in premium within BDCs (See Table 1). To formalize the need for fixed effects, I used a joint test on an OLS regression with a dummy variable for each fund. The p-value of the joint test was 0.000, indicating that one can reject the hypothesis that the dummy variable coefficients are all jointly equal to 0. From this one might conclude that the addition of fixed effects terms to the model would be an appropriate change. Incorporating a fixed effect term v_i gives us the regression

$$\Delta \log \text{premium}_{i,t+1} = \alpha + v_i + \beta \log \text{premium}_{i,t} + \epsilon_{i,t+1} \tag{9}$$

where α is the constant and β is the relationship between $\Delta \log \text{premium}_{i,t+1}$ and $\log \text{premium}_{i,t}$ after having removed the effect of time-invariant fund-specific characteristics. $\epsilon_{i,t+1}$ is the error term, still assumed to be a normally distributed random variable with mean 0. A constraint must be imposed on either α or v_i in order to prevent the infinite solutions that can arise, for instance, by adding an amount to α and subtracting that same amount from all of v_i. The choice of constraint should be informed by the preferred interpretation of α and v_i. To give α a similar meaning in this fixed effects model as in the previous model, I can impose the constraint: $\sum_{i=1}^{N} v_i = 0$. This constraint has no effect on β, but it makes it such that α must take the average value of the fixed effects v_i. Therefore, α can be interpreted as the amount by which one can expect log premium to change for the average BDC should log premium be 0.

Table 4: Regression Table (Fixed-Effects)

	(1) $\Delta \log \text{premium}_{i,t+1}$	(2) $\Delta \log \text{price}_{i,t+1}$	(3) $\Delta \log \text{NAV}_{i,t+1}$
$\log \text{premium}_t$	-0.242***	-0.215***	0.0269
	(0.0485)	(0.0407)	(0.0143)
Constant	-0.0426***	-0.0498***	-0.00719***
	(0.00692)	(0.00580)	(0.00203)
Observations	903	903	903

Note: Standard errors clustered at fund level in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

The results of the regression in Equation 9 can be found in the first column of Table 3. The relationship between $\log \text{premium}_{i,t}$ and $\Delta \log \text{premium}_{i,t+1}$ remains both negative and significant. However, with the BDC fixed effects, the magnitude of the coefficient is much larger. This means that when the equilibrium $\overline{\text{premium}}$ is allowed to differ between funds, the premium returns to the equilibrium faster than if all funds use the same equilibrium as in OLS.

Equilibrium discount for the a given BDC is still reached at the point when the change in log premium is set equal to 0, however because of the estimated fixed effects term v_i, the equilibrium $\overline{\text{premium}_i}$ will be distinct for each fund:

$$\log \overline{\text{premium}_{i,t}} = -\frac{\alpha + v_i}{\beta} \tag{10}$$

Using the above formula to determine the distribution of equilibrium premia according to this fixed effects model yields the distribution of equilibrium premia represented in Figure 5. The distribution is noticeably wider than that of the observed historical premia shown in Figure 3. This is likely due to outlier estimations, a result of the large number of free parameters in comparison to the small size of the dataset.

Now that mean reversion has been shown in the fixed effects model of BDC premium behavior, the next step is to isolate the effects of independent changes in log price and log NAV. To do so, I similarly split up the original fixed effects premium regression represented in Equation 9 into two component regressions

$$\Delta \log \text{price}_{i,t+1} = \alpha + v_i + \beta \log \text{premium}_{i,t} + \epsilon_{i,t+1} \tag{11}$$

and

$$\Delta \log \text{NAV}_{i,t+1} = \alpha + v_i + \beta \log \text{premium}_{i,t} + \epsilon_{i,t+1}. \tag{12}$$

The results of the fixed effects component regressions can be found in Table 4. Similar

Figure 5: Distribution of FE Equilibrium Premium

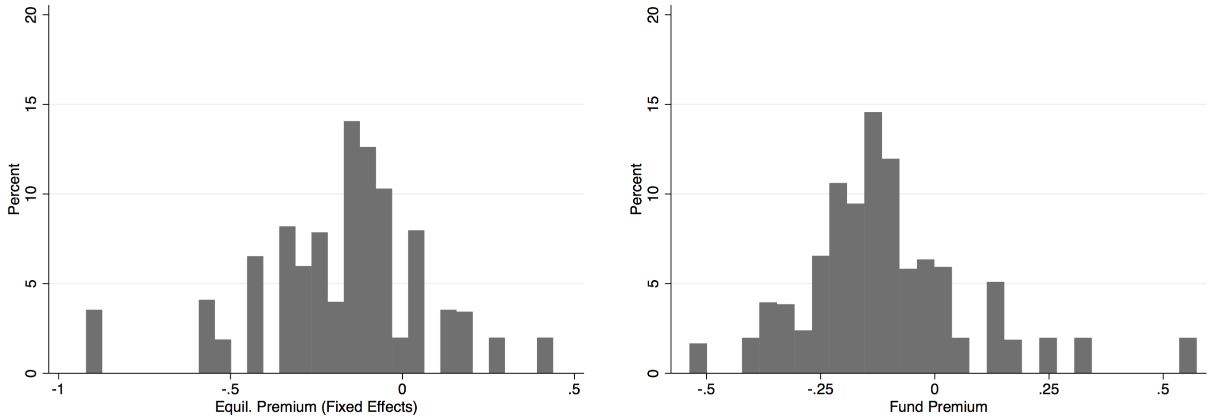

Note: Figure 3 reproduced on the right for ease of comparison

to with the OLS model, the results show that in this model both changes in price and changes in NAV play a role in discount mean reversion. However, after controlling for BDC fixed effects, price plays a much larger role than before, and NAV is no longer significant at an alpha of 0.05, meaning one can no longer reject the null hypothesis that $\log \text{premium}_{i,t}$ is unrelated to $\Delta \log \text{NAV}_{i,t+1}$. Such strong evidence for changes in price being the driving component of BDC discount mean reversion suggests perhaps the presence of noise/sentiment trading in the BDC market. Furthermore, as before, if the premium of a BDC relative to that BDC's equilibrium $\overline{\text{premium}_i}$ is known, one can reliably predict that it will mean revert back to the equilibrium and this mean reversion is driven primarily by price. This knowledge could be used as a trading strategy.

However, both of the models discussed so far suffer from look-ahead bias in the sense that the coefficients and constants incorporate information from over the whole sample, and using the model on any specific date would be impossible without knowing future information. The fixed effects model in particular, due to its large number of free variables and the small size of the dataset, is susceptible to incorporating a large amount of future information in the fixed effects terms. To get around this issue, I can again split the data into a training sample and testing sample and attempt to predict $\Delta \log \text{price}_{i,t+1}$ and $\Delta \log \text{NAV}_{i,t+1}$, storing the estimates in $\Delta \widehat{\log \text{price}}_{i,t+1}$ and $\Delta \widehat{\log \text{NAV}}_{i,t+1}$. To evaluate the accuracy of the predictions, I regress them against the observed values for those same rows.

