How do private equity fees vary across public pensions?

Internet Appendix

Juliane Begenau* Emil Siriwardane†

Abstract

In this appendix, we investigate other sources of within-fund return variation and discuss how they would impact our main conclusions. We then provide more detail about the contracting environment in private-market funds. We also present several supplementary results, including: (i) robustness tests using other unsupervised learning techniques to identify return clusters; (ii) whether net-of-fee returns clusters within funds are present using alternative measures of returns; (iii) an analysis of how commitment size determines assignment to management-fee tiers; (iv) tests showing that pension effects are predictable out-of-sample and exist with alternative measures of returns; and (v) an analysis of whether characteristic-adjusted pension effects are sensitive to how returns are measured.

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

†Siriwardane: Harvard Business School and NBER. E-mail: esiriwardane@hbs.edu.

A Potential Sources of Within-Fund Return Variation

In this section, we provide a supplemental analysis of different sources of within-fund return dispersion. Section A.6 reproduces our main analyses in a subsample of funds where we think these alternative sources are less likely to occur.

A.1 Measurement Error

A.1.1 Audit Results

As discussed in Section 3.1, measurement error could arise when Preqin tabulates the data provided to them by pensions. To gauge the size of these potential errors, we directly filed FOIA requests with a subset of the pensions in the Preqin data. We chose the pensions for the audit based on one of two criteria. First, we ranked pensions based on their forgone capital due to fee dispersion and chose the top and bottom 15 pensions according to this metric.[1] Second, we selected the funds with the widest range of within-fund variation in returns, and then chose the pensions with the largest and smallest returns in these funds.[2] Both selection methods are designed to ensure that we examine the pensions whose reported returns deviate the most from others in their respective funds. In total, we sent FOIA requests to 65 pensions and received responses from 48 of them.

We use the following measure to quantify differences between our audited dataset and Preqin for a given variable v:

$$D_{if}^{v} = \frac{\left| v_{if}^{P} - v_{if}^{A} \right|}{v_{if}^{P}}$$

where v_{if}^{P} is the value of variable v for pension i in fund f and v_{if}^{A} is the value from our direct FOIA, both rounded to three decimal places. When $v_{if}^{P} = v_{if}^{A}$, we set $D_{if}^{v} = 0$ and when $v_{if}^{P} = 0$ we define D_{if}^{v} by scaling absolute deviations between the two datasets by v_{if}^{A}.

[1] For each pension in this set, we selected the quarter q between 2010 and 2018 with the maximum number of observations in the Preqin data. We then filed a FOIA for the pension's data in quarter q. One of the pensions in this list erroneously responded with flows for their chosen quarter, as opposed to cumulative distributions and contributions. This pension did not respond to our second request for data, so we have dropped them from this analysis.

[2] For each pension in this set, we filed a FOIA for data in the quarter that is used in the core sample.

Table IA1 summarizes the results of our audit. The first three rows of the table show that the vast majority (94-98%) of the data in Preqin *perfectly* matches the audited data. The remaining rows break out the distribution of D^v_{if} for the limited number of entries where the two datasets do not perfectly align (i.e., $v^P_{if} \neq v^A_{if}$). For most of these entries, the audited data and the Preqin data are still close to each other. For example, there are only 22 observations where TVPI – our main measure of returns – is not identical in the audited and Preqin data. However, three-quarters of these imperfect matches are still within 6% of each other. In addition, we have manually confirmed that most of the large outliers where D^v_{if} is near 100% are concentrated in just a handful of investor-fund observations. Thus, taken together, these audit results suggest that Preqin is accurately reports the data provided to them by pensions. This does not rule out the possibility that pensions are the source of measurement error. However, several patterns in the data (e.g., return clustering) suggest that truly random errors are not the primary source of within-fund return variation.

A.2 Multiple Commitment Rounds

As noted in the main text, performance may mechanically vary across investors in the same fund who commit at different rounds. This form of dispersion results from differences in the timing of contributions and is distinct from dispersion in management fees. We now present several pieces of evidence that suggest our main conclusions are robust to this potential source of dispersion.

A.2.1 Contribution Tiers and Fund Age

In practice, there is an institutional feature that would equalize or "true up" investors who commit at different rounds. These equalization methods generally require investors who join at later rounds to contribute such that all investors appear as if they had joined at the same time. This process also ensures that future distributions are dispersed as if all investors had joined at the same time. Thus, in principle, equalization makes it more plausible that investors will indeed earn equal gross distributions in a fund. In practice, small and permanent differences in contributions may arise because late-closing investors often make a one-time interest payment to investors who committed earlier (e.g., LIBOR plus a spread).

Because equalization takes time and is imperfect, some of the dispersion in the data may be driven by differential commitment times and not fees. This channel should be particularly strong early in a fund's life. We explore this possibility as follows. First, we calculate the number of clusters N_{ft} for contribution rates in fund f on each date t. We then take the average of N_{ft} over all $t \geq q$ and define its rounded value $N_f(q)$ as the number of contribution clusters in f. It is useful to note that we set $q = 4$ as our baseline explicitly to mitigate the impact of timing-induced dispersion.

We can then use variation in $N_f(q)$ to look for evidence that dispersion is driven by staggered commitments and the subsequent equalization process. Specifically, we compute the fraction of funds for whom $N_f(u) = N_f(q = 4)$. Figure IA2 plots this fraction for different $u \geq 4$. If the staggered commitments are the dominant source of contribution clusters, then this fraction should decline rapidly as funds age and the equalization process runs its course. In contrast, the figure shows that the number of contribution tiers is very stable for most funds. For example, for 93% of funds, the number of contribution tiers is identical if we exclude the first four quarters versus the first twenty quarters of data.

The imperfect nature of the equalization mechanism should also generate a distinct pattern in contribution rates. Namely, after equalization payments are made, any dispersion in contribution rates should remain constant as the fund ages thereafter. However, Panel A of Figure 6 shows a strong and positive linear relationship between within-fund dispersion in contribution rates and age. As we argue in Section 3.2, this pattern is exactly what one would expect if management fees differed across investors. Indeed, the slope of the line provides an estimate of the average within-fund dispersion in management fees.

Overall, the relationship between tiering rates, contribution dispersion, and age suggests that commitment timing is not a major issue in our data. Moreover, to the extent that equalization gives investors equal claims to gross distributions, commitment timing cannot account for the large dispersion in distributions that we observe empirically. Instead, the observed patterns of dispersion strike us as more consistent with some form of fee dispersion. Next, we develop this argument further using data on fundraising rounds.

A.2.2 Single versus Multi-Round Funds

We also obtained an indicator variable from Preqin's research team that equals one if a fund had multiple fundraising rounds and zero otherwise.[3] Single-round funds compromise 63% of our sample. Even though timing-induced is not possible in these funds, 57% have multiple DVPI tiers. This simple fact cuts against the idea that differences in entry dates due to multiple funding rounds confound the interpretation of our results. Later, in Section A.6, we also reproduce our main analyses in the sample of single-round funds that also are unlikely to have investor-specific mandates. Return dispersion in these funds is more likely attributable to fees and we find that all of our main conclusions hold in this sample.

A.3 Secondary Market Transactions

In principle, secondary sales or purchases could also generate performance dispersion. To see why, consider the following stylized example of a three-period fund that has two investors (A and B), no fees of any kind, and grows at a rate of 20% per year. Each investor initially contributes \$50 in period 1, meaning the value of the fund is \$100 in period 1, \$120 in period 2, and \$144 in period 3. Further suppose that investor A sells its stake for 90% of its NAV in period 2 to a third investor (C). The cash flows and returns (TVPI) to each investor evolve as follows:

	Investor A			Investor B			Investor C		
	CF	NAV	TVPI	CF	NAV	TVPI	CF	NAV	TVPI
$t = 1$	-50	50	1	-50	50	1	-	-	-
$t = 2$	54	0	1.08	0	60	1.2	-54	60	1.11
$t = 3$	-	-	-	72	0	1.44	72	0	1.33

[3] The fund-details data from WRDS contains fundraising launch-date, fundraising status, last-close-date, final-close-date, and first-close-date. However, because WRDS only provides a recent snapshot of the funding round status for funds, we can not easily infer the number of rounds for each fund. This prompted us to reach out to Preqin directly to inquire about the time-series of fundraising status for each fund and they kindly provided us with the data.

In this example, TVPI clearly differs across all investors at period 2 and again at period 3, yet by construction there is no dispersion in fees.

While secondary sales and purchases are certainly possible, we have several reasons to think that these types of transactions occur infrequently in our data. First, it seems unlikely that direct secondary purchases are common for U.S. public pensions. Moreover, most secondary purchases occur through dedicated secondary funds as opposed to direct purchases (Nadauld et al., 2019), and we exclude these funds from our sample. As one example, we analyzed the composition of the historical private equity portfolio of CalSTRS, the second largest U.S. pension.[4] As of September 2020, only 2.6% of CalSTRS' historical capital commitments in private equity went to secondary buyouts. Given that the largest U.S. pensions are relatively inactive in the direct secondary market, it seems reasonable to us to conclude that smaller pensions are as well.

We also attempt to identify potential secondary transactions directly in the data. Our approach is based on the idea that an investor who sells their stake in the secondary market should presumably report cash flows to Preqin at the beginning of the fund's life and then abruptly stop after the sale. We operationalize this logic by searching for instances where an investor stopped reporting for at least eight quarters prior to the last observation in a fund. For example, if a fund's last observation was in 2018Q1, we flag investors whose last observation was prior to 2016Q1 as those who potentially sold in the secondary market. We further require that these investors have data available for 90% of the potential quarters up until their last observation (i.e., 90% of quarters since the fund's final close). This screen is designed to identify LPs who were regularly reporting to Preqin prior to a potential sale event.[5] These cases are quite rare, as only 5% of investor-fund pairs are flagged as potential secondary sales. We view this as an upper bound on the prevalence of secondary sales in our data, given that an LP who simply stops reporting data on a fund to Preqin will appear as if they sold their share in the secondary market.

[4] https://resources.calstrs.com/publicdocs/Page/CommonPage.aspx?PageName=DocumentDownload&Id=400737b3-d226-462e-ba22-5a1e5cd470ca

[5] We also exclude funds whose vintage is prior to 2000 and those with less than 8 quarters of total data. These screens ensure that we do not confuse potential secondary sales with sporadic reporting.

In a secondary sale, a natural set of buyers are the existing LPs in the fund. Thus, another way to identify potential secondary activity is to check whether commitment size changes over time for a given investor in a fund. This approach should also identify instances where an LP sells only a partial stake in a fund. We find that commit size changes in only 8% of investor-fund pairs. This is again likely an upper bound on secondary market activity, since commitment size could change for many reasons other than a secondary sale (e.g., increasing an initial commitment).

In sum, secondary transactions could in principle generate return dispersion, though there is also good reason to believe that public pensions do not *directly* participate in this market very often.[6] We discuss how secondary transactions would impact our main conclusions later in Section A.7.

A.4 Recallable Capital

We discuss in Section 5.3 how the accounting of recallable capital could generate within-fund variation in returns. In this subsection, we use data from a single large LP to develop a sense of how much recallable capital accounting could explain the observed data. For this particular LP on date t, we were able to obtain cash flow data that includes recallable capital in both contributions and distributions (i.e., following GIPS standards), as well as data that nets recallable capital from both variables. We then compare the observed dispersion within fund f on date t against the dispersion that would arise if LPs in fund f use both accounting approaches when reporting returns to Preqin.

To start, denote the TVPI in fund f that includes recallable capital as r_f^G and the TVPI that nets recallable capital r_f^{NG}. The G and NG in the superscripts stand for GIPS and No-GIPS, respectively. We suppress the time stamp t to keep the notation simple. Next, assume that all dispersion in our data is due to the difference between these two variables. Under this assumption, we would expect the dispersion in fund f to be $\sigma_f^{RC} \equiv |r_f^G - r_f^{NG}|$.

Figure IA1 plots the observed range of returns in fund f, denoted $\hat{\sigma}_f$, against σ_f^{RC}. Under the null hypothesis that recallable capital accounting fully drives within-fund variation in TVPI, we

[6]Note that we exclude PE funds that specialize in secondary markets from all of our analysis.

would expect all of the points in the plot to lie on the 45-degree line, which appears in orange. The plot shows clearly that we can easily reject this null hypothesis, as a regression of $\hat{\sigma}_f$ on σ_f^{RC} delivers an intercept of 0.08 ($t = 6.32$) and a slope of 0.71 ($t = 8.94$).

About 20% of funds satisfy the condition that $|\hat{\sigma}_f - \sigma_f^{RC}| < 0.01$. We view this as a reasonable upper bound on the fraction of funds whose observed dispersion data could be plausibly explained by differences in recallable capital accounting. The ability to generalize this analysis depends on whether this particular sample of funds suffers from selection bias, though we see no obvious reason why this would be the case. Thus, while accounting biases relating to recallable capital (or other LP-specific adjustments) are surely present in the data, this analysis supports our argument in Section 5.3 that any such biases are likely to be a small source of within-fund variation in returns. In addition, accounting biases cannot explain our finding that some funds or GPs are more likely to have return clusters – hence our conclusion that fees vary within many funds – and should not bias our estimation of contract differences in Section 3.2.

