Skill and Luck in Private Equity Performance

Korteweg, Arthur; Sørensen, Morten

Document Version

Accepted author manuscript

Published in:

Journal of Financial Economics

DOI:

[10.1016/j.jfineco.2017.03.006](https://doi.org/10.1016/j.jfineco.2017.03.006)

Publication date:

2017

License

CC BY-NC-ND

Citation for published version (APA):

Korteweg, A., & Sørensen, M. (2017). Skill and Luck in Private Equity Performance. *Journal of Financial Economics*, 124(3), 535-562. <https://doi.org/10.1016/j.jfineco.2017.03.006>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 26. Apr. 2023

Skill and Luck in Private Equity Performance

Arthur Korteweg and Morten Sørensen

Journal article (Accepted version)

CITE: Skill and Luck in Private Equity Performance. / Korteweg, Arthur; Sørensen, Morten. In: *Journal of Financial Economics*, Vol. 124, No. 3, 2017, p. 535-562.

DOI: [10.1016/j.jfineco.2017.03.006](https://doi.org/10.1016/j.jfineco.2017.03.006)

Uploaded to [Research@CBS](https://research.cbs.dk): January 2018

© 2017. This manuscript version is made available under the CC-BY-NC-ND 4.0 license
<http://creativecommons.org/licenses/by-nc-nd/4.0/>

Skill and luck in private equity performance[☆]

Arthur Korteweg^{a,*}, Morten Sorensen^b

^a*Marshall School of Business, University of Southern California, USA*

^b*Copenhagen Business School, Solbjerg Plads 3, 2000 Frederiksberg, Denmark*

A B S T R A C T

Private equity (PE) performance is persistent, with PE firms consistently producing high (or low) net-of-fees returns. We use a new variance decomposition model to isolate three components of persistence. We find high long-term persistence: the spread in expected net-of-fee future returns between top and bottom quartile PE firms is 7 to 8 percentage points annually. This spread is estimated controlling for spurious persistence, which arises mechanically from the overlap of contemporaneous funds. Performance is noisy, however, making it difficult for investors to identify the PE funds with top quartile expected future performance and leaving little investable persistence.

JEL classification: C11 D83 G11 G14 G23 G24

Keywords: Persistence, Private equity, Venture capital, Skill, Learning

[☆]We are grateful to the editor (Bill Schwert), an anonymous referee, Ulf Axelson, Andrea Buraschi, Jung Hyun Choi, Peter Cornelius, Wayne Ferson, Ulrich Hege, Chris Jones, Bing Liang, Yun Ling, Georgios Magkotsios, John Matsusaka, Andrew Metrick, Paul Pfleiderer, Matt Rhodes-Kropf, Berk Sensoy, Peter Shepard, Per Stromberg, Irene Yi, and seminar participants at Arizona State University, Columbia Business School, Cornerstone Research, HEC Paris, Hong Kong University, National University of Singapore, Syddansk Universitet, University of California at San Diego, University of Alberta, University of Virginia Darden School of Business, the 2013 Spring Journal of Investment Management (JOIM) conference on Private Equity, the 7th Private Equity Findings Symposium at London Business School, the Hong Kong University of Science and Technology Conference on Entrepreneurship and Finance, the 2014 University of North Carolina Private Equity Conference, the 2014 Argentum Private Equity Symposium, the 2015 American Finance Association meetings, the 2015 Financial Intermediation Research Society (FIRS) conference, the 2015 Four Nations Cup, and the Conference on Entrepreneurship and Finance in memory of Ola Bengtsson for helpful comments and feedback. Morten Sorensen gratefully acknowledges funding by the Danish Council for Independent Research (“Det Frie Forskningsrd”) under the Sapere Aude program through grant number: DFF-4003-00095.

*Corresponding author. Tel.: +1 213 740 0567; fax: +1 213 740 6650.

Email address: korteweg@marshall.usc.edu (A. Korteweg)

1. Introduction

A central question in finance is whether some investment managers persistently outperform and, if so, how rents are shared between these managers and their investors. Most empirical studies focus on mutual funds and hedge funds, with mixed results.¹ To the extent that mutual funds generate rents, Berk and Green (2004) argue that they should accrue fully to the manager. For private equity investments, e.g., in venture capital (VC) or buyout (BO) funds, there are several reasons why investors could extract a share of the rents. A private equity (PE) firm typically manages a sequence of funds that are raised several years apart, so capital cannot flow elastically across firms. Furthermore, realized returns are known only at the end of a fund's life, and the environment of high uncertainty and large information asymmetries could enable some investors in these funds (called limited partners or LPs) to learn about the PE firm's skill, resulting in persistence in these LPs' net-of-fee returns.

The seminal study by Kaplan and Schoar (2005) finds evidence of persistence in net-of-fee PE performance. They estimate the regression

$$y_{i,N} = \alpha + \beta \cdot y_{i,N-1} + \varepsilon_{i,N}, \quad (1)$$

where $y_{i,N}$ is the net-of-fee performance of fund number N managed by PE firm i . They find a positive and statistically significant β coefficient, which means that the performance of fund number $N - 1$ predicts the performance of fund N of the same PE firm. Their interpretation is that PE firms differ in their

¹Lack of persistence in investor returns is shown, for example, by Jensen (1968), Malkiel (1995), Gruber (1996), and Carhart (1997) for mutual funds, by Brown, Goetzmann, and Ibbotson (1999) and Griffin and Xu (2009) for hedge funds, by Timmermann and Blake (2005) and Busse, Goyal, and Wahal (2010) for institutional trading desks, by Graham and Harvey (1996) for investment newsletters, and by Barber and Odean (2000) for individual investors. The evidence against persistence in mutual and hedge funds is not unequivocal, however, and substantial heterogeneity exists across managers. See, for example, Baks, Metrick, and Wachter (2001), Kacperczyk, Sialm, and Zheng (2005), Busse and Irvine (2006), and Kojen (2014) for mutual funds, and Titman and Tiu (2011) and Jagannathan, Malakhov, and Novikov (2010) for hedge funds. For comprehensive reviews of the performance measurement literature for mutual and hedge funds, see Ferson (2010) and Wermers (2011).

skills and abilities and that LPs invested in funds that are managed by skilled PE firms persistently outperform, even-net-of-fees.²

Regression (1) is motivated by a cross-sectional intuition: Some funds outperform, and when these funds' predecessors also outperformed, this is evidence of performance persistence. Eq. (1) is a time series AR(1) model, however, and a tension exists between this cross-sectional intuition and the model's time series properties. In the AR(1) model, persistence is short term. The expected return of fund N depends only on the performance of fund $N - 1$, regardless of whether this past performance was due to skill or luck, and in the long run all PE funds have the same expected performance, $E[y] = \frac{\alpha}{1-\beta}$. Hence, the AR(1) model is a model of performance persistence that, by construction, does not allow for long-term performance differences, which is clearly undesirable.

To better capture the cross-sectional intuition, we analyze a variance decomposition model of PE performance. Taking the LP's perspective, performance persistence has three components. First, long-term persistence refers to the possibility that some PE firms generate consistently higher (or lower) expected returns (net of fees). LPs can outperform by investing in these skilled PE firms with high expected returns. Second, investable persistence reflects the difficulty of identifying the PE firms with higher expected returns. When performance is noisy, top quartile past performance could be due to luck and does not necessarily predict future top quartile performance. This noise makes it difficult for LPs to identify skilled PE firms, implying low investable

²Following Kaplan and Schoar (2005), several studies of PE performance persistence have developed different variations of the AR(1) model. Phalippou and Gottschalg (2009) consider persistence after correcting for biases in the reported interim net asset values (NAVs). Phalippou (2010) and Chung (2012) find weaker evidence for persistence when regressing $y_{i,N}$ on $y_{i,N-2}$ and argue that persistence is short-lived. Robinson and Sensoy (2016) find persistence in a more recent sample than the original Kaplan and Schoar study. Harris, Jenkinson, Kaplan, and Stucke (2014b) and Braun, Jenkinson, and Stoff (2016) find that persistence has declined for buyout firms post-2000. Li (2014) finds evidence of stronger persistence in buyout firms compared with VC. Hochberg, Ljungqvist, and Vissing-Jorgensen (2014) also add interim performance of the most recent fund to the AR(1) regression, and find that this helps predict performance of the next fund.

persistence. The third component, spurious persistence, arises from the partial overlap of consecutive funds that are managed by the same PE firm. Partially overlapping funds are exposed to the same market conditions during the overlap period. Even though these contemporaneous exposures are transitory, so that past performance does not predict future performance, they induce a positive correlation in the performance of subsequent overlapping funds that shows up as spurious persistence in the AR(1) model. The AR(1) model is constrained to capture the three persistence components with its single beta coefficient, which is too restrictive given their variation in the data. To the best of our knowledge, our analysis is the first to formally distinguish between these components of persistence and to quantify them separately.