The results of the above regressions can be found in Table 5. Overall, the results are very similar; the fixed effect estimates have a similar coefficient and constant in the regression with observed values as the OLS estimates. However, the one difference between the performance of the models is that the fixed effects model has a significantly higher RMSE in the prediction of $\Delta \log \text{price}_{i,t+1}$. The OLS model, despite its shortcomings in accommodating persistent differences in BDCs, outperformed the fixed effects model when taken out-of-sample.

Table 5: Out-of-Sample Predicted vs. Observed Regression (FE)

	(1) $\Delta \log \text{price}_{i,t+1}$	(2) $\Delta \log \text{NAV}_{i,t+1}$
$\Delta \widehat{\log \text{price}}_{i,t+1}$	0.162	
	(0.100)	
$\Delta \widehat{\log \text{NAV}}_{i,t+1}$		0.279
		(0.193)
Constant	0.00713	0.000399
	(0.00649)	(0.00275)
Observations	428	428
RMSE	0.1333	0.0432

Note: Standard errors clustered at the fund level in parentheses

The fixed effects model, while a good theoretical representation of how one would expect BDC premia to behave, is less useful in a practical setting. As a model to be used for prediction, it has limited usefulness due to the many free parameters it estimates which can lead to overfitting, especially if the data is as small and noisy as in this case. Consequently, the barebones OLS model outperformed it out-of-sample. As a model built to improve understanding of BDC premium dynamics, the fixed effects model is similarly unhelpful as it gives us no insight as to what are the time-invariant characteristics of the different funds. Since certain, measurable, characteristics of BDCs are potentially important in explaining BDC premia, the next step ought to be to include those measurements as exogenous variables in the original OLS mean-reversion model. This way, I realize the benefit of a more realistic model with more explanatory power than the original OLS model while retaining more predictive power than with fixed effects.

4.3 OLS with Market Capitalization (OLS w/ MC)

Much care must be taken when selecting which BDC characteristics to put into the OLS premium model. If BDC characteristics are added haphazardly and without theoretical justification, issues like collinearity and overfitting come into play. I chose four BDC characteristics that I both had data for and believed had the strongest theoretical case for consideration: expense ratio, average loan size ($mil), market cap, and institutional ownership. For each, I describe the theoretical case for inclusion as well as test whether the data support the theoretical relationship.

Expense ratio is defined as the total operating expenses of the fund divided by the BDC's net assets. Expense ratio encompasses selling, general, and administrative costs as well as management fees. In neoclassical theory, management fees and fund expenses represent costs that, when taken in perpetuity, slightly diminish future returns such that the fundamental value of the fund, as measured by the present value of future cash flows, is lower than the fundamental value as measured by the sum of the market value of the underlying assets. Theoretically, therefore, funds with higher expense ratios should have steeper discounts. Figure 6, however, shows that premium and expense ratio seem to have no obvious relationship despite the rational justification for one.

Institutional ownership, defined as the fraction of outstanding shares owned by institutional investors, plays a theoretical role through a behavioral model rather than a neoclassical one. In the behavioral model, institutional capital is controlled by the rational investors while individual investors are considered irrational noise traders. Noise traders in the market for BDC shares can induce volatility in BDC prices but not to BDC underlying assets. Therefore, low levels of institutional ownership should be related to higher price volatility, with more exaggerated swings around an equilibrium discount level.

Market capitalization is defined as the total market value of all of BDCs shares. Market cap, while potentially an important variable itself, is interesting because of its likely relationship with manager reputation and perceived skill. In the neoclassical model, investors are willing to pay a premium for manager skill as an effective manager will increase the future cash flows of the fund through security selection, and in the case of BDCs, operational assistance to portfolio assets. As a result, the fundamental value of the managed portfolio as measured by the present value of future cash flows, is higher than the fundamental value of the underlying securities on their own. As discussed in previous sections, anecdotal evidence suggests that reputation plays an important role in investment decisions in the BDC market, and the fund size is a strong proxy for reputation. This strong theoretical relationship between market cap and premium is supported by the data, as Figure 6 shows that BDC premium and market cap show a positive relationship.

Finally, average loan size ($mil) is defined as the total fair market value of a BDC's portfolio divided by the number of loans. Loan facility size, while potentially important in and of itself, was selected due to its likely high correlation with liquidity and the lack of a better proxy for portfolio asset liquidity. Smaller loans trade less frequently and have a smaller market of buyers and sellers. Liquidity is an important determinant of premium in the neoclassical model, either as a driver of discount from barriers to arbitrage, or as a driver of premium due to clientele effects from a fund holding in-demand but illiquid portfolio securities. As a result, one would expect an inverse relationship between premium and average loan size, as smaller, less-liquid loans would demand a higher premium. However, the data shows the opposite relationship (see Figure 6). This is likely a result of log average loan size and log market cap having a positive relationship ($R = 0.769$), potentially interpreted as a tendency for larger funds to invest in larger credit facilities.

To determine whether average loan size has any informativeness outside of its relationship to market cap, I regress log premium on log market cap and see whether any variation in the residuals can be explained by average loan size. Doing so quantifies the significant relationship between log market cap and log premium (See Figure 9), but shows average loan size has no ability to explain the residuals of the relationship ($p = 0.306$).

The above discussion of the BDC characteristics data supports moving forwards in two directions. The first is inclusion of only market cap in the OLS regression. Market cap has one of the stronger theoretical relationships with premium and has clearly the strongest empirical relationship with premium (see Figure 6). Only average loan size comes close to showing any relationship with premium, however I determined portfolio average loan size has no relationship with premium outside of its relationship with market cap (as larger funds buy larger loans). Furthermore, market cap is the variable with the most cross-entity variation between BDCs relative to its within-entity variation and fund average variation over time (see Table 1). This means that market cap behaves the most like the time-invariant, between-entity-variant factors that fixed effects are meant to pick up. Because market cap stands out above the other characteristics in the theoretical justification for its inclusion, its empirical relationship with premium, and its variety between funds, there is a strong case for its inclusion. I will not include the other variables in order to stay far from the pitfalls of collinearity and overfitting through including too many variables with individually less predictive power.

The second analysis justified by the above discussion of BDC characteristic data is categorizing funds based on institutional ownership levels and determining whether less institutional ownership (and subsequently higher noise trader ownership) is related with steeper discounts and more price-based mean reversion. The rest of this section discusses the inclusion of market cap in the OLS model, and the next section will discuss the evaluation of the effects of institutional ownership.