A.5 Dispersion Excluding Large LPs and Prior to 2010

Differences in gross-returns across LPs could also lead to net-of-fee return dispersion within a fund. Gross returns could differ due to the commingling of co-investment or other special purpose vehicles with main fund returns, though this is unlikely to be a large issue given that pensions do not appear to commingle these vehicles when reporting to Preqin (Section 5.4). Gross returns could also differ due to investor-specific mandates like ESG restrictions. One way to assess the potential bias introduced by special purpose vehicles or LP-specific mandates is to measure within-fund return dispersion in smaller LPs prior to 2010. The logic of this analysis is that smaller LPs are less likely to invest in special purpose vehicles or have the bargaining power to impose mandates on GPs. In addition, it is our understanding that these types of custom exposures were far less common prior to 2010. Panel A of Figure IA3 shows the cross-sectional average within-fund standard deviation in DVPI for vintages through 2010. To construct the plot, we compute the standard deviation of DVPI, σ_f, within each fund in the core sample, excluding all LPs whose

AUM is ever over \$100 billion for our sample period. We then take the average of σ_f in each fund vintage and plot it in the figure. Panel B repeats the calculation using TVPI. In both cases, the level of dispersion excluding large LPs is comparable, albeit slightly lower than in the full sample. This supports the notion that special purpose vehicles and other LP-specific mandates are not the primary source of return variation within funds.

A.6 Single-round funds that are unlikely to have investor-specific mandates

For robustness, we now probe how any timing- or mandate-induced dispersion would bias the interpretation of our results. To handle the latter type of dispersion, we create an indicator variable for whether a fund is likely to have investor-specific mandates using its industry focus. Specifically, we identify multi-industry funds that also invest in the oil and gas industry. LP-specific mandates are arguably more likely in these funds since LPs may want to restrict their exposure to oil and gas portfolio companies while still retaining exposure to the other industries in which the fund invests. Based on conversations with Preqin, this is a more precise way to identify funds with potential LP-specific mandates compared to using the ESG variable that is available on WRDS ("firm ethos"). This is because the ESG variable on WRDS is based on whether a GP has made any public statements regarding their view on ESG – it may not reflect any actual investment behavior. Moreover, the ESG variable from WRDS is based on recent public statements. In contrast, our classification based on industry focus reflects the structure of the fund at the time of actual investment.

We then construct a sample of single-round funds that do not invest in oil and gas. Henceforth, we refer to this set of funds as green single-round funds. Because timing- and mandate-induced dispersion should be minimal in this subset, we use it to probe the robustness the following sets of results: (i) estimates of how much management and carry vary in the typical fund; (ii) characteristics of funds that tier; (iii) which investors are more likely to be top-tier; and (iv) how much money investors have left on the table.

In Section 3.2, we use dispersion in call and distribution rates to estimate within-fund dispersion in management and carry.Table IA2 reports the estimates for green single-round funds. The estimates from the this sample are generally comparable to our full-sample estimates for both types of fees, though average carry dispersion is somewhat lower for green real estate funds that were raised in a single round. In all cases, the standard errors are such that we cannot reject the null that the estimates are different across the two samples.

In Section 3.3, we analyze whether some funds or GPs are more likely than others to tier their investors. Our outcome variable is an indicator variable for whether a fund has tiers in call rates, distribution rates, or DVPI. Columns (1)-(5) of Table IA3 present a series of regressions of this indicator variable on GP characteristics using the full sample. This corresponds to Table 3 of the paper. Columns (6)-(10) repeat the regressions using only the sample of green single-round funds. The main thing to notice is that the regression coefficients are comparable across the two samples, as is the overall model fit. In other words, our characterization of which funds tier does not appear to be materially biased by multiple fundraising rounds or investor-specific mandates.

Section 4.1 uses fixed-effects regressions to test the null hypothesis that pensions are randomly assigned to fee tiers in their funds. This null is strongly rejected in the full sample and we interpret this result as evidence of pension effects – some pensions are consistently top-tier investors in all of their funds. An alternative interpretation is that some pensions consistently commit in earlier or later rounds when investing in private funds. Absent fee dispersion, these pensions may then consistently (and mechanically) under or over perform due to timing-related issues. A similar story can be told regarding LP-specific mandates. However, we continue to find pension effects ($p < 0.01$) in green single-round funds, suggesting that these alternative interpretations are unlikely.

In Section 4.2, we characterize the types of LPs that are more likely to be top-tier. One concern with this analysis is that we are simply picking up on the characteristics that predict which LPs are likely to commit later or earlier in a fund. Columns (1) and (2) of Table IA4 show that this is unlikely to be the case. In column (1), we regress an indicator variable for top-tier status in fund f on a host of LP characteristics. In column (2), we repeat the regression in the subsample of green single-round funds. In both cases, large pensions who commit more to the fund and have relationships with the GP are more likely to be top-tier. The coefficient on the track record of the LP is also comparable in green single-round funds, though it is estimated with less precision.[7] Hence, it does not appear that timing- or mandate-induced dispersion biases our conclusions regarding the types of LPs who are likely to be top-tier.

Finally, Figure IA4 shows forgone capital curves for clean single-round funds versus all other funds. If anything, it appears that forgone capital is higher in clean single-round funds. This is the opposite that one would expect if timing- or mandate-induced dispersion swamped dispersion induced by fees.

To summarize, several patterns in the data suggest that neither staggered commitment dates nor LP-specific mandates confound our analysis. Moreover, our main conclusions are robust if we focus on a subset of funds for which this is not an issue.

A.7 How Would Other Sources of Return Dispersion Impact Our Results?

A concern with the preceding argument is that fees may not be the only source of return dispersion even within clean single-round funds. For example, differences in LP-accounting practices could still generate return dispersion in these funds. Studying non-oil funds also does not guarantee that gross returns are equal across investors, since LPs could have other mandates that prevent them from investing in specific portfolio companies that are outside of the oil and gas industry. Because we cannot perfectly eliminate or rule out all non-fee sources of return dispersion, we now discuss how their presence would impact our main conclusions.

A.7.1 Estimates of Dispersion in Management Fees and Carry

In Section 3.2 we use the full time-series of commitments and distributions to estimate the size of within-fund dispersion in both management fees and carry. The logic of our approach is as follows: dispersion in management fees should generate a distinct linear relationship between within-fund

[7] LPs who were early investors in PE do not appear as likely to be a top-tier investor in green single-round funds, though this is likely an issue with power since these funds tend to be older and our definition of an early PE investor is based on if an LP appears in the data before 2006.

dispersion in contributions (per dollar of commitment) and age. Similarly, dispersion in carry rates should generate a call-option like relationship between dispersion in distributions (per dollar of commitment) and fund performance. Intuitively, dispersion in carry should not be detectable in unprofitable funds.[8] These patterns bear out clearly in Figure 6, suggesting that dispersion in management and carry in the average fund are 91 bps and 5.8%, respectively.

As we discuss in Section 5.1, measurement error could in principle generate the linear relationship between dispersion in contributions and age. To see why, suppose that in each period LPs input contributions with random error. In this case, dispersion in *cumulative* contributions will grow linearly with age and inflate our estimate of management fee dispersion. This explanation is, however, inconsistent with the presence of clusters in contributions (see Internet Appendix C.2). Moreover, measurement error would not generate the call-option like relationship between dispersion in distributions and performance.

Similarly, it seems implausible that variation in LP-specific accounting practices or differences in gross return exposure could generate the patterns seen in Figure 6. Secondary transactions should also not bias our estimates of management fee and carry dispersion. To see why, consider a fund with three investors and no differences in carry. Further suppose that one investor sells after T periods. Prior to the sale, dispersion in distribution rates should be insensitive to performance since there are no differences in carry. In the period of the sale, distribution rates may differ across the investors. Nonetheless, after the sale, there should be no dispersion in distribution rates for the remaining investors. Thus, our estimation approach would correctly conclude dispersion in distribution rates is not sensitive to fund performance (i.e. no carry differences), even though there was a one-time difference in returns. A similar logic applies to our estimates of management fee dispersion.

[8] See Section 3.2 for a more formal argument.

A.7.2 Characterization of the Funds that Use Multiple Fee Schedules

In Section 3.3, we characterize the funds that are likely employ multiple fee schedules. Our analysis revolves around regressions of a fund-level indicator for multiple return tiers on fund characteristics. The indicator variable is based on the return clusters that are estimated in Section 3.1. The coefficient estimates from these regressions are not biased by the presence of measurement error. This is because the response variable will inherit any measurement error in returns, which does not bias regression estimates.

It is also unlikely to us that LP-specific accounting, gross-return differences, or secondary market activity would change our broad interpretation. For example, we find that VC funds are far less likely to have multiple-return tiers compared to buyout funds. Moreover, we find that law-firm fixed effects drive heterogeneity in the likelihood that a fund has multiple return clusters, as does the use of placement agents and subscription status. For variation in LP accounting practices to generate these results, it would have to be that such practices vary less across investors in VC funds, funds that use a particular law firm, or undersubscribed funds. The same logic applies when considering if gross-return differences or secondary market activity can explain these findings. On the other hand, it seems more plausible that fee practices (e.g., side letter usage) would depend on fund type or the identity of law firms.

A.7.3 Pension Effects

Section 4.1 contains evidence rejecting the null hypothesis that pensions are randomly assigned returns in each fund (so-called pension effects). This immediately suggests that measurement error cannot fully account for within-fund return variation. It is also consistent with the idea that some LPs consistently pay lower fees relative to others in their funds. However the interpretation would change if return variation is driven by heterogeneity in LP accounting practices. To illustrate, suppose there is a subset of LPs that do not follow GIPS standards in their accounting of recycled capital. Even with uniform fees, DVPI (and TVPI) will differ within funds, and regression tests will detect pension effects. Still, as discussed in the main text and in Section C.6.2, pension effects are still present when returns are measured using IRRs, yet IRRs do not depend on the accounting of recallable capital. Moreover, in Begenau et al. (2020), we asked a subset of LPs in our sample (via FOIA) whether they followed GIPS accounting standards and 100% of respondents answered yes.

Gross-return differences could also generate pension effects. For instance, some LPs may impose investment restrictions in all of their funds (e.g., ESG) and thus earn consistently different net-of-fee returns. Though possible, this alternative is unlikely because investor-specific ESG mandates were not popular for most our sample period (Section 5.4). As mentioned above, we also find pension effects in green single-round funds, ($p < 0.01$), a set of funds where ESG mandates are less likely.

By the same token, pension effects would arise if some LPs consistently transact in the secondary market. Again, for reasons discussed in Section A.3, this is unlikely because public pensions were not active in the secondary market during our sample.

A.7.4 Characterization of Top-Tier LPs

In Section 4.2, we characterize the types of LPs that are more likely to be in the top-tier of a given fund. Our main findings are that large and experienced LPs with high past performance are more likely to outperform other LPs in their funds. As with pension effects, it is immediately clear that pure measurement error cannot generate these results. LP-specific accounting practices could do so if, for example, large and experienced pensions use different accounting practices than small and inexperienced ones. This is certainly possible, though again the FOIA results in Begenau et al. (2020) suggest such heterogeneity is small.

Analogously, for gross-return differences to explain the results, it would have to be that large or experienced LPs consistently restrict investments such that they earn higher gross returns in their funds. One way that this could occur is through co-investments; however, we exclude co-investment vehicles in our analysis and also find that co-investments are a relatively small portion of public pensions' current PE portfolio (see Section 5.4.1). Secondaries could also explain the results in Section 4.2 if some investor types are more likely to transact in the secondary market. Though we cannot definitively rule out this possibility, we think it is unlikely given the overall low rates of secondary activity discussed above.

A.7.5 Forgone Capital and the Bias in Aggregate Performance Metrics

Section 6 presents two supplementary pieces of analysis. First, we compute the amount of forgone capital due to fee dispersion. Second, we calculate aggregate performance using different LPs in each fund, which provides a gauge for the bias that fees induce on aggregate performance data. Importantly, both sets of analyses are potentially biased by non-fee sources of return dispersion. This is why we view both as more suggestive in nature and have relegated them to the end of the paper.

To summarize, measurement error and other sources of within-fund return dispersion should not materially bias our characterization of the tiering behavior of funds and LPs, nor should it influence our measurement of within-fund dispersion in management and carry. The existence of pension effects could in principle be driven by some of these other sources, though there are good reasons to think fees are their primary source. The same holds true for our characterization of the types of LPs who are likely to be top-tier in terms of fees. The measurement of forgone capital and aggregate performance bias in Section 6 is indeed sensitive to all forms of return dispersion.

B The Contracting Environment in Private Market Funds

In this section, we provide more institutional details about the contracting environment in private market funds. We also discuss theoretical predictions that arise from this contracting environment.

