How do private equity fees vary across public pensions?*

Juliane Begenau† Emil Siriwardane‡

February 2021

Abstract

We provide evidence that investment fees vary within private equity funds. Net-of-fee return clustering suggests that 70% of funds group investors into two fee-tiers that vary along both fixed and variable components. Managers of venture capital funds and those with high past performance are less likely to tier their investors. Some investors consistently earn higher net-of-fee returns relative to others within their funds. Investor size, experience, and past performance explain some but not all of this effect, suggesting that unobserved traits like negotiation skill or bargaining power materially impact the fees investors pay in private equity.

*Special thanks go to our discussants Arthur Korteweg, Yael Hochberg, and Stijn Van Niewerburgh. We are grateful to Jonathan Berk, Ted Berk, John Campbell, Josh Coval, Peter DeMarzo, Mark Egan, Steve Grenadier, Debbie Gregory, Paul Gompers, Daniel Green, Robin Greenwood, Ben Hebert, Victoria Ivashina, Josh Lerner, Hanno Lustig, Chris Malloy, Ludovic Phalippou, Monika Piazzesi, Josh Rauh, Claudia Robles-Garcia, Amit Seru, Andrei Shleifer, Ilya Strebulaev, Per Stromberg, Adi Sunderam, Boris Vallee, and Luis Viceira, as well as seminar participants at Harvard, the NBER CF SI 2020, the NBER LTAM 2021, the Red Rock Conference, the Virtual Corporate Finance Conference, the University of Washington, the University of Southern California, Stanford, and the Swedish House of Finance for their comments and discussions. Sid Beaumaster, Ruying Gao, Francesca Guiso, Rick Kaser, and Yuan Tian provided excellent research assistance. We are also indebted to Maeve McHugh and Michael Smyth at Preqin for their help compiling the dataset for our analysis. First draft is from November 2019.

†Begenau: Stanford GSB and NBER and CEPR. E-mail: begenau@stanford.edu

‡Siriwardane: Harvard Business School. E-mail: esiriwardane@hbs.edu

1 Introduction

Over the last twenty years, state and local defined-benefit pensions have increasingly shifted capital out of traditional asset classes like fixed income and into private-market investment vehicles like private equity and venture capital (Ivashina and Lerner, 2018). While there is a general consensus that fees in private-market funds are large (Gompers and Lerner, 2010; Metrick and Yasuda, 2010; Phalippou et al., 2018), there is virtually no systematic and large-sample analysis of how fees are determined in this asset class. The main reason why is that investment terms are privately negotiated and are thus rarely observed by outsiders. In this paper, we exploit a novel dataset to overcome this empirical challenge and provide evidence that pensions often pay different fees when investing in the *same* private-market fund.

Our empirical design is based on the simple idea that two investors who invest at the same time in the same fund should earn the same gross-of-fee return. In our data, we observe net-of-fee returns for multiple investors in the same fund. Thus, variation in net-of-fee returns across these investors should be informative about within-fund variation in fees.[1] By fees, we mean any management and performance fees, fund expenses, or other costs that are borne by investors (i.e., the fund's limited partners or LPs). As an illustration of our approach, Figure 1 shows the cumulative return earned by two investors in the same fund, where returns at each point in time equal cumulative distributions received per dollar of investment (DVPI). At the beginning of the fund's life, both investors earn identical returns, though after five years the orange investor has earned \$1.7 per dollar invested compared to the \$1.4 for the blue investor. We argue that this return gap can be used to understand differences in fees paid by the two investors.

Our main findings are as follows. We start by showing that the intuition of Figure 1 generalizes to a sample of over 2,400 funds managed by 857 different fund managers (i.e., the general partners or GPs). Returns vary considerably across investors in the average fund. For instance, for funds between 8 to 12 years of age, the average within-fund volatility of DVPI is 0.05 compared to an

[1] The structure of private market funds means that investors generally invest at the same time (see Section 2). We discuss alternative sources of within-fund return variation (e.g., measurement error) below and in Section 3.

average level of 1.19. Moreover, we observe similar patterns in within-fund volatility regardless of how we measure returns (e.g., including unliquidated investments or IRRs) and across funds of all vintages.

We then show that net-of-fee returns within a fund tend to cluster, as opposed to being continuously distributed across investors. Using textbook machine learning techniques (Jain et al., 1999), we find that 70% of the funds in our sample have two clusters of returns, suggesting that investors in most funds are grouped into one of two fee-tiers. The majority of remaining funds (23% of total) appear to use a single fee structure for all investors, though GPs differ in their propensity to do so. For example, venture capital funds are far more likely – roughly 32 percentage points – to use a single fee structure, as are successful GPs with a long track record of high performance. The latter finding is consistent with a model where high-performing GPs can charge a single high fee to all LPs without losing capital commitments, whereas low-performing GPs have to negotiate more to attract investors (e.g., Berk and Green, 2004).

Fee structures within a fund could differ along several dimensions, such as performance-based fees, management fees, fund expenses, or portfolio company fees and associated fee offsets. Though we do not observe these exact contractual splits, the panel nature of our data allows us to decompose differences across investor tiers into two dimensions: a fixed fee component that scales with commitment size and time and a variable fee component that scales with performance. Intuitively, any within-fund dispersion in fixed fees should grow linearly with fund age and any dispersion in variable fees should grow linearly with performance. Building on this logic, we estimate that the average within-fund volatility of fixed and variable fees are 84 basis points and 5.3%, respectively.[2] These estimates vary strongly across asset classes. Venture capital funds have the smallest within-fund dispersion in both fee components whereas private debt and infrastructure funds have the widest dispersion. We also develop a placebo test by exploiting the fact that performance fees are contingent on a minimum level of fund performance. Accordingly, in all asset classes, we confirm that dispersion in variable fees is not detectable in unprofitable funds.

[2]These estimates are in line with the menu-model that Bain Capital has offered to its investors in recent years (Zuckerman and Or, 2011; Markham, 2017).

In the last part of the paper, we document that some pension investors consistently earn higher net-of-fee returns relative to others in their funds. Within a given fund, we categorize an investor as being top-tier in terms of fees if it earns above-median net-of-fee returns for the majority of the fund's life. An F-test from a fixed-effects regression comfortably rejects the null hypothesis that top-tier fee status is randomly assigned to investors in each fund. The rejection of random tier assignment is driven by the wide observed distribution of "pension effects": some investors in our sample are in the top performance tier for over 70% of their funds, whereas others are in the top tier for less than 15% of their funds. This finding supports the idea that some investors consistently select or are offered the best fee structure in their respective funds, at least in terms of ex-post performance. We provide further evidence that part of these pension effects are driven by selective matching between LPs and GPs (e.g., relationships).

There are several possible reasons why some pensions could consistently pay lower fees than others when investing in private markets. For instance, GPs could offer fee reductions to pensions who lower the cost of raising a fund, perhaps by drawing in other investors or by providing larger amounts of capital. Consistent with this intuition, pensions that are large in overall size are roughly 20 percentage points more likely to be in the lowest-fee tier for the average fund. Similarly, those that contribute more capital to a fund are more likely to be in the lowest fee tier. We also find that proxies for investor sophistication correlate with tier assignment, as pensions in low-fee tiers tend to be better governed, more experienced, and have high past performance. Nonetheless, even after controlling for all of these observable characteristics, there are still a subset of pensions who consistently outperform others within their respective funds. We interpret this as evidence that unobservable traits related to negotiation and contracting skill materially impact the fees that investors pay in private-market funds.

Our empirical analysis relies on within-fund variation in returns. In Section 3, we investigate several alternative sources of this variation, including measurement error, pension-specific accounting practices, or within-fund differences in portfolio composition.[3] While these alterna-

[3] Some funds allow LPs to deploy additional capital to specific portfolio companies in a fund at lower or no cost (so-called co-investment rights, Fang et al. (2015)). These types of special-purpose vehicles are a small part of public

tives could all be present in the data, none can fully explain the empirical patterns discussed above, nor would they bias our estimates of within-fund dispersion in fixed and variable fees. For example, any LP-specific differences in accounting (e.g., reinvested distributions) cannot account for why some funds or GPs are more likely to have a single return cluster than others. Thus, our results collectively suggest that fees are a central reason why investors in the typical private-market fund appear to earn different returns.

Fee dispersion is a natural outcome in private-market funds because investment terms are often negotiated bilaterally between LPs and GPs (see Section 3.4). This means that the fee-setting mechanism can be understood through search and bargaining models (Burdett and Judd, 1983; Bester, 1988; Duffie et al., 2005), which generally predict that consumers or investors will pay different prices for the same product in equilibrium. This prediction has been confirmed in many market settings, including those for health care (Sorensen, 2000; Grennan, 2013), automobiles (Goldberg, 1996), financial securities (Eisfeldt et al., 2020), residential mortgages (Allen et al., 2019), and mutual funds (Hortaçsu and Syverson, 2004). Our results extend these empirical studies to the $5.8 trillion private-capital market (McKinsey, 2019). In addition, our estimates of within-fund dispersion in fixed fees are comparable to Hortaçsu and Syverson (2004), who show that management fees for S&P 500 index funds range from 10 to 268 basis points. Private-market funds are far more complex and opaque than S&P 500 index products, so if anything, one would expect higher levels dispersion in our setting (Salop and Stiglitz, 1977; Gabaix and Laibson, 2006).

This study also contributes to prior research on private equity. Much of the previous work in this setting has focused on across-fund variation in fees (Robinson and Sensoy, 2013), whereas we are among the first to study within-fund variation. Our results reveal important features of the fund-formation process (e.g., investor tiering) and suggest that GPs vary considerably in how they set investment terms with their LPs. The notion that GPs group investors into fee tiers is also consistent with recent studies showing that GPs differentiate among investors through co-investments and other special purpose vehicles (Lerner, Mao, Schoar, and Zhang, 2018; Fang, Ivashina, and Lerner,

pensions' portfolios during our sample and are listed as separate entities, which allows us to exclude them entirely from our analysis. See Section 3.3 for a complete discussion.

2015; Braun, Jenkinson, and Schemmerl, 2019). Additionally, the observation that some pensions consistently receive better terms aligns with the findings of Lerner et al. (2018), who show that GPs offer certain special purposes vehicles only to a select set of investors.

From a methodological perspective, we also contribute to the ongoing policy debate on fee transparency in private-market funds (State Comptroller SEC letter, 2015). As discussed above, the opacity of private-market funds has made it difficult to systematically study fee determination in a large-sample setting. We are able to sidestep many of these issues by studying ex-post returns of multiple investors in the same fund. An added advantage of this approach is that net-of-fee returns reflect all costs borne by investors, not just management and performance fees. For example, our analysis of within-fund return clustering will detect differences in fund expenses or fee-offsets related to portfolio company fees, the latter of which has been shown to be large in private-market funds (Phalippou et al., 2018).

Finally, our results contribute to the literature on public pensions investment decisions (Hochberg and Rauh, 2012; Andonov et al., 2018). Consistent with this literature, we show that pensions with more elected board members tend to outperform others in their funds. In addition, our analysis of pension effects complements prior work documenting frictions in the labor market for public pension managers and investment staff (Dyck et al., 2018). Specifically, we provide evidence that the most talented negotiators in terms of fees do not perfectly match with the largest pensions or private equity investors, as one might expect in a frictionless model of talent mobility (e.g., Gabaix and Landier, 2008).

The paper is organized as follows. In Section 2, we discuss the data and return measures and show how net-returns vary within fund across fund vintages. Section 3 discusses the potential drivers of fee dispersion. Section 4 documents that net-returns cluster at the fund level and characterizes dispersion in two fee categories at the fund level. In Section 5, we show that some pension investors systematically outperform others in all their funds and describe their characteristics. Section 6 concludes. Additional details and results are available in an online appendix.

2 Data and Motivating Evidence

2.1 Background and Data Description

We study public pension investments into private market vehicles, namely private equity (PE). A typical PE fund has two types of investors, the general partner (GP) and the limited partners (LPs). The GP manages the fund and usually contributes about 1-5% of its own capital to the fund. The bulk of the fund's capital therefore comes from LPs, who are entities like pensions, endowments, and family offices. At the beginning of a fund's life-cycle, GPs secure capital commitments from LPs, after which capital is formally "called" from the LPs. Some of this called capital is invested by the GP, while the rest is used to pay management fees and other fund expenses that are borne by LPs.[4] In most cases, each LP has the same pro rata claim on the investments made by the fund, meaning gross-of-fee returns are equal across LPs. These investments are held for several years before they are liquidated. The GP then withholds a portion of the investment proceeds as a performance fee (or "carry") before issuing distributions back to the LPs. From start to finish, most funds have a total lifespan of ten to fifteen years. This structure makes it reasonable to compare returns across investors in the same fund.

We obtain investment performance data from Preqin, a data provider that specializes in alternative assets markets. Preqin's data on private market investments is sourced primarily from Freedom of Information Acts (FOIA) requests of public pensions and legally-required annual reports.[5] The Preqin data covers funds from vintage year 1990 onward and contains cash-flow data on LP-level investment into individual funds. We specifically observe the amount of committed capital by the investor in the fund, the amount of capital that has been "called" from the investor (i.e., actual contribution amounts), and the amount of capital that has been distributed back to the investor by the fund. These variables are all reported in cumulative terms. Importantly, distributions are re-

[4] We discuss the contracting environment and different fees charged by GPs at length in Section 3.4.