Incorporating a term for log market cap into the OLS regression yields

$$\Delta \log \text{premium}_{i,t+1} = \alpha + \beta_1 \log \text{premium}_{i,t} + \beta_2 \log \text{marketcap} + \epsilon_{i,t+1}, \tag{13}$$

Figure 6: Log premium vs. Fund Characteristics (Mean by Fund)

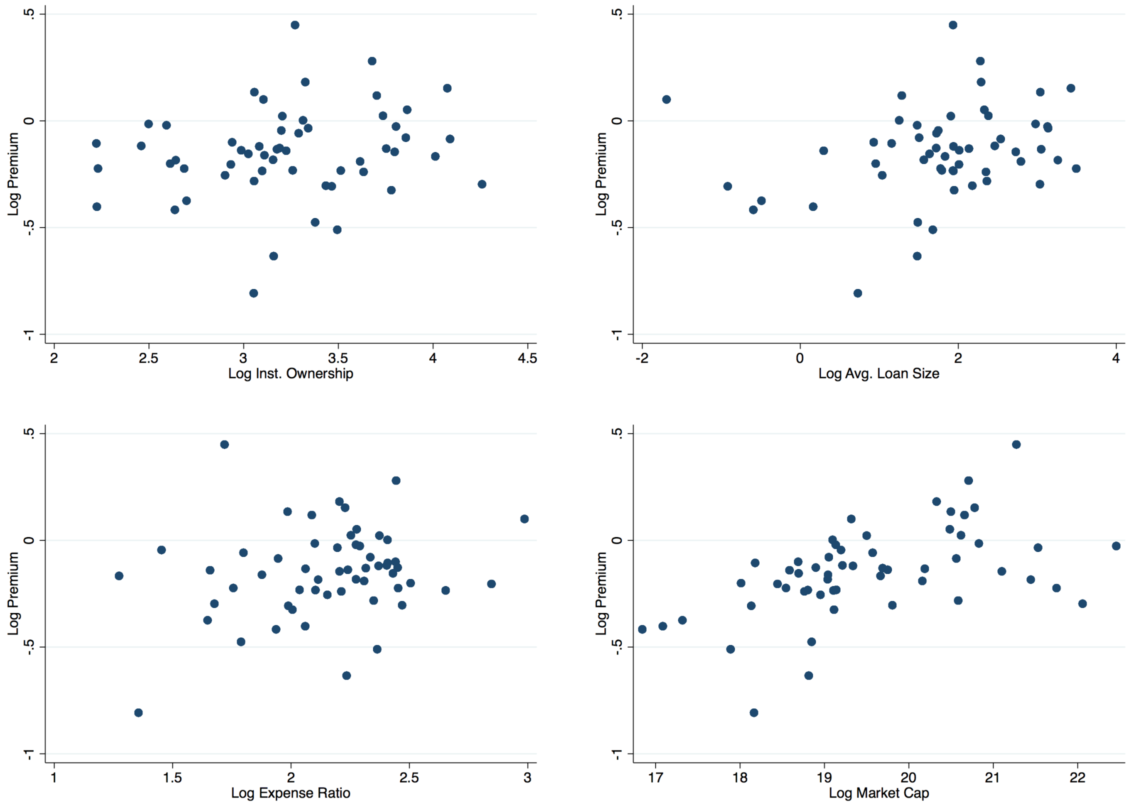

where β_1 and β_2 are the coefficients representing the relationship between their respective explanatory variables ($\log \text{premium}_{i,t}$ and $\log \text{marketcap}_{i,t}$) and the endogenous variable $\Delta \log \text{premium}_{i,t+1}$. The error term is $\epsilon_{i,t+1}$. An important note is that this model reflects an assumption that the left over error is random with mean 0, implying that that only a fund's market cap and current premium affect change in premium.

The results of the OLS regression with market cap can be found in the first column of Table 6. Mean reversion is present, as the log premium coefficient β_1 is both significant and negative. Furthermore, log marketcap has a significant effect on the mean reversion behavior, as the log marketcap coefficient β_2 is positive and significant. This means that, all else being equal, a larger BDC will have its changes in premium be slightly more positive than the change in premium of smaller BDC. There are implications of this when it comes to determining the model's equilibrium $\overline{\text{premium}_i}$, mainly that now market cap directly

Table 6: Regression Table (OLS w/ Market Cap)

	(1) $\Delta \log \text{premium}_{i,t+1}$	(2) $\Delta \log \text{price}_{i,t+1}$	(3) $\Delta \log \text{NAV}_{i,t+1}$
$\log \text{premium}_{i,t}$	-0.0840***	-0.0563***	0.0277*
	(0.0160)	(0.0121)	(0.0119)
$\log \text{marketcap}_t$	0.00630**	0.00718*	0.000885
	(0.00235)	(0.00311)	(0.00234)
Constant	-0.143**	-0.167**	-0.0244
	(0.0464)	(0.0623)	(0.0468)
Observations	903	903	903

Note: Standard errors clustered at fund level in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

influences the equilibrium $\overline{\text{premium}}_i$ that a BDC reverts to, now represented as

$$\log \overline{\text{premium}}_{i,t} = -\frac{\alpha + \beta_2 \log \text{marketcap}}{\beta_1} \tag{14}$$

where $\epsilon_{i,t+1}$ disappears as it is randomly distributed with a mean of 0. Figure 7 shows the distribution of equilibrium premia according to the OLS model with market cap as an additional explanatory variable. The distribution is more closely packed and smoother than the distribution of equilibrium premia produced by the fixed effects model and seems to indicate a more reasonable distribution than the fixed effects model. However, one strange aspect is that there are no equilibrium premia above 0, despite β_1 being negative and canceling with the negative sign out front in Equation 16.

Next, I look at the components of premium, price and NAV, through the lens of the OLS model with market cap and determine the role in premium mean reversion of their changes through the regressions

$$\Delta \log \text{price}_{i,t+1} = \alpha + \beta_1 \log \text{premium}_{i,t} + \beta_2 \log \text{marketcap}_{i,t} + \epsilon_{i,t+1} \tag{15}$$

Figure 7: Distribution of OLS w/ Market Cap Equilibrium Premium

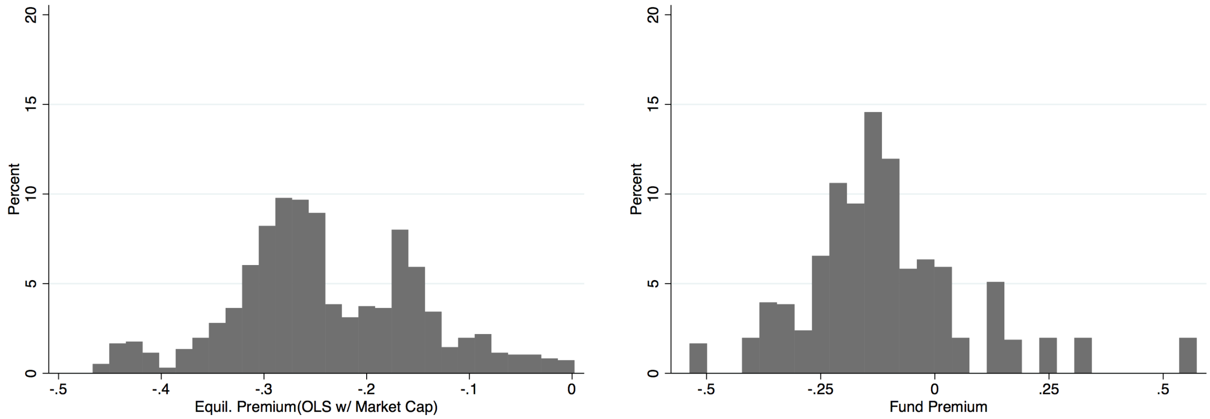

Note: Figure 3 reproduced on the right for ease of comparison

and

$$\Delta \log \text{NAV}_{i,t+1} = \alpha + \beta_1 \log \text{premium}_{i,t} + \beta_2 \log \text{marketcap}_{i,t} + \epsilon_{i,t+1}, \tag{16}$$

the results of which can be found in Columns 2 and 3 in Table 6. All of log premium, log marketcap, and the constant are significant for determining change in log price, while only log premium is significant for determining change in log NAV. This indicates that log marketcap plays a role in the price component of mean reversion, but not that of NAV. Specifically, market cap has a positive coefficient and therefore dampens the price mean reversion effect represented by the negative coefficient on $\log \text{premium}_{i,t}$; all else being equal, a larger BDC will have a more subdued the price component of mean reversion.