B.1 The Limited Partnership Agreement

The limited partnership agreement (LPA) dictates all investment terms between the investment manager (GP) and the investors (LP). Examples of terms that are contained in the LPA include, but are not limited to: (i) the definition of the "entire agreement", which typically consists of the LPA and any associated amendments or side letters; (ii) rules governing contributions and distributions; (iii) fees, expenses, and taxes that are paid by LPs, (iv) GP liability and the extent to which the partnership is liable for litigation involving the GP; and (v) confidentiality of the agreement, particularly as it pertains to public entities that are subject to FOIA request.[9]

LPAs are privately negotiated between LPs and their GPs and therefore vary widely from fund to fund.[10] As one simple example of this heterogeneity, in our analysis of 91 confidential LPAs in Section 2.1.2, the lowest observed word count was roughly 16,000 and the highest was 70,000. In recent years, the Institutional Limited Partnership Association (ILPA) has pushed for the adoption of a homogenous and standardized LPA. The goal of this initiative is to lower the costs and complexity of negotiating an LPA, thereby increasing transparency in private-market funds.

B.2 Side Letters

As discussed in Section 2.1.2, LPAs are often superseded by so-called side letters. From a contracting perspective, side letters are simply a private contract between the GP and a individual LP that sits alongside the LPA (main agreement). In principle, side letters can substantially change any aspect of the LPA for a specific LP or subset of LPs, though they cannot modify the terms of other LPs in the fund unless explicitly stated by the LPA. The legal use of side letters is granted in the LPA through an "enabling clause" that allows the GP to use and extend side letters. These clauses can vary in specificity about when, for what, and with whom side letters can be used. Below, we present one example of an enabling clause that we obtained from a public LPA (from article 9. Section 9.1. Amendment of Partnership Agreement and Co-Investors (A) Partnership Agreement (b)):

"Notwithstanding the provisions of this Agreement, including Section 9.1(a), it is

[9] See this model LPA as an example.

[10] According to the London-based law firm MJ Hudson, management fees are among the most negotiated terms in an LPA. Their full report can be found here. It is also worth noting that not all investors are able or choose to negotiate their LPAs. According to Da Rin and Phalippou (2017), only 59% of all LPs and only 36% of small LPs "always negotiate contract terms".

hereby acknowledged and agreed that the General Partner on its own behalf or on behalf of the Partnership without the approval of any Limited Partner or any other Person may enter into one or more side letters or similar agreements with one or more Limited Partners which have the effect of establishing rights under, or altering or supplementing the terms of this Agreement. The parties hereto agree that any terms contained in a side letter or similar agreement with one or more Limited Partners shall govern with respect to such Limited Partner or Limited Partners notwithstanding the provisions of this Agreement." (source)

Furthermore, Morgan Lewis, a global top 10 law firm based on employees and revenue, states that side letters are commonplace and often used...

"to grant special rights and privileges to important investors (e.g., seed investors, strategic investors, those with large commitments, and employees, friends, and family) or to those subject to government regulation (e.g., ERISA, the Bank Holding Company Act, or public records laws)."(Morgan Lewis, 2015)

These special privileges can include fees, exceptions from certain fund expenses like legal fees of the partnership or placement agent fees, and management-fee offsets. For example, the following quote from the public LPA mentioned above stipulates that the allocation of fund profits (i.e, carry) can be altered by a side letter:

"Each Limited Partner's rights and entitlements as a Limited Partner are limited to the rights to receive allocations and distributions of Capital Profit and Operating Profit expressly conferred by this Agreement and any side letter or similar agreement entered into pursuant to Section 9.1(b) and the other rights expressly conferred by this Agreement and any such side letter or similar agreement or required by the Act..." (source)

In general, there is no standard for whether all LPs are notified of any side letters or whether all LPs can obtain the terms contained in side letters. Based on several conversations with SEC lawyers and former commissioners, it is our understanding that side letters can be legally used to confer economic benefits to a subset of LPs so long as their use is explicitly stated in the LPA. Importantly, GPs are not required to state the exact contents of side letters (though they may choose to do so), but instead must only state that side letters could be used to confer some LPs economic benefits. By explicitly stating that some LPs may receive better terms, fund managers are legally protected from any future litigation from LPs. From an economic perspective, this contracting environment means that investment terms between GPs and LPs are effectively negotiated on a bilateral basis. As we discuss in Section B.6, this is one important way through which fee dispersion can arise in equilibrium.

The use of side letters also appears to be widespread. For instance, the law firm Wrick Robbins states on its website:

> *While the limited partnership agreement ("LPA") of any investment fund will typically function as the figurative spine, general partners will almost invariably enter into "side letter" agreements with limited partners that branch off from the LPA.* (source)

The emphasis in the quote is ours. In a limited sample of LPs, we have confirmed that side letters are commonplace in private-market funds (see Section 2.1.2). Specifically, we find that 75% of LPAs that we analyzed under confidentially agreements contain language permitting the use of side letter by GPs.

B.3 Most-favored-nation (MFN) clauses

Given side letters may confer some economic benefits to a subset of LPs, investors will often try to negotiate for most favored nation (MFN) status (Morgan Lewis, 2015). Generally speaking, an investor with MFN status can view and select the terms of all side letters that are offered by the GP. In some cases, MFNs are written such that LPs will be automatically given any economic benefits. In other cases, LPs with MFN status must opt into any benefits they choose within a fixed period of time after the fund's close, usually 20-60 days.

In practice, MFNs do not unconditionally guarantee that LPs will receive the best terms offered by the fund. For instance, many MFNs state that LPs can expect equal treatment vis-a-vis a "similarly situated" investor. In some cases, the notion of "similarly situated" is precisely defined based on commitment size, but in other cases it can be quite vague. According to our own conversations with several large law firms that specialize in private-market fund formation, there are also many ways in which MFNs can be "carved out." For instance, LPs with MFNs are not typically able to receive terms offered to LPs that are designed by the GP as "friends and family." Moreover, GPs usually have full discretion over who they designate as friends and family.

According to an industry survey of GPs, nearly two-thirds of PE buyout funds have MFN clauses with some of their LPs (Toll and Centopani, 2017, Chart 2.32). Furthermore, only about half of GPs that grant MFNs do so based on commitment size (Toll and Centopani, 2017, Chart 2.33), meaning other investor characteristics determine the assignment of MFN status in many funds. Da Rin and Phalippou (2017) further report that 55% of the large LPs in their survey "always obtain a MFN clause" but only 13% of small LPs do.

For our purposes, the fact that MFNs are not ubiquitously given to all investors makes it legally possible for economic terms to differ across LPs in the same fund. Indeed, it would be a somewhat puzzling equilibrium outcome if all investors received MFNs, since this would eliminate the need for MFNs in the first place.

B.4 Registered Investment Advisor

Since the enactment of the Dodd-Frank Wall Street Reform and Consumer Protection Act, GPs of private-market funds are required to register with the SEC as investment advisers. As a registered investment advisor, GPs have more reporting requirements, increased fiduciary burdens to their LPs, required compliance policies, and periodic examination by the SEC.

In some cases, GPs may ask for their fiduciary responsibility to be waived by LPs (Appelbaum et al., 2016). As an illustration of this dynamic, Appelbaum et al. (2016) cites a private placement memo sent by TPG Partners VII's to potential investors:

> *"Since the amount of carried interest payable to the General Partner depends on the Partnership's performance, we may have an incentive to approve and cause the Partnership to make more speculative investments than it would otherwise make in the absence of such performance-based compensation. We may also have an incentive to dispose of the Partnership's investments at a time and in a sequence that would generate the most carried interest, even if it would not be in the Partnership's interest to dispose of the investments in that manner."*
>
> *"The General Partner may take its own interests into account in the exercise of such discretion. The exercise of such discretion may negatively impact the Limited Partners generally or may impact some Limited Partners disproportionately."*

In the context of our analysis of within-fund fee dispersion, this quote highlights that the ability of the GP to put its interest above others may disproportionally affect some LPs. The memo also makes clear that carry is one dimension over which the GP is permitted to favor its own interests. While we have no way of knowing how widespread this type of language is in the industry, the example demonstrates how GPs can have wide discretion over how carried interest or other return sources are allocated within the fund.

B.5 Channels for Fee Dispersion

The fee structure of private market funds is complex. In addition to the usual management and carry fee (e.g., the famous 2% and 20%), there are several other fee and expense categories that could vary across investors in the same fund. We discuss all of these sources now.

Management fees Management fees are charged as a percentage fee on an LP's commitment to the fund, typically in the range of 1%-3% of total committed capital during the first three to five years of the funds (i.e., the investment period). In our data, they are added to contributions. Management fees are used to cover the fund's overhead such as the salary and benefits of the investment manager and rent. After the end of the fund's investment period, the management fee is reduced to a percentage of actual invested capital, or a reduced percentage of the committed capital amounts. As discussed above, different investors in the same fund can be charged different management fees. For example, large investors or “friends and family” of the GP may be offered management fee reductions. These reductions can be accomplished via side letters (see for example page 7 of the Duane Morris’ report about private equity fees) or through a menu model, the latter of which has been offered by Bain Capital funds in recent years (Markham, 2017; Zuckerman and Or, 2011). Management fee waivers or offsets for some LPs are another possible source for within-fund fee dispersion (see the paragraph on fee offsets below).

Performance-based fees The performance-based fee, also called the carried interest, determines the profit share the GP retains out of the fund’s distribution. In our data, performance fees are netted out of distributions. It is the most important component of the GP’s expected incentive compensation and often viewed to align incentives between LPs and GPs (Marrs et al., 2009). Carried interest is subordinate to the contributions of investors, meaning that they can only be charged once the fund has returned the contributions back to LPs and earned a preferred return (e.g., hurdle rate). Once that has occurred, the GP typically receives all additional profits until it achieved its carried (catch-up period) interest, after which each additional dollar is split between the LP and the GP according to the pre-determined profit-sharing agreement. Side letters can change the carried interest agreement between an LP and a GP, as evidenced by the publicly available LPA cited above (see SEC exhibit).

Organizational fees and fund expenses Not all fund or organizational expenses are covered by the management fees. The LPA may stipulate that some out-of-pocket expenses related to the formation of the fund, audits, taxes, and some fund administration expenses are not covered by the management fees and must therefore be covered by the limited partners. This can include expenses for office supplies, travel, legal and accounting fees, fees involved with analyzing and researching potential portfolio companies, or placement agent fees. Many LPAs have indemnification clauses that ensure that the GP is not liable for legal action or associated expenses that may be brought against the fund. Some LPs can negotiate exemptions from certain fund expenses such as placement agent fees or legal fees in side letters, providing further scope for within-fund dispersion in net-of-fee return. Table 3 shows that funds with placement agents are 13 percentage points more likely to have more than investor-fee tier compared to the average fund. A public example of an LP that has a policy in place to not pay any fee for placement agents comes from the Commonwealth of Pennsylvania Public School Employee's Retirement System (PSERS). PSERS (2020) stipulates "*No Placement Agents shall be used and no payments from or on behalf of PSERS to Placement Agents shall be made in connection with PSERS' investments in or through Investment Managers.*" Unless LPs with placement agent policies pay more than their pro-rata share for all non-placement agent related fees, LP policies of this kind introduce variation in pro-rata share fund expenses across LPs within the same fund.

Fund expenses have come under much scrutiny after SEC investigations uncovered expenses that should have been borne by the fund manager (and thus the management fee) were instead charged to the fund (see Andrew Bowden's "Sunshine" speech). Andrew Bowden's speech highlighted that GPs increasingly charge LPs for operating partner expenses, consultants, and other expenses that management fees were designed to cover. For example:

> *"Another similar observation is that there appears to be a trend of advisers shifting expenses from themselves to their clients during the middle of a fund's life — without disclosure to limited partners. In some egregious instances, we've observed individuals presented to investors as employees of the adviser during the fundraising stage who have subsequently being terminated and hired back as so-called "consultants" by the funds or portfolio companies. The only client of one of these "consultants" is the fund or portfolio company that he or she covered while employed by the adviser. We've also seen advisers bill their funds separately for various back-office functions that have traditionally been included as a service provided in exchange for the management fee, including compliance, legal, and accounting — without proper disclosure that these costs are being shifted to investors. "*

Given the lack of transparency related to fund expenses, it seems natural to think that investors in the same fund could differ in terms of fund expenses, particularly if they differ in terms of their information or contract-processing costs (Gabaix and Laibson, 2006; Salop and Stiglitz, 1977). The lack of transparency may also make it difficult for some LPs to verify or enforce the terms of LPAs. Consistent with this idea, side letters often require more transparent disclosures about fund expenses.

Fee offsets and portfolio company fees Some special investors may also have management fee waivers. For instance, LPs that are on the fund's advisory may pay lower management fees. GPs charge also fees to the portfolio companies in which the partnership invests (Phalippou, Rauch, and Umber, 2018). LPAs may stipulate that this income to the fund should offset their management fees. SEC (2020) has pointed out that in some instances these offsets may not occur based on an LP's pro-rata share in the fund.