[5] To encourage the same reporting standards across investors and funds, Preqin provides detailed guidelines on submitting performance data in their FOIA requests. After data is submitted to Preqin, the information is reviewed internally and, when possible, is cross-referenced against as many different sources as possible. Further details on Preqin's collection process can be found in Preqin's Private Capital Performance Data Guide. https://docs.preqin.com/reports/Preqin-Private-Capital-Performance-Data-Guide.pdf

ported net of performance-based fees that are withheld by the GP (e.g., carry), and contributions are inclusive of fixed fees such as management fees that are calculated as a percentage of an LP's committed capital. This means that investment multiples are net of fees We also observe the net asset value (NAV) of each investor's current investments in the fund. For a given investor in a fund, the NAV reflects the market value of investments that have not yet been liquidated.

2.2 Sample definition

The sample and variables that we use are taken directly from Begenau et al. (2020). In that companion paper, we discuss a variety of quality control filters that we apply to the raw Preqin data in order to ensure that the resulting cash flow variables are comparable across investors in the same fund.[6] To keep the current paper self-contained, we summarize the main features of our approach below.

The raw data file by Preqin has roughly 750,000 observations and is unique at the level of data source, LP, fund, and date. To be included into our sample, we require a complete set of non-missing identifiers in terms of investor, fund, fund manager, date and fund vintage, as well as non-missing information regarding an LP's contribution, distribution, commitment size, and fund net-asset-value. In addition, we require cash flows to be denominated in USD and focus on LPs who are U.S. public pension funds.[7] This choice eliminates an potential issues that currency conversion may have on our analysis of within-fund returns. There are 376,394 observations that remain after applying these filters and deleting duplicates.

In addition to these basic sample filters, we drop any source-investor-fund cell in which a negative contribution or distribution occurs and is too large to plausibly reflect a fee offset.[8] These cases are incredibly rare and only affect 0.28% of observations. To be conservative, we also drop

[6]The latest version of the companion paper can be found here. To keep our analysis as transparent as possible, we have also posted the code we use to clean the Preqin data here.

[7]The vast majority of investors in our data are U.S. public pension funds (83%) and UK public pension funds (7%). Other investor types in our dataset include public university endowments, government agencies, insurance companies, foundations, and private sector pensions. Throughout the paper, we only use data on U.S. public pensions.

[8]Fee offsets may reflect, among other things, monitoring fees that are passed from portfolio companies back to LPs. See Appendix Section B.5 for more details on the types of fund income that can lead to fee offsets.

any LP from the sample if more than 2.7% of their observations have any potential quality issue (e.g., large negative distribution), which only affects 15 LPs and 0.86% of total commitments in the sample. We discuss the choice of these cutoffs in detail in Begenau et al. (2020).

After applying these additional filters, we retain fund-quarter cells in which there are at least two LPs reporting cash flows, since our focus is on within-fund variation in returns. We also drop all funds that are related to multi-strategy investment (only 2 funds), co-investment, or secondary sales.[9] This leaves us with 233,907 observations that are unique at the investor-fund-quarter level (p, f, t). For some of our subsequent analysis, we condense the data further so that it is unique at the investor-fund (p, f) level. To do so, we find all observations for a fund that are within 20 quarters of the last observed date. For each fund, we then pick the quarter with the largest number of investors reporting returns. This approach allows for any fee differences within the fund to play out over a long enough horizon, while still including as many investors as possible. We refer to this condensed data as the core sample.

2.3 Definition of Returns

Together, contributions, distributions, and remaining net asset values allow us to calculate standard industry return multiples. Specifically, we define the realized cash multiple for investor p in fund f at time t as:

$$r_{pft} \equiv \frac{\text{Cumulative Distribution}_{pft}}{\text{Cumulative Contribution}_{pft}}. \tag{1}$$

We refer to this a a realized multiple because it only reflects distributions that have been paid by the fund to LPs. In practice, it is commonly referred to as the distributed value to paid-in capital ratio or DVPI. Similarly, we define the total multiple on invested capital as:

$$r^T_{pft} \equiv \frac{\text{NAV}_{pft} + \text{Cumulative Distribution}_{pft}}{\text{Cumulative Contribution}_{pft}}. \tag{2}$$

[9] We identify co-investment funds based on their category type in Preqin and if their listed name includes “Co-”. We identify secondary transactions in a similar manner.

Compared to DVPI, this measure reflects both remaining net asset value (unrealized value) and realized distributions. It is commonly referred to as the total value to paid-in capital ratio or TVPI.

LPs also report internal rates of return (IRRs) to Preqin. However, we primarily analyze return multiples and instead use IRRs only for robustness tests. The main reason is that reported IRRs are missing for 20% of the observations in our data. We have also computed IRRs ourselves based on observed cash flows, though in most cases these cannot be used to measure within-fund return variation. This is because most funds do not have a fully balanced panel across investors and IRRs are very sensitive to the timing of cash flows.

2.4 Dispersion in net-of-fee returns

Table 1 describes the core sample. Panel A shows that we have 9,847 investor-fund (p, f) level observations, covering 219 unique pension funds (LPs), 857 unique fund managers (GPs), and 2,407 funds. Roughly half of our observations are investments in Private Equity funds, 1,956 are in Venture Ca[ital, 1,895 are in Real Estate, 1,215 are in Private Debt, and 306 are in Infrastructure.

Panel B of Table 1 presents summary statistics of the core sample. The average fund is 9 years old (25th percentile and the 75th percentile are 4 years and 12 years, respectively). Age denotes the years since the final close date. The average investment size (i.e., commitment amount) is \$55 millions (the 25th percentile and the 75th percentile are \$12 millions and \$74 millions, respectively). The average pension fund investment represents 5% of total commitments (the 25 percentile and 75th percentile are 1% and 6.5%, respectively). The average investor invests in 4 funds (the 25 percentile and 75th percentile are 2 and 5, respectively). Pensions in our sample have invested (i.e., assets under management or AUM) roughly \$24 billions into the funds we cover. We do see some large funds with over \$300 billion in assets.

Panel C presents the average level and dispersion of three different net-of-fee return measures at the fund level, broken out by fund age. DVPI denotes the realized multiple, TVPI the total multiple including the fund's net-asset-value from the non-liquidated portfolio companies. The reported IRR is the internal rate of return calculated using the same cash flow variables as for

TVPI. Since the profile of cash flow matters for IRR, we only compare IRRs of investors that present with an identical timing of cash flows. Across measures, performance tends to increase with fund age, while net-of-fee return dispersion is sizable regardless of the measure and the fund's age. We measure dispersion as the within-fund net-of-fee return standard deviation. It is fairly stable across fund age, with the exception of IRR according to which the youngest funds have the highest dispersion. Measuring the average range as two times the standard deviation, the within-fund range of DVPI is 10% relative to the mean of DVPI for 4 to 8 year old funds, 8% for 8 to 12 year old funds, and 6% for funds older than 12 years. The TVPI range is similar. The within-fund IRR range is 15% relative to the within-fund mean IRR for 4 to 8 year old funds, 23% for 8 to 12 year old funds, and 17% for funds older than 12 years.

Figure 2 presents an overview of the within-fund net-of-fee return dispersion in our core sample data, using the realized multiple DVPI (Panel A), the total multiple TVPI (Panel B), and the reported IRR (Panel C) from Preqin. To construct this graph, we calculate the within-fund standard deviation of net-of-fee returns for each fund and show its distribution in a box plot for fund groups organized by fund vintage. Regardless of fund age and how returns are measured, there is sizable net-of-fee return variation within the same fund. This fact is the motivation for the remainder of the paper.

3 Potential Sources of Within-Fund Return Dispersion

In this section, we discuss and evaluate several potential sources of variation in within-fund returns. We focus on four broad channels: (i) measurement error; (ii) accounting practices that vary across LPs; (iii) differences in gross-of-fee returns across investors in the same fund; and (iv) variation in fees or other contract terms that impact LP performance. Table 2 summarizes the extent to which each can explain different patterns in the data that we document in Sections 4 and 5. While none of these channels are mutually exclusive, there are several patterns in the data (e.g., return clustering) that suggest fees play a central role in generating the dispersion observed in Figure 2.

3.1 Measurement Error

Measurement errors are one simple reason why returns in our data could differ across LPs in the same fund. Given that the data is sourced primarily via FOIA requests, these errors could occur when Preqin transcribes the FOIA data that they receive from LPs. To gauge the size of this channel, we created our own dataset by filing FOIAs directly with a sixty-five of the pensions in our sample. We chose these pensions based on the funds with the most observed dispersion and the LPs whose performance was the most extreme relative to others in their respective funds. In the vast majority of cases (~97%), the data from our direct FOIA was identical to the Preqin data. For the small number of cases where the data did not perfectly match, the size of the deviations was economically small. Section A.1 of the online appendix contains the full results of this audit, including the exact language of our FOIA requests.

Measurement error could also occur if LPs report erroneous data in their FOIA replies to Preqin. While this is certainly possible and even probable, three patterns in the data cannot be generated by true measurement error. First, net-of-fee returns are clustered or bunched at the fund level (Section 4.1). Second, some funds and GPs are more likely to exhibit return dispersion than others (Section 4.2). Third, some LPs are more likely to outperform others when investing in the same fund (Section 5.1). Most public pensions are also audited annually, which in principle should reduce the occurrence of reporting errors over time.

3.2 Accounting Practices

Public pensions that invest in private capital vehicles have no legally mandated accounting standards, which could lead some LPs in our sample to report cash flows or NAVs differently than others. Any such differences could cause returns to systematically differ across LPs in the same fund. Nonetheless, there is an simple institutional reason why accounting differences are unlikely to be a large source of within-fund return variation. GPs typically send their LPs a quarterly report that summarizes the current state of their investments in the fund. The detail of these reports varies substantially across GPs, but all generally provide a running total of distributions and an estimate of the market value of unliquidated investments. According to several large LPs and Preqin, the content of these reports is generally used to satisfy any FOIA requests. In our data, this means that accounting practices should vary across funds, not within funds. With that said, it is still possible for specific LPs to adjust the cash flows and NAVs contained in the investment reports for FOIA requests. We now discuss two specific variables for which any such adjustments are most likely to occur.

Net asset values (NAVs) Fund NAVs measure the estimated value of each LPs share in the fund in the event of an orderly liquidation. LPs could systematically differ in how they report NAVs if some deduct expected performance fees (carry) that would be charged by the GP. If this were the case, these LPs would consistently report lower TVPIs in their respective funds and generate the dispersion observed in Panel B of Figure 2. However, the presence of sizable within-fund variation in DVPI (Panel A, Figure 2) suggests that NAV-accounting is not the primary source of net-of-fee return variation within funds.

Recallable (or Recyclable) Capital Within-fund dispersion in DVPI (or TVPI) could also arise due to differences in how LPs account for recallable capital. Recallable (or recyclable) capital refers to proceeds from liquidated investments that can be reinvested by the GP. The specific terms of this reinvestment are prescribed by so-called recycling provisions, which prescribe the amount and horizon over which recallable capital can be deployed. According to the CFA Institute's Global Investment Performance Standards (GIPS), LPs should account for recallable capital by recording a new distribution and a new contribution equal to amount that is being recalled. To the extent that LPs do not follow GIPS standards, they could instead net out recallable capital, recording no new distribution and no new contribution. In a given fund, these two approaches could lead to the appearance of sizable variation in DVPI (and TVPI).

Recallable capital accounting is unlikely to be the primary source of within-fund return variation for at least three reasons. First, in Begenau et al. (2020), we directly FOIA'd a subset of pensions about their accounting of recallable capital and 100% of respondents conformed to GIPS standards. Second, IRRs are not sensitive to the accounting of recallable capital and we still observe meaningful within-fund variation in IRRs (Panel C, Figure 2). Third, we show in Section 4.2 that certain GPs and investment styles (e.g., venture capital) are more likely to exhibit within-fund return dispersion, which cuts against the view that dispersion is primarily driven by differences in LP-specific accounting practices.[10]

3.3 Differences in the gross-return exposure

Net-of-fee returns could differ across investors in the same fund if they have different gross (or pre-cost) exposure to the fund. There are two mechanisms through which this could occur in practice: (i) co-investment vehicles and (ii) LP-specific restrictions on investment.

3.3.1 Co-investment

Co-investment vehicles allow LPs to augment their exposure to the "main" fund by allocating additional capital towards a particular deal or set of deals (see Fang et al. (2015) for an in-depth discussion of co-investment). These structures are related to so-called side-car or parallel fund vehicles. To see why co-investments structures could generate net-of-fee return dispersion, consider a fund in which only investor A has the ability to co-invest. LP in the fund are otherwise equal in terms of their commitment size and all investment terms, namely fees. Further suppose that investor A combines the returns on its co-investment portfolio and the main fund when responding to FOIA requests and reporting to Preqin. If the co-investment vehicle tilts more towards certain portfolio companies relative to the main fund, then investor A's reported net-of-fee return will differ from other LPs. The resulting dispersion would be even larger if the co-investment vehicle had reduced cost structure compared to the main fund, as it often does in practice.

There are several reasons why co-investment vehicles are not the primary source of within-fund return variation that we observe empirically. First, and most importantly, we exclude any

[10] Appendix Section A.2 also contains an exercise where we bound the fraction of funds whose dispersion could be plausibly attributed to accounting differences.

funds that Preqin classifies as a co-investment vehicle from our analysis. We expect this classification to be relatively accurate because LPs generally list co-investment vehicles as a separate fund when reporting performance to Preqin. For example, “Fortress Investment Fund IV” and “Fortress Investment Fund IV - co-investment” appear as two separate funds and we drop the latter. Moreover, for several of the largest LPs in our data, we have manually compared the co-investments that are reported on their websites and annual reports against the data in Preqin. In all cases, we found that cash flows from co-investments were indeed listed separately in the Preqin data.