Table 7 contains the performance of the out-of-sample predictions. One can expect this model to perform better than other models in out-of-sample predictions because through only including one additional variable it is parsimonious enough to avoid the overfitting issues that befall the fixed effects model, while allowing for fund-specific time-invariant differences (market cap) and their accompanying effects on equilibrium discount. Consistent with this reasoning, the OLS with market cap estimator performed better than both the OLS and the fixed effects estimators when taken out-of-sample (RMSE was the lowest of the three).

Table 7: OOS Predicted vs. Observed Regression (OLS w/ Market Cap)

	(1) $\Delta \log \text{price}_{i,t+1}$	(2) $\Delta \log \text{NAV}_{i,t+1}$
$\Delta \widehat{\log \text{price}}_{i,t+1}$	1.240	
	(0.416)	
$\Delta \widehat{\log \text{NAV}}_{i,t+1}$		0.562
		(0.295)
Constant	0.0515	0.00610
	(0.0181)	(0.00434)
Observations	428	428
RMSE	0.1275	0.0395

Note: Standard errors clustered at the fund level in parentheses

4.4 OLS with Institutional Ownership (IO)

Along with market cap, the other BDC characteristic that I determined to be worth investigating was institutional ownership. Institutional ownership is valuable as it can be interpreted as a proxy for the proportion of rational investors vs. noise traders. For this section, I operate off the assumption that institutional investors behave more like the rational investors described by the behavioral model than do retail investors. Under this assumption, knowing the proportion of institutional vs. individual investors yields a proxy that allows us to investigate the merits of the behavior model discussed in De Long, Shleifer, Summers, and Waldmann (1990) as it applies to BDCs.

The De Long et al. model suggests that low institutional investor ownership of BDCs (and therefore large presence of noise traders) induces stochastic, sentiment-driven volatility in BDC prices. Meanwhile, BDC NAV values are not be affected by the presence of noise traders in the BDC market, as NAV is determined by the leveraged loan market whose investor base consists of only institutional investors due to barriers to entry. The additional volatility induced by noise traders in BDC prices, but not NAV, yields three hypotheses testable using the BDC data: first, noise trader presence will result in higher variation in premia in funds with lower institutional ownership. Second, the additional risk of investing in BDCs with lower institutional ownership requires a higher expected return and therefore a steeper discount. Third, the additional variability posed by noise trader presence should make price the dominant component in premium mean reversions of BDCs with higher noise trader presence, and NAV the dominant component in premium mean reversions of BDCs with higher institutional investor presence.

To test the first prediction, I divide BDCs into a high average institutional ownership (IO) group and a low IO group using BDCs' average institutional ownership level over the course of the sample. Those with average IO above or equal to the median get placed in the high IO group, and those with IO less than the median get placed in the low IO group. The mean within-fund standard deviations of premium within each group are in Table 8. While the low IO group does have a higher variation in premium, it is not a statistically significant difference ($p = 0.298$).

Table 8: Comparison of Mean Std of Prem. between High IO and Low IO BDCs

	Obs	$\overline{\sigma_i(\text{prem}_{i,t})}$
Low IO	29	0.0156
		(0.0047)
High IO	29	0.0127
		(0.0028)
Diff		0.0029
		(0.0055)

Note: Standard errors in parentheses

To test the final two predictions, I adjust the high IO and low IO cohort definitions as follows. The high IO group is now composed of the funds that have IO above or equal to the median level for a given quarter. The low IO group is composed of funds with IO below the median level for a given quarter. This allows funds to change group should their IO rise or fall such that their original group assignment no longer accurately describes them. To designate this distinction, I define the indicator variable $\mathbf{I}_{i,t}$ as

$$\mathbf{I}_{i,t} \equiv \begin{cases} 1 & \text{IO}_{i,t} \geq \text{median}_t\{\text{IO}_{i,t}\} \\ 0 & \text{IO}_{i,t} < \text{median}_t\{\text{IO}_{i,t}\} \end{cases}, \quad (17)$$

where $\text{IO}_{i,t}$ is the percent institutional ownership for fund i during quarter t. To test the theoretical prediction that low IO funds trade at a steeper discount than high IO funds, I calculate the mean premium and standard error for the high IO and low IO group. These values are shown in Table 9, and seem to indicate that funds with relatively high levels of noise trader ownership (low IO) on average trade at a steeper discount to NAV than do high IO funds. An unequal variances t-test supports that the low IO BDC group has a significantly deeper mean discount than does the high IO group ($p = 0.001$).

Table 9: Comparison of Mean Premium of High IO and Low IO BDCs

	$\mathbf{I}_{i,t}$	Obs	Mean Premium
Low IO	0	398	-0.1444
			(0.0118)
High IO	1	415	-0.0982
			(0.0100)
Diff			-0.0462
			(0.0154)

Note: Standard errors in parentheses

The last prediction from the De Long et al. behavioral model testable using the BDC dataset is that high levels of noise trader presence will cause BDC premium mean reversion to be driven more by changes in price, and less by changes in NAV, due to the irrational sentiment risk present in the market for BDC shares but not in the markets for the underlying leveraged loans. If a BDC investor base is mostly noise, prices would reflect sentiment-driven expectations of future value and would fluctuate around a relatively stable NAV. Mean reversions, therefore, would be predominantly driven by price in a market with high noise trader presence. If a BDC investor base is mostly rational investors, then changes in price reflect reasonable expectations about unpublished changes in NAV; contractions in premium would be driven by published NAV adjusting towards the price to reflecting the unpublished information about NAV that the rational price had already picked up.