We find substantial long-term persistence in expected net-of-fee returns. Top quartile PE firms have annual expected net-of-fee returns that are 7 to 8 percentage points higher than bottom quartile firms, on average, across all fund types. Performance is noisy, though, and we find little investable persistence, particularly for VC firms. VC performance is mostly driven by luck, and LPs need to observe an excessive number of past funds (25 or more) to identify VC firms with top quartile expected returns with reasonable certainty. Furthermore, we find a substantial amount of spurious persistence, accounting for 44% of the observed autocorrelation in PE firms' returns, on average across model specifications.

Comparing subsamples, smaller funds have greater long-term persistence and more investable persistence (higher signal-to-noise ratios) than larger funds, especially for VC firms. We find less long-term persistence for PE firms in the US, followed by Europe, and greatest persistence for PE firms in the rest of the world (ROW), which also have more volatile performance. We confirm the findings in Braun, Jenkinson, and Stoff (2016) and Harris, Jenkinson, Kaplan, and Stucke (2014b) that long-term persistence has declined in the 2000s, relative to the 1990s. This decline is largest for VC firms. Substantial persistence still exists for BO and other firms post-2000.

If LPs are risk neutral, so that they simply maximize expected returns

net-of-fees, then our results are directly applicable to their portfolio allocation decisions. For risk-averse LPs, we estimate a second set of risk-adjusted performance specifications that control for common risk exposures, such as systematic market risks. These specifications assume identical exposures in the cross section of PE funds of the same type and vintage year, but they allow for time-varying risk exposures and risk premiums across vintages. This follows the general PE literature [e.g., Driessen, Lin, and Phalippou (2012) and Hochberg and Rauh (2013)], in which the limitations of fund-level data do not allow for estimation of fund or firm-specific risk exposures. Our results suggest that there is long-term persistence even in risk-adjusted returns.

Our findings of large long-term persistence but little investable persistence have important implications. They can explain LPs' recent focus on obtaining more detailed information about PE firms and their past funds (such as firms' internal organization and culture, compensation structure, incentives, alignment of interests, internal processes, and deal sourcing) to help analyze past performance [see Ewens and Rhodes-Kropf (2015); Jenkinson, Jones, and Martinez (2016); Kuckertz, Kollmann, Rohm, and Middelberg (2015)]. The additional information is necessary for LPs to identify top PE firms, as past fund performance by itself is insufficient. This lends support to a new explanation for the puzzle that net-of-fee outperformance is not competed away, contrary to the prediction by Berk and Green (2004). Skilled PE firms are scarce, but LPs with the ability to identify these skilled firms can also be scarce [for example, Lerner, Schoar, and Wongsunwai (2007); Hochberg and Rauh (2013); Sensoy, Wang, and Weisbach (2014) present evidence that LPs have heterogeneous investment skills].³ Skilled LPs should therefore earn rents. Our results are also consistent with the model in Hochberg, Ljungqvist, and Vissing-Jorgensen (2014), in which

³As the mutual fund literature argues, manager skill is measured by gross-of-fee performance, possibly adjusted for the amount of capital under management [e.g., Jensen (1968); Malkiel (1995); Berk and Green (2004); Berk and Van Binsbergen (2015)]. Such skill is a necessary condition [though not sufficient, as Berk and Green (2004) argue] for persistence in net-of-fee returns. PE firms report only net-of-fee performance, and therefore we cannot quantify the existence of managerial skill directly.

net-of-fee performance persistence arises from a hold-up problem created by the asymmetric information between LPs and the PE firm, and the model by Garleanu and Pedersen (2016), in which search frictions generate net-of-fee persistence.⁴ Furthermore, our evidence helps explain empirical results in the recent literature on the strength of flow-performance relations in PE. For example, the low investable persistence of VC firms is consistent with their lower flow-performance sensitivity relative to buyout firms, as reported by Kaplan and Schoar (2005) and Chung, Sensoy, Stern, and Weisbach (2012). We also find that learning effects are strongest for smaller size funds of a given type, which tend to be funds of younger firms, consistent with the result in Chung, Sensoy, Stern, and Weisbach (2012) that funds with lower sequence numbers have stronger flow-performance sensitivity than later funds of the same firm.

Our empirical variance decomposition model, or hierarchical linear model, generalizes the classical analysis of variance (ANOVA) method to capture the particular features of PE performance. The model allows us to separate the three different forms of persistence (long-term, investable, and spurious persistence), and it differs from the AR(1) model in several ways. First, it explicitly models the timing of funds without relying on their numbering, which is important when funds are simultaneous, so that it is arbitrary which one is labeled N or $N - 1$. Second, it accounts for the overlap between funds and distinguishes situations in which fund N follows quickly after fund $N - 1$ from cases in which funds are many years apart. Third, it is robust to missing data for intermediate funds. Fourth, unlike the AR(1) regressions, our model can accommodate PE firms with only a single fund, which do not have a fund $N - 1$. These firms are typically worse firms, and they are excluded in the AR(1) model, which could introduce systematic biases. Our method also captures estimation error in the model parameters and the effects of parameter uncertainty on the LPs' inference problem of identifying skilled PE firms.

⁴Glode and Green (2011) present a similar model for hedge funds in which the source of asymmetry is knowledge about the PE manager's investment strategy. Marquez, Nanda, and Yavuz (2015) consider the information asymmetry between entrepreneurs and PE firms.

Econometrically, unlike standard ANOVA models, our model is semi-parametric, and it does not require fund returns to follow a normal distribution. We use a mixture-of-normals distribution and a Bayes factor test to determine the appropriate number of mixtures. This generalization is especially important for VC funds, whose returns are highly skewed. Our Bayesian estimation approach, which is described in detail in the Appendix, is computationally efficient, and it provides accurate small sample inference for the estimated parameters, which is important because the parameters of interest are variances and ratios of variances (in the case of signal-to-noise ratios), which have nonstandard asymptotic distributions.

One limitation of our analysis is that it focuses on cross-sectional heterogeneity in PE performance. We do not consider aggregate skill and performance, and we do not investigate whether PE in aggregate outperforms the market or other types of investments. For studies of aggregate PE performance, see, for example, Korteweg and Sorensen (2010), Ang, Chen, Goetzmann, and Phalippou (2014), Harris, Jenkinson, and Kaplan (2014a), and Korteweg and Nagel (2016).

Our analysis is related to the recent literature on performance persistence in mutual funds and hedge funds. Kosowski, Timmermann, Wermers, and White (2006) and Fama and French (2010) estimate manager skill fixed effects using a bootstrap method, which allows for flexible distributions of the return residuals, like our mixture-of-normals distribution. Barras, Scaillet, and Wermers (2010) and Ferson and Chen (2015) extend the bootstrap approach to the null hypothesis that only a fraction of managers have no skill, a generalization that is similar to our prior distribution of managers' skill. However, fixed effects skill estimates and bootstrapping are more applicable in settings with long panel data sets (i.e., with many observations per individual), such as mutual and hedge funds.

A main benefit of the Bayesian approach is its application to learning and investability, but it also has technical benefits for estimating variance decomposition models, for example, in enforcing non-negative variance

estimates. Bayesian approaches to skill assessment in mutual and hedge funds have been applied, amongst others, by Baks, Metrick, and Wachter (2001), Pastor and Stambaugh (2002a, 2002b), Jones and Shanken (2005), Avramov and Wermers (2006), and Busse and Irvine (2006). In particular, Jones and Shanken (2005) use a random effect prior on skill that is close to ours. However, these studies do not allow for non-normal residuals, and they do not explore the speed of learning about skill over time. To our knowledge, the variance decomposition approach we pursue here is new to the literature.