Second, while co-investment vehicles have been increasing in popularity in recent years, they have not been a large part of public pension PE investments for most of our sample (1990-2018). Based on data from CEM Benchmarking, a provider of benchmarking services for thousands of global pensions, Beath et al. (2014) find that less than 5% of U.S. public pensions had any co-investments in PE as of 2014.[11] Smaller pensions may be less able or inclined to co-invest because it requires the internal infrastructure to evaluate individual portfolio companies and then deploy capital on relatively short notice. Even for larger pensions, co-investments are not yet a large portion of their portfolios. For instance, in 2019, CalSTRS – the second largest pension fund in the U.S. – reported that less than 5% of its PE portfolio was through co-investments (CalSTRS, 2019).[12]

Third, as part of our data-quality audit (Internet Appendix Section A.1), we asked pensions via FOIA if they utilized any special investment arrangements such as a side-car deals or co-investments. The vast majority responded that they had no such arrangement. For the few cases that affirmed co-investment arrangements, we confirmed that these co-investment relationships were reported separately and therefore not included in our analysis.

[11] (Preqin, 2014) finds that “relatively few LPs are being offered co-investment rights by GPs in the Limited Partnership Agreement,” despite strong interest from LPs for such rights. The survey further states that “there seems to be some contradiction between the attitudes towards and the actual co-investment activity occurring.”

[12] Co-investment by CalPERS – the largest U.S. pension fund – was relatively infrequent prior to 2011, when it launched a dedicated co-investment program (CalSTRS, 2019). The program was suspended in 2016.

3.3.2 Other Investor-Specific Mandates

Another reason why the gross-return may deviate for some investors in a fund is what we call investor-specific mandates. One prominent example that has boomed in popularity in recent years are so-called environmental, social, and governance (ESG) restrictions. These restrictions mean that one investor might restrict investment into portfolio companies based on ESG criteria (e.g., firms with large carbon footprints). Any such restrictions will naturally cause returns to differ across investors in the same fund.

To the best of our knowledge, data on investor-specific restrictions are not available for the private-market funds in our sample. However, the National Association of State Retirement Administrators (NASRA) reports that relatively few U.S. pension plans incorporate ESG in their investment process, though some of the larger U.S. pensions have started to do so more in recent years (NASRA, 2018). Motivated by this evidence, we exclude large LPs (those with AUMs over \$100 bn) and compute within-fund return volatility for funds launched prior to 2010. Figure IA2 in the internet appendix shows that average level of dispersion is marginally lower for this sample of funds and LPs. Assuming this sample is less biased by investor-specific mandates or co-investment, the figure therefore suggests that gross-return differences are not the primary source of within-fund return variation that we observe empirically.

3.4 Differences in Contract Terms

Differences in fees (investment costs) are a natural reason that net-of-fee returns might vary within a fund. We now outline two features of the contracting environment in private markets that provide a legal and theoretical foundation for costs to differ across LPs: (i) private and bilateral negotiation, (ii) complexity. We discuss the contracting environment at length in Internet Appendix B.

3.4.1 Bilateral Contracting

Investment into private market vehicles is governed by a private contract, the limited partnership agreement (LPA), between the investment manager (GP) and the investors (LPs). Generally speaking, the GP and the LPs privately negotiate the terms of the LPA, including the expenses borne by LPs, tax treatment of fund income, the ability of the GP to unilaterally amend the LPA, and the degree to which the GP is indemnified through the partnership.

The LPA dictates and governs four broad types of expenses that ultimately determine the returns of LPs: (i) management fees, which are typically a percentage of committed capital; (ii) performance-contingent fees or carry; (iii) fund and organizational expenses; and (iv) portfolio company fees. Portfolio company fees are paid to the GP by the firm in which the partnership invests and in many cases the LPA stipulates that the GP is supposed to share this income to LPs in the form of a fee offset or reduction (Phalippou et al., 2018). We provide more detail on the nature of all of these expenses in Internet Appendix B.5.

LPAs can also be used to create multiple investor classes. For instance, tax-exempt investors like pensions can opt to be separated from taxable investors in order to minimize tax burdens for both groups.[13] Funds may also allow LPs to choose from a menu of fees, with each choice representing a different investor class. Bain Capital is one notable example of a GP who recently shifted to a menu-model, offering investors a choice to pay 1% management fee and 30% carry or 2% management fee and 20% carry (Zuckerman and Or, 2011; Markham, 2017).

Though the LPA is visible and agreed upon by all LPs in the fund, its terms are often superseded by additional agreements (so-called "side letters") that are negotiated bilaterally between the GP and individual LPs.[14] Side letters can alter many aspects of the original LPA, such as reporting requirements by the GP, explicit modifications of fees, or exemptions from paying certain fund expenses (e.g., placement agent fees). They can also establish provisions for "most favored nation"

[13]LPAs for tax-exempt investors can allow capital to flow through blocker corporations that improve tax efficiency. According to several large LPs and our read of LPAs, public investors may have to opt in or negotiate for these types of tax optimization services because they are not always treated as tax-exempt by default.

[14]Not all LPs get to negotiate with contract terms. According to Da Rin and Phalippou (2017), only 59% of all LPs and only 36% of small LPs "always negotiate contract terms", implying that a substantial fraction of LPs does not.

(MFN) status, which under certain conditions allow an LP to view and select the terms of side letters that have been offered to other LPs.[15]

The use of side letters to modify LPAs is widespread. Under confidentiality agreements, we obtained a subset of LPAs for 91 funds in our sample. Within this subset, 76% have language indicating that the GP has the sole discretion to enter into a side letter with any LP. Furthermore, the LPAs explicitly state that any such side letter supersedes the LPA and may confer rights and benefits to an LP that are not granted to others. In addition, the LPAs that we studied varied in terms of how much side letters were required to be made visible to other LPs: 46% had no language requiring GPs to notify LPs of side letters, 41% required GPs to notify LPs but did not require the terms of side letters to be granted to all LPs, and 13% stated all LPs would be notified of any side letters and could opt into their terms within 20-60 days of the fund's close.

In principle, side letters or similar LPA amendments are therefore one mechanism through which economic terms could differ across investors in the same fund. In an ideal world, we could directly explore this channel if we observed the side letters associated with each fund in our data, though this is difficult in practice due to the private nature of these contracts. Nonetheless, industry surveys of GPs indicate that nearly 50% use side letters to offer some investors more favorable terms (Toll and Centopani, 2017, Chart 2.31). This is likely a lower bound because such surveys are self-reported and GPs have little incentive to reveal their contracting strategy. We provide further background on how the contracting environment (e.g., LPAs, side letters, MFNs) could give rise to fee dispersion in Internet Appendix B.

3.4.2 Contract Complexity

The preceding discussion highlights that most LPAs allow for investment terms to be bilaterally negotiated between LPs and GPs. It is thus natural to think of this contracting environment through the lens of search and bargaining models (e.g., Burdett and Judd, 1983; Bester, 1988; Hortaçsu and

[15] The nature of MFNs and carve outs that apply to them vary across LPAs. In some cases, MFN clauses will automatically confer the benefits of all other side letters. In others, MFNs give LPs the ability to opt into side letter provisions granted to LPs of a similar size and within a fixed window (e.g., 30 days after close), see Toll and Centopani (2017, Chart 2.32).

Syverson, 2004; Duffie et al., 2005; Allen et al., 2019). Price dispersion is a ubiquitous feature of these models due to costs associated with search and bargaining. In our setting, these costs depend for example on the ability of LPs to obtain information on different fee structures offered in a fund (e.g., through side letters or negotiation) as well as in other funds (i.e., outside option), or the speed at which LPs can evaluate the economic content of different contracts. For many pensions, evaluating the terms of an LPA is time consuming because LPAs are long and complex contracts: within the set that we analyzed, the average LPA was 75 pages long and contained 41,000 words, though some contained as many as 70,000. By comparison, the prospectus of a typical Vanguard mutual fund is 8 pages long and contains less than 2,000 words.

Complexity can also cause and amplify price dispersion in models where investors differ in information-processing costs or sophistication. For example, in the classic model of Salop and Stiglitz (1977), price dispersion occurs in equilibrium because consumers differ in their ability to discern the true price of a good. Analogously, LPA complexity could make it difficult for some LPs to accurately estimate the cost of investing in a private funds. As one example of this friction, a group of state comptrollers recently filed a complaint to the U.S. Securities and Exchange Commission (SEC) about fee disclosure practices in private market vehicles (State Comptroller SEC Letter, 2015).[16] Relatedly, in Gabaix and Laibson (2006), complexity leads to effective price dispersion in equilibrium because it allows producers and sophisticated consumers to benefit from consumers who myopically evaluate complex contracts. Finally, heterogeneity in LP sophistication could lead to ex-post dispersion in fees if LPA adherence is difficult to enforce or verify. Indeed, the former SEC compliance office director, Andrew J. Bowden, stated in 2014 that a review of LPAs by the agency revealed "what we believe are violations of law or material weaknesses in controls" in over 50% of cases.[17]

In sum, given that contracts between LPs and GPs are often negotiated on a bilateral basis, fees

[16]In response, the Institutional Limited Partnership Association (ILPA) has developed a "standard Model LPA" specifically designed to reduce "the cost, time and complexity of negotiating the terms of investment." With regard to side letters, the ILPA's model LPA would: (i) deliver full transparency of all side letters by default; and (ii) provide all LPs with the more favorable rights of any side letter by default. See Section Article 20.6.2 of the model LPA.

[17]See the "sunshine" speech by Andrew J. Bowden in 2014.

may naturally differ across LPs within the same fund. Fee dispersion could arise in equilibrium through simple cost-based pricing, search and negotiation frictions, contract complexity, or heterogeneity in LP sophistication. In the next two sections, we present several pieces of evidence that are consistent with the existence of fee dispersion within funds and characterize this dispersion further. Furthermore, we show that many of these patterns are inconsistent with the other sources of return variation discussed above (see Table 2 for a summary).

4 Characterizing Fee Dispersion

In this section, we present evidence that fees are an important source of within-fund variation in net-of-fee returns. We then characterize which funds are more likely to employ multiple fee structures and estimate how much specific contract features vary across LPs. Our analysis proceeds in three steps. First, we show that returns typically cluster together. This clustering (or tiering) is the main source of within-fund return variation and is consistent with a model where GPs offer LPs a limited number of contracts with different fee structures. Second, we show that certain funds and GPs are more likely to offer multiple contracts. Third, we propose a method to estimate contractual differences across investor tiers, exploiting both the cross-sectional and time-series variation in LP contributions and distributions in the same fund. Although we do not observe specific contract terms, our approach allows us to differentiate between fees and expenses that are contingent on performance (e.g., carry) and those that scale with commitments and time (e.g., accumulated management fees).

4.1 Within-fund clusters of net-returns

LPAs and associated side letters are time consuming for GPs and LPs to negotiate. To streamline the process, GPs often standardize the language of their side letters to address overlapping concerns and limit the number of provisions that most favored nation clauses could elect (Morgan Lewis, 2015). Contract standardization should therefore lead to clustering of net-of-fee returns within a fund. Panel A of Figure 3 plots the distribution of within-fund net-returns measured by DVPI for an anonymized fund at a single date. This fund has 16 investors for whom we have a full panel of returns that all start at the same date. The distribution of DVPI exhibits two distinct clusters, with a majority of investors earning \$1.14 per dollar invested and a smaller subset earning \$1.95.

Panel A of Figure 4 shows that this pattern of clustering is common within most funds. To construct this graph, we use machine learning techniques to partition observations within each fund into distinct groups. A textbook approach to this problem is called k-means clustering, which assigns each return observation in a fund to one of k clusters based on their distance to the clusters (see e.g., Steinhaus, 1956; MacQueen, 1967; Jain et al., 1999). The k clusters are themselves chosen to minimize the total distance of observations to their respective clusters. We select the optimal number k of clusters based on Silhouette scores, as is common in the machine learning literature, and define $Tiers_{ft}$ as the number of clusters in fund f at time t. The number of clusters $Tiers_f$ in fund f is then defined as the time-series average of $Tiers_{ft}$, rounded to the nearest integer. Panel A of Figure 4 plots the distribution of $Tiers_{ft}$ for funds with at least two years since the final close. According to this procedure, 23% of funds in our sample have one DVPI tier, 74% have two tiers, and 3% have three or more tiers. Hence, most funds appear to group their investors into contract tiers in terms of fees.[18]

We next analyze clustering in (capital) call rates, defined as the ratio of cumulative contributions per dollar of commitments. Call rates are useful for understanding fee differences across investors because contributions include capital that is invested by the GP, management fees, fund expenses, and portfolio-company fee offsets (Phalippou et al., 2018). This observation means that call rates for investor p in fund f at time t can be decomposed into three terms:

$$\text{call-rate}_{p,f,t} = i_{p,f,t} + m_{p,f} \times t + \varepsilon^{m}_{p,f,t}. \tag{3}$$

$i_{p,f,t}$ is defined as the cumulative amount of capital that has been invested into the fund per dollar of commitments. $m_{p,f}$ denotes any fees that are charged on an annual basis as a percentage of

[18] We observe similar clustering patterns using TVPI and IRR. See Figure IA3 in the Internet Appendix.

investor p's commitment size (e.g., management fees). $\varepsilon^{m}_{p,f,t}$ is a residual term that captures any fund expenses, portfolio company fees, or measurement error. $i_{p,f,t} = i_{f,t}$ will be constant across LPs under the assumption that capital is invested based on the pro rata share of commitments. Under this assumption, cross-sectional variation in call rates will reflect variation in $m_{p,f}$ or $\varepsilon^{m}_{p,f,t}$. Moreover, as we exploit in Section 4.3, any variation in $m_{p,f}$ will generate a linear relationship between call rate variation and age.