To evaluate this prediction, I expand upon the OLS model by introducing the high IO indicator $\mathbf{I}_{i,t}$ and an interaction term between $\mathbf{I}_{i,t}$ and $\log \text{premium}_{i,t}$. Similar to in the previous models, I regress all of log premium, log price, and log NAV on the set of predictor variables. The first will confirm that mean reversion is present in the OLS model with IO, the second and third will allow us to attribute the mean reversion to either changes in log price or log NAV as well as determine if IO has any impact on the role of each. The new regressions can be written as

$$\Delta \log \text{premium}_{i,t+1} = \alpha + \beta_1 \log \text{premium}_{i,t} + \beta_2 \mathbf{I}_{i,t} + \beta_3 \mathbf{I}_{i,t} \times \log \text{premium}_{i,t} + \epsilon_{i,t+1}, \quad (18)$$

$$\Delta \log \text{price}_{i,t+1} = \alpha + \beta_1 \log \text{premium}_{i,t} + \beta_2 \mathbf{I}_{i,t} + \beta_3 \mathbf{I}_{i,t} \times \log \text{premium}_{i,t} + \epsilon_{i,t+1}, \quad (19)$$

and

$$\Delta \log \text{NAV}_{i,t+1} = \alpha + \beta_1 \log \text{premium}_{i,t} + \beta_2 \mathbf{I}_{i,t} + \beta_3 \mathbf{I}_{i,t} \times \log \text{premium}_{i,t} + \epsilon_{i,t+1}, \quad (20)$$

where β_2 is the coefficient on the indicator variable for high IO and β_3 is the coefficient on the interaction term between the indicator and $\log \text{premium}_{i,t}$. These additional terms function to conditionally modify the original OLS intercept α and slope β_1, such that if $\mathbf{I}_{i,t} = 1$ then the intercept is $\alpha + \beta_2$ and slope on $\log \text{premium}_{i,t}$ is $\beta_1 + \beta_3$. Otherwise, if $\mathbf{I}_{i,t} = 0$, then the intercept and slope are represented by α and β_1 respectively. Should β_2 or β_3 be significant, then one can conclude that slope or intercept of the regression is significantly affected by high institutional ownership $\mathbf{I}_{i,t}$ to an extent described by the magnitude and sign of β_2 or β_3 respectively. The results of thee regressions are displayed in Table 10.

The estimates of the regression described by Equation 20 are in Column 1. The coefficient on $\log \text{premium}_{i,t}$ is both negative, significant, and close in magnitude to the OLS

Table 10: Regression Table (OLS w/ Inst. Ownership)

	(1) $\Delta \log \text{premium}_{i,t+1}$	(2) $\Delta \log \text{price}_{i,t+1}$	(3) $\Delta \log \text{NAV}_{i,t+1}$
$\log \text{premium}_{i,t}$	-0.0682***	-0.0555***	0.0127
	(0.0186)	(0.0150)	(0.0113)
$\mathbf{I}_{i,t}$	-0.00373	-0.00432	-0.000584
	(0.00715)	(0.00866)	(0.00394)
$\log \text{premium}_{i,t} \times \mathbf{I}_{i,t}$	0.0208	0.0684*	0.0476*
	(0.0268)	(0.0292)	(0.0204)
Constant	-0.0126*	-0.0197**	-0.00709**
	(0.00509)	(0.00693)	(0.00259)
Observations	756	756	756

Note: Standard errors clustered at fund level in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

estimates in Table 2, indicating the presence of premium mean reversion behavior in this model. Neither the indicator nor interaction coefficient for high IO are significant, meaning the premium mean reversion tendency is not affected by which of the high or low IO group a fund is in. Subsequently, the implied equilibrium premium is identical for all funds and similar to that of the original OLS model ($\overline{\text{premium}} = -16.8\%$). Columns 2 and 3 of Table 10 show the estimates from Equations 21 and 22. The indicator coefficient β_2 is not significant in either the $\Delta \log \text{price}_{i,t+1}$ or the $\Delta \log \text{NAV}_{i,t+1}$ model, meaning $\mathbf{I}_{i,t}$ has no significant effect in either model. However, the interaction coefficient β_3, is significant. This means that if $\mathbf{I}_{i,t} = 1$ then the "slope" of the reversion for $\Delta \log \text{price}_{i,t+1}$ and $\Delta \log \text{NAV}_{i,t+1}$ is 0.0684 and 0.0476 higher respectively, as shown in Table 11.

The slopes in Table 11 represent the magnitude of mean reversion movement in each component conditional on whether or not the fund has high institutional ownership. These results support the prediction of the De Long et al. behavioral model. When $\mathbf{I}_{i,t} = 0$, noise trader levels are high, and subsequently the magnitude of the slope of the price reversion is much larger than that of the NAV reversion. Therefore, the data shows a BDC with low institutional ownership and a large premium (discount) will have its premium eliminated

Table 11: Conditional Slope Estimates

$\mathbf{I}_{i,t}$	Eqn	Δ log price$_{i,t}$	Δ log NAV$_{i,t}$
0	β_1	-0.0555***	0.0127
		(0.0150)	(0.0113)
1	$\beta_1 + \beta_3$	0.0129	0.0603**
		(0.0276)	(0.0166)

Note: Standard errors clustered at the fund level in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

primarily through a decrease (increase) in price. When $\mathbf{I}_{i,t} = 1$, noise trader levels are relatively low, and as a result of the of the significant conditional modifier β_3, the slope of the NAV component is much steeper than that of price. Therefore, the data shows a BDC with high institutional ownership and a large premium (discount) will have its premium eliminated primarily through an increase (decrease) in NAV.

5 Results and Discussion

In the last section I detailed the methodology behind identifying and decomposing mean reversion in BDC premia and briefly reviewed the results of the strategies. In this section, I summarize the main findings of these methods and discuss their implications with respect to the closed-end fund puzzle as well as regulatory oversight of the small- and medium- enterprise loan markets in which BDCs operate.

5.1 Summary of Findings

Of the four features of the closed-end fund puzzle, the presence of only two (moderate long run discounts and mean reverting fluctuations) were testable within the limitations of the BDC data. The data showed that most BDCs trade around a mid single-digit discount in the long run ($\overline{\overline{\text{prem}}}_{MCW} = -4.6\%$), that these discounts fluctuate significantly within funds over time ($\overline{\sigma_i(\text{prem}_{i,t})} = 11.9\%$), and that these fluctuations are mean reverting (all models discussed in Section 4 revealed a significant negative relationship between the premium of a BDC and its subsequent change in premium, see the first row of Table 12). These results confirm the presence of the second and third features of the CEF puzzle, establishing BDCs as a relevant source of new data on the issue.

The next step, and the main focus of this paper, was decomposing the BDC premium mean reversions into changes in price vs. NAV. Doing so would help isolate the extent to which fluctuations in BDC premia are determined by the behavioral vs. the neoclassical explanations discussed in the Background section. All regression models presented in Section 4 support change in price as the dominant factor in BDC premium reversions (See Table 12). The four models, each presupposing unique economic assumptions about BDC discounts, yielded distinct but broadly consistent findings. The OLS model, a summary of which is in Column 1 of Table 12, revealed mean reversion behavior driven to a larger extent by price, but not exclusively so, as the NAV component was also significant. The second model incorporated fund fixed effects, economically substantiated as a reflection of the impact of the time-invariant differences between funds on premium behavior and possessed much stronger in-sample performance. This boost in explanatory power from the addition of fixed effects justified the further investigation of time-invariant BDC characteristics in order to be able to attribute the fixed effects to specific BDC features.