Taxes Some funds offer tax-exempt investors to invest through blocker corporations in order to minimize tax liabilities. LPAs may stipulate that any costs of setting up blocker vehicles are charged to the LPs that invest through them. In a fund with taxable and tax-exempt investor, tax-exempt investors are not always treated as such by default. The LPA or a side letter may allow LPs to opt into tax minimization services. Since there is not always default treatment of tax-exempt investors, taxes are another reason why effective carry rates could differ across investors in the same fund. For an extensive discussion of tax issues in private equity refer to Appelbaum et al. (2016).

B.6 Connection to Search and Bargaining Models

The fact that the majority of funds have two return tiers suggests to us that many GPs group investors into one of two fee tiers. This is a predictable equilibrium outcome given the contracting environment in private market investment vehicles. Side letters allow GPs and some LPs to negotiate individualized investment terms. 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 Syverson, 2004; Duffie et al., 2005; Allen et al., 2019), wherein price (or fee) dispersion is driven by costs that affect the search and bargaining process. 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 arduous 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. This friction has been cited in complaints made by state comptrollers to the U.S. Securities and Exchange Commission (SEC) about fee disclosure practices in private market vehicles (State Comptroller SEC Letter, 2015).[11] 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.[12]

[11] 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.

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

C Supplementary Analysis and Details

In this section, we present the several pieces of analysis that supplement our results in the main text.

C.1 Cluster Counts using the Gap Statistic

C.1.1 Description of the Gap Statistic

In Section 3.1, we use a various machine learning techniques to gauge the degree of return clustering in each fund. Our baseline results use Silhouette scores (Rousseeuw, 1987) to select the optimal number of clusters in a fund. For robustness, we also use the Gap statistic method developed in Tibshirani et al. (2001). The broad intuition of Tibshirani et al. (2001) is to determine the number of clusters based on what one would observe if the data were randomly chosen from a single-cluster distribution. Specifically, denote observed returns r_p for $p = 1, ..., P$ investors in a fund. Further denote the minimum and maximum observed return as r_{min} and r_{max}, respectively. The Gap statistic approach then proceeds as follows:

1. Cluster the observed returns using k-means clustering and varying the number of clusters from $k = 1, ..., K$. For each k, compute the sum of squared distances W_k of each return r_p to its nearest cluster.

2. Generate S sets of simulated data based on draws from a null distribution with a single cluster. For each simulation s, Tibshirani et al. (2001) recommend drawing P values from the uniform distribution over the range $[r_{min}, r_{max}]$.

3. For each simulation s, cluster the data using k-means clustering, again varying the number of clusters from $k = 1, ..., K$. Denote the sum of squared distances for simulation s and cluster number k by W^*_{sk}.

4. Compute the Gap statistic for k as:

$$Gap(k) = S^{-1} \sum_{s} log(W^*_{sk}) - log(W_k) \tag{1}$$

5. Denote $\bar{l} = S^{-1} \sum_s log(W^*_{sk})$ and define s_k as follows:

$$s_k = \left[1 + S^{-1}\right]^{1/2} \times \left[S^{-1} \sum_{s} \left(log(W^*_{sk}) - \bar{l}\right)^2\right]^{1/2}$$

   which is essentially the standard deviation of $log(W^*_{sk})$ across the S simulations, after a degree of freedom adjustment.

6. Choose the optimal number of clusters k via the following rule:

$$\hat{k} = \text{Minimum } k \text{ such that } Gap(k) \geq Gap(k+1) - s_{k+1}$$

As described in detail in Tibshirani et al. (2001), the Gap statistic is based on the null of a single-cluster model and rejects it in favor of the most likely multi-cluster model. Thus, in the event that the data is truly noise, the Gap statistic will conclude that there is no return tiering at all and hence fees are constant within a fund. As we discuss in Section 3.1, for the purposes of classifying whether funds as have a single return tier or not, the Gap statistic and Silhouette score approaches agree for 96% of funds.

C.1.2 Robustness of Main Results

Tables 2-5 form the basis of our main conclusions regarding: (i) the size of ex-ante fee dispersion; (ii) the types of funds that employ multiple fee schedules; (iii) whether some pensions consistently pay lower fees than others; and (iv) the types of LPs that are likely to pay lower fees in their funds. Thus, for completeness, we reproduce these tables using the Gap statistic to determine whether a fund uses multiple fee schedules (i.e., has multiple return clusters). These results are contained in

Tables IA6-IA9. These tables show that our main conclusions are robust to using the Gap statistic instead of the Silhouette score (Rousseeuw, 1987). This should not come as surprise given the two approaches largely agree on which funds use multiple fee schedules (96% agreement).

C.2 Clustering in Contributions and Distributions

We show in Section 3.1 that the majority of funds have two DVPI tiers. DVPI can be broken out into two components: contributions and distributions. Thus, we now 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:

$$\text{call-rate}_{p,f,t} = i_{p,f,t} + m_{p,f} \times t + \varepsilon^{m}_{p,f,t}. \tag{2}$$

$i_{p,f,t}$ 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 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 3.2, any variation in $m_{p,f}$ will generate a linear relationship between call rate variation and age.

Panel A of Figure IA6 analyzes clustering in call rates based on Silhouette scores as described in Section 3.1. According to this procedure, 38% of funds have one cluster, 59% of funds have two clusters, and 3% 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 IA5 Panel A, 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 B 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:

$$d_{pft} = g_{pft}\left[1 - c_{pf} \times h_{pft}\right] + \nu_{pft}$$

where g_{pft} is the unobserved gross-of-fee distribution process, c_{pf} is the rate of net-of-tax performance fees, h_{pft} is an indicator variable for whether fund f has cleared p's hurdle rate, and ν_{pft} is a residual term (e.g., measurement error). Under the assumption that gross returns ($g_{pft} = g_{ft}$) and hurdle rates ($h_{pft} = h_{ft}$) are equal across investors, then any clustering in distribution rates will reflect clustering in $c_{p,f}$ or $\varepsilon^{c}_{p,f,t}$.

Panel B of Figure 4 shows distribution of clustering across funds based on distribution rates, again using Silhouette scores for cluster selection. We find that 38% of funds have one cluster, 59% have two clusters, and 2% 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 B of Figure IA5 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 $c_{p,f}$.

C.3 Clustering Using TVPI and IRR

Panel A of Figure IA7 repeats the *k*-means clustering analysis from Section 4.1 using TVPI to measure returns. Panel B of the figure uses reported IRRs to measure returns. In both cases, around 80% of funds have 2 clusters of returns. These are largely consistent with our analysis of DVPI in Section 4.1, though the fraction of funds with a single cluster is lower when using TVPI and IRR.

C.4 Characterizing Funds and GPs that Tier Investors

C.4.1 GP Track Record

In Section 3.3 we explore the hypothesis that GPs without an established track record have a harder time raising capital and are thus more willing to negotiate fees. Our measurement of track record is based on the number of funds previously raised by the GP. When measuring previous fund counts, we only include funds that are in the same strategy.[13] Column (1) of Table IA5 shows that there is a negative relationship between our fund-level tiering indicator on fund series number. By putting fund number directly into the regression, this specification assumes a linear relationship between tiering propensity and fund number. However, this assumption may not be appropriate if inexperienced GPs are particularly willing to negotiate in order to attract capital.

In column (2), we probe the assumption of linearity by including dummy variables based on fund number in a given strategy. Funds numbered 6 or higher are omitted from the regression, which means the reported coefficients are all relative to this set of funds. We refer to funds in this omitted category as being run by "seasoned GPs". First-time funds appear to have a relatively higher propensity to tier. The coefficient in the regression indicates first-time GPs are 15 pp ($t =$ 4.58) more likely to tier than seasoned GPs. The coefficient on 2nd-time funds is similarly high but then quickly declines and becomes insignificant for 4th-time funds. This pattern of decay is also very clearly non-linear, suggesting that the specification in column (1) is indeed misspecified.

[13] In the Preqin data, this is the variable that designates the fund series number.

These findings are generally consistent with the hypothesis that is challenging for GPs with a limited track record to raise a fund. Consequently, these GPs appear more willing to concede to some LPs in order to secure capital commitments. The coefficients in column (2) also suggests that it may take two or three funds for a GP to truly establish a track record, which makes sense given that the true performance of a GP's first fund may not be known for several years after its launch. As a result of this analysis, we use an indicator variable for first, second, or third-time fund in the main text.

C.5 Characterizing Top-Tier Investors

This subsection presents additional analysis of the types of investors who are most likely to be top-tier in terms of fees.

C.5.1 Commitment Size

One potential hypothesis regarding tier assignment is that GPs offer management fee discounts to all LPs who commit over a certain amount. To investigate this hypothesis, we analyze the set of funds that have multiple clusters of contribution rates (percent of commitment called), since management fees should appear in contributions.[14] We label p as a low management fee ("low-fee") investor in fund f if it has below-median contribution rates for the entirety of the fund's life. This definition captures the idea that an investor who pays relatively low management fees should have consistently lower contribution rates relative to others in the fund. We define $C_f^* = min_{p \in T_f} C_{pf}$ as the minimum commitment among low-fee investors, where T_f is the set of low-fee investors and C_p is investor p's commitment. For each fund f, we then define the following

[14]We classify a fund has having multiple tiers in contribution rates if it has multiple clusters for the majority of its dates. We determine the number of clusters on a given date using Silhouette scores and a k-means clustering algorithm. This approach mimics the one we use in the main text.

indicator variable:

$$S_f = \begin{cases} 1 & \text{if } C_p \leq C_f^* \quad \forall p \notin T_f \\ 0 & \text{otherwise} \end{cases}$$

Thus, S_f equals one if all investors in the high management-fee tier – those with higher contribution rates – have committed less to the fund than everyone in the low-fee tier. If size is the only factor that determines assignment to management-fee tiers, then $S_f = 1$ for all of the funds in our sample.

In our full sample of funds, we find that 63% have $S_f = 1$. In other words, in 63% of funds, all investors in the low management-fee tier have higher commitment sizes than those in the high management-fee tier. We think this is a relatively high rate, especially considering that any classification errors of high and low-fee investors will bias this rate down from 100%.

There are also several economic reasons why we should not expect a rate of 100%. One reason is that investors may trade off contracts with high management fees and low carry against contracts with low management fees and high carry. This sort of tradeoff accords with the contract offering that has been famously used by recent Bain Capital funds (Markham, 2017). In this case, it is not clear whether size should perfectly sort investors into low and high-management fee tiers.

A second and more important reason why we should not expect S_f to equal one for all funds is that size is one of only several factors that could determine tier assignment. For example, GPs could give fee breaks to investors who are thought to be skilled in order to entice other LPs to commit to the fund, though skill may not be perfectly captured by commitment size. To the extent that these types of other factors influence tier assignment, then we would not expect S_f to equal one in all funds. And empirically, our regression evidence in Table (5) suggests that these other factors do have explanatory power for tier assignment over and above what is contained in commitment size.

From a legal perspective, our analysis of limited partnership agreements (LPAs) also suggests that tier assignment may depend on more than just commitment size. For example, 54% of the LPAs that we analyzed either stated that the GP was not required to offer the terms of any side letter to other LPs or contained no language regarding any such restrictions. 32% of LPAs did indicate that the GP would offer terms to other LPs who it deemed "similar" to those that had received side letters. However, in these cases, the notion of similarity was left vague – likely intentionally in our view – and size was mentioned as one of several factors that decided similarly. In other words, to the extent that it is representative, our sample of LPAs does not imply that it is standard for all LPs above a certain commitment threshold to be legally guaranteed the same fee terms. This view is supported Toll and Centopani (2017, Chart 2.33), who report that only about half of GPs that grant MFNs do so based on commitment size.

This is not to argue that commitment size does not play a large role in tier assignment, but rather that it is one of a few factors that could play a role. Indeed, the fact that $S_f = 1$ for a large fraction of funds suggests that size is an importing sorting mechanism and accords with anecdotal evidence from practitioners regarding "check size" and fee breaks. In many ways, the goal of our analysis is to let the data speak to which factors determine tier assignment, as opposed to pre-committing to a specific hypothesis (e.g., that check size is the *only* factor).

C.6 Pension Effects

C.6.1 Out-of-Sample Tests

We interpret the existence of pension effects in Section 4 as evidence that some pensions are able to consistently pay lower fees in their respective funds. As one way to validate this view, we test whether a pension's relative performance in a fund is predictable based on its relative performance in other funds. Specifically, for each pension p in fund f, we estimate pension effects in the sample *excluding* fund f, after which we apply the Bayesian shrinkage procedure described in Section 4.1.1 We denote this Bayesian-adjusted pension effect estimate as $\widehat{\alpha}_{p,-f}$, where the subscript $-f$ indicates that it is based on the sample of funds excluding f. We then use this estimate to predict whether p is a top-tier investor in fund f as follows:

$$y_{pf} = \lambda_a + \beta \widehat{\alpha}_{p,-f} + \varepsilon_{pf}$$

where y_{pf} is an indicator for whether p is top-tier in fund f and λ_a are fixed effects based on vigintiles of fund f's age.