Panel B of Figure 4 analyzes clustering in call rates in our data using the k-means clustering method described above. According to this procedure, 30% of funds have one cluster, 66% of funds have two clusters, and 4% of funds have three or more clusters. Thus, in the typical fund, investors appear to be grouped into two tiers in terms of fees that are included in contributions. For example, for the sample fund in Figure 3 Panel B, the call rates of 16 investors cluster around two values of 0.99 and 1.04. To understand the magnitude of this dispersion, suppose it is fully driven by $m_{p,f}$ and that $i_{f,t} = 90\%$. In this case, these call rates would imply 1.80% and 1.90% for the two values of $m_{p,f}$.

Panel C of Figure 4 analyzes clustering in distribution rates, defined as the ratio of cumulative distributions over commitments. Much like call rates, distribution rates are useful to study because they are net of carry and any other performance-contingent expense that the GP would deduct before returning capital. We can decompose distribution rates into two components as follows:

$$\text{dist-rate}_{p,f,t} = d_{p,f,t} \times \left(1 - c_{p,f}\right) + \varepsilon^{c}_{p,f,t}, \tag{4}$$

where $d_{p,f,t}$ is the gross-of-fee distribution per dollar of committed capital for investor p in fund f at time t, $c_{p,f}$ is the carry rate or performance fee charged by the GP, and $\varepsilon^{c}_{p,f,t}$ is a residual that reflects any additional expenses that come out of distributions or measurement error. $d_{p,f,t} = d_{f,t}$ will not vary if each investor's share of gross distributions is based on its commitment size, as is common in most funds. Under this assumption any clustering in distribution rates will reflect clustering in $c_{p,f}$ or $\varepsilon^{c}_{p,f,t}$.

Panel C of Figure 4 shows distribution of clustering across funds based on distribution rates.

We find that 27% of funds have one cluster, 69% have two clusters, and 4% have three or more clusters of distribution rates. Thus, two values of $c_{p,f}$ – or more generally, fees that are charged on distributions – are typical for most funds. Panel C of Figure 3 shows that our example fund features two clusters of distribution rates, one at 1.13 and the other at 1.17. To get a sense of magnitude, assume $d_{f,t} = 1.5$ and all variation in call rates is driven by performance fees. In this scenario, the clustering in distribution rates implies 25% and 22% for the two values of $c_{p,f}$.

In summary, we find that net-of-fee returns within the typical fund are clustered or tiered together, as opposed to being continuously distributed across investors. Most funds (around 70%) have two investor tiers in terms of fees that are charged prior to investment and those that are deducted prior to distributions. Around a quarter of funds have only a single tier of returns, which is consistent with a single fee structure across investors. Importantly, this pattern of clustering at the fund level cuts against the idea that measurement errors alone drive within-fund variation in returns.

4.2 Are some funds more likely to use contract tiers?

The preceding analysis showed that about one-quarter of funds exhibit no dispersion in net-of-fee returns, which is an indication that all investors are charged the same fee. This section investigates what distinguishes these funds and GPs from those that use multiple fee structures. The main variable that we analyze is a dummy variable that equals 1 if fund f has more than one tier in terms of DVPI, call rates, or distribution rates and 0 otherwise. We define tiers in each category as in Section 4.1. We restrict our analysis to funds that are at least one year old to give time for potential fee differences to materialize in the data.

Table 3 presents a set of OLS regressions of a fund's propensity to tier investors, $Tiers_f > 1$, on various correlates. In column (1), we regress the tier-indicator on a full set of GP fixed effects. The R^2 in this regression is just under 30% and an F-test of the null of no GP effects is strongly rejected. This suggests that some GPs consistently tier investors across all their funds, whereas others consistently offer only one contract to their pension LPs. The existence of GP effects also implies that LP-specific accounting or differences in gross-return exposure cannot fully account for within-fund return variation.

We explore covariates next. In all regressions, we include fixed effects for the number of investors in the fund. This eliminates any bias that measurement error might introduce to our machine-learning classification of tiers. For example, measurement error could lead to an upward bias in the number of tiers for funds with many investors.

In column (2) of Table 3, we regress the fund level dummy, $Tiers_f > 1$, on past and current performance of the GP. Past performance is the average quartile ranking of all funds raised by the GP raised prior to fund f's final close date. Current performance is fund f's realized quartile ranking. In both cases, quartile rankings are measured as of 2020Q4.[19] Funds in quartile 1 are the best performing fund and funds in quartile 4 are the worst performing. To see the motivation for this regression, consider a model where capital commitments reflect manager skill (Berk and Green (2004)) and contract negotiations are costly. In this case, a high performing GP could charge a high fee to all LPs without losing capital commitments, while low performing GPs may have to negotiate more to attract investors (Toll and Centopani, 2017, p. 29). Consistent with this intuition, column (2) indicates that GPs with one extra quartile of past performance are four percentage points more likely to use a single fee structure. This effect is large when considering that the unconditional likelihood of charging a single fee is about 20%. The logic of Berk and Green (2004) further suggests that the number of funds previously raised by a GP is also a proxy for its skill. Accordingly, column (3) shows that GPs who raise one extra fund are one percentage point more likely to uses a single fee tier.

In column (4), we explore whether a fund's propensity to tier investors varies across asset classes by including indicator variables for private debt, private equity, real estate, and venture capital. The indicator for infrastructure funds is omitted, so all point estimates are relative to that asset class. Venture capital funds are far less likely to use multiple investor tiers than other asset classes – 32 percentage points less likely than infrastructure and 23 percentage points less

[19] See this report by Preqin for the methodology.

likely than private equity.[20] This finding is consistent with the observation that the venture capital industry was an earlier adopter of standardized LPA and side letter provisions (Robbins, 2019). Private equity funds are 8 percentage points less likely to tier than infrastructure funds, though the point estimate is statistically significant only at the 10% level. The point estimates on private debt and real estate funds are not statistically different from zero. Overall, the fact that the propensity to tier investors varies by asset class cuts against the idea that LP-specific accounting conventions are the primary driver of within-fund return variation.

Because fund size and contracting terms are determined endogenously, in columns (5-7) we rerun the regressions in columns (2-4) with a fixed effect based on the deciles for the fund's size. The point estimates are comparable across these specifications and lead to the same conclusions.

4.3 Within-fund dispersion in Fees

Most funds in our sample appear to tier investors into one of two fee structures. Fee terms could differ along several dimensions, including management fees, fund expenses, portfolio company fees and associated offsets, taxes, or performance fees. Though we do not observe exact contractual differences, the panel nature of our data allows us to decompose differences in fees into "fixed fees" that scale with commitment and time, and variable fees that scale with performance. To see why, recall from (3) that capital call rates can be written as the sum of three components: capital used to fund investment, fixed fees $m_{p,f}$ that scale with time, and a residual term (e.g., fund expenses that do not scale with time). Let p^{σ}_{ft} be the within-fund standard deviation of call rates in fund f at time t. Under the assumption that investment rates (per dollar of commitments) are identical across investors $\left(\text{i.e., } i_{p,f,t} = i_{f,t}\right)$, from (3) we can write:

$$p^{\sigma}_{ft} = m^{\sigma}_{f} \times age_{f,t} + \varepsilon^{\sigma}_{f,t}, \tag{5}$$

[20] These results are consistent with our analysis of LPAs from Section 3.4. In this limited subset, funds whose LPAs contain side letter language have effectively no chance of having a single return tier. Moreover, venture capital (VC) funds are 7 pp less likely to include side letter language in their LPAs. Within the funds whose LPAs have side letter language, VC funds are 15 pp more likely to state that all LPs can view any side letters and 9 pp more likely to allow for LPs to opt in (either conditionally or guaranteed) to any side letter provisions.

where m_f^σ is the within-fund dispersion in fixed fees, $age_{f,t}$ is the fund age measured as the years elapsed since the fund's close, and $\varepsilon_{f,t}^\sigma$ is the dispersion in the residual. This equation says that dispersion in call rates should be linearly related to age with a slope of m_f^σ. Panel A of Figure 5 confirms a strong linear relationship using a binned scatter plot. The plot pools data on all funds within their first five years of life, as this is the period when management fees are typically charged as a percent of committed capital. The slope of 84 basis points (standard error of 10) reflects the average within-fund dispersion in fixed fees across all funds. We view this as evidence that fixed fees vary within the average fund, in part because it is hard to imagine how LP-specific accounting, measurement error, or within-fund differences in gross returns could generate such a strong linear pattern.

Table 4 repeats this analysis within each asset class. To increase power, we exclude all funds that we categorize as having a single fee tier. The first column of the table shows the point estimates of m_f^σ and the second column is the associated standard error. Within-fund dispersion in fixed fees varies strongly across asset classes. Fixed fees within infrastructure funds vary on average by 100 basis points, whereas they vary by only 37 basis points in venture capital funds. The low dispersion in venture capital funds is consistent with our finding that those funds are also less likely to tier investors in the first place (Section 4.2).

A similar logic allows us to uncover the average within-fund variation in variable fees $c_{p,f}$, which could arise from differences in carry or differences in the tax-treatment of carry. Specifically, dispersion in performance fees should be linearly related to performance, conditional on performance fees being charged. From (4), we can write d_{ft}^σ the within-fund standard deviation of distribution rates as:

$$d_{ft}^\sigma = c_f^\sigma \times \max(\widetilde{r}_{ft} - r_f^h, 0) + \nu_{f,t}^\sigma, \tag{6}$$

where $\widetilde{r}_{ft}$ is the gross return of the fund and r_f^h is a minimum return or hurdle rate the fund must achieve before it can charge performance fees. c_f^σ is the within-fund standard deviation in performance fees, and $\nu_{f,t}^\sigma$ is the dispersion in the residual term that reflects non-performance related fees that are deducted from distributions (or measurement error). Equation (6) implies that the relationship between dispersion in distribution rates should have a specific functional form, one that resembles the payoff of a call option. In the region where funds are below their hurdle rate, d_{ft}^{σ} should be insensitive to performance. In the region where funds are above their hurdle rate, d_{ft}^{σ} should be linearly increasing in performance with a slope equal to c_f^{σ}.

Panel B of Figure 5 shows that this call-option pattern is precisely what we observe in the data. To construct the binned scatter plot, we pool over all funds and proxy for each fund's gross fund performance $\widetilde{r}_{ft}$ using its maximum TVPI of fund f at time t. We define funds as profitable enough to charge performance fees if they have a TVPI of at least 1.09 and an IRR of at least 9%.[21] The binned scatter plot also partials out vintage fixed effects to control for any potential age-specific effects on performance fees. In the region where funds are defined as unprofitable, the slope estimate is 0.1 and is not statistically different from zero. In the region where funds are profitable, the slope of $c = 5.3\%$ (standard error equals 0.7%) indicates that variable fees vary by 5.3% in the average fund. Once again, it seems hard to imagine that anything other than fees could generate this precise pattern.

The middle columns of Table 4 present estimates of c^{σ} by asset class and their associated standard error. These estimates are based only on funds who we define as profitable enough to charge performance fees. Real estate and private debt funds display the largest volatility in performance-based fees at 6.2% and 6%, respectively. Much like our dispersion estimates of fixed fees, venture capital funds have the smallest average dispersion in performance variable fees at 0.6%. Private equity funds lie in the middle of these two extremes, with an average within-fund dispersion of 3.5%. In all asset classes, we reject the null that the average within-fund dispersion in variable fees equals zero.

As discussed above, we should not be able to estimate dispersion in variable fees in the region where funds are unprofitable. We exploit this observation to implement a series of placebo tests in the last two columns of Table 4. The table reports the point estimate of a regression of dispersion in distribution rates on fund performance in the subsample of unprofitable firms. As expected, in

[21] Performance fee or carry is typically charged once the IRR is above 8-10%.

all asset classes, the estimated slopes are not statistically different from zero.

To summarize, venture capital funds have the lowest average dispersion in fixed and variable fees. Within private equity funds, the average dispersion in fixed and variable fees are around 80 bps and 3.5%, respectively, which is consistent with public reports on the menu-model that Bain Capital has used in recent years (Markham, 2017). Infrastructure, private debt, and real estate funds tend to have larger within-fund dispersion in fixed and variable fees. As further validation for our estimation approach, we find that unprofitable funds do not exhibit dispersion in performance variable fees.

5 Do some investors consistently get better terms?

In this section, we show that LPs who outperform in one fund are more likely to outperform in their other funds. This pattern is consistent with the idea that some LPs are able to consistently select or obtain the best terms in their respective funds, at least ex-post. Furthermore, matching between GPs and LPs appears to be an important component of this persistent outperformance. We then analyze the characteristics of the LPs that are most likely to outperform and estimate how much these traits explain their persistent outperformance.

5.1 Pension-effects

5.1.1 Baseline estimates

In light of our evidence on within-fund tiering of investors, we now analyze whether some investors are more likely to be in the top tier of performance in all of their funds. Specifically, within each fund, we define an indicator variable y_{pf} based on whether investor p has above-median TVPI for the majority of fund f's life. We construct y_{pf} using medians based on our finding that most funds have two tiers of investors. In this sense, y_{pf} can be interpreted as a measure of whether p is a top-tier investor in fund f. We measure performance using TVPI, since it should capture any advantageous contract terms in either contributions (e.g., fixed fees) or distributions (e.g., performance-based fees). To the extent that a fund's net-asset-value (NAV) excludes management fees, TVPI also offers us some more power to detect differences in management fees. We then assess the persistence of within-fund performance across funds using the following regression:

$$y_{pf} = \lambda_a + \alpha_p + \varepsilon_{pf}, \tag{7}$$

where λ_a are age fixed effects as measured by vigintiles of fund f age and α_p denotes an investor fixed effect. Controlling for fund age allows us to better isolate pensions who truly have dominated fee contracts from those who trade off high fixed fees for low variable fees or vice versa. For instance, pensions who trade off the two fee components may outperform early in a fund's life but not later. When estimating (7), we exclude funds that we classified as having a single fee tier in Section 4.1.