Subsequently, I explored the extent to which market capitalization, expense ratio, average portfolio loan size, and institutional ownership are related to premium. While all exhibited significant between-fund heterogeneity and all had strong theoretical cases for inclusion (other than institutional ownership) under the neoclassical rationale, only market cap showed any material empirical relationship to premium ($R = 0.538$). These findings provided the motivation for the third model, OLS with market cap. The OLS with market cap model revealed both a significant positive relationship between market cap and equilibrium premium and a significant inverse relationship between BDC valuation and the magnitude of the price component of premium mean reversion (See Table 6).

Table 12: Summary of Coefficients on $\log \text{premium}_t$

Dep Var:	(1) OLS	(2) FE	(3) OLS w/ MC	(4) OLS ($\mathbf{I}_{i,t} = 0$)	(5) OLS ($\mathbf{I}_{i,t} = 1$)
$\Delta \log \text{premium}_{i,t+1}$	-0.0693***	-0.242***	-0.0840***	-0.0682***	-0.0682***
	(0.0150)	(0.0485)	(0.0160)	(0.0186)	(0.0186)
$\Delta \log \text{price}_{i,t+1}$	-0.0395**	-0.215***	-0.0563***	-0.0555***	0.0129
	(0.0138)	(0.0407)	(0.0121)	(0.0150)	(0.0276)
$\Delta \log \text{NAV}_{i,t+1}$	0.0298*	0.0269	0.0277*	0.0127	0.0603**
	(0.0124)	(0.0143)	(0.0119)	(0.0113)	(0.0166)

Note: Standard errors clustered at fund level in parentheses. Table entries reflect coefficients on $\log \text{premium}_t$ when used to predict the dependent variable listed in the row titles along the regression models listed in the column titles. $\mathbf{I}_{i,t} = 0$ or 1 indicates low or high institutional ownership, respectively. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

The final model discussed in Section 4 was OLS with institutional ownership. Institutional ownership was included via an indicator variable (indicating whether fund institutional ownership was greater than or less than median for respective quarter) and an interaction term with log premium. This setup is economically justified through the behavioral model, as noise trader presence is a fundamental distinction in the behavioral setup and the institutional ownership indicator variable allows us to toggle regression estimates conditional on this proxy for noise trader levels. Doing so reveals that funds with low institutional ownership (and therefore high estimated levels of noise traders) exhibit premium mean reversion driven primarily by price, while premium reversions in funds with high institutional ownership are driven primarily by NAV (See the final two columns of Table 12). Furthermore, the low IO group had a significantly steeper mean discount ($p = 0.001$) and higher, but not significantly so, mean within-fund variation in discount ($p = 0.298$).

5.2 Discussion

Implications for CEF Puzzle

The features of the closed-end fund puzzle seem to contradict classical finance theory, and as a result, the principal question is why these features are consistently observed in CEF data. The literature on the topic can be divided into two schools of thought: the neoclassical theory believes CEF discounts and their puzzling pricing characteristics are the result of legitimate reasons for fundamental value to deviate from NAV, while the behavioral school attributes the CEF puzzle features to the presence of noise traders who make trading decisions based on irrational expectations of future performance. The BDC market provided a data set unstudied with respect to the above question, and the findings of this paper support the behavioral model as the primary explanation of BDC discounts.

One prediction of the behavioral model supported by the findings of this paper is that premium fluctuations are primarily driven by price. The presence of noise traders, who mistake public or inconsequential information as an investing edge and overvalue securities in some periods and undervalue them in others, will induce sentiment-driven volatility to the prices of BDCs. BDC net asset values, on the other hand, are unaffected by noise traders as the markets for the underlying leveraged loans are institutional. The behavioral model attributes discounts to the irrationality of the noise traders, and predicts that price, not NAV, will drive fluctuations. The findings from the four regression models discussed in Section 4 showed that price is the dominant component in BDC discount fluctuations.

Another finding in support of the behavioral model is that expense ratio and average loan size (as a proxy for liquidity), key determinants of discount in the neoclassical model, are uncorrelated with premium. The neoclassical model would predict that the perpetual draw on future cash flows of high expenses would cause a BDC to trade at a price below its NAV, and that a fund holding in-demand but illiquid portfolio securities would earn a liquidity premium. However, my findings show these expected relationships are not observed.

Finally, these conclusions in support of the behavioral model are consistent with relatively low presence of institutional capital in the BDC market. To the extent institutional capital serves as a proxy for rational investors, a low institutional presence would elicit more of the behaviors predicted by the behavioral model. Furthermore, I found that the funds with lower institutional ownership possessed more price-driven premium fluctuations and a significantly steeper discount ($p = 0.001$) than did those with higher institutional ownership, showing that, even within the BDC market, changes in noise trader presence have the effect predicted by the behavioral model.

Implications for Regulators

As summarized in the Background section, as a tougher regulatory environment after the global financial crisis has constrained the ability of traditional bank lenders to finance small- and medium- enterprises (SME), new, non-bank lenders have filled in. Loan mutual funds, collateralize loan obligations, and BDCs have rapidly grown to occupy the unfulfilled credit demand left by the exit of traditional lenders (Beltratti and Bock 2018). The rapid growth in the market for leveraged loans combined with a lack of visibility or oversight has many regulatory bodies concerned. Blunt warnings have been issued over the past 6 months by institutions including the US Federal Reserve, the International Monetary Fund, and the Bank of International Settlements (Warnings over Leveraged Loans 2019).

Publicly-traded BDC shares offer a real-time indirect market valuation of portfolios of some of the most opaque and illiquid leveraged loans. Should fluctuations these market valuations reflect a reasonable expectation of changes in the value of respective BDC portfolios, BDC prices could serve as an additional perspective from which to oversee the lower end of the leveraged loan market. The findings of this paper, however, imply that BDC prices do not serve as a reasonable representation of the SME leveraged loan market. Their prices relative to NAV do not reflect rational expectations of future changes in value of BDC portfolio facilities, and instead consist primarily of noise. This does not mean BDC premia are useless as they still might serve other purposes, for example they could function as a proxy for loan market sentiment, but such uses would require further research.

5.3 Limitations

Certain important qualifications bear mentioning on findings and implications presented above. The net asset value data reported by BDCs in their quarterly earnings is unaudited. As a result, it can be the case that the reported NAV does not accurately reflect the fair market value of portfolio loans. Occasionally, this distinction is made clear; as recently as December 2018 BDC operator Fifth Street Management settled with the SEC for failing to reasonably control the quality of its valuation models "for illiquid assets whose values could not be determined by reference to market prices or quotes," among other violations (McElhaney 2018).

Furthermore, untested assumptions underlie the use of certain measurements as proxies for variables for which I had no data or could not measure. For example, institutional ownership percentage was used as a proxy for rational investor ownership. While a reasonable assumption that institutional capital represents rational investors, in reality there might exist institutions who invest irrationally and retail investors make rational investment decisions. Another manifestation of this limitation is the assumed correlation between liquidity and loan size underlying my use of average loan size as a proxy for liquidity, while in reality there might exist large loan facilities that are illiquid and smaller ones that are more liquid.