Figure IA8 depicts the estimated relationship using a binned scatter plot, after controlling for the aforementioned fixed effects. The plot shows a clear linear relationship between actual and predicted tier assignment. The point estimate on $\widehat{\alpha}_{p,-f}$ is $\beta = 1.10$. The standard error, which is adjusted for heteroskedasticity, equals 0.06 and indicates that we can not reject the null that $\beta = 1$. In other words, the Bayesian-adjusted pension effect based on all funds except f is an unbiased forecast of tier assignment in fund f. These results are consistent with the idea that pensions are endowed with a persistent form of negotiation skill or bargaining power that impacts the fees they pay across all of their funds. Alternatively, they could be driven by differences in accounting that cause some LPs to overperform in all of their funds, though we argue that this unlikely in Section A.4.

C.6.2 Pension Effects Using DVPI and IRR

One natural concern with our analysis of pension effects is that they rely on TVPI as a measure of returns. For each investor p in fund f, TVPI reflects actual cash distributions that have been received by p as well as the reported net asset value of p's non-liquidated positions in the fund. Consequently, if some pensions consistently report net asset values differently than others, then one might be concerned that this drives our estimate of pension effects. To alleviate these concerns, we recalculate Table IA10 using DVPI to measure returns instead of TVPI. DVPI is based only on cash distributions and is therefore not influenced by how pensions report net asset values. When using DVPI to measure returns, we continue to find evidence of pension effects and LP-GP effects.

C.7 The Impact of Characteristics

For completeness, Figures IA9 and IA10 replicate Figure 6 using DVPI and IRR, respectively, to measure returns. Panel A of Figures IA9 and IA10 illustrate that the observed distribution of pension effects is much wider than what would be observed if contracts (i.e., returns) were randomly assigned to LPs in each fund. Panel B of Figures IA9 and IA10 show that adjusting returns for characteristics as in Section 5.2 accounts for some of the observed pension effects, yet some pensions continue to persistently over- or underperform in their respective funds even controlling for observables. This latter finding reinforces our analysis using TVPI and is consistent with the existence of unobservable contracting skills that vary across public pensions in the United States.

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.

Appelbaum, E., R. Batt, et al. (2016). Fees, fees and more fees: how private equity abuses its limited partners and us taxpayers. Technical report, Center for Economic and Policy Research (CEPR).

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*.

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

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

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.

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.

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.

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

Marrs, N. M., L. D. Hellebusch, and K. Das (2009). Variations in structuring" whole fund" and" deal by deal" carried interest or promote in real estate funds and joint ventures. *Real Estate Finance Journal 24*(4), 5.

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

Nadauld, T. D., B. A. Sensoy, K. Vorkink, and M. S. Weisbach (2019). The liquidity cost of private equity investments: Evidence from secondary market transactions. *Journal of Financial Economics 132*(3), 158–181.

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

PSERS (2020, March). Placement agent policy.

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics 20*, 53–65.

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.

SEC (2020). Observations from examinations of investment advisers managing private funds.

Tibshirani, R., G. Walther, and T. Hastie (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 63*(2), 411–423.

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 IA1: Audit Results

	Contributions	Distributions	Fund Value	DVPI	TVPI
Total Number Matched	932	875	841	814	558
Number of Perfect Matches	908	853	809	799	536
Perfect Match Rate (%)	97.4	97.5	96.2	98.2	96.1
D for Non-Matches					
p25	0.1	1.1	0.6	1.7	0.3
p50	9.4	15.9	18.5	10.7	2.0
p75	19.2	50.0	62.2	28.0	5.9
Max	53.7	100.0	101.7	100.0	64.7

Notes: This table presents the results from our audit of the Preqin data. We sent direct FOIA requests to 65 pensions and received data responses from 48, which we then compared to the Preqin data. For each variable v, we compute the percentage deviation between the value in our direct FOIA and Preqin:

$$D^v_{if} = \frac{|v^P_{if} - v^A_{if}|}{v^P_{if}}$$

where v^P_{if} is the value for pension i invested in fund f in Preqin and v^A_{if} is the associated value from our direct FOIA. When $v^P_{if} = v^A_{if}$ we set $D^v_{if} = 0$ and when $v^P_{if} = 0$ we scale absolute deviations by v^A_{if}. The first row of the table shows the number of non-missing observations that are in Preqin and our manual audit data. The second row shows the number of exact matches for each variables (e.g., $v^P_{if} = v^A_{if}$) and the third row shows the perfect match rate. The remaining rows in the table show the distribution of D^v_{if}, expressed in percentage terms, for observations that are not perfectly matched across the two datasets ($v^A \neq v^P$).

Table IA2: Within-Fund Dispersion in Fixed and Performance Fees in Green Single-Round Funds

	Mgmt (bps)		Carry (%)		Carry Placebo		
	m	$se(m)$	c	$se(c)$	u	$se(u)$	p-value
Private Debt	77	23	7.7	1.9	1.1	1.7	0.50
Private Equity	93	10	2.7	1.1	-0.0	0.7	0.95
Real Estate	81	14	3.1	2.8	-1.5	1.5	0.30
Venture Capital	33	8	0.4	0.3	-0.1	1.0	0.89

Notes: This table presents estimates of the standard deviation 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 the within-fund standard deviation in fund f's capital call rate and age_{ft} is its age in years at time t. The call rate equals the fraction of committed capital that has been called for investment. We estimate the regression for funds that are at most 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 standard deviation 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 carry dispersion should not be detectable ($u = 0$). 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 under the "Carry Placebo" header report u, the standard error of u, and the p-value from the test of the null that $u = 0$. In all cases, standard errors are clustered by fund. The estimation sample is for single-round funds that do not invest in oil and gas. Infrastructure funds are excluded due to lack of data. We also only analyze funds in the master sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Silhouette score approach (Rousseeuw, 1987) described in Section 3.1.

Table IA3: Robustness: Characteristics of Funds that Use Multiple Fee Structures

	Dependent Variable: $100 \times 1(\text{Tiers}_f > 1)$									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Quartile of GP's Prior Funds	4.21**					2.22				
	(2.24)					(0.74)				
Undersubscribed Fund		12.04**					10.64**			
		(4.83)					(2.42)			
Fund Number 1-3			12.95**					15.84**		
			(6.63)					(6.07)		
Uses Placement Agent				12.76**					13.01**	
				(6.83)					(4.77)	
Infrastructure					15.05**					17.54*
					(4.45)					(1.83)
Private Debt					5.00*					3.90
					(1.86)					(1.01)
Real Estate					5.98**					7.22**
					(2.75)					(2.24)
Venture Capital					-33.04**					-33.48**
					(-13.36)					(-11.36)
1-Round, No Oil and Gas						x	x	x	x	x
Adjusted R^2	0.05	0.05	0.04	0.04	0.15	0.04	0.03	0.04	0.03	0.14
N	1,163	1,498	2,260	2,256	2,260	564	695	1,393	1,393	1,393

Notes: This table reports OLS estimates from a linear probability model of the likelihood that a fund has multiple investor tiers. We define an indicator variable for whether a fund has multiple tiers based on the k-means clustering analysis and Silhouette score approach (Rousseeuw, 1987) described in Section 3.1. The dependent variable in the regression is this indicator variable multiplied by 100. For each fund f run by GP g, the quartile of g's prior funds is defined as the average performance quartile of all g's funds that were raised at least four years before f's final close. Quartile rankings are from Preqin and measured as of the time of close and higher values correspond to worse performance, i.e., the bottom quartile is coded as 4 while the top quartile is coded as 1. Funds are undersubscribed if their final close size is below their target fund size. Number of funds raised by the GP is measured as of the time of f's close and includes the current fund. Use of placement agent is based on information from Preqin. We assume missing entries mean no placement agent is being used. Asset class designations are also from Preqin. Columns (1)-(5) are estimated on the full sample of funds and columns (6)-(10) are estimated for funds raised in single-round funds that do not invest in oil and gas firms. All regressions include: (i) a fixed effect based on the average number of investors in the fund over its lifetime, rounded to the nearest integer; and (ii) a fixed effect based on the decile of the fund's size. In columns (5) and (10), buyout funds are the omitted asset class in the regression. Heteroskedasticity-robust t-statistics are reported below point estimates.

Table IA4: Robustness: The Determinants of Top-Tier LPs

Dependent Variable:	$100 \times 1(p$ is Top-Tier in $f)$	
	(1)	(2)
Percent of Fund	0.96**	0.74**
	(4.80)	(2.39)
Large Pension (AUM)	13.97**	8.66*
	(4.97)	(1.93)
LP-GP Fund Count	0.71**	1.50**
	(2.25)	(2.70)
Quartile of LP's Prior Funds	-9.51**	-8.66**
	(-3.31)	(-2.02)
Early PE Investor	8.44**	4.76
	(3.44)	(1.39)
Elected Board Members (%)	0.12**	0.15**
	(2.79)	(2.42)
Board Size	0.61**	0.51
	(2.80)	(1.57)
1-Round, No Oil and Gas		x
R^2	0.11	0.11
N	5,050	1,843

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. An investor is defined as top-tier if it has above-median returns (TVPI) in fund f for the majority of its life. Percent of fund is defined as p's commitment relative to total fund 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 performance quartile ranking of p's funds that were active at the time of fund f's close, conditional on at least four years of performance history. Quartile rankings come from Preqin and are measured as of each fund's close date. The bottom quartile is coded as 4, meaning higher values correspond to worse performance. We define an indicator variable for whether p was an early investor in private markets if its 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 vigintiles of fund age. Column (1) is based on all funds in the core sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Silhouette score approach (Rousseeuw, 1987) described in Section 3.1. Column (2) restricts the sample to include funds that were raised in a single round and do not invest in oil and gas firms. In both cases, we require that funds are at least one year old. t-statistics are reported below point estimates and are based on standard errors that are clustered within each investor-vintage cell.

Table IA5: First-Time Funds and the Likelihood of Using Multiple Fee Structures

	Dep. Variable: $100 \times 1(\text{Tiers}_f > 1)$		
	(1)	(2)	(3)
Number of Funds Raised by GP	-1.70**		
	(-5.35)		
Fund Number 1		15.11**	
		(4.58)	
Fund Number 2		15.23**	
		(5.44)	
Fund Number 3		10.82**	
		(3.78)	
Fund Number 4		3.40	
		(1.06)	
Fund Number 5		-0.99	
		(-0.28)	
Fund Number 1-3			12.95**
			(6.63)
Adjusted R^2	0.04	0.04	0.04
N	2,260	2,260	2,260

Notes: This table reports OLS estimates from a linear probability model of the likelihood that a fund has multiple investor tiers. We define an indicator variable for whether a fund has multiple tiers based on the k-means clustering analysis and Silhouette score approach (Rousseeuw, 1987) described in Section 3.1. The dependent variable in the regression is this indicator variable multiplied by 100. The indicator for first-time funds is based on whether the fund is the first launched by a GP in a given strategy. In column (3), we include indicator variables for the number of fund and the omitted category is fund whose number is 6 or higher. Thus, all coefficients in column (2) are relative to this omitted category of funds. All regressions include a fixed effect based on the decile of the fund's size and the average number of investors in the fund over its lifetime, rounded to the nearest integer. Heteroskedasticity-robust t-statistics are reported below point estimates.

Table IA6: Within-Fund Dispersion in Fixed and Performance Fees Using the Gap Statistic

	Mgmt (bps)		Carry (%)		Carry Placebo		
	m	$se(m)$	c	$se(c)$	u	$se(u)$	p-value
Infrastructure	96	23	5.1	3.6	5.4	3.3	0.11
Private Debt	98	19	6.5	1.8	1.5	1.7	0.37
Private Equity	85	7	3.6	0.9	0.8	0.5	0.13
Real Estate	70	11	7.0	1.7	1.2	1.1	0.27
Venture Capital	43	8	0.5	0.3	0.0	0.7	0.95

Notes: This table presents estimates of the standard deviation 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 the within-fund standard deviation 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 at most 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 standard deviation 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 carry dispersion should not be detectable ($u = 0$). 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 under the "Carry Placebo" header report u, the standard error of u, and the p-value from the test of the null that $u = 0$. In all cases, standard errors are clustered by fund. All estimates are based on funds in the master sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Gap statistic approach (Tibshirani et al., 2001) described in Section 3.1.