The problem of producing a skill estimate for a manager who simultaneously manages investments across several funds has not received much attention in the mutual and hedge fund literatures. Wu, Wermers, and Zechner (2016) show that multi-fund managers are common and that the average mutual fund manager manages 2.2 funds at any given time. Our method for controlling for overlap can therefore also be useful for mutual and hedge fund studies.

More generally, evaluating persistence and separating skill from luck are fundamental issues in economics and finance. For example, other applications of our method include the performance persistence of serial entrepreneurs [e.g., Gompers, Kovner, Lerner, and Scharfstein (2010) and Bengtsson (2013)] and studies of management style, such as Bertrand and Schoar (2003), which address the closely related issue of persistence in managers' business practices.

The paper proceeds as follows. In Section 2, we present the data. Section 3 describes our empirical model. Section 4 discusses our results and the evidence for long-term persistence in private equity performance. Section 5 evaluates investable persistence. Section 6 analyzes various subsamples of the data. Section 7 discusses robustness to using an alternative measure of fund returns, and Section 8 concludes.

2. Data

The analysis uses an extensive data set with information about PE firms and the funds they manage. The data are obtained from Preqin and include

fund-level information such as performance, size, type (e.g., VC or BO), and geography. Preqin is a commercial data provider that started collecting performance data using Freedom of Information Act requests to public investors and later extended the scope of its data collection to other public filings and to voluntary reports by some general partners (GPs) and LPs. For each fund, Preqin reports aggregate net-of-fees fund performance, such as the internal rate of return (IRR) and the total value to paid-in capital multiple (TVPI). No information about gross-of-fees performance is available. We also cannot calculate the public market equivalent (PME) measure of fund performance, which has some advantages when evaluating PE performance (Korteweg and Nagel (2016); Sorensen and Jagannathan (2015)), because the data do not contain cash flows between the GP and LPs.

We focus on the IRR, which is the annualized return to LPs net of performance fees (carried interest) and management fees. While it has well-known limitations, IRR is the most widely available fund performance measure, and it is commonly used to analyze PE performance.⁵ Though the IRR is an absolute performance measure, our model controls for general market performance and systematic risks. The analysis in Section 7 confirms our results using TVPI as an alternative measure of performance.

Harris, Jenkinson, and Kaplan (2014a) compare several data sets with PE fund performance. Most of these data are from commercial data providers (Preqin, Burgiss, and Cambridge Associates) and one data set is from a large anonymous LP that was also studied by Robinson and Sensoy (2016). For VC funds, Preqin has slightly weaker coverage for the 1980s and 1990s, but it provides the most comprehensive coverage for the 2000s. For BO funds, Harris,

⁵Other papers that use IRR to measure PE performance include Ljungqvist and Richardson (2003), Kaplan and Schoar (2005), Lerner, Schoar, and Wongsunwai (2007), Ljungqvist, Richardson, and Wolfenzon (2008), Kaplan and Stromberg (2009), Chung (2012), Franzoni, Nowak, and Phalippou (2012), Higson and Stucke (2012), Acharya, Gottschalg, Hahn, and Kehoe (2013), Caselli, Garcia-Appendini, and Ippolito (2013), Hochberg and Rauh (2013), Jenkinson, Sousa, and Stucke (2013), Lopez-de Silanes, Phalippou, and Gottschalg (2015), Robinson and Sensoy (2016), Harris, Jenkinson, and Kaplan (2014a), Hochberg, Ljungqvist, and Vissing-Jorgensen (2014), Li (2014), and Sensoy, Wang, and Weisbach (2014).

Jenkinson, and Kaplan find that Preqin contains the largest total number of funds for the 1990s and 2000s (but not for the 1980s). Importantly, the Preqin data have performance information for the largest number of both BO and VC funds. Moreover, they find no evidence that Preqin’s IRR performance data are biased relative to the performance data from the other data sources. Phalippou (2014) finds that Preqin is representative of the data sets from Burgiss and Cambridge Associates, though Phalippou considers only BO funds. Hence, when analyzing the performance and persistence of PE funds, Preqin is one of the best data sets available.

The two main fund types are VC and BO funds, but Preqin also classifies some funds as real estate, fund-of-funds, infrastructure, turnaround, special situations, co-investment, and venture debt funds, which we collectively refer to as Other funds. The majority of Other funds are real estate and funds-of-funds, and though the investments by these two fund types are very different, they have (surprisingly) similar performance and persistence characteristics. Therefore, we combine all Other funds for most of our analysis.

We define a fund’s location by the geographical location of its GP. This location can sometimes differ from the locations of its portfolio companies, but we obtain very similar results when we define location in terms of the fund’s geographical investment focus.

2.1. Descriptive statistics

We restrict our sample to funds with available performance information.⁶ Our model does not require a fund to be preceded by another fund, and

⁶Our estimates remain valid with randomly missing fund performance data. Moreover, survivorship does not bias our parameter estimates when each PE firm’s survival depends only on past observed data (i.e., performance), but not on the true parameter values, so observed data are a sufficient statistic for survivorship. Similar (or stronger) assumptions are common in studies of mutual fund performance persistence, such as, for example, Baks, Metrick, and Wachter (2001) and Pástor and Stambaugh (2002). Formally, it always holds that $p(\theta|data, survival) = [p(survival|data, \theta)/p(survival|data)] \cdot p(\theta|data)$, where θ contains the model parameters. The θ estimates are then valid when $p(\theta|data, survival) = p(\theta|data)$, which holds when the fraction in brackets equals one.

we include PE firms with only a single fund.⁷ To avoid concerns about funds' self-reported intermediate IRRs and NAVs [see Phalippou and Gottschalg (2009); Brown, Gredil, and Kaplan (2016); Jenkinson, Sousa, and Stucke (2013); Barber and Yasuda (2016)], we restrict our sample to fully liquidated funds. We eliminate funds with less than \$5 million in committed capital (in 1990 US dollars) to exclude smaller, idiosyncratic funds.

Table 1 shows descriptive statistics for our final sample. The sample contains 1,924 funds, raised between 1969 and 2001 and managed by 831 unique PE firms. Of these funds, 842 are VC funds (managed by 409 firms), 562 are BO funds (285 firms), and the remaining 518 funds (197 firms) are classified as Other funds.⁸ Fig. 1 shows the distribution of the number of funds per PE firm, by firm type. The average VC and BO firm manages about two funds during our sample period (2.6 for Other firms), but the median VC and BO firm manages just a single fund. Table 1 shows that the average (median) overlap between two funds managed by the same PE firm, which is important when assessing persistence, is 5.8 years (6.0 years) for VC and BO and 6.8 years (7.0 years) for Other.⁹

** INSERT TABLE 1 AND FIGURE 1 NEAR HERE **

The average (median) VC fund has \$207 million (\$110 million) of committed capital, compared with \$694 million (\$300 million) for BO and \$373 million (\$207 million) for Other funds. For VC and Other funds, the subclassifications in Panel B of Table 1 show that late-stage VC funds and distressed debt funds

⁷Firms with a single fund do not help identify persistence parameters directly. With a sample of only single fund firms, it would not be possible to distinguish whether the dispersion in fund performance were due to dispersion in long-term persistence, idiosyncratic variation, or some combination. Still, single fund firms are informative about the remaining model parameters. They contain information about time fixed effects and the total variance of performance. That said, we have confirmed that the results are robust to excluding single fund firms from the sample.

⁸The number of firms by type add up to more than the 831 unique firms in our sample, because some PE firms manage funds of several types.

⁹We report overlaps for fund pairs, and there can be more pairs than individual funds, as is the case for VC funds. To illustrate, a PE firm that raises a fund every second year manages five partially overlapping funds (raised in years 0, 2, 4, 6, and 8). These five funds form ten fund pairs.

tend to be larger, whereas early stage VC and natural resource funds tend to be smaller. There are no subclassifications of BO funds.