6 Conclusions

This paper has shown that, to the extent they are testable, the features of the closed-end fund puzzle are present in the BDC market. Furthermore, the qualities of BDC discounts generally corroborate the predictions of the behavioral explanation of the closed-end fund puzzle. Features both predicted by the behavioral model and observed in the data include:

price-driven fluctuations in discount, little to no relationship between observed discounts and commonly cited neoclassical determinants of discounts, and a relationship between proxy noise trader prevalence and sentiment-driven discount behavior.

One further direction for research involves comparisons across fund structures (how do discounts compare not just within BDCs or CEFs, but between the two markets). Doing so and isolating the fund structure-specific features, such as BDC operational assistance and CEF liquidity of underlying, can test other hypothesis generated from past literature include the importance of manager ability as well as Cherkes, Sagi, and Stanton's (2009) liquidity gap premium theory. In addition, although this paper determined BDCs not to be a valuable regulatory tool for overseeing middle-market loan valuations, further study of the noise in the market could reveal a relevant proxy measure of market sentiment.

Appendix

Figure 8: Premia by Quarter

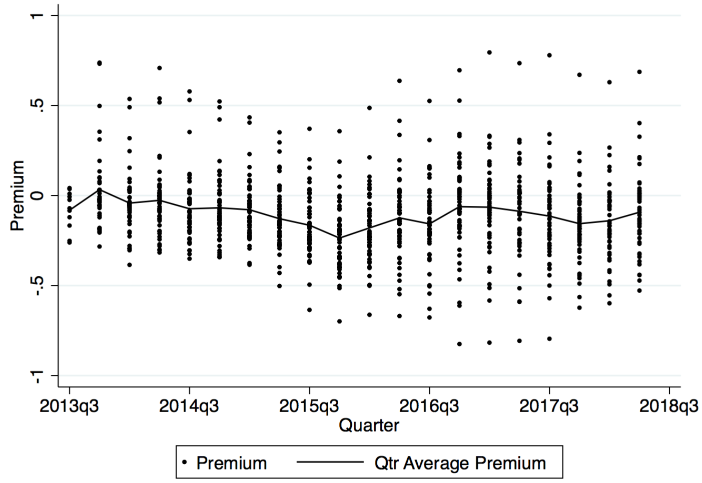

Figure 9: Log Market Cap vs. Log Premium (Avg by Fund)

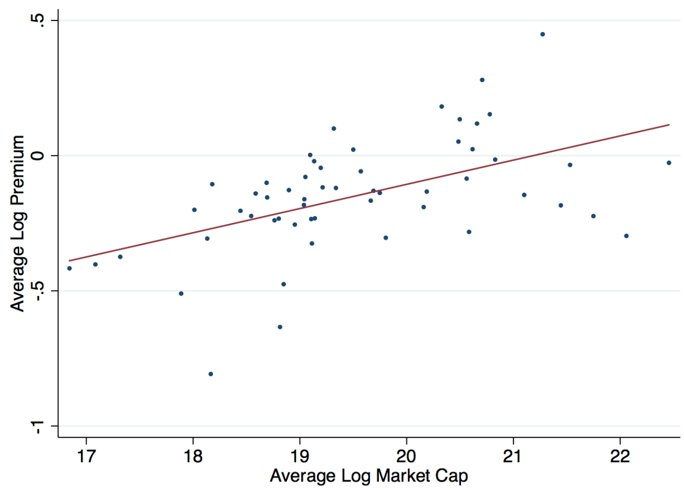

OLS Estimates Overlayed: $\beta = .0895^{***}, \alpha = -1.897^{***}, R^2 = 0.2738$

Table 13: Summary of BDC Premium Data

Ticker	Data Start	Prem Avg	Prem Std	Mkt Cap	Firm Actions
ABDC	2014Q2	-0.1932	0.1513	151.52 $M	Listed 2014Q2
ACAS	2013Q4	-0.2547	0.0575	3814.17	Acq. by ARCC 2016Q4
ACSF	2014Q1	-0.1284	0.063	119.04	Liquidated 2018Q3
AINV	2013Q4	-0.1298	0.0942	1487.89	
ARCC	2013Q4	-0.0241	0.0633	5755.26	
BBDC	2013Q4	0.2322	0.286	682.53	PKA TCAP
BKCC	2013Q4	-0.123	0.0512	593.79	
CCT	2017Q4	-0.1679	0.0144	2056.67	Listed 2017Q4
CGBD	2017Q2	-0.0142	0.0231	1112	Founded 2017Q2
CMFN	2014Q1	-0.2079	0.0897	142.94	Listed 2014Q1
CPTA	2013Q3	-0.1951	0.1455	203.95	
CSWC	2013Q4	-0.1486	0.0912	385.84	
EQS	2014Q1	-0.3227	0.1101	26.8	
FDUS	2013Q4	0.0303	0.1304	299.89	
FSIC	2014Q2	-0.0298	0.0862	2256.67	Listed 2014Q2
FULL	2013Q4	-0.1668	0.167	67.28	Acq. by GECC 2016Q4
GAIN	2013Q4	-0.1087	0.0867	256.79	
GARS	2013Q3	-0.1599	0.1098	192	
GBDC	2013Q4	0.1283	0.0737	948.68	
GECC	2016Q4	-0.1998	0.0195	114	Listed 2016Q4
GLAD	2013Q4	-0.0114	0.1328	207.37	
GSBD	2014Q4	0.1486	0.1074	801.4	Listed 2015Q1
GSVC	2013Q3	-0.3649	0.1419	159.62	
HCAP	2013Q3	-0.0966	0.0824	78.94	
HRZN	2013Q3	-0.0923	0.0757	131.26	
HTGC	2013Q4	0.3299	0.1367	991.53	
KCAP	2013Q3	-0.2096	0.164	180	
MAIN	2013Q4	0.5734	0.1459	1773.68	
MCC	2013Q4	-0.2461	0.1619	424.74	
MCGC	2013Q4	-0.1469	0.0618	190.43	Acq. by PFLT 2015Q3
MFIN	2013Q4	-0.3522	0.3765	179.59	PKA TAXI
MRCC	2013Q4	0.0038	0.052	204.89	
Averages		-0.114	0.119	641	
MCW		-0.046	0.088		

Source: Closed-End Fund Advisors. Note: Data start is the first quarter of data used for this paper for a particular fund. Prem Avg is the average premium for a fund over the sample. Prem Std is the standard deviation of the premium of a fund over the sample. Mkt Cap is average market cap (in $Mil) over sample. MCW are market cap weighted averages. PKA (previously known as) indicates data for part of sample was pulled using different ticker and then merged.