Table IA7: Characteristics of Funds that Use Multiple Fee Structures Using the Gap Statistic

	Dependent Variable: $100 \times 1(\text{Tiers}_f > 1)$						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Quartile of GP's Prior Funds			5.71**				
			(2.92)				
Undersubscribed Fund				12.52**			
				(4.75)			
Fund Number 1-3					11.14**		
					(5.43)		
Uses Placement Agent						10.87**	
						(5.49)	
Infrastructure							14.43**
							(3.33)
Private Debt							2.25
							(0.76)
Real Estate							7.77**
							(3.35)
Venture Capital							-31.53**
							(-12.48)
GP FE	x						
Law-Firm FE		x					
p(FEs = 0)	0.00	0.00					
Adjusted R^2	0.27	0.05	0.04	0.04	0.03	0.03	0.12
N	1,916	950	1,163	1,498	2,260	2,256	2,260

Notes: This table reports OLS estimates from a linear probability model of the likelihood that a fund has multiple investor tiers. We define an indicator variable for whether a fund has multiple tiers based on the k-means clustering analysis and Gap statistic approach (Tibshirani et al., 2001) described in Section 3.1. The dependent variable in the regression is this indicator variable multiplied by 100. For each fund f run by GP g, the quartile of g's prior funds is defined as the average performance quartile of all g's funds that were raised at least four years before f's final close. Quartile rankings are from Preqin and measured as of the time of close and higher values correspond to worse performance, i.e., the bottom quartile is coded as 4 while the top quartile is coded as 1. Funds are undersubscribed if their final close size is below their target fund size. Number of funds raised by the GP is measured as of the time of f's close and includes the current fund. Use of placement agent is based on information from Preqin. We assume missing entries mean no placement agent is being used. Asset class designations are also from Preqin. Columns (1) and (2) respectively include only a GP fixed effect and a law-firm fixed effect. The regressions in column (3)-(7) include: (i) a fixed effect based on the average number of investors in the fund over its lifetime, rounded to the nearest integer; and (ii) a fixed effect based on the decile of the fund's size. In column (7), buyout funds are the omitted asset class in the regression. Heteroskedasticity-robust t-statistics are reported below point estimates.

Table IA8: Do Some LPs Consistently Pay Lower Fees? Robustness Using the Gap Statistic

	LP Effects					LP-GP Effects				
Age Min.	F	p	p^*	K	N	F	p	p^*	K	N
1	3.51	<0.01	<0.01	181	6,731	1.62	<0.01	<0.01	1,480	4,086
4	3.36	<0.01	<0.01	158	5,116	1.67	<0.01	<0.01	1,061	2,818
8	2.66	<0.01	<0.01	132	3,304	1.74	<0.01	<0.01	635	1,658

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. 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. λ_a are fixed effects based on vigintiles of fund f's age. $\alpha_k = \alpha_p$ are fixed effects for LPs (p) for the results in the columns on the left and $\alpha_k = \alpha_{gp}$ are LP-GP fixed effects for the results in the columns on the right. 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 500 times then generate p^* by comparing the actual F-statistic to the simulated distribution of F-statistics. We repeat the analysis using DVPI and IRR to measure returns in Internet Appendix C.6.2. All estimates are based on funds in the core sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Gap statistic approach (Tibshirani et al., 2001) described in Section 3.1.

Table IA9: The Determinants of Top-Tier LPs Using the Gap Statistic

Dependent Variable:	$100 \times 1(p$ is Top-Tier in $f)$		
	(1)	(2)	(3)
Percent of Fund	1.06**	0.84**	1.01**
	(4.76)	(3.10)	(2.89)
Large Pension (AUM)	13.79**	16.65**	15.04**
	(4.93)	(5.04)	(3.58)
LP-GP Fund Count	0.63*	0.55	1.05
	(1.80)	(1.24)	(1.56)
Quartile of LP's Prior Funds	-9.91**	-11.95**	-13.61**
	(-3.33)	(-3.57)	(-3.15)
Early PE Investor	9.60**	10.76**	6.65*
	(3.82)	(3.59)	(1.74)
Elected Board Members (%)	0.11**	0.16**	0.26**
	(2.59)	(2.96)	(3.51)
Board Size	0.65**	0.38	0.08
	(2.96)	(1.42)	(0.23)
Fund Age Min. (yrs)	1	4	8
R^2	0.11	0.12	0.13
N	4,677	3,171	1,574

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. An investor is defined as top-tier if it has above-median returns (TVPI) in fund f for the majority of its life. This indicator variable is multiplied by 100 in the regression. Percent of fund is defined as p's commitment relative to total fund 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 performance quartile ranking of p's funds that were active at the time of fund f's close, conditional on at least four years of performance history. Quartile rankings come from Preqin and are measured as of each fund's close date. The bottom quartile is coded as 4, meaning higher values correspond to worse performance. We define an indicator variable for whether p was an early investor in private markets if its 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 vigintiles of fund age. t-statistics are reported below point estimates and are based on standard errors that are clustered within each investor-vintage cell. All estimates are based on funds in the core sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Gap statistic approach (Tibshirani et al., 2001) described in Section 3.1.

Table IA10: Do Some LPs Consistently Pay Lower Fees? Robustness using DVPI

	LP Effects					LP-GP Effects				
Age Min.	F	p	p^*	K	N	F	p	p^*	K	N
1	2.43	<0.01	<0.01	182	7,210	1.32	<0.01	<0.01	1,580	4,460
4	2.27	<0.01	<0.01	159	5,487	1.35	<0.01	<0.01	1,142	3,084
8	1.93	<0.01	<0.01	135	3,518	1.20	<0.01	<0.01	678	1,789

Notes: This table is based on the following regression: $y_{pf} = \lambda_a + \alpha_k + \varepsilon_{pf}$, where y_{pf} is an indicator variable that equals 1 if p has above median returns in fund f. 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 DVPI. λ_a are fixed effects based on vigintiles of fund f's age. $\alpha_k = \alpha_p$ are fixed effects for LPs (p) for the results in the columns on the left and $\alpha_k = \alpha_{gp}$ are LP-GP fixed effects for the results in the columns on the right. 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 500 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. All estimates are based on funds in the core sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Silhouette score approach (Rousseeuw, 1987) described in Section 3.1.

Table IA11: Do Some LPs Consistently Pay Lower Fees? Robustness using IRR

	LP Effects					LP-GP Effects				
Age Min.	F	p	p^*	K	N	F	p	p^*	K	N
1	2.34	<0.01	<0.01	158	6,310	1.33	<0.01	<0.01	1,378	3,851
4	2.45	<0.01	<0.01	140	4,957	1.25	<0.01	<0.01	1,028	2,764
8	2.34	<0.01	<0.01	119	3,196	1.32	<0.01	<0.01	616	1,619

Notes: This table is based on the following regression: $y_{pf} = \lambda_a + \alpha_k + \varepsilon_{pf}$, where y_{pf} is an indicator variable that equals 1 if p has above median returns in fund f. 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 IRR. λ_a are fixed effects based on vigintiles of fund f's age. $\alpha_k = \alpha_p$ are fixed effects for LPs (p) for the results in the columns on the left and $\alpha_k = \alpha_{gp}$ are LP-GP fixed effects for the results in the columns on the right. 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 500 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. All estimates are based on funds in the core sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Silhouette score approach (Rousseeuw, 1987) described in Section 3.1.

Figure IA1: Dispersion Implied by Recallable Capital Accounting versus Actual Dispersion for a Sample LP

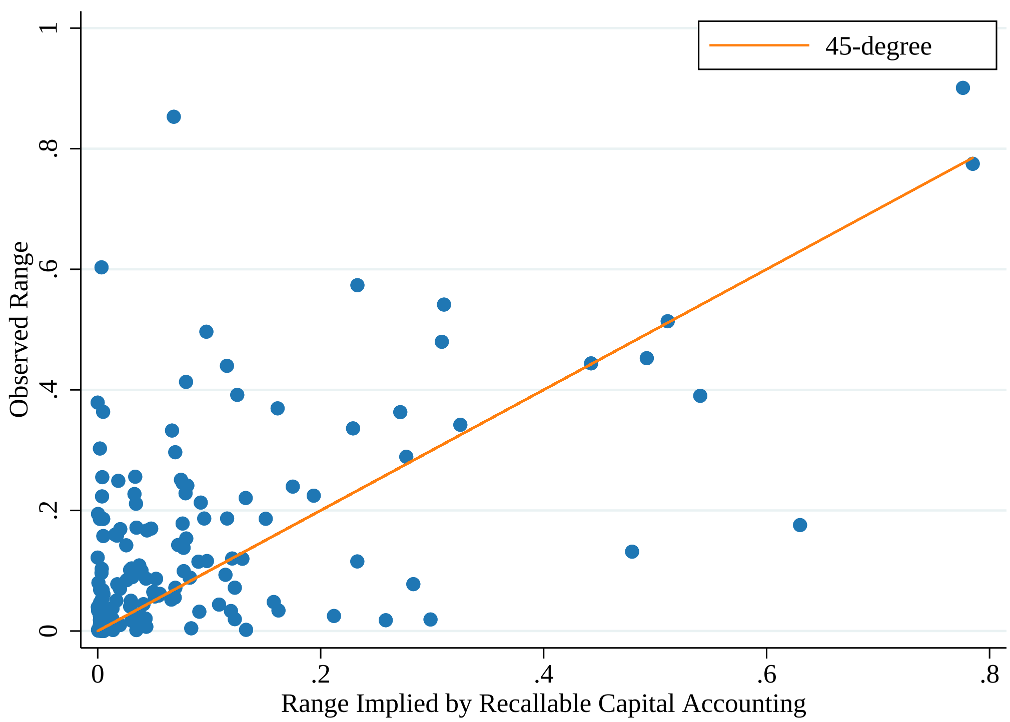

Notes: This x-axis of this plot corresponds to the range of returns in fund f that we would expect to see from accounting differences in recallable capital. The y-axis is the actual range observed in the data. Each dot in the scatter plot corresponds to a fund. The orange line is the 45-degree line. The data for the plot comes from a sample LP for whom we were able to obtain distributions and contributions when recallable capital was included, as well as when it is not. See Section A.4 for more details.

Figure IA2: The Impact of Excluding Initial Quarters on Contribution Tiering Rates

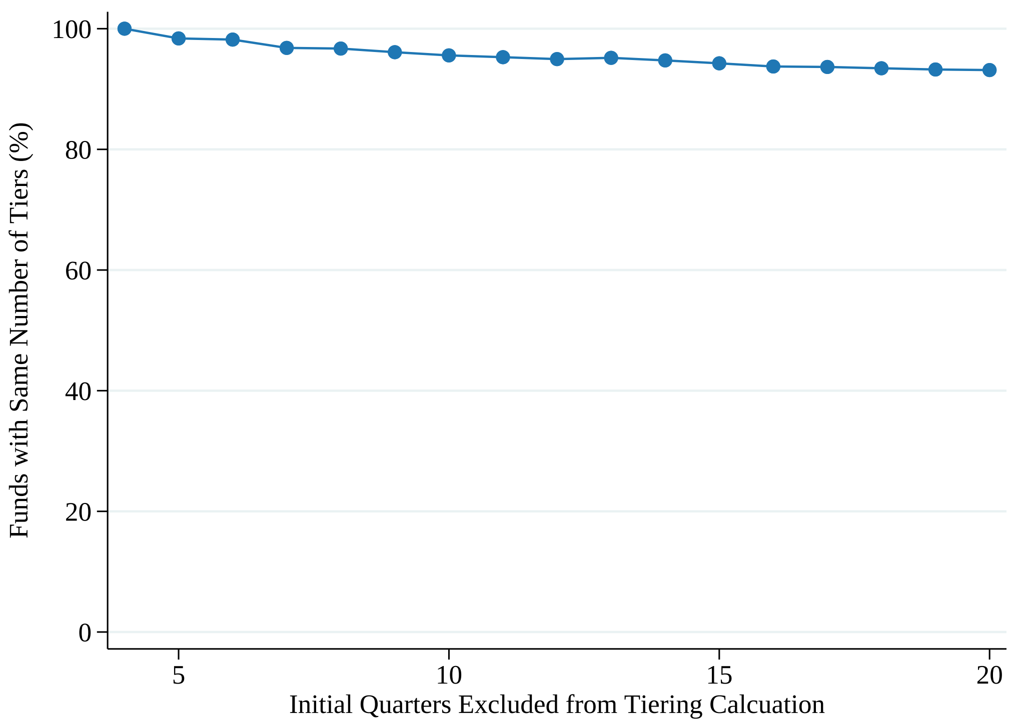

Notes: For each fund f and date t, we compute the number of clusters N_{ft} for contribution rates using a k-means clustering analysis, where the number of clusters is chosen based on Silhouette scores (Rousseeuw, 1987). We then take the average of N_{ft} over all $t \geq q$ and define its rounded value $N_f(q)$ as the number of contribution clusters in f. The plot shows the percent of funds for which $N_f(u) = N_f(q = 4)$ as a function of u. We use $q = 4$ as our baseline when computing cluster counts for all variables throughout the paper. See Section 3.1 for more details on the clustering approach.