Turning to performance, Panel A of Table 2 shows that the average (median) fund IRR is 17.7% (8.6%) for VC funds, 16.9% (14.9%) for BO funds, and 13.9% (11.9%) for Other funds. Table 2 shows IRRs by vintage year and fund type, and Fig. 2 plots average IRRs. For VC funds, very strong performance is evident during the dot-com bubble in the late 1990s, with average IRRs as high as 45.2% annually, followed by a sharp drop after the dot-com bubble burst. Funds have ten-year lives, so funds with vintage years well before 2000 were exposed to this bubble and have lower performance. BO performance is more stable and shows a recovery toward the end of the sample period, relative to VC and Other funds. The performance of Other funds is even more stable, showing an earlier but more modest decline in the late 1990s followed by a comparable modest recovery.

** INSERT FIGURE 2 NEAR HERE **

** INSERT TABLE 2 NEAR HERE **

Our empirical analysis uses total log returns (i.e., continuously compounded returns) instead of annualized IRRs. The total log return for fund u of firm i is denoted y_{iu} . It is calculated by compounding the fund's IRR over its ten-year life:

$$y_{iu} = 10 \cdot \ln(1 + IRR_{iu}). \quad (2)$$

The natural logarithm reduces the skewness of the IRRs, and it allows us to decompose the total ten-year return into a sum of annual returns. The transformation fails for the two funds in our sample with IRRs of -100% (one is a vintage 2001 VC fund and the other is a 1998 BO fund), and these two funds are excluded in the empirical analysis, but our results are robust to including them with IRRs set equal to the first (lowest) percentile of the IRR distribution.

3. Variance decomposition model

Our empirical model is a hierarchical linear model, which generalizes the classical analysis-of-variance (ANOVA) decomposition. Hierarchical linear models have been widely used, for example, in education research, because they capture the hierarchical structure that arises when, for example, one observes individual students, who are grouped into classrooms, which are grouped into different schools, which are located in different districts, etc. Such a hierarchical structure also arises for PE when individual PE funds are managed by different PE firms and span different time periods. Although not pursued here, our model could be extended to include data at additional levels, for example, data for individual deals, as in Braun, Jenkinson, and Stoff (2016), or data with LPs' holdings of PE funds, as in Sensoy, Wang, and Weisbach (2014). General overviews and discussions of hierarchical models are in Raudenbush and Bryk (2002) and De Leeuw, Meijer, and Goldstein (2008).

3.1. *Economic intuition*

To illustrate the intuition behind the variance decomposition in short panels, consider a set of 60 PE firms. Each firm makes just two investments (or manages two funds), each of which succeeds or fails. Of the resulting 120 investments, we observe that one-half fails and the other half succeeds, and the unconditional success probability is thus 50%. If the investments were statistically independent, each firm would have a 25% probability of zero, a 50% probability of a single, and a 25% probability of two successful investments, and we should see 15 of the 60 PE firms with no successes, 30 with a single success, and the remaining 15 firms with two successful investments. But imagine instead that the successes are evenly distributed among the 60 PE firms, so 20 have zero, 20 have one, and 20 PE firms have two successful investments. In this case, the performance variation between PE firms exceeds the variation that is implied by the variation within PE firms if they were

statistically independent, so the investments cannot be independent.¹⁰ Some PE firms have higher (and some have lower) success probabilities. In this example, the even distribution of successes among PE firms is consistent with each firm's success probability being drawn from the uniform distribution on $[0, 1]$. If $p_i \sim U[0, 1]$ denotes firm i 's success probability, then the expected probability of two successes is $E[p_i^2] = 33\%$. Following this intuition, our model defines and measures persistence by comparing performance variation within PE firms with performance variation between firms. Excess variation between firms, as in this example, implies persistence.

Excess variation, and hence persistence, also implies that the outcomes of the investments made by a PE firm are positively correlated. In the example, let s_1 and s_2 be indicators for the outcomes of the two investments by a given PE firm. When $p_i \sim U[0, 1]$, then $\text{cov}(s_1, s_2) = E[s_1 s_2] - E[s_1] E[s_2] = 1/12$.

This intuition leads to a natural distinction between PE firms' past performance and expected future performance. With $p_i \sim U[0, 1]$ and using Bayes's rule, the posterior density of p_i conditional on observing two past successes is $f(p_i | s_1 = 1, s_2 = 1) = 3p_i^2$. Some firms with two successes are mediocre firms that have been lucky, and the probability that a firm with top tercile past performance (two successes) has top tercile expected future performance is only $\Pr(p_i \in [0.66, 1] | s_1 = 1, s_2 = 1) = 70\%$.¹¹ The success probability of a new investment by firms with top tercile past performance is $E[p_i | s_1 = 1, s_2 = 1] = 75\%$, and the success probability for firms with actual top tercile expected performance is $E[p_i | p_i > 66\%] = 83\%$. To summarize, in short panels, it is not possible to determine precisely which firms have high or

¹⁰Here we ignore the possibility that the number of successes are evenly distributed across firms due to random sampling. Our argument is an identification argument, and it should be interpreted in the limit as the number of firms tends to infinity. Put more simply, nothing changes if one replaces our 60 firms making a total of 120 investments with, e.g., 60 million firms making a total of 120 million investments. What is important is that each firm makes only a small number of investments (funds).

¹¹In this particular example, in which a third of the funds have two successes, a third have a single success, and a third have none, it is natural to consider top tercile performance. In practice, it is more common to consider top quartile performance.

low p_i , but the dispersion in p_i can still be determined.

3.2. Comparing random and fixed effects

The previous example essentially describes a random effects model in which the random effect is each firm’s success probability, p_i . The example illustrates some advantages of the random effects model over the more typical fixed effects model in short panels (i.e., panels with few observed outcomes for each individual firm) with large idiosyncratic error terms. In these short panels, fixed effects are estimated with substantial standard errors, making it difficult to interpret the dispersion in the estimated fixed effects, because it confounds the true variation in the fixed effects, which is due to skill differences between firms, with the standard errors of the fixed effects. To illustrate, continuing the example, let f_i be the estimated fixed effect for PE firm i , so PE firms with two successes have $f_i = 1$, firms with a single success have $f_i = \frac{1}{2}$, and firms with no successes have $f_i = 0$. Because p_i varies uniformly between zero and one, while there are equal numbers of firms with $f_i = 0, \frac{1}{2}, 1$, these fixed effects are clearly noisy estimates of the true underlying p_i . The variance in the estimated fixed effects, $\text{var}[f_i] = \frac{1}{6}$, is twice the variance due to actual long-term persistence, which is $\text{var}[p_i] = \frac{1}{12}$.¹² Random effects models avoid this problem by explicitly parameterizing and estimating $\text{var}[p_i]$, at the cost of more restrictive assumptions (primarily distributional assumptions and independence between the random effects and the covariates).

A second advantage of the random effects model is that the model and its estimates are directly applicable to the LP’s inference problem and, hence, provide a consistent way to evaluate the speed of learning and investable persistence.

¹²The literature about this problem [dating back to Stein (1956) and Lindley (1962)] has developed a number of “shrinkage” methods that reduce the estimated dispersion in the fixed effects, under various assumptions, to isolate the dispersion that is due to actual cross-sectional differences. Shrinkage methods are effectively ways to pull estimates from fixed effect models closer to the corresponding random effects estimates [see, for example, Efron and Morris (1975)].

3.3. Full model

Our full model generalizes the above example. The main outcome variable is fund performance, not simply success or failure. The model allows for overlapping funds and includes explanatory covariates. The basic intuition, however, remains the same.

The model, which is estimated separately for each fund type (VC, BO, Other) is set up as follows. Let PE firms be indexed by i . Each firm manages a sequence of funds indexed by u . Each observation contains the performance of a fund and characteristics of the fund and firm. The ten-year total log return of fund u managed by firm i is

$$y_{iu} = X'_{iu}\beta + \sum_{\tau=t_{iu}}^{t_{iu}+9} (\gamma_i + \eta_{i\tau}) + \varepsilon_{iu}. \quad (3)$$

The summation runs over the fund's ten-year life, with year t_{iu} denoting the fund's first year of operation (its vintage year). X_{iu} is a vector of observed fund-specific covariates and can include fund or firm characteristics, vintage year fixed effects, etc.