Under the null hypothesis of no pension effects, the estimated α's should not be statistically distinguishable from each other. In other words, if contract terms in a given fund are randomly assigned to pensions, then we should not be able to reject an F-test that the α's are jointly equal to each other. In Table 5, we report the number of pension effects K, the F-tests and their associated p-values based on the core sample. When moving from rows (1) to (3), we conduct the F-test for whether the α's are jointly equal based on funds that are at least one, four, and eight years old, respectively. In all cases, the estimated F-statistic is large enough that we reject a null of no pension effects with a p-value of less than 0.01.

The standard approach to conducting F-tests like those in Table 5 rely on parametric assumptions to test the null of no pension effects. As a robustness check, we calculate non-parametric p-values based on a permutation test where we: (i) randomly assign return paths to investors within each fund f; (ii) calculate the simulated value of y_{pf}; (iii) run regression (7); and (iv) recalculate the F-statistic from the test of equality across α's. We repeat this procedure 1,000 times to generate an simulated distribution of F-statistics, after which we compute a non-parametric p-value based on where the actual F-statistic falls in this distribution. We denote the p-values based on these permutation tests as p^*. Reassuringly, we again reject the null of no pension effects.

Though the preceding F-tests provide a statistical sense of the size of pension-effects in our data, they do not easily convey the economic magnitude of such effects. To get a better sense, we compute the distribution of the estimated pension effects and compare it to the distribution implied by random assignment of returns (i.e., fees) within each fund. Because the true distribution of pension effects $\alpha's$ differs from the estimated distribution due to sampling error, we adjust the estimated pension effects using an Empirical Bayes method (see Chetty et al. (2014) and Egan, Matvos, and Seru (2018) for examples). Let $\hat{\alpha}$ denote the vector of estimated α's based on regression (7). Using Casella (1992, Eqs. 7.11 and 7.13), we can calculate the empirical Bayes estimate $\widetilde{\alpha}$ as

$$\widetilde{\alpha} = \bar{\alpha} + \max\ (1 - B, 0) \times (\hat{\alpha} - \bar{\alpha}),$$

where $\overline{\alpha}$ is the average of the estimated fixed effect vector $\hat{\alpha}$ and

$$B = \frac{1}{F}\left(\frac{K-1-2}{K-1}\right)$$

is a shrinkage coefficient. The F in the formula for the shrinkage coefficient B corresponds to the F-statistic from the joint test that the $\hat{\alpha}$ are equal (as reported in Table 5).[22] Intuitively, the F-statistic is larger (and B smaller) when the pension-effects are estimated with more precision, and in turn, the Bayes estimate does not shrink $\hat{\theta}$ as much towards its mean.

Panel A of Figure 6 visualizes the rejection of the F-test. The orange line shows the distribution of pension effects under the random assignment of return paths (i.e. contracts) within each fund. The distribution of observed pension effects $\widetilde{\alpha}$ (in blue) has much fatter tails, hence why we reject the null of random assignment. The distribution of observed pension effects shows that the 95% percentile outperforms in 73% of its funds while the 10% percentile pension outperforms in only 13% of its funds.

Pension effects could in principal be driven by LP specific accounting conventions. However,

[22]See Morris (1983) for a formal description of Empirical Bayes inference.

we observe similar patterns when using DVPI and IRR to measure returns (Panel A of Figures IA4 and IA5 in the internet appendix), meaning NAV reporting or recallable capital accounting are unlikely to drive our results. Moreover, LP-specific accounting is inconsistent with our finding that some funds and GPs are more likely to tier investors than others (Section 4.2). We therefore interpret this evidence as showing that some pensions consistently obtain or are offered the best contract in all their funds, at least on the basis of ex-post performance.

5.1.2 LP-GP Effects

Search and bargaining models suggest that contracting can depend on the specific match between LPs and GPs. Empirically, Hochberg, Ljungqvist, and Lu (2007) show that relationships between LPs and GPs matter for fund performance. We investigate whether matching between GPs and LPs relates to an LPs net-of-fee return outperformance by augmenting regression (7) as follows:

$$y_{pfg} = \lambda_a + \eta_{pg} + \varepsilon_{pfg} \tag{8}$$

where y_{pfg} is the indicator of investor's p relative outperformance in fund f managed by fund manager g. We define the indicator as above. As before, λ_a are the vigintiles of fund f age. The new term in the regression is η_{pg}, which are LP-GP fixed effects. The LP-GP effects η_{pg} measure the outperformance of investor p in funds managed by GP g. If, for instance, some investors receive better terms than others in funds managed by a specific set of GPs, then we should reject an F-test of the joint significance of the η's.

Table 5 reports F-statistics and their associated p-values from testing whether the η's are jointly equal. We reject the null of no LP-GP effects (η's) when using parametric p-values and non-parametric p-values based on the permutation tests described in Section (5.1.1). This evidence suggests that matching between LPs and GPs is important for understanding why some pensions consistently outperform others when investing in the same fund.

5.2 Observable Pension Characteristics

In this section, we map the pension effects to observable pension characteristics to better understand why some pensions consistently outperform others in their respective funds. We proceed in two steps. First, we replace the pension effects in regression (7) with characteristics related to size and investor sophistication. Second, we show that characteristics explain some, but not all of the observed pension effects. This suggests that unobservable traits like negotiation skill or bargaining power materially impact the fees pensions pay in private-market funds.

5.2.1 Observable Characteristics X_{pf}

Pension effects are silent on why some pensions consistently outperform other pensions ex-post. To better understand the economic determinants of these pension effects, we replace pension fixed effects with observable characteristics X_{pf} in regression (7) as follows:

$$y_{pf} = \mu_f + \beta \mathbf{X}_{pf} + \varepsilon_{pf}, \tag{9}$$

where y_{pf} is the indicator of investor's p relative outperformance in fund f as described above, and μ_f is a fund fixed effect and λ_a are the vigintiles for fund age as before. We include fund fixed effects in the regression to ensure that β is identified using within-fund variation.

We consider the following set of observable characteristics X_{pf}. For each investor p in fund f, we compute p's share of the total fund as their commitment amount divided by the total fund size. We include each investor's share of the fund to account for potential returns to scale when raising capital. For example, one might expect that GPs might reduce fees for investors that account for a larger fraction of the fund, as this would then free up the GP to focus on optimizing the investment portfolio instead of raising capital.

Due to information asymmetries about manager skill, signaling effects are likely to be important for GPs when they raise a fund. For example, if a GP secures a capital commitment from a large and well-known pension, then other pensions may be more willing to commit capital to the fund. We code investor p as "Large" if its total assets under management are over \$100 billion at the time of fund f's launch, a designation that is reserved for easily recognizable pensions in our data. In addition to the potential signaling effect that they may have on fund raising, large investors are also more likely to possess the ability to deploy large amounts of capital quickly, so size is likely related to the economies to scale in fund raising discussed above.

We include three variables that capture the experience and potential negotiation skill of each investor in private markets. A priori, it is plausible to think that skill in fee negotiation improves as investors become more experienced in the nature of private market investment vehicles. Motivated by our finding of LP-GP effects in Section 5.1.2, we include the LP-GP prior fund count. Specifically, for each LP-GP pair, we count the number of funds that are managed by general partner g in which p has invested. We use the full dataset to compute this measure because we want to capture settings where a GP reduces fees for investor p in fund f in expectation that the investor will invest in future funds raised by the GP. The second variable captures how well an investor's prior funds have performed. Arguably, LPs skilled at manager selection are also skilled in contract negotiations. We measure each investor p's past fund performance as the average quartile ranking of its active funds at the time of fund f's close, where quartile rankings are the same used in Section (4.2). The third variable we use is an indicator for whether investor p was an early private equity investor based on having invested in PE prior to 2008.

The last set of variables that we include are related to pension governance.[23] We include board size to account for any potential coordination problems that may cause larger boards to sub-optimally negotiate fee contracts. In addition, Andonov et al. (2018) find that pension boards with more state officials are more likely to make poor investment decisions in private equity, likely due to distortions from political considerations. Motivated by that finding, we include the percent of pension p's board that is made up of elected, as opposed to appointed, members. Intuitively, boards that have more elected officials are more likely to focus on the pension plan beneficiaries when determining fee schedules with their general partners. For each pension and fund pair (p, f),

[23] We are grateful to Josh Rauh for sharing the data on pension board composition from Andonov et al. (2018). The data ends in 2013, so we extend it to 2018 to better match our sample.

the total number of board members and the percent of elected members are both measured as of fund f's vintage year.

Table 6 presents the results from regression. Columns (1) to (3) focus on funds that meet a certain minimum age requirement (i.e., 1, 4, and 8 years). Across all subsamples, both measures of size substantially impact the likelihood of being a top-tier investor in a fund. Moving above the threshold of $100 billion dollars in assets under management (AUM) increases the probability of being a top-tier investor by around 20 percentage points. One percentage point increase in an investor's check size, i.e., percent of fund, increases the probability of being a top tier investor by about 30 to 69 basis points.[24] The point estimate on the percent of the fund is measured imprecisely in the sample of funds that are 8 years old. This is largely because the investor's share of the fund is correlated with other covariates, namely size. If we regress y_{pf} on percent of the fund excluding all other covariates, the slope coefficient is 0.63 and has a t-statistic of 2.92.

Measures of investor sophistication and experience also appear to play an important role in determining an investor's net-of-fee return performance. For instance, each additional fund between the LP and the GP increases the likelihood of being a top-tier investor by 0.75%, though the point estimate is not always precisely measured. This finding relates to our evidence of LP-GP matching in Section 5.1.2. Investors with past success in manager selection are 10 to 16 percentage points more likely to be top-tier investors in their funds. Recall that we measure past success for each investor based on the average quartile ranking of its previous funds. Similarly, pensions who were early investors in private equity are 10 percentage points more likely to outperform others in their funds. Finally, there is also some evidence that better pension governance improves the probability of being a top-tier investor. Pensions whose boards have 10 percent more elected members are roughly 2 percent more likely to be top-tier investors. For context, the standard deviation of percent of elected board members is 26% in our sample. On the other hand, there does not appear to be a stable relationship between board size and top-tier investor status.

[24] This is consistent with tiered MFN clauses that make access to better terms conditional on commitment size (source).

5.2.2 How much of pension-effects are due to observables?

In this section, we assess how much observable characteristics account for the observed pension effects documented in Section 5.1. There are three steps to adjust pension effects for observable characteristics. First, we take the full panel of raw returns, as measured by TVPI, and regress it on fund-by-quarter fixed effects and the full set of covariates from Section 5.1.1 interacted with vigintiles for fund age. We use the raw level of returns because it allows us to flexibly control for covariates. Second, we use the residuals from the regression to compute whether investor p is a top-tier investor in fund f, denoted by $\tilde{y}_{pf}$, as in Section 5.1. Third, we estimate pension effects based on this characteristic-adjusted tier assignment $\tilde{y}_{pf}$ and apply the empirical Bayes procedure described in Section 5.1.1.

Panel B of Figure 6 plots the distribution of the resulting characteristic-adjusted pension effects as a dashed-green line, as well as the observed pension effects (characteristic unadjusted) as a blue line. The blue line is identical to Panel A. The plot shows that characteristics do account for some of the observed pension effects. The characteristic-adjusted distribution is shifted to the left because some of the differences in contract terms (taken as given in Panel A) are absorbed by observable characteristics like size. If all pensions earned the same within-fund return after characteristic adjustments, then no pension would outperform any other and the entire mass of the green-dashed line would be at zero. Nonetheless, the figure also shows that there are still many LPs who consistently over- or underperform after controlling for characteristics. For example, the far right tail of the characteristic-adjusted pension effects is at 80% and the far left tail is at 10%. Based on the set of covariates that we consider, this evidence suggests that a subset of pensions consistently outperform others in their funds for reasons that are orthogonal to pension size, share of a fund's commitments, or past experience in private-market funds. We interpret this as evidence that unobserved traits like negotiation skill or bargaining power meaningfully impact the fees that pensions pay in private-market funds.

6 Conclusion

Research on the costs of private-market funds has confronted two important empirical challenges. First, investment terms are privately negotiated and rarely observed by outsiders. Second, fee structures in private-market funds are complex, so much so that investors often struggle to calculate how much they have paid in total fees. This paper has developed a method of addressing these challenges by using variation in net-of-fee returns across investors in the same private-market fund.

We apply this method to a sample of nearly $440 billion of investments made by 219 public pensions into 2,407 funds managed by 857 GPs. Within this relatively large sample, we document sizable variation in within-fund net-of-fee returns, regardless of how returns are measured. Our analysis suggests a significant component of this variation is driven by fees. Most funds feature two clusters of net-of-fee returns, which suggests that the typical fund groups investors into one of two fee-tiers. We find that some investment managers and funds appear more likely to tier their investors than others. While there are several dimensions through which fees could differ across investor tiers, we develop a methodology to estimate the average within-fund dispersion of fees that accumulate linearly with age and those that depend on performance. We then provide evidence that some pensions consistently select or are offered the lowest fee structures in their funds, at least in terms of ex-post performance. Characteristics like size explain some of these pension effects, though unobservable traits like negotiation skill also appear to play an important role.