Figure IA3: Within-Fund Dispersion of Returns Excluding Large LPs Prior to 2010

Panel A: DVPI Dispersion

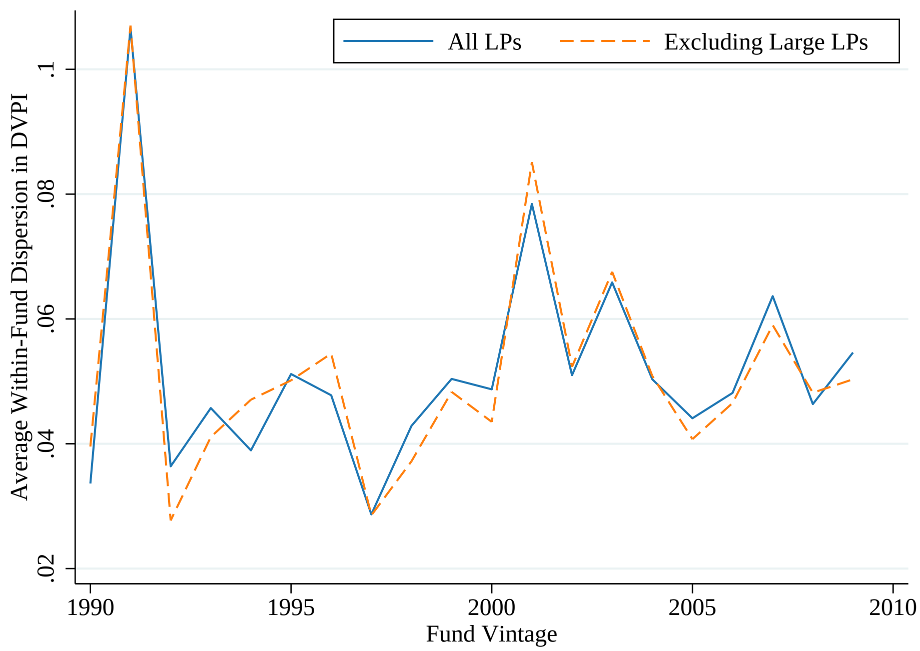

Panel B: TVPI Dispersion

Notes: This plot shows the cross-sectional average within-fund standard deviation of returns for funds of a given vintage. The blue line in both plots shows dispersion when including all LPs and the orange line shows dispersion when excluding large LPs. Panel A of the figure measures returns using DVPI and Panel B measures them using TVPI. We define large LPs as those with over $100 billion in AUM.

Figure IA4: Forgone Return Curves by Fund Types

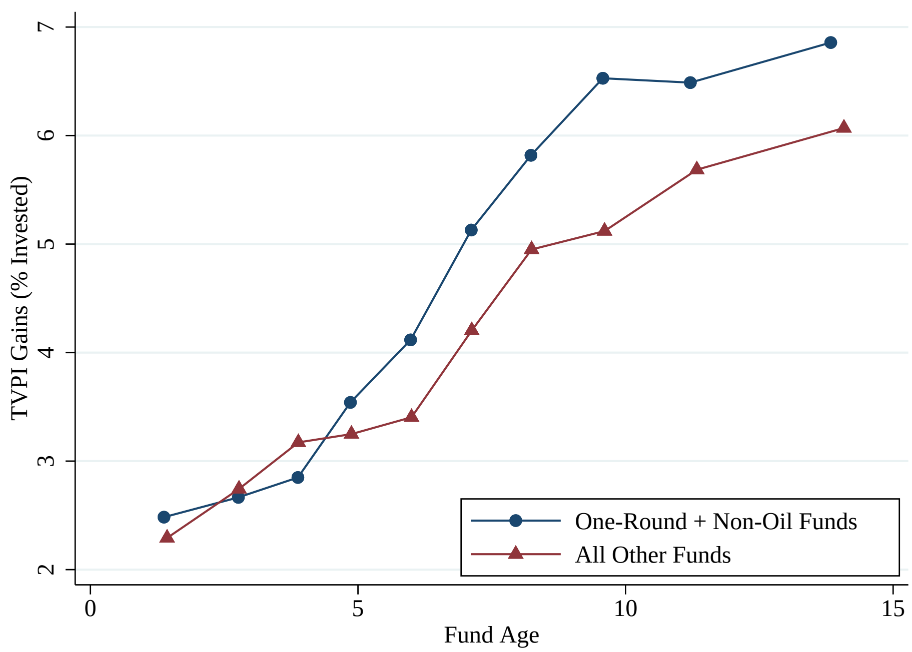

Notes: This plot shows forgone capital per dollar of contributions over the fund lifecycle. Formally, let r_{pfa} equal the TVPI of investor p in fund f at age a and define $\widetilde{r}_{fa}$ as the average return of LPs who are in the top-tier of fees. As in Section 4.1, top-tier LPs are defined as those with above-median returns for the majority of a fund's life. Forgone capital for each investor is defined as $\Gamma_{pfa} = K_{pfa} \max(\widetilde{r}_{fa} - r_{pfa}, 0)$, where K_{pfa} equals cumulative contributions. To aggregate, we group observations based on declines of fund age a, sum the total amount of forgone capital Γ in each group, and scale the total by the sum of contributions in the group. We then plot the resulting amount of forgone capital per dollar of contributions against the average fund age in each decile. The plot shows forgone capital based on whether or not funds were raised in a single round and do not invest in oil and gas firms.

Figure IA5: Example of Clustering in Call and Distribution Rates

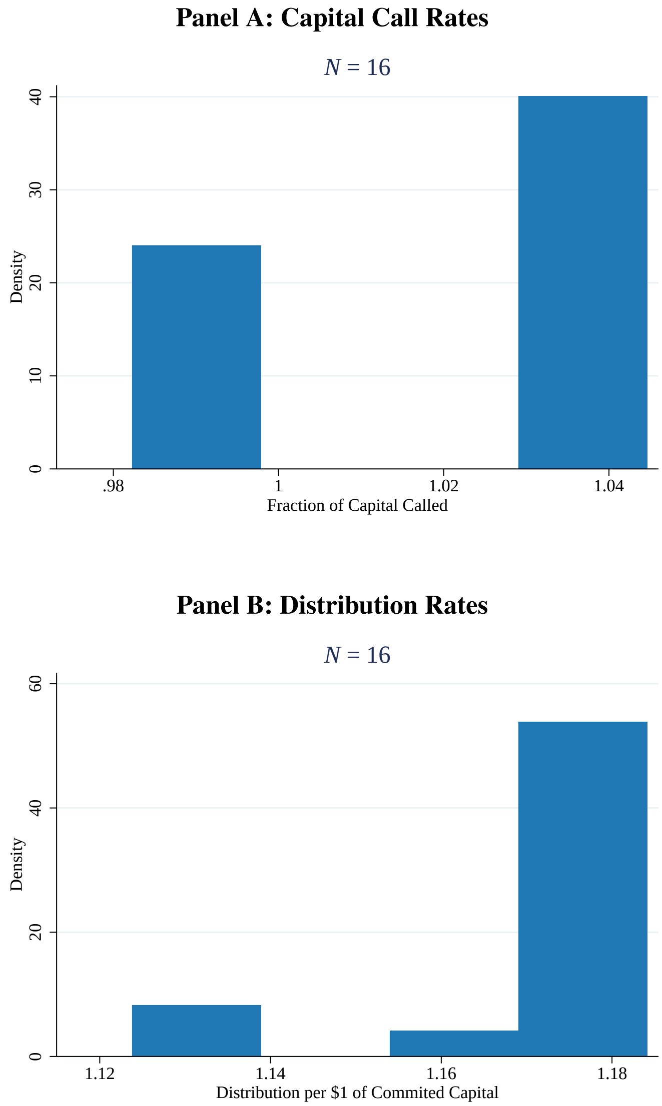

Notes: Panel A shows capital call rates (contributions divided by committed capital) for 16 investors in the same fund at a fixed point in time. Similarly, Panel B shows distribution rates (distributions per dollar of committed capital) for these investors in the same fund at the same time. We are not able to identify individual funds or investors per our data-sharing agreement with Preqin. The example fund and date used in this graph are the same as the one used in Figure 3. The fund was closed in a single fundraising round, which means all investors entered it at the same time.

Figure IA6: Distribution of Call and Distribution Cluster Counts Across Funds

Panel A: Contribution Rates

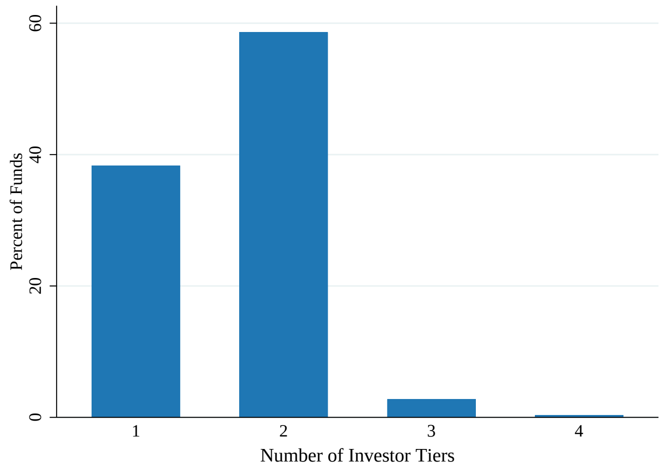

Panel B: Distribution Rates

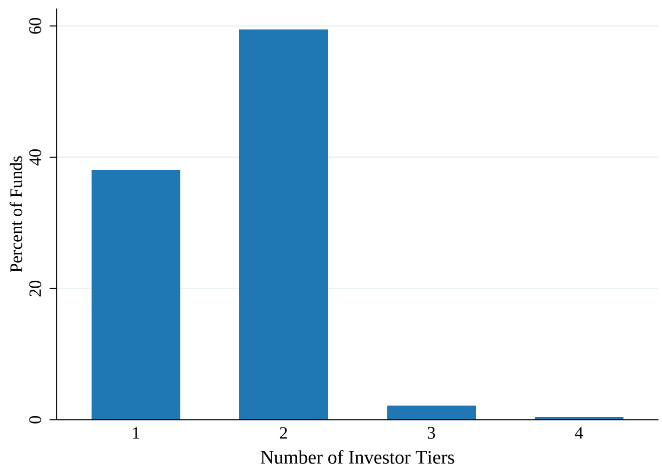

Notes: This plot shows the results of the clustering analysis across funds. For each fund f and date t, we compute the number of clusters in a given variable using a k-means clustering analysis, where the number of clusters is chosen based on Silhouette scores (Rousseeuw, 1987). The number of clusters at the fund level is defined as the average number of clusters over each fund's life, rounded to the nearest integer. **Panel A** summarizes the cluster analysis for contribution rates (call rates), defined as 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 B** 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. See Section 3.1 for more details on the clustering approach.

Figure IA7: Distribution of TVPI and IRR Cluster Counts Across Funds

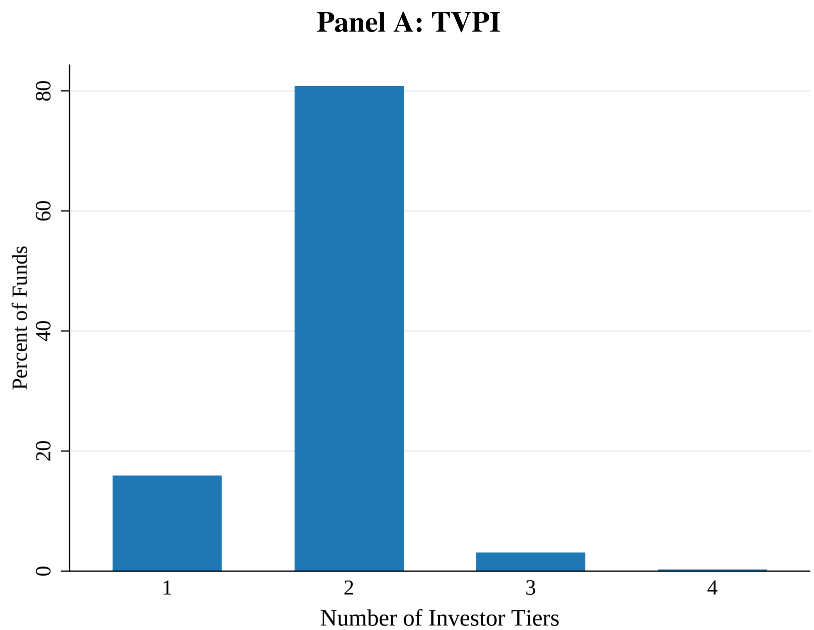

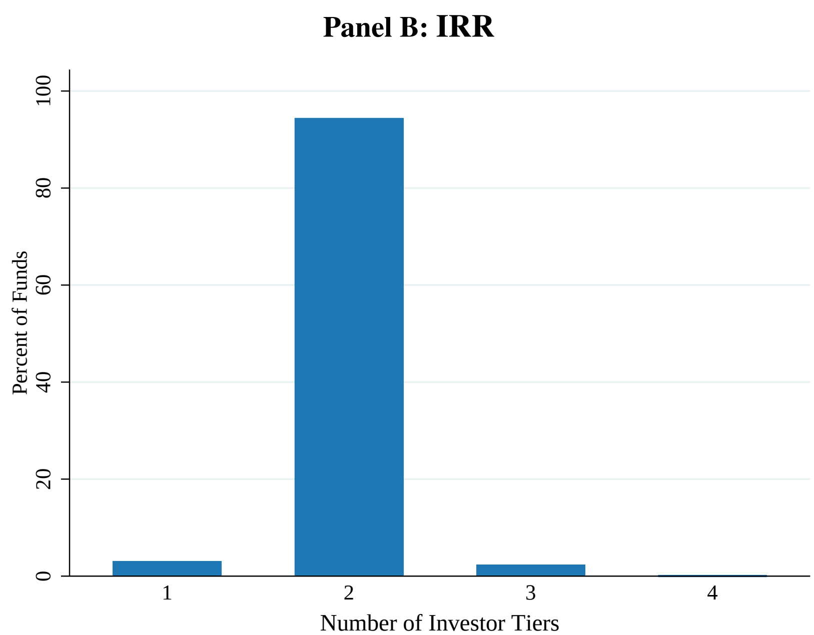

Notes: This plot shows the distribution of within-fund clusters (investor tiers) in net-of-fee returns across funds. For each fund f and date t, we compute the number of clusters in either TVPI or IRR using a k-means clustering analysis, where the number of clusters is chosen based on Silhouette scores (Rousseeuw, 1987). 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 A shows the distribution of cluster counts across funds for TVPI and Panel B repeats its for IRR. See Section 3.1 for more details on the clustering approach.