The three random effects that determine the covariance structure are γ_i , $\eta_{i\tau}$, and ε_{iu} . The γ_i term captures long-term persistence, similar to the success probability, p_i , in the example above. At birth, each PE firm receives an independent draw of γ_i , distributed $\gamma_i \sim \mathcal{N}(0, \sigma_\gamma^2)$, which remains constant for all funds managed throughout the life of the firm.¹³ Funds of a PE firm with higher γ_i have persistently higher expected returns (corresponding to a higher success probability, p_i , in the example). The variation in γ_i across PE firms reflects differences in expected returns across PE firms. When there is little variation in γ_i (i.e., σ_γ^2 is small), then PE firms are similar, and there are few

¹³In PE there are too few funds for each firm to allow γ to vary throughout the life of a PE firm. In Section 6, we divide the sample period into early and late subsamples, which permits γ to vary across these subsamples and somewhat mitigates this problem. A constant γ is the best case for the LP's inference problem of determining γ , and this case thus represents an upper bound on the investable persistence in PE.

long-term persistent differences in their performance. When σ_γ^2 is larger, more of the performance differences are due to heterogeneity in expected returns across PE firms. Without loss of generality, the mean of γ_i (along with the means of the other random effects) is normalized to zero, and the industry level of expected returns is captured by the intercept term in X_{iu} . The model is parameterized with γ_i inside the sum in Eq. (3), so γ_i is the annualized abnormal performance for firm i relative to its peers, and each fund earns γ_i ten times during its life.

The covariance in the returns of partially overlapping PE funds is captured by the PE firm-year-specific effects, independent and identically distributed (i.i.d.) $\eta_{i\tau} \sim \mathcal{N}(0, \sigma_\eta^2)$. Two overlapping funds that are managed by the same PE firm share an $\eta_{i\tau}$ term for each year of overlap, and these terms generate short-term correlations in the overlapping funds' performances. Economically, such correlations can arise due to common strategies, common risk exposures, or common investments in the same portfolio companies.¹⁴ To illustrate, suppose a PE firm manages two funds with vintage years 1999 and 2001, both focused on investments in emerging markets. From 2001 to 2009, these two funds would then have similar emerging market exposures. Due to these common exposures, an AR(1) regression of $y_{i,N}$ on $y_{i,N-1}$ would yield a positive coefficient, but this coefficient would not be evidence of actual persistence, and it does not imply that past performance predicts future performance. When these spurious correlations are large, the estimated σ_η^2 is large. Conversely, when there is no effect of shared exposures, the variance tends to zero. The model accommodates all of these cases.¹⁵

The error term ε_{iu} captures fund-specific idiosyncratic performance shocks. It is i.i.d. across funds, across firms, and over time. Because fund performance

¹⁴Braun, Jenkinson, and Stoff (2016) find that an average BO fund with 15.6 investments contains 1.0 deals that are common with another fund that is managed by the same PE firm. They do not have data for VC or Other firms.

¹⁵We also estimate specifications that account for overlap only during the first five years of the funds' lives, to focus on the correlation that is due to the investment decisions that are made in these initial years. These specifications give similar estimates, and the main results and conclusions remain unchanged.

is skewed, ε_{iu} is modeled using a mixture-of-normals distribution, which is considerably more general than the normal distribution. This flexibility is particularly important for VC performance, which we find requires a mixture of three normals, whereas the performance of BO and Other funds is captured by mixtures of just one or two normal distributions. Details are in the Appendix.

The total variance of y_{iu} is the sum of the variances of the three random effects. The summation in Eq. (3) contains the same γ_i term ten times, and it contains ten i.i.d. $\eta_{i\tau}$ terms, so the total variance is

$$\sigma_y^2 = 100\sigma_\gamma^2 + 10\sigma_\eta^2 + \sigma_\varepsilon^2. \quad (4)$$

The covariance between two funds that are managed by the same PE firm and have N years of overlap is

$$\text{cov}(y_{iu}, y_{iv}) = 100\sigma_\gamma^2 + N\sigma_\eta^2. \quad (5)$$

This covariance relation is plotted in Fig. 3, and this figure also illustrates the identification of the model. The parameters of interest are the variances of the three random effects, σ_γ^2 , σ_η^2 , and σ_ε^2 . In Fig. 3, the intercept is σ_γ^2 and the slope is σ_η^2 , so these two variances are identified by comparing the covariances of funds with increasing amounts of overlap. Given σ_γ^2 and σ_η^2 and observing the total variance, σ_y^2 , the remaining σ_ε^2 is identified as the residual variance in Eq. (4). Thus, although the model cannot determine when exactly returns are earned during a fund's life (as we use only each fund's total return, when liquidated), it can determine how much of the variation in funds' total performance is due to each of the three random effects.

**** FIGURE 3: OVERLAP AND COVARIANCE ****

With respect to the vector of covariates, X_{iu} , we estimate two specifications. In Specification I, X_{iu} contains only a constant term. All correlations in contemporaneous performance, including correlations due to common exposures to systematic risk factors that are shared across all firms, are then captured

by the $\eta_{i\tau}$ terms. To explicitly control for these risk exposures, Specification II adds vintage year fixed effects to X_{iu} (the model is then formally a mixed-effects model). All funds with the same vintage year experience the same factor returns (e.g., total market returns) over their ten-year lives. If these funds have the same exposures to these factors [e.g., the same beta in the capital asset pricing model (CAPM) or the same loadings on the Fama and French risk factors], then vintage year fixed effects capture the common performance component that is due to these exposures. Under this assumption, γ_i then represents the PE firm’s risk-adjusted annualized expected return relative to other PE firms (as it is centered at zero). It does not, however, show performance relative to the market overall, and it cannot tell whether PE outperforms in the aggregate.

Apart from the assumption that PE firms have the same contemporaneous factor exposures, this risk-adjustment is general.¹⁶ Each vintage year has a separate fixed effect, so risk exposures and factor premiums can vary over time, for example, due to trends in leverage and credit market conditions. The exposures can also vary by fund type, because we estimate the model separately for VC, BO, and Other funds. In Section 6, we also estimate the model separately for finer subsamples.

3.4. *Benefits of the Bayesian estimator*

We estimate our model using a Bayesian estimator that exploits recent advances in numerical computing [Markov chain Monte Carlo (MCMC), Gibbs sampling, and posterior augmentation]. The estimation procedure is described in detail in the Appendix. Essentially, our random effects model imposes a particular structure on the covariance matrix. While this model could, in principle, be estimated by classical maximum likelihood estimation (MLE) or the minimum-norm quadratic unbiased estimation (MINQUE) of Rao (1972), the Bayesian estimator offers several advantages. When the parameters of

¹⁶We cannot estimate fund-level risk loadings because we observe only one return per fund. Even firm-level risk loadings are not feasible given the typical number of funds per firm. The risk-adjustment issue is pervasive in the PE literature, and vintage-year fixed effects are also used to control for systematic risk, for example, by Hochberg and Rauh (2013).

interest are variances and functions of variances, classical estimators have difficulty imposing non-negativity constraints on these parameters. It is also difficult to derive the asymptotic distributions of variance parameters and (nonlinear) functions of these parameters. The Bayesian estimator solves these problems. In addition, it allows for correct small sample inference, and it explicitly accounts for parameter uncertainty in the estimated parameters. Finally, allowing for non-normal distributions is important in PE data, but mixtures of normals models are difficult to estimate using MLE due to the label-switching problem. The standard solution is to use an expectation maximization (EM) algorithm, which uses data augmentation in a similar fashion as the Bayesian approach. However, the EM algorithm is complex and does not resolve the other issues with MLE-type estimators.