Our analysis can be extended in several dimensions, particularly to address the welfare and policy implications of our results. For example, we find that some pensions consistently pay higher fees ex-post than others in their funds. Our results do not allow us to determine whether these pensions make suboptimal decisions ex-ante. The most direct way to address this question is to refine our method of estimating within-fund dispersion in fixed and variable fees, specifically by applying it at the fund-level. With fund-level estimates of specific contract parameters, one could determine which pension appears to invest under ex-ante dominated investment terms. A related question is whether some pensions optimally pay higher fees in funds managed by more skilled

GPs. Overall, this type of analysis would allow for a more careful decomposition of within-fund fee dispersion into supply-side (e.g., cost-based pricing) and demand-side (e.g., LP search frictions) factors, which is critical for measuring the welfare implications of fee dispersion in this setting.

Though the ex-ante welfare implications of within-fund fee variation are challenging to pin down, it is clear that this variation has large consequences for ex-post performance. To illustrate this point, we compute how much extra wealth each pension would have generated had it earned the best observed return (i.e., paid the lowest fee) in its respective funds. By this simple metric, public pensions in the U.S. would have earned $43 billion more – or $9.7 more per $100 invested – over our entire sample with better contract terms.

References

Allen, J., R. Clark, and J.-F. Houde (2019). Search frictions and market power in negotiated-price markets. *Journal of Political Economy 127*(4), 1550–1598.

Andonov, A., Y. V. Hochberg, and J. D. Rauh (2018). Political Representation and Governance: Evidence from the Investment Decisions of Public Pension Funds. *The Journal of Finance*.

Beath, A., C. Flynn, and J. MacIntosh (2014). How implementation style and costs affect private equity performance. *Rotman International Journal of Pension Management 7*(1), 1–5.

Begenau, J., C. Robles-Garcia, E. N. Siriwardane, and L. Wang (2020). An empirical guide to investor-level private equity data from preqin. *Working paper*.

Berk, J. B. and R. C. Green (2004). Mutual fund flows and performance in rational markets. *Journal of political economy 112*(6), 1269–1295.

Bester, H. (1988). Bargaining, search costs and equilibrium price distributions. *The Review of Economic Studies 55*(2), 201–214.

Braun, R., T. Jenkinson, and C. Schemmerl (2019). Adverse selection and the performance of private equity co-investments. *Journal of Financial Economics*.

Burdett, K. and K. L. Judd (1983). Equilibrium price dispersion. *Econometrica: Journal of the Econometric Society*, 955–969.

CalSTRS (2019). *CalSTRS Investment Committee Report*. CalSTRS.

Casella, G. (1992). Illustrating empirical bayes methods. *Chemometrics and intelligent laboratory systems 16*(2), 107–125.

Chetty, R., J. N. Friedman, and J. E. Rockoff (2014, September). Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American Economic Review 104*(9), 2593–2632.

Da Rin, M. and L. Phalippou (2017). The importance of size in private equity: Evidence from a survey of limited partners. *Journal of Financial Intermediation 31*, 64–76.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

Dyck, A., P. Manoel, and A. Morse (2018). Outrage and compensation in public pension performance. Working Paper.

Egan, M. L., G. Matvos, and A. Seru (2018). Arbitration with uninformed consumers. Technical report, National Bureau of Economic Research.

Eisfeldt, A., B. Herskovic, S. Rajan, and E. N. Siriwardane (2020). Otc intermediaries. *Working paper*.

Fang, L., V. Ivashina, and J. Lerner (2015). The disintermediation of financial markets: Direct investing in private equity. *Journal of Financial Economics 116*(1), 160 – 178.

Gabaix, X. and D. Laibson (2006). Shrouded attributes, consumer myopia, and information suppression in competitive markets. *The Quarterly Journal of Economics 121*(2), 505–540.

Gabaix, X. and A. Landier (2008, 02). Why has CEO Pay Increased So Much?*. *The Quarterly Journal of Economics 123*(1), 49–100.

Goldberg, P. K. (1996). Dealer price discrimination in new car purchases: Evidence from the consumer expenditure survey. *Journal of Political Economy 104*(3), 622–654.

Gompers, P. A. and J. Lerner (2010). What drives venture capital fundraising? *Brookings Papers on Economic Activity*, 149.

Grennan, M. (2013). Price discrimination and bargaining: Empirical evidence from medical devices. *American Economic Review 103*(1), 145–77.

Hochberg, Y. V., A. Ljungqvist, and Y. Lu (2007). Whom you know matters: Venture capital networks and investment performance. *The Journal of Finance 62*(1), 251–301.

Hochberg, Y. V. and J. D. Rauh (2012, 12). Local Overweighting and Underperformance: Evidence from Limited Partner Private Equity Investments. *The Review of Financial Studies 26*(2), 403–451.

Hortaçsu, A. and C. Syverson (2004). Product differentiation, search costs, and competition in the mutual fund industry: A case study of s&p 500 index funds. *The Quarterly Journal of Economics 119*(2), 403–456.

Ivashina, V. and J. Lerner (2018). Looking for alternatives: Pension investments around the world, 2008 to 2017. Working Paper.

Jain, A. K., M. N. Murty, and P. J. Flynn (1999). Data clustering: a review. *ACM computing surveys (CSUR) 31*(3), 264–323.

Lerner, J., J. Mao, A. Schoar, and N. R. Zhang (2018). Investing outside the box: Evidence from alternative vehicles in private capital. MIT Sloan working paper.

MacQueen, J. (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, Volume 1, pp. 281–297. Oakland, CA, USA.

Markham, I. (2017). *Bain keeps two investor class structure for Fund XII.* Private Funds CFO.

McKinsey (2019). Mckinsey global private markets review 2019. Technical report, McKinsey & Company.

Metrick, A. and A. Yasuda (2010). The economics of private equity funds. *The Review of Financial Studies 23*(6), 2303–2341.

Morgan Lewis (2015). *Venture Capital & Private Equity Funds, Deskbook Series*. Morgan Lewis.

Morris, C. N. (1983). Parametric empirical bayes inference: theory and applications. *Journal of the American statistical Association 78*(381), 47–55.

NASRA (2018). *ESG - Environmental, Social And Governance*. NASRA.

Phalippou, L., C. Rauch, and M. Umber (2018). Private equity portfolio company fees. *Journal of Financial Economics 129*(3), 559–585.

Preqin (2014). *The State of Co-Investments*. Preqin.

Robbins, W. (2019, November). Client alert: Ilpa releases model limited partnership agreement. Online.

Robinson, D. T. and B. A. Sensoy (2013, 11). Do Private Equity Fund Managers Earn Their Fees? Compensation, Ownership, and Cash Flow Performance. *The Review of Financial Studies 26*(11), 2760–2797.

Salop, S. and J. Stiglitz (1977). Bargains and ripoffs: A model of monopolistically competitive price dispersion. *The Review of Economic Studies 44*(3), 493–510.

SEC (2015). *Standardized Private Equity Fee Disclosure*. SEC.

Sorensen, A. T. (2000). Equilibrium price dispersion in retail markets for prescription drugs. *Journal of Political Economy 108*(4), 833–850.

Steinhaus, H. (1956). Sur la division des corps matériels en parties. *Bull. Acad. Polon. Sci 1*(804), 801.

Toll, D. M. and P. Centopani (2017). Pe/vc partnership agreements study, 2018-2019.

Zuckerman, G. and A. Or (2011). *Bain Capital Lowers Its Fees*. WSJ.

Table 1: Summary Statistics for Core Sample

Panel A: Funds, GPs, and LPs by Asset Class

	Full Sample	Infrastructure	Private Debt	Private Equity	Real Estate	Venture Capital
Funds	2,407	70	272	968	499	598
GPs	857	36	108	349	195	248
LPs	219	82	127	194	148	117
N	9,847	306	1,215	4,475	1,895	1,956

Panel B: Core Sample Characteristics

	Mean	Stdev	Min	p25	p50	p75	Max
Fund Age (years)	9	6	-1	4	8	12	27
Commitment ($ mm)	55	82	0	12	30	74	1,600
Percent of Fund	5.1	7.2	0.0	1.0	2.7	6.5	100.0
Investors per Fund	4	3	2	2	3	5	29
AUM ($ bn)	23.81	39.09	0.05	2.47	8.80	28.20	354.00

Panel C: Level and Dispersion of Performance

	Fund Age (years)			
	< 4	4-8	8-12	12+
DVPI Level	0.13	0.79	1.19	1.54
DVPI Dispersion	0.03	0.04	0.05	0.05
TVPI Level	1.09	1.56	1.55	1.62
TVPI Dispersion	0.05	0.06	0.06	0.05
IRR Level	6.48	12.99	7.63	10.10
IRR Dispersion	3.35	0.99	0.87	0.87
Funds	676	517	563	651

Notes: Panel A of this table contains reports the unique funds, GPs, LPs, and total sample size by asset class in the core sample (Section 2.1).. Panel B reports summary statistics on the core sample. Fund age is defined based on the final close date of each fund and negative values indicate that an investor enters the data prior to final close. AUM measures the total assets under management of each LP as of each fund's vintage year. In Panel C, we first compute the median and standard deviation of performance within funds, denoted by p_f and s_f, respectively. The table then reports the average p_f and s_f across funds, conditional on age. DVPI is defined as cumulative distributions divided by contributions. TVPI equals DVPI plus the reported liquidation value of any remaining investments in the fund, scaled by cumluative contributions. We use DVPI, TVPI, and the percentage IRR that is reported in Preqin to measure performance. IRRs are not fully populated and are missing for 23% of the 10,397 observations in the core sample. The number of funds is based on the available sample of funds for which we can compute return multiples.

Table 2: Potential Sources of Dispersion in Net-of-Fee Returns

		Measurement	Accounting Explanations:		Gross Return Differences:	
Explains?	Fees	Error	Recallable Capital	NAVs	Co-Investment	LP-Specific Mandates
Observed Dispersion:						
Using TVPI	x	x	x	x	x	x
Using DVPI	x	x	x		x	x
Using IRR	x	x		x	x	x
Old Funds with Small LPs	x	x	x	x		
Direct FOIA:						
Cash Flows + NAVs	x		x	x	x	x
Accounting of Recallable Capital	x				x	x
Within-Fund Return Clustering	x				x	x
GP and Fund-Specific Dispersion	x					
Dispersion Estimates: Mgmt + Carry	x					
Pension Effects	x		x	x		
GP Survey Evidence	x					
Additional Notes					Small % of portfolios and excluded from analysis	

Notes: This table presents different sources of within-fund dispersion in returns and evidence for or against each explanation. "x" denotes that the explanation in the column is consistent with the evidence in each row. The row "GP and Fund-Specific Dispersion" refers to the evidence in Section 4.2 showing that some GPs or funds are more likely to have multiple investor tiers. The row "Dispersion Estimates: Mgmt + Carry" refers to the patterns in call rates and distribution rates that allow us to estimate within-fund differences in effective management and performance contingent fees (see Section 4.3). "Pension effects" refers to the evidence in Section 5 that some pensions are consistently more likely to outperform other investors with whom they invest. GP Survey Evidence is based on charts 2.31 and 2.32 of (Toll and Centopani, 2017) showing that the majority of GPs use side letters to confer certain investors additional economic benefits and that the majority of LPs in a given fund are not given most favored nation status (see Section 3.4).

Table 3: Fund Characteristics and the Likelihood of Multiple Fee Structures

	Dependent Variable: $100 \times 1(\text{Tiers}_f > 1)$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Quartile of GP's Prior Funds		3.91^{**}			4.30^{**}		
		(3.25)			(3.56)		
Quartile of Current Fund		-0.01			0.03		
		(-0.02)			(0.04)		
Number of Funds Raised by GP			-0.91^{**}			-1.20^{**}	
			(-3.61)			(-4.63)	
Private Debt				-3.78			-3.79
				(-0.80)			(-0.80)
Private Equity				-8.24^{*}			-8.29^{*}
				(-1.88)			(-1.90)
Real Estate				-5.26			-5.33
				(-1.16)			(-1.18)
Venture Capital				-31.67^{**}			-31.71^{**}
				(-7.06)			(-7.08)
GP FE	x						
Size FE					x	x	x
Adjusted R^2	0.27	0.04	0.02	0.11	0.04	0.03	0.11
N	2,053	1,778	2,407	2,407	1,778	2,407	2,407

Notes: This table reports OLS estimates from a linear probability model of the likelihood that a fund has multiple investor tiers. For each fund-quarter, we first use a k-means clustering algorithm to classify the number of investor tiers based on capital call rates (contributions over commitments), distribution rates (distributions over commitments), and net-returns (DVPI). As discussed in Section 4.1, capital call and distribution rates can, respectively, be used to identify within-fund differences in fixed fees (e.g., management) and performance-based fees (e.g., carry). The number of fixed fee tiers is computed by taking the time-series average of each fund's data after its first year, rounded to the nearest integer. Performance-based fee and net-return tiers are determined analogously. The number of tiers is defined based on their within-fund average over time, rounded to the nearest integer. We then create an indicator variable that equals one if fund f has multiple tiers for all categories. The dependent variable in the regression is this indicator variable multiplied by 100. For each fund f run by GP g, the previous quartile of the GP's funds is the average performance quartile of g's funds that were raised before f's final close date. Number of funds raised by the GP is is measured as of the time of f's close and includes the current fund. Quartile rankings and asset class designations come from Preqin and are current as of 2020Q1. Higher quartiles correspond to better performance. Column (1) includes only a GP fixed effect. All of the remaining regressions include a fixed effect based on the average number of investors that we observe in the fund over its lifetime, rounded to the nearest integer. Columns (5) through (7) also include a fixed effect based on the decile of the fund's size. In columns (4) and (7), infrastructure funds are omitted from the regression, so the coefficients should be interpreted as the percentage likelihood of a fund having one investor tier relative to infrastructure funds.