Figure IA8: Pension Effects: Actual vs Predicted

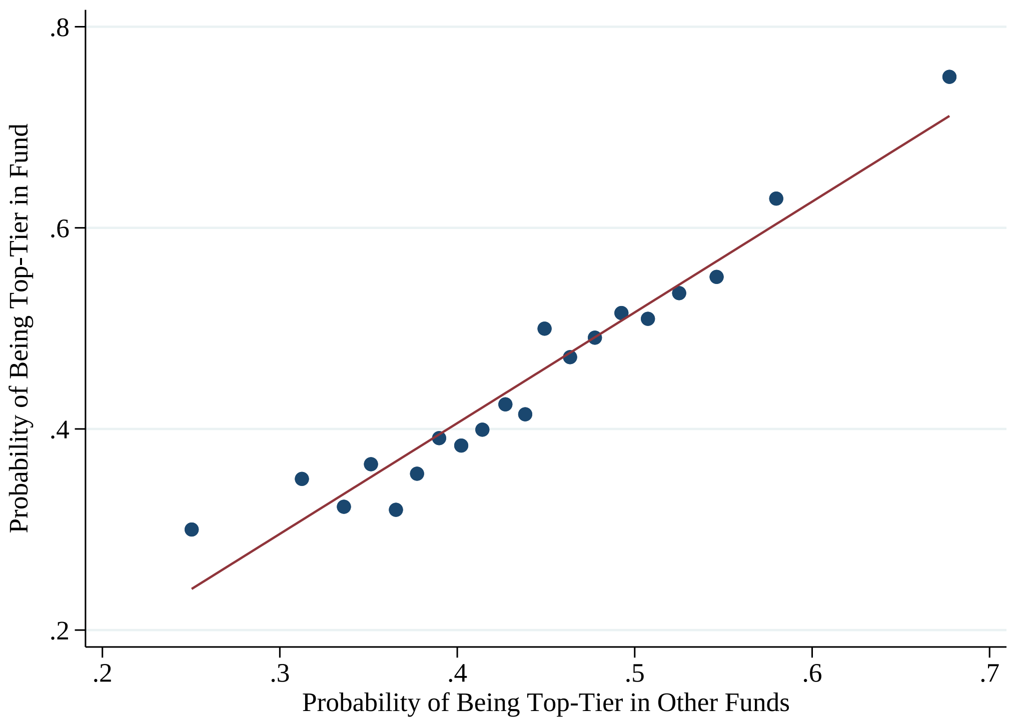

Notes: For each pension p in fund f, we estimate pension effects using all funds except f according to Equation (6), after which we shrink the estimates using the empirical Bayes method described in Section 4.1.1. These estimates are denoted by $\widehat{\alpha}_{p,-f}$. The figure then shows a binned scatter plot of an indicator y_{pf} for whether p is top-tier in fund f against $\widehat{\alpha}_{p,-f}$, after controlling for fixed effects for fund and those based on vigintiles of fund age. The plot shows that a pension's relative performance in one fund is predictable using its relative performance in its other funds. All estimates are based on funds in the core sample (see Section 2.2.2) that were determined to have multiple investor tiers by the Silhouette score approach (Rousseeuw, 1987) described in 3.1.

Figure IA9: The Distribution of Pension Effects based on DVPI

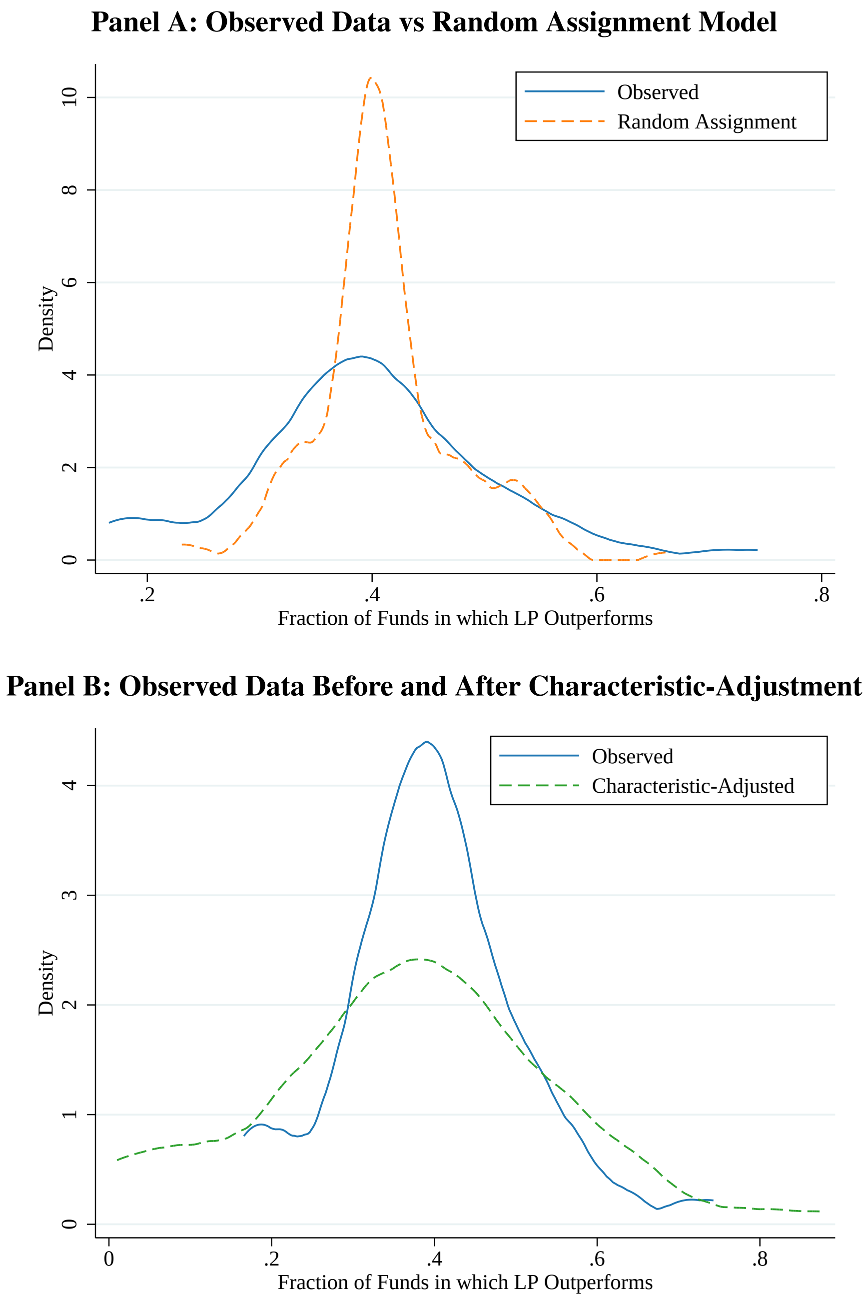

Notes: This plot shows the likelihood of being a top-tier investor across pensions ("pension effects"). In Panel A, pension effects are estimated by regressing pension p's status in fund f, y_{pf}, on pension fixed effects and fixed effects based on the vingintiles of fund age. The estimated pension fixed effects are then shrunk towards their mean using an empirical Bayes estimate. The blue line in the plot shows the results of this procedure. The orange line in the top plot shows the pension effects that we obtain when simulating random assignment of contracts to pensions in each fund. 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 orange, alongside the pension effects before adjustments. Returns are measured using DVPI. See Sections C.7, 4.1.1 and 4.2.2 for more details.

Figure IA10: The Distribution of Pension Effects based on IRR

Panel A: Observed Data vs Random Assignment Model

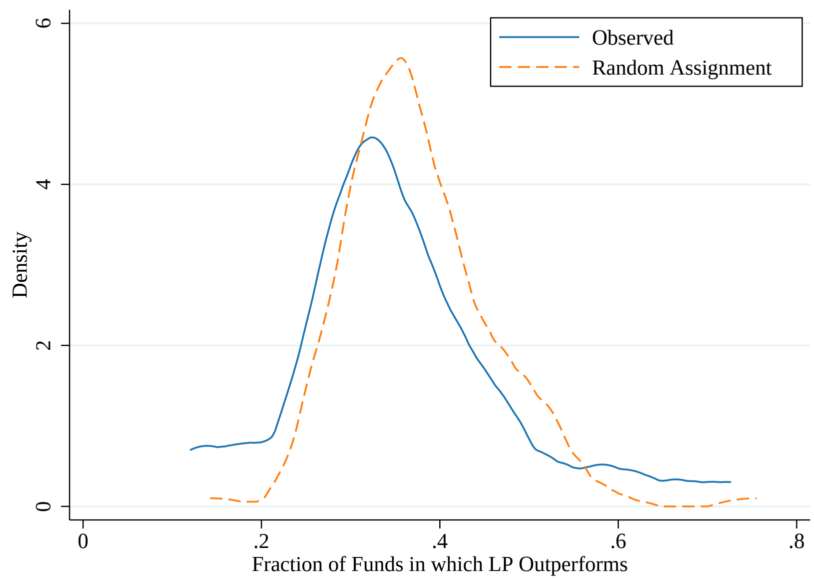

Panel B: Observed Data Before and After Characteristic-Adjustment

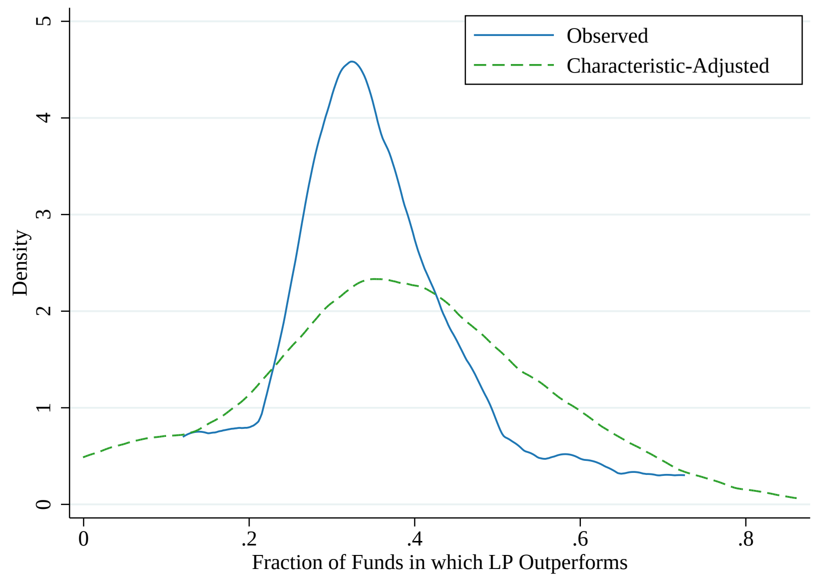

Notes: This plot shows the likelihood of being a top-tier investor across pensions ("pension effects"). In Panel A, pension effects are estimated by regressing pension p's status in fund f, y_{pf}, on pension fixed effects and fixed effects based on the vingintiles of fund age. The estimated pension fixed effects are then shrunk towards their mean using an empirical Bayes estimate. The blue line in the plot shows the results of this procedure. The orange line in the top plot shows the pension effects that we obtain when simulating random assignment of contracts to pensions in each fund. 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 orange, alongside the pension effects before adjustments. Returns are measured using the reported IRRs in Preqin. See Sections C.7, 4.1.1 and 4.2.2 for more details.