4. Results

4.1. *IRR regressions*

We first confirm the findings by Kaplan and Schoar (2005) using our data. Table 3 (which follows the layout of Table VII in Kaplan and Schoar) shows ordinary least squares (OLS) regressions of $IRR_{i,N}$ on $IRR_{i,N-1}$ with various controls, including further performance lags in some specifications. All regressions have vintage year fixed effects and standard errors are clustered at the firm level. In most specifications, the previous fund’s performance strongly predicts the performance of the next fund. For example, the coefficient of 0.162 in the first specification shows that a fund with a 1% higher IRR predicts a 0.162% higher IRR for the next fund. The second specification suggests that this effect is even stronger when controlling for the performance of fund $N - 2$, although the coefficient on this second fund’s performance is negative (albeit insignificant). These results are robust to controlling for the fund’s (log) size and sequence number. The final six specifications in Table 3 show separate estimates of the AR(1) model for each fund type. When estimated separately, the VC results are similar to the full sample results, but the BO effects become

slightly stronger. For Other funds, the coefficient is positive and significant in Specification I but becomes smaller and insignificant when fund $N - 2$ is included, although this weaker result could be due to the smaller sample size.

** TABLE 3: IRR REGRESSIONS **

In these specifications, even fund $N - 2$ can overlap with fund N , and the positive coefficients can partly reflect the overlap instead of actual performance persistence. Panel B of Table 3 shows estimates using a sample of only the funds that are entirely nonoverlapping. This restriction further reduces the sample size and leaves no remaining signs of persistence. However, this weaker result could be due to the lower statistical power of the smaller sample.

A natural interpretation of the AR(1) results in Table 3 is that BO funds have more persistence (they have the largest coefficients and R^2 , and the coefficient remains statistically significant with fund $N - 2$), followed by VC funds (smaller coefficients and R^2 than BO funds, but still significant with fund $N - 2$), and that Other funds have the least persistence (the smallest coefficients and R^2 , and insignificant with fund $N - 2$). This ranking of BO, VC, and Other funds becomes more nuanced when we distinguish between the different forms of persistence using our model

4.2. Long-term persistence

Table 4 reports estimates of our model for VC, BO, and Other funds. Panel A shows the magnitudes of the three random effects as measured by their standard deviations (σ_γ , σ_η , and σ_ε).¹⁷ The variances ($100 \times \sigma_\gamma^2$, $10 \times \sigma_\eta^2$, σ_ε^2 and σ_y^2) are easier to interpret, and they are reported in Panel B. For each fund type,

¹⁷Though we use a Bayesian estimator, we discuss results using standard frequentist terminology: our point estimate is the mean of the posterior distribution, and our standard error is the standard deviation of the posterior distribution. A parameter is statistically significant, at a given level, when zero is not contained in the corresponding symmetric credible interval, as usually defined in Bayesian statistics. Our Bayesian estimator produces exact small-sample inference, even for nonlinear transformations of the estimated parameters, and all reported inference is calculated without asymptotic approximations.

we use two specifications. In Specification I, X_{iu} contains only an intercept. Specification II includes vintage year fixed effects.

** TABLE 4: PARAMETER ESTIMATES **

For VC funds, Specification I in Table 4 shows a total unconditional variance (σ_y^2) of 6.933. This variance is decomposed into three components, with 0.243 due to long-term persistence ($100 \times \sigma_\gamma^2$), 0.675 due to the overlap effect ($10 \times \sigma_\eta^2$), and the remaining 6.015 due to idiosyncratic variance (σ_ε^2).

We find that σ_γ^2 is statistically significant, consistent with the findings from the AR(1) regressions.¹⁸ Denote the x th percentile of the γ_i distribution by $q_\gamma(x)$. With the point estimate of σ_γ of 0.049, $q_\gamma(75\%) = 3.30\%$, implying that the marginal top quartile VC firm has an expected annual return that is 3.30% higher than the average VC firm (net-of-fees). The spread in the expected returns of the marginal top and bottom quartile firms is $q_\gamma(75\%) - q_\gamma(25\%) = 6.60\%$, annually. These percentiles are calculated from the point estimate of the standard deviation, assuming it is perfectly estimated, but Specification I in Table 4 shows that the estimate of σ_γ for VC funds has a standard error of 0.007. Because our estimation procedure gives the full posterior distribution of σ_γ , we can generate the corresponding posterior distribution of the spread $q_\gamma(75\%) - q_\gamma(25\%)$, accounting for parameter uncertainty. Table 4 reports that this spread estimate is 6.59%, close to the 6.60% spread calculated from the point estimate of σ_γ , which suggests that parameter uncertainty is a relatively small problem for these spreads. Nevertheless, because the calculation is simple, all reported gamma spreads in Panel C of Table 4 account for parameter uncertainty. The table also reports the spread in expected returns between the median top and bottom quartile PE firms, $q_\gamma(87.5\%) - q_\gamma(12.5\%)$, which is estimated to be 11.24% annually.

These gamma spreads cannot be calculated by simply subtracting the

¹⁸Testing statistical significance of variance parameters is complicated by the one-sided alternative hypothesis. We use a Bayes factor test to test $H_0 : \sigma_\gamma^2 = 0$ against $H_A : \sigma_\gamma^2 > 0$, as discussed in the Appendix.

observed IRRs of bottom quartile funds from those of top quartile funds. Top quartile performance does not imply top quartile expected returns, and this empirical difference confounds long-term persistence and noise. To illustrate, consider the case in which performance is noisy, so σ_ε^2 is large, but little difference exists between PE firms, so σ_γ^2 is small. In this case, there is little long-term persistence and the spread $q_\gamma(75\%) - q_\gamma(25\%)$ is small. Noisy performance, however, still leads to large differences in observed fund IRRs, so the empirically observed difference between top and bottom quartile funds can still be large, but this is entirely due to noise. Conversely, the empirical difference could understate actual long-term persistence. In periods when many particularly high-quality (or low-quality) PE firms are active, the empirical difference could be too small, because it is calculated from funds in a narrower range of the γ_i distribution.

Panel B in Table 4 also shows that for BO funds the variance that is due to long-term persistence ($100 \times \sigma_\gamma^2$) is 0.361, which is higher than for VC funds. Hence, the gamma spreads are also higher, and for BO firms $q_\gamma(75\%) - q_\gamma(25\%) = 8.03\%$ and $q_\gamma(87.5\%) - q_\gamma(12.5\%) = 13.70\%$ annually. The variance due to the overlap ($10 \times \sigma_\eta^2$) is 0.216, which is smaller than for VC funds, but the difference disappears when adding vintage year fixed effects to the model. Importantly, BO funds have substantially smaller idiosyncratic risk (σ_ε^2) and are less noisy than VC funds. The smaller noise in the returns of BO funds can also explain the stronger persistence results for BO in the AR(1) regressions, as less noisy returns give stronger statistical power in the AR(1) model.

For Other funds, the overlap and long-term persistence effects are largely similar to those for VC and BO funds. Comparing Specification I for the various fund types, the σ_γ estimates for Other and VC funds are similar, so their gamma spreads are similar as well. Idiosyncratic volatility of Other funds, however, is significantly less than the volatilities of both VC and BO funds, resulting in a better signal-to-noise ratio for Other funds.

To summarize, long-term persistence, as measured by the gamma spreads,

is greatest for BO funds and slightly lower for VC and Other funds, but the difference in the gamma spread between the three fund types is modest. These gamma spreads are calculated from the population distribution across PE firms. To earn this spread, an LP must perfectly identify PE firms with the highest and lowest gammas. Hence, these gamma spreads are upper bounds on the outperformance LPs can earn by identifying PE firms with higher expected returns. We return to this issue in Section 5.

4.3. Overlap effects

The returns of overlapping funds are correlated. In Table 4, Panel B, Specifications I and II for VC funds show overlap effects of 0.675 and 0.386, respectively. Without vintage year fixed effects, the $\eta_{i\tau}$ terms capture all correlations between contemporaneous funds, including correlations arising from common exposures to the market (and other common risk factors) during the overlap period. To control for these shared exposures, Specification II includes vintage year fixed effects, which capture market-wide movements and imply that $\eta_{i\tau}$ terms capture deviations only from these market-wide movements that are specific to a PE firm (e.g., due to the PE firm making similar decisions across its contemporaneous funds). The resulting overlap effects are largest for BO funds and smallest for Other funds. Generally, the variation in performance that is due to the overlap effect exceeds the variation that is due to the difference in expected returns.