Table 13 (Cont.): Summary of BDC Premium Data

Ticker	Data Start	Prem Avg	Prem Std	Mkt Cap	Firm Actions
MVC	2016Q1	-0.2753	0.0564	200.5 $M	
NEWT	2014Q4	0.1139	0.1403	254.2	
NMFC	2013Q4	0.025	0.0429	908.11	
OCSI	2013Q3	-0.2019	0.0889	221.38	PKA FSFR
OCSL	2013Q4	-0.2363	0.1246	905.42	PKA FSC
OFS	2013Q3	-0.1414	0.0568	133.2	
OHAI	2013Q4	-0.3852	0.133	72.13	PKA NGPC
OXSQ	2013Q3	-0.1237	0.0958	390.9	PKA TICC
PFLT	2013Q4	-0.0544	0.0606	337.79	
PNNT	2013Q4	-0.1653	0.1153	580	
PSEC	2013Q4	-0.1912	0.1241	2826.84	
RAND	2013Q3	-0.3327	0.1045	20.93	
SAR	2013Q4	-0.1745	0.1344	104.65	
SCM	2013Q4	-0.114	0.1008	163.95	
SLRC	2013Q4	-0.0789	0.0694	853.89	
SUNS	2013Q4	-0.0413	0.0731	222.05	
SVVC	2014Q3	-0.5372	0.1389	84.71	
TCPC	2013Q4	0.0554	0.0653	796.16	
TCRD	2013Q4	-0.118	0.0859	361.74	
TINY	2013Q3	-0.2535	0.1263	78.85	
TPVG	2014Q1	-0.0693	0.1098	191.11	Listed 2014Q1
TSLX	2014Q1	0.1677	0.0754	1065.44	Listed 2014Q1
WHF	2013Q4	-0.1069	0.0825	224.11	
XRDC	2014Q1	-0.3023	0.1164	36.71	PKA BDCV, KIPO
Averages		-0.114	0.119	641	
MCW		-0.046	0.088		

Source: Closed-End Fund Advisors. Note: Data start is the first quarter of data used for this paper for a particular fund. Prem Avg is the average premium for a fund over the sample. Prem Std is the standard deviation of the premium of a fund over the sample. Mkt Cap is average market cap (in $Mil) over sample. MCW are market cap weighted averages. PKA (previously known as) indicates data for part of sample was pulled using different ticker and then merged.

References

Alexander GJ, Peterson MA. (2016). Short selling and the pricing of closed-end funds. *Journal of Financial Markets, 33,* 124–142

Beltratti A, Bock J. (2018). BDCs: The most important commercial lenders you've never heard about. *Journal of Alternative Investments, 20,* 8–20

Black F. (1986). Noise. *Journal of Finance, 41,* 529–543

Boehm SB, Friedberg HL. (2008). Everyman a Venture Capitalist: Taking Private Equity Public. *Business Law Today, 17,* 12–17

Berk JB, Stanton RH. (2007). Managerial ability, compensation, and the closed-end fund discount. *Journal of Finance, 62,* 529–556

Brauer GA. (1984). 'Open-ending' closed-end funds. *Journal of Financial Economics, 13,* 491–507

Brickley J, Schallheim J. (1985). Lifting the lid on closed-end investment companies: a case of abnormal returns. *Journal of Financial and Quantitative Analysis, 20,* 107–117

Brown GW. (1999). Volatility, Sentiment, and Noise Traders. *Financial Analysts Journal, 55,* 82–90

Cherkes M. (2003). A positive theory of closed-end funds as an investment vehicle. Working Paper, Department of Economics, Princeton University. **http://papers.ssrn.com/sol3/papers.cfm?abstract_id=567084**

Cherkes M. (2012). Closed-End Funds: A Survey. *Annual Review of Financial Economics, 4,* 431–445

Cherkes M, Sagi J, Stanton R. (2009). A liquidity-based theory of closed-end funds. *Review of Financial Studies, 22,* 257–297

Closed-End Fund Advisors. BDC Database. Retrieved October 10, 2018.

De Long JB, Shleifer A, Summers L, Waldman R. (1990). Noise trader risk in financial markets. *Journal of Political Economy, 98,* 703–738

Gasbarro D, Johnson RD, Zumwalt JK. (2003). Evidence on the Mean-Reverting Tendencies of Closed-End Fund Discounts. *Financial Review, 38,* 273–291

Gemmill G, Thomas DC. (2002). Noise trading, costly arbitrage, and asset prices: evidence from closed-end funds. *Journal of Finance, 57,* 2571–2594

Goel T. (2018, September). The rise of leveraged loans: a risky resurgence? *BIS Quarterly Review.* **https://www.bis.org/publ/qtrpdf/r_qt1809u.html**

Hanley KW, Lee CM, Seguin PJ. (1995). The Marketing of Closed-end Fund IPOs: Evidence from Transactions Data. *Journal of Financial Intermediation, 5,* 127–159

Haunss, K. (2018). Regulators sound the alarm about leveraged loan market. *Reuters.* Retrieved from **https://www.reuters.com/article/regulators-sound-the-alarm-about-leverag/regulators-sound-the-alarm-about-leveraged-loan-market-idUSL8N1XA7MH/**

Investment Company Institute. (2018, April). Retrieved from **https://www.ici.org/cef/background/bro_g2_ce**.

Lee CMC, Shleifer A, Thaler RH. (1991). Investor sentiment and the closed-end fund puzzle. *Journal of Finance, 46,* 76–110

Malkiel BG. (1977). The valuation of closed-end investment-company shares. *Journal of Finance, 32,* 847–858

Malkiel BG. (1995). Returns from Investing in Equity Mutual Funds 1971 to 1991. *Journal of Finance, 50,* 549–572

Malkiel BG, Xu Y. (2005). The persistence and predictability of closed-end fund discounts. Working Paper, Department of Economics, Princeton University. **http://www-stat.wharton.upenn.edu/ steele/Courses/956/Resource/CEF/MalkielXu.pdf**

McElhaney A. (2018, December). Former BDC Giant Fifth Street Management Settles With SEC. *Institutional Investor.*

Mcgrath C. (2017, April). 80% of equity market cap held by institutions. *Pensions & Investments.*

Orr L. (2017, August). Be Damned Careful. *Institutional Investor.*

Peavy JW. (1988). Closed-end fund new issues: Pricing and aftermarket trading considerations. Working Paper, Southern Methodist University.

Pontiff J. (1995). Closed-end fund premia and returns implications for financial market equilibrium. *Journal of Financial Economics, 37,* 341–370

Pratt EJ. (1966). Myths associated with closed-end investment company discounts. *Financial Analysts Journal, 22,* 79–82

Ross SA. (2002). Neoclassical finance, alternative finance and the closed-end fund puzzle. *European Financial Management, 8,* 129–137

SEC EDGAR System. Company filings. Retrieved February 21, 2019.

Sharpe WF, Sosin HB. (1975). Closed-end investment companies in the United States: Risk and return. *European Finance Association 1974 Proceedings*, North Holland, Amsterdam.

Sias RW, Starks LT, Tinic SM. (2001). Is noise trader risk priced? *Journal of Financial Research, 24,* 311–329

Thompson R. (1978). The information content of discounts and premiums on closed-end fund shares. *Journal of Financial Economics, 6,* 151–186

Warnings over Leveraged Loans are Not Mere Fear-Mongering. (2019, February). *Financial Times*. Retreived from **https://www.ft.com/content/b8264050-2ee5-11e9-ba00-0251022932c8**.

Weiss K. (1989). The post-offering price performance of closed-end funds. *Financial Management, 18,* 57–67