Table 4: Average Within-Fund Dispersion in Fee Parameters

	Mgmt (bps)		Carry (%)		Carry Placebo		
	m	$se(m)$	c	$se(c)$	u	$se(u)$	p-value
Infrastructure	101	26	5.5	2.7	4.2	2.6	0.12
Private Debt	97	31	6.0	1.5	1.8	1.5	0.22
Private Equity	79	9	3.5	0.7	0.3	0.5	0.53
Real Estate	68	12	6.2	1.6	1.0	1.1	0.34
Venture Capital	37	7	0.6	0.3	-0.5	0.4	0.25

Notes: This table reports the dispersion in fixed fees (e.g., management fee, m) and performance-based fees (e.g., carry rates c) for the average fund, along with a placebo test (u) of our carry estimator in a subsample of unprofitable funds. Dispersion in fixed fees (m) is estimated via the following regression: $p^{\sigma}_{ft} = a + m \times age_{ft} + \varepsilon_{ft}$, where p^{σ}_{ft} is dispersion in fund f's capital call rate (contribution/commitment) and age_{ft} is its age in years at time t. We estimate the regression for funds that are less than five years old. Within the set of profitable funds, dispersion in performance-based fees (c) is estimated via the following regression: $d^{\sigma}_{ft} = \alpha + c \times \widetilde{r}_{ft} + \varepsilon_{ft}$, where d^{σ}_{ft} is the within-fund volatility of distributions-to-commitments and $\widetilde{r}_{ft}$ is the within-fund maximum of TVPI for fund f at time t. α is a set of fixed effects for fund vintage. The placebo test for carry estimates the same regression $d^{\sigma}_{ft} = \alpha + u \times \widetilde{r}_{ft} + \varepsilon_{ft}$ in the subset of unprofitable funds, where we should not be able to detect carry dispersion ($u = 0$). We define profitable funds as those with: (i) a TVPI above 1.09 and (ii) an IRR above 9%. All regressions are weighted by the average number of investors in each fund. The columns number the Carry Placebo headers report u, the standard error of u, and the p-value from the test of the null that $u = 0$. Standard errors are clustered by GP and fund. We exclude funds that have a single tier of investors based on a k-means clustering analysis.

Table 5: The Persistence of Relative Within-Fund Performance using TVPI

	LP Effects					LP-GP Effects				
Age Min.	F	p	p^*	K	N	F	p	p^*	K	N
1	3.87	<0.01	<0.01	174	6,427	1.74	<0.01	<0.01	1,398	3,893
4	3.71	<0.01	<0.01	152	4,917	1.79	<0.01	<0.01	1,017	2,706
8	3.02	<0.01	<0.01	121	3,112	1.37	<0.01	<0.01	594	1,532

Notes: This table is based on the following regression: $y_{pf} = \lambda_a + \alpha_p + \varepsilon_{pf}$, where y_{pf} is an indicator variable that equals 1 if p has above median returns in fund f. λ_a are fixed effects based on vigintiles of fund f's age. α_p are fixed effects for LPs (p). The table shows the F-statistic, the p-value, and a nonparametric p-value of the null hypothesis that the α_p jointly equal zero. To generate the nonparametric p-value (p^*), we randomly assign return paths within each fund, compute y, run the regression, and retain the F-statistic. We do so 1,000 times then generate p^* by comparing the actual F-statistic to the simulated distribution of F-statistics. We determine whether p has above median returns in fund f based on whether it is above median on average over the life of the fund. Returns are measured using TVPI. We report the same analysis using DVPI to measure returns in the internet appendix.

Table 6: Investor Characteristics and the Likelihood of Within-Fund Outperformance

Dependent Variable:	$100 \times 1(p$ Outperforms in $f)$		
	(1)	(2)	(3)
Percent of Fund	0.69**	0.45**	0.32
	(3.64)	(2.18)	(1.30)
Large Pension (AUM)	17.39**	19.35**	20.11**
	(6.20)	(6.03)	(5.30)
LP-GP Fund Count	0.71**	0.63*	0.75
	(2.25)	(1.69)	(1.42)
Quartile of LP's Prior Funds	-10.36**	-10.83**	-15.79**
	(-2.59)	(-2.36)	(-2.64)
Early PE Investor	9.83**	9.78**	7.77**
	(4.06)	(3.52)	(2.34)
Elected Board Members (%)	0.13**	0.19**	0.27**
	(3.32)	(4.08)	(4.46)
Board Size	0.42*	0.16	-0.20
	(1.91)	(0.62)	(-0.62)
R^2	0.10	0.10	0.11
Fund Age Min. (yrs)	1	4	8
N	5,382	3,930	2,220

Notes: This table reports OLS estimates from a linear probability model of the likelihood that an investor p is a top-tier investor in fund f. The dependent variable is an indicator for whether p has above-median returns (TVPI) in fund f for the majority of its life. Percent of fund is based on the$p$'s commitment relative to the fund's size. Large Pension (AUM) is an indicator if p has AUM over \$100 billion. LP-GP Fund Count is the number of funds between p and the manager of fund f, measured over our full sample. The variable Quartile of LP's Prior Funds measures the average quartile ranking p's funds that were active at the time of fund f's close. Quartile rankings come from Preqin and reflect performance as of 2020Q1. Higher quartiles correspond to better performance. We define an indicator variable for whether p was an early investor in private markets if it's first entry into the dataset is before 2008. Elected Board Members (%) is the percent of p's board that is elected by members or the general public, measured at the time of f's close. Board size equals the number of board members at the same point in time. All regressions include a fund fixed effect and fixed effects based on fund age vigintiles. Standard errors are clustered within each investor-vintage cell.

Figure 1: Example Cashflow Profiles of Investors in the Same Fund

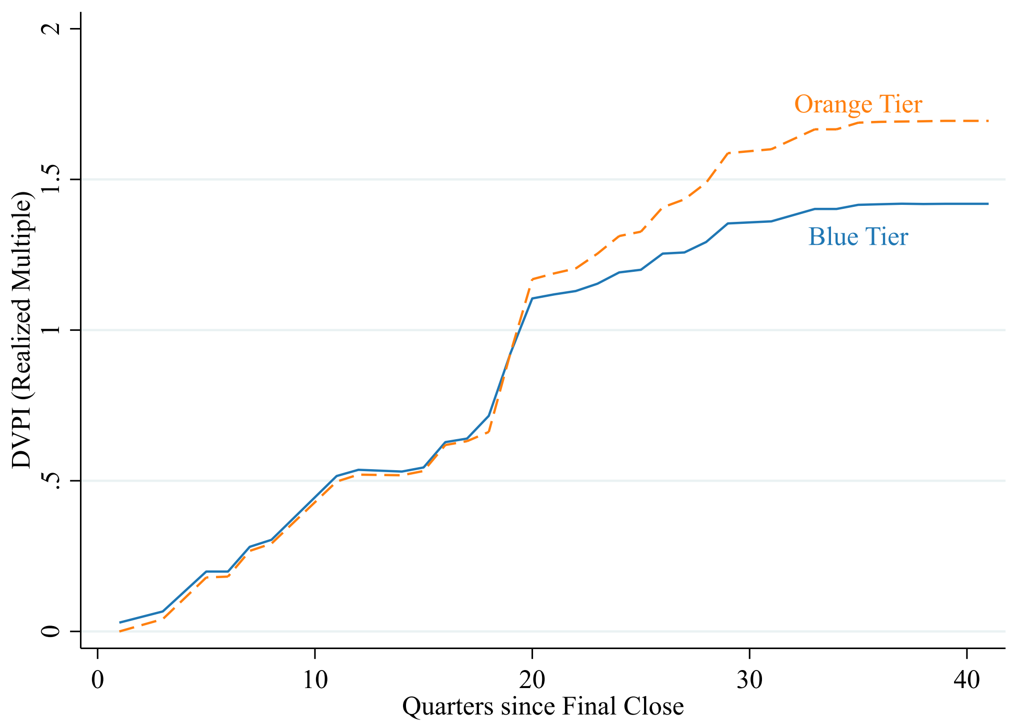

Notes: This plot shows the time-series evolution of DVPI for two anonymous investors in the same fund. DVPI is defined as the cumulative amount of distributions, scaled by the cumulative amount of contributions. We are not able to identify individual funds or investors per our data-sharing agreement with Preqin.

Figure 2: Within-Fund Dispersion in Net-of-fee Returns

Panel A: DVPI (Realized Multiple)

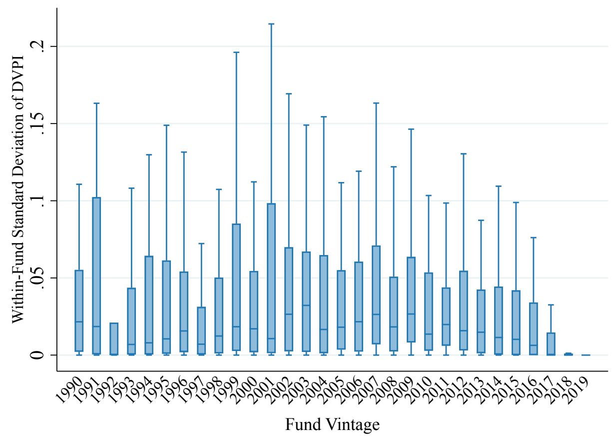

Panel B: TVPI (Total Multiple)

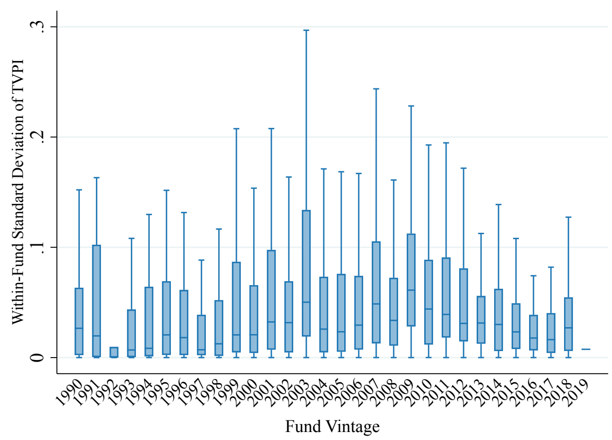

Panel C: Reported IRR

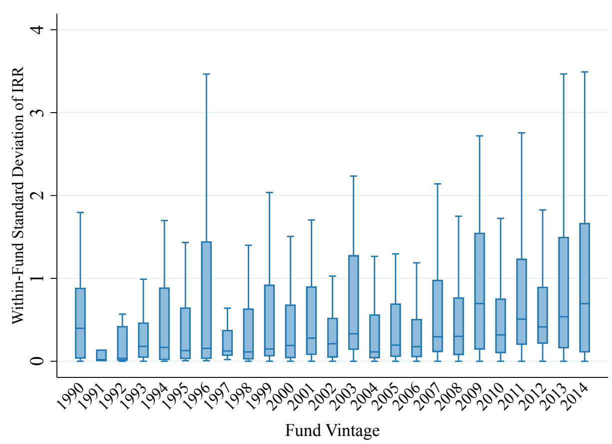

Notes: This plot shows the distribution of within-fund return dispersion across funds. Dispersion is defined as the within-fund standard deviation of returns, σ_f. The boxplot summarize the distribution of σ_f across funds, broken down by the vintage year of f. See Section 2.1 for more details on our sample construction. Panel A and B of the figure show σ_f when measuring returns using DVPI and TVPI, respectively. DVPI is defined as the cumulative amount of distributions, scaled by the cumulative amount of contributions. TVPI is defined as the cumulative amount of distributions plus any remaining net-asset-value, scaled by the cumulative amount of contributions. Panel C uses the IRRs reported in Preqin, which are expressed in percentage points. For IRRs, we exclude vintage years after 2015 since IRRs are typically unstable early in a fund's life-cycle.

Figure 3: Example of Clustering in Returns, Call Rates, and Distribution Rates

Panel A: Clustering in Net-of-Fee Returns (DVPI)

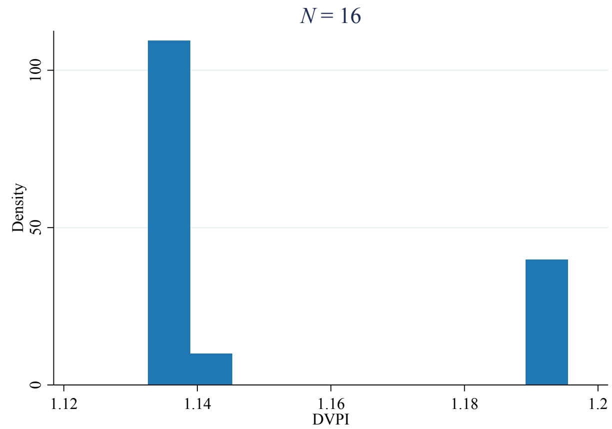

Panel B: Capital Call Rates

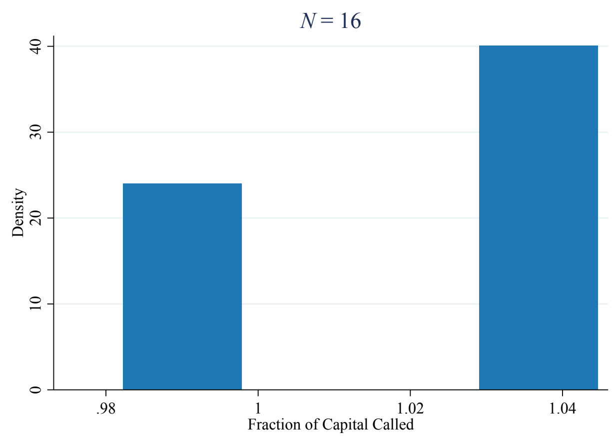

Panel C: Distribution Rates

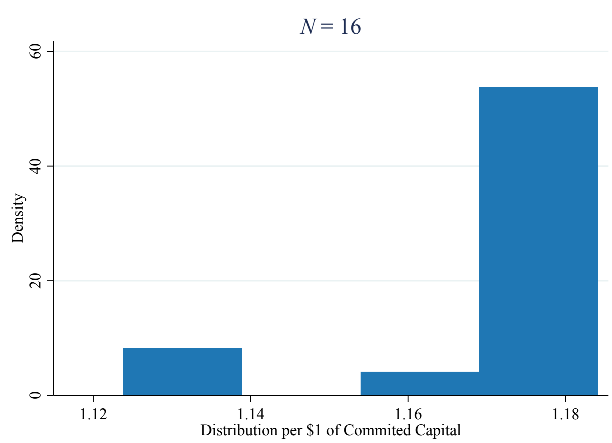

Notes: Panel A shows the distribution of net-of-fee returns as measured by DVPI for 16 investors in the same fund at a fixed point in time. Panel B shows capital call rates (contributions divided by committed capital) for the same investors, in the same fund at the same time. Similarly, Panel C shows distribution rates (distributions per dollar of committed capital) for these investors in the same fund at the same time. We are not able to identify individual funds or investors per our data-sharing agreement with Preqin.