The estimates in Table 4 allow us to quantify the effect of overlap on the AR(1) regression coefficient for fund N on $N - 1$. This coefficient can be positive due to the overlap effect, even when there is no actual long-term persistence and past performance does not predict future performance. Table 1 shows an average overlap of subsequent funds of 5.8 to 6.8 years. Using Eq. (5), the estimates in Table 4 imply that funds with such average overlaps have total covariances of 0.37 to 0.64. But 25.8% to 61.7% of this covariance (with an average across specifications of 43.6%) is due to the overlap, suggesting that the AR(1) coefficients as a measure of performance predictability is biased upward

by 34% to 168%. The overlap effect for Other funds is smaller, implying a smaller upward bias in the AR(1) coefficients. Overall, when evaluating two overlapping funds from the same PE firm that have both outperformed, their shared outperformance is more likely due to their overlap than to the firm generating persistently higher returns.

5. Learning and investable persistence

The estimates from section 4 show substantial differences in long-term persistence across PE firms, but they do not show how difficult it is for LPs to identify PE firms with higher expected returns. We quantify this investable persistence in two ways. First, we estimate the speed of learning about γ_i using the signal-to-noise ratio. This ratio is simple to calculate, it allows for a direct comparison of different firm types, and it has a simple economic interpretation based on the updating of beliefs about γ_i . The drawback is that it relies on normal distributions. Second, we use the full mixtures-of-normals model to evaluate how many past funds an LP must observe to estimate a PE firm's expected return with reasonable certainty. We find that the signal-to-noise ratio is low, and it is difficult for LPs to identify PE firms with high expected returns based on their past performance. An LP needs to observe an excessive number of past funds to evaluate a firm's expected return with reasonable certainty. In practice, LPs therefore need additional information, such as detailed information about individual deals and individual partners associated with these deals or a firm's internal organization and culture, to inform their investment decisions (see Ewens and Rhodes-Kropf (2015); Jenkinson, Jones, and Martinez (2016); Kuckertz, Kollmann, Rohm, and Middelberg (2015)).¹⁹

¹⁹An interesting issue is that estimating the hierarchical model using individual deal performance (instead of aggregate fund performance) is not necessarily more informative. Expected returns (i.e., differences in means) are no better measured with higher frequency observations, as pointed out by Merton (1980). Moreover, the overlap problem still exists for individual deals. The optimal aggregation of information about fund managers' performance, combining high and low frequency performance measures, is an interesting topic to pursue.

5.1. Signal-to-noise

Our model contains two types of shocks. Transitory shocks are drawn independently each period and given by the $\eta_{i\tau}$ and ε_{iu} terms. Persistent shocks, given by γ_i , generate long-term differences in expected returns across PE firms. The signal-to-noise ratio, s_γ , is defined as the ratio of the variance of the persistent shock relative to total variance.²⁰

$$s_\gamma = \frac{100\sigma_\gamma^2}{\sigma_y^2} \quad (6)$$

This signal-to-noise ratio, which is bounded between zero and one, has a simple economic interpretation. In a Gaussian learning model, an LP updates beliefs about γ_i as follows. Let the LP's beliefs about γ_i after observing N funds be $\mathcal{N}(\gamma_{i,N}, \sigma_{i,N}^2)$. After observing the performance of one additional (non-overlapping) fund, the LP's updated beliefs are $\mathcal{N}(\gamma_{i,N+1}, \sigma_{i,N+1}^2)$, where

$$\gamma_{i,N+1} = s_\gamma \cdot \frac{y_{i,N+1} - X'_{i,N+1}\beta}{10} + (1 - s_\gamma) \cdot \gamma_{i,N} \quad (7)$$

and

$$\sigma_{i,N+1}^2 = (1 - s_\gamma) \cdot \sigma_{i,N}^2. \quad (8)$$

Eq. (7) gives the mean of the LP's updated beliefs, which is a combination of two terms, weighted by the signal-to-noise ratio. The fraction in the first term is the new information, specifically, the surprise performance of the new fund relative to its expected performance. The greater the signal-to-noise ratio, the more weight is placed on the new information and the faster the LP updates beliefs. This also follows from Eq. (8), which shows the dispersion in the LP's beliefs. A larger signal-to-noise ratio means that the dispersion declines faster as new information arrives, i.e., the LP learns faster. When the signal-to-noise ratio is low, new performance is largely uninformative about the PE firm's expected

²⁰For example, Cochrane (1988) uses a similar variance ratio to evaluate the persistence of GDP shocks.

return. Because it is difficult for the LP to infer γ_i , more weight is put on the prior belief, $\gamma_{i,N}$ [the second term in Eq. (7)], and the dispersion in beliefs in Eq. (8) declines only slowly.

Point estimates of s_γ are reported in Panel B of Table 4. Fig. 4 plots the posterior distributions of s_γ for VC, BO, and Other firms, with and without vintage year fixed effects. The signal-to-noise ratio is lowest for VC funds and is larger for BO and Other funds. For VC funds, the large amount of noise makes performance less informative. Other funds have the least amount of noise and the highest signal-to-noise ratio, making it is easier for LPs to identify which Other firms have higher expected returns.

** FIGURE 4: ESTIMATES OF SIGNAL-TO-NOISE RATIO **

5.2. Identifying PE firms

Using the full mixtures-of-normals model, Fig. 5 plots the probability that top quartile performance implies top quartile expected returns. This probability is $\Pr \left[\gamma_i \geq q_\gamma(75\%) \mid \frac{1}{N} \sum_{n=1}^N y_{i,n} \geq Q_N \right]$, where Q_N is the average observed performance of the marginal top quartile firm with N past funds. To interpret Fig. 5, consider the limit with little long-term persistence, where σ_γ^2 converges to zero. Top quartile performance is entirely due to luck, and the probability that a PE firm with top quartile performance also has top quartile expected returns is simply 25%. As long-term persistence increases and performance becomes informative, this probability increases. At the other limit, when σ_γ^2 becomes very large (relative to σ_y^2), persistence dominates. In this case, top quartile performance is entirely due to skill and perfectly identifies PE firms with top quartile expected returns, and the probability converges to 100%. In practice, for a realistic number of funds, the probability remains well below this limit. For example, Fig. 5 shows that for VC firms with five past funds, having top quartile past performance implies a 37% probability of also having top quartile expected future returns, which is slightly better than the 25% probability in the uninformative case. For BO and Other firms, with five past funds and top

quartile performance, this probability increases to 47% and 51%, respectively. These estimates are consistent with the signal-to-noise ratios, which also show that Other funds have the most informative performance, followed by BO and VC funds.

** FIGURE 5: LEARNING SPEED **

Fig. 5 shows the probabilities for up to 50 past funds. As Fig. 1 shows, no PE firm has managed even close to 50 funds, so this is an upper bound on the ability of LPs to discriminate between PE firms based on their past performance. Even at this upper bound, just 53% of the VC firms with top quartile past performance have top quartile expected future returns. Put differently, an LP that invests in a portfolio of funds selected only on PE firms with historical top quartile performance finds that only 53% of these firms continue to deliver top quartile performance in the future.

5.3. Investable persistence

The previous analysis shows that LPs find identifying Other funds with higher expected returns easier, but it does not account for the value of identifying these funds. For example, Table 4 shows that Other funds have less long-term persistence than BO funds, as measured by σ_γ^2 , which means that although it is more difficult to identify the best BO funds, there is more to gain from identifying them. Fig. 6 shows the combination of these two effects. This figure plots the expected gamma for PE firms with top quartile past performance, calculated among all PE firms with N past funds or, formally, $E\left[\gamma_i \mid \frac{1}{N} \sum_{n=1}^N y_{i,n} \geq Q_N\right]$. Initially, it is easier to identify Other firms, but the value of identifying these firms is limited by their lower gamma spreads. For BO firms, the gamma spread is larger, and after observing four to five funds, BO firms are sufficiently well identified that, relative to Other firms, the benefit of their greater gamma spreads outweighs the difficulty of identifying them. For PE firms with five or more funds, BO funds have the greatest investable persistence.