Figure 4: Number of Clusters in Returns, Call Rates, and Distribution Rates

Panel A: Within Fund DVPI Clusters across funds

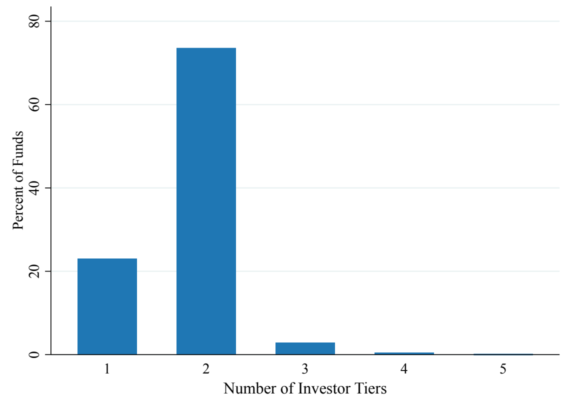

Panel B: Clusters of Fees Included in Contributions

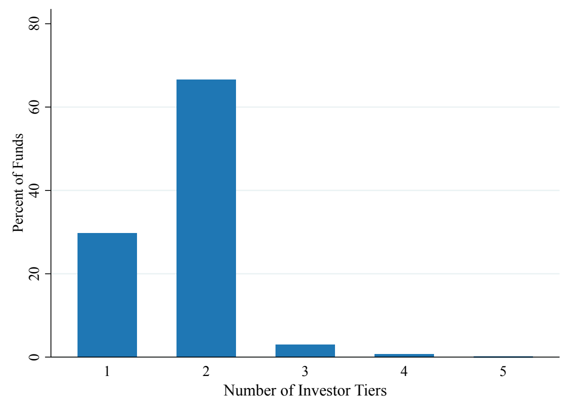

Panel C: Clusters of Fees Deducted from Distributions

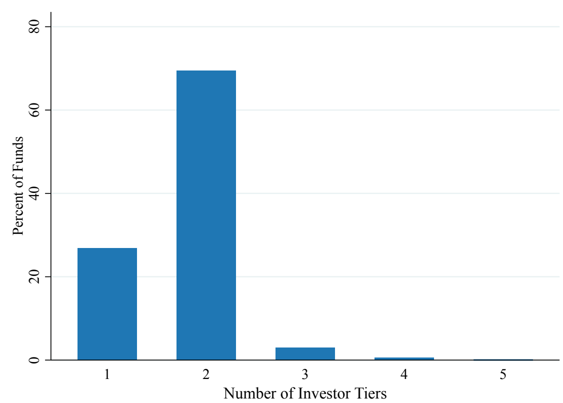

Notes: **Panel A** shows the number of within-fund clusters (investor tiers) in net-of-fee returns (DVPI) across funds. For each fund f and date t, we compute the number of clusters in DVPI using a k-means clustering analysis, where the number of clusters is chosen based on Silhouette scores. The number of clusters at the fund level is defined as the average number of clusters over each fund's life, rounded to the nearest integer. **Panel B** repeats the cluster analysis for capital call rates, cumulative contributions per dollar of committed capital. Contributions by LPs into a fund include organizational expenses, management fees, and capital that is ultimately invested by the GP. **Panel C** repeats the cluster analysis for distribution rates, defined as cumulative distributions per dollar of committed capital. Distributions out of funds are net of fees, including carry.

Figure 5: Estimating Within-Fund Fee Dispersion

Panel A: Dispersion in Fixed Fees (e.g,. management fees)

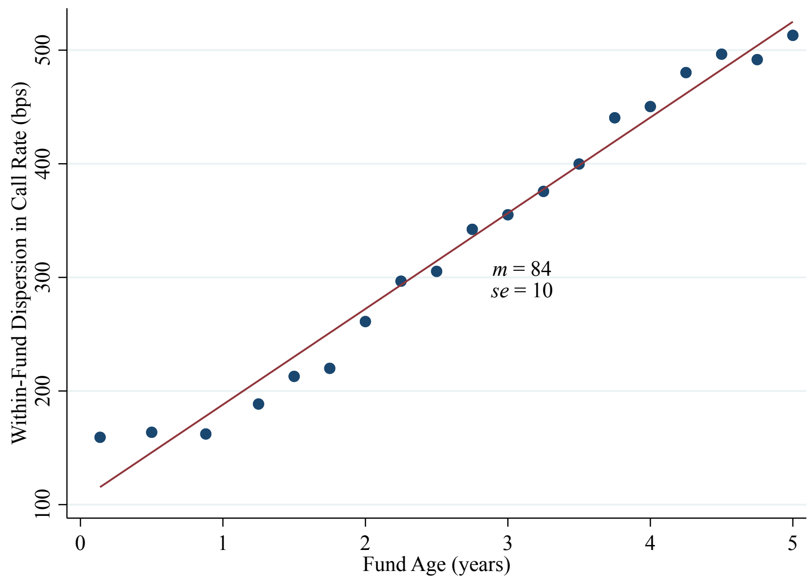

Panel B: Dispersion in Performance Fees (e.g., carry)

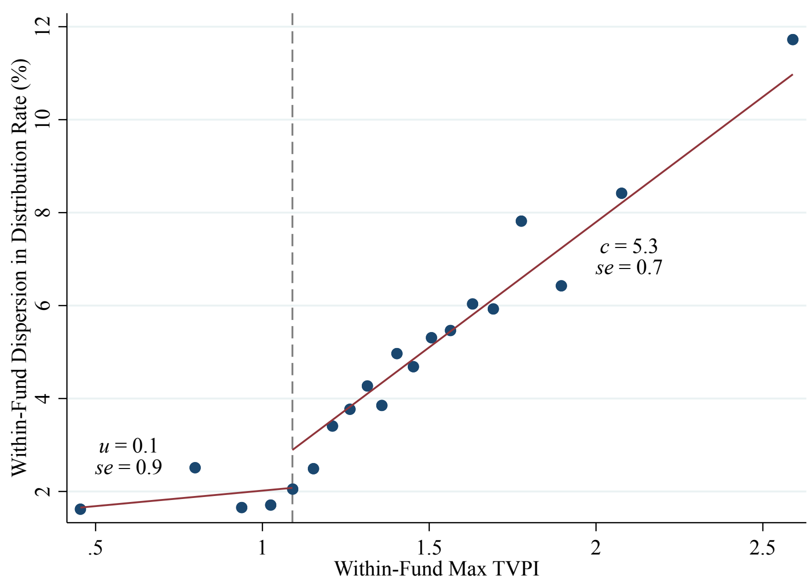

Notes: This plot depicts how we estimate dispersion in fixed fees (e.g., management, m) and performance-based fees (e.g., carry, c) within the average fund. **Panel A** shows a binned scatter plot of the within-fund volatility of call rates for fund f at time t against f's age at the same date. The call rate equals the fraction of committed capital that has been called for investment. The plot is made using all data on funds whose age is less than five years. **Panel B** is a binned scatter plot where the y-axis is the within-fund volatility of distribution rates for f at time t. Distribution rate is the percent of distributions relative to commitment amount. The x-axis of the plot is f's maximum TVPI at time t. Both variables are shown after partialing out vintage fixed effects. Funds must have at least 5 observed investors at t to be included. The vertical dotted line in Panel B marks the boundary of funds with a TVPI of 1.09, a proxy for those that are profitable enough to charge carry. On either side of the boundary, we report the slope of the line of best fit and its standard error, which is clustered by fund and GP.

Figure 6: Within-Fund Outperformance Rates by LP

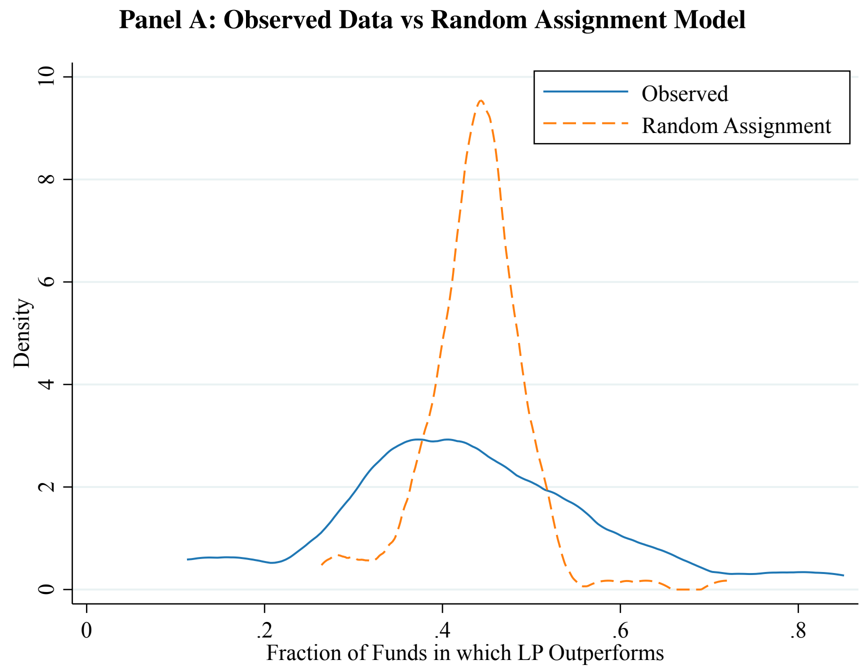

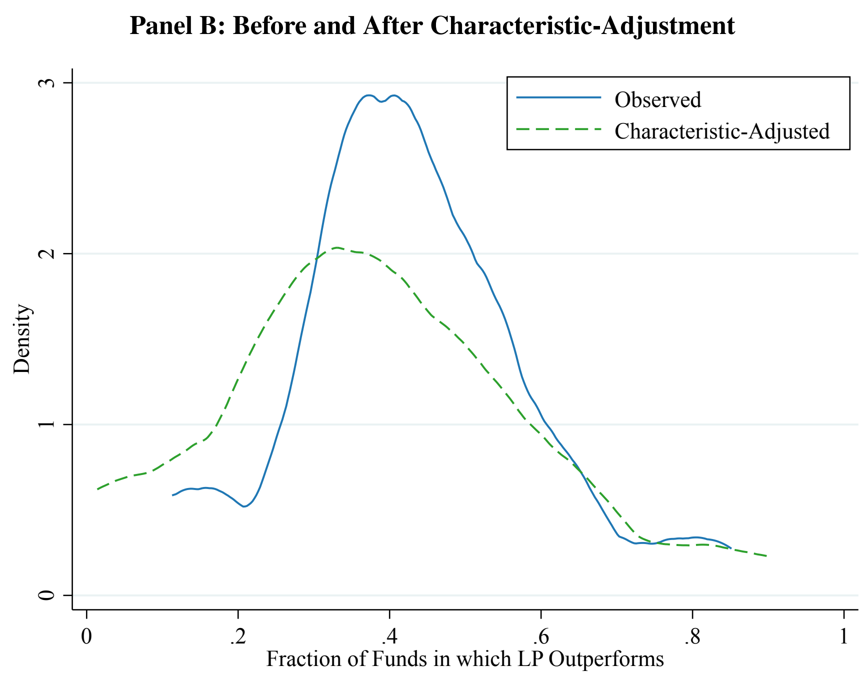

Notes: This plot shows the fraction of funds in which a pension outperforms other investors ("pension effects"). The blue line in Panel A is created by regressing an indicator for whether pension p earns above-median returns in fund f, y_{pf}, on fixed effects for pension and age vigintile. The estimated pension fixed effects are then shrunk towards their mean using an empirical Bayes estimate and shown in blue. The orange line shows simulated pension effects based on the random assignment of contracts to pensions in each fund (see Section 5.1.1). In Panel B, we evaluate how observable characteristics account for the observed pension effects. To do so, we regress returns r_{pft} of pension p in fund f at time t on a vector of characteristics and fixed effects for fund-date. We use the residuals from the regression to determine characteristic-adjusted status in each fund $\tilde{y}_{pf}$, re-estimate pension effects, and apply the empirical Bayes procedure. The resulting characteristic-adjusted pension effects are plotted in green, alongside the pension effects before adjustments. Returns are measured using TVPI. See Sections 5.1.1 and 5.2.2 for more details.