**** FIGURE 6: INVESTABLE PERSISTENCE ****

VC firms have less investable persistence. Their signal-to-noise ratio is low, so it is difficult to identify better VC firms. Moreover, the long-term persistence and gamma spreads are also lower. They are similar to Other firms and below the spread of BO firms, and the gains from identifying better VC firms are therefore also lower. This overall lower investable persistence of VC firms shows in Fig. 5 and 6, in which VC firms are below BO and Other firms. This result is also consistent with the empirical evidence that VC funds have lower flow performance sensitivities relative to BO funds [Kaplan and Schoar (2005) and Chung, Sensoy, Stern, and Weisbach (2012)], as good VC fund performance is more likely due to luck compared with BO and Other funds and therefore induces lower capital flows.

The reported probabilities and results in this section were calculated based on the assumption of nonoverlapping funds. All else equal, learning is slower and investable persistence is lower with overlapping funds, as some fraction of the correlation in subsequent fund returns is due to the overlap (η) not γ . This reinforces our main conclusion that investable persistence is low, based on past fund returns alone.

6. Subsamples

Table 5 reports model estimates for subsamples of the data. We estimate each subsample independently. For example, when we compare the persistence of small and large funds in Panel A of Table 5, the model is estimated twice, once for large funds and once for small funds. PE firms that manage both small and large funds are represented in both sets of estimates, possibly with different gammas. These narrower samples help alleviate concerns about heterogeneity in σ_γ^2 , σ_η^2 , σ_ε^2 , and risk factor loadings.²¹ The subsamples also provide a more

²¹The full sample results in Table 4 are robust to including fund size, sequence number, and indicators for location and fund type as covariates in X_{iu} . Even when these covariates

nuanced picture of persistence. Our general finding is that PE firms in more developed and competitive markets, such as larger and more recent funds located in the US, show less performance persistence.

** TABLE 5: SUBSAMPLES **

6.1. Fund size

Table 5 reports the long-term persistence of small and large funds. Panel A shows that small funds have more long-term persistence and greater gamma spreads than large funds, across VC, BO, and Other firms.²² This performance difference is not simply due to their greater volatility, which shows in the larger σ_γ^2 for small funds. Panel B of Table 5 also shows that their signal-to-noise ratios are higher, implying that the performance of small funds is more informative about the PE firm’s expected returns. Because fund size is strongly correlated with firm age within each fund type, this evidence is also consistent with the results in Chung, Sensoy, Stern, and Weisbach (2012) that funds with lower sequence numbers have a stronger flow-performance sensitivity than later funds of a given PE firm.

6.2. GP location

Table 5, Panel A, shows that PE firms located in the rest of the world have more long-term persistence, followed by firms in Europe, while US-based firms have the least persistence. Total volatility follows a different pattern, with ROW being the most volatile, followed by the US, and then European funds. Panel B shows that for VC and BO funds, the signal-to-noise ratio is substantially higher for European funds. For Other funds, those in ROW have the most informative performance. The performance of US-based funds is relatively less informative, which is consistent with the US PE industry being more mature.

are included in the model, the variances are still constant across the entire sample, implying that the various persistence measures are constant as well. To evaluate different degrees of persistence for different fund types, the variances must be allowed to vary across fund types, which is most easily achieved by estimating the model separately on subsamples of fund types.

²²The number of funds for the small and large fund subsamples does not add up to the total number of funds in Table 1, because fund size is not reported for 224 funds in the data set.

6.3. Investment style

VC and Other funds can be further categorized by their investment styles, as classified by Prequin. Table 5 shows that early stage VC funds have lower long-term performance persistence and the least informative performance. Generalist VC funds have the most long-term persistence, but late-stage funds have the most informative performance.

For Other funds, we separately analyze the two major classes of funds: real estate funds and fund-of-funds. These fund types are very different but have surprisingly similar persistence characteristics. Fund-of-funds have slightly greater long-term persistence than real estate funds, although their long-term persistence is still well below the levels of VC and BO funds. Real estate funds have more informative performance, however. In fact, the performance of real estate is more informative than the performance of both VC and BO funds. Consistent with this result, Chung, Sensoy, Stern, and Weisbach (2012) find that the flow-performance relation is stronger for real estate funds than for both VC and BO funds.

6.4. Time period

We confirm the findings by Braun, Jenkinson, and Stoff (2016) and Harris, Jenkinson, Kaplan, and Stucke (2014b) that persistence has been declining. Table 5 shows estimates for the early and late halves of our sample period. Panel A shows that long-term persistence has declined substantially across all fund types. Panel B shows that fund performance has also become less informative. This decline is particularly pronounced for VC funds, and it is more marginal for BO and Other funds.

7. Robustness: total value to paid-in capital

Another commonly used measure of fund performance is TVPI, which is the total capital returned to LPs (net of fees) divided by the capital invested in the

fund, without accounting for the timing of the cash flows.²³ The TVPI is less prone to produce outliers in situations in which capital is quickly returned to LPs, which can distort the IRR calculation. Another advantage is that, unlike the IRR, the TVPI always exists and is unique. TVPI is widely used in practice, and it is available for a slightly larger sample of funds than the IRR. In this section, we confirm that our results hold for the TVPI.

The distribution of TVPI for VC, BO, and Other funds is not surprising. VC and BO funds have an average TVPI of 2.0, and BO has a higher median (1.7, compared with 1.4 for VC). Other funds have lower TVPI, with an average (median) of 1.7 (1.5). VC funds have the highest standard deviation of TVPI, at 2.9, followed by BO (1.4) and Other (0.9) funds. The rank ordering of TVPI means, medians, and standard deviations across types is the same as for IRR. Our empirical specifications use $\ln(\text{TVPI})$ as the dependent variable. For brevity, we refer to this log measure as simply TVPI. The natural logarithm reduces the skewness of the TVPI and makes it more consistent with our model and easier to interpret. These are the same reasons for using the log IRR above.

Table 6 shows the model estimates using TVPI. Because the IRR and TVPI are scaled differently, the magnitudes of the estimated coefficients are not directly comparable, but the general patterns are broadly similar to the IRR results in Table 4. As before, σ_γ^2 is largely the same for VC, BO, and Other firms, meaning that the cross-sectional variance in performance due to long-term persistence is similar for these three types of firms. The most notable difference to the IRR results is that the TVPI estimates suggest that VC firms have slightly more long-term persistence than BO or Other firms. Spurious persistence estimates are similar across the three fund types (as they are in the IRR results), with the largest σ_η^2 in Specification I for VC and the smallest σ_η^2 in Specification II for Other firms. Finally, as for the IRR estimates, the

²³For nonliquidated funds, the remaining NAV is typically added to the returned capital in the numerator of the TVPI, which raises some issues about adjusting the remaining NAV (which is gross of fees) for expected fees, to get the net-of-fee TVPI. Because we use only fully liquidated funds, this adjustment is irrelevant.

idiosyncratic variance is highest for VC funds, with σ_ε^2 about 2.3 times larger than for BO funds and about 4.6 times larger than for Other funds.

** INSERT TABLE 6 NEAR HERE **

The TVPI and IRR are available for slightly different samples of funds. The results are not affected by differences across these samples. We confirm in the Appendix that the TVPI and IRR models give similar results for the common subsample in which both measures are available.

Fig. 7 shows the speed of learning and investable persistence calculated from the TVPI coefficients. The speed of learning (in Panel A) is unchanged when performance is measured using TVPI, with Other firms having the fastest speed, followed by BO and VC firms. The results on investable persistence (Panel B) change slightly. Due to the higher long-term persistence for VC firms when calculated using TVPI, it is more valuable, albeit still slower, to learn about these firms. For about 15 VC funds, the precision in the LP’s inference about the firms’ gamma coefficient becomes sufficiently precise, and the investable persistence of VC overtakes that of BO and Other firms. In our sample, about 95% of VC firms manage five or fewer funds, so this reversal is of little practical relevance.

** INSERT FIGURE 7 NEAR HERE **

8. Conclusion

We decompose the persistence of private equity performance into long-term, investable, and spurious persistence. Across all types of PE firms, we find large amounts of long-term persistence. The spread in expected returns between top and bottom quartile PE firms is 7 to 8 percentage points, annually. In contrast, we find low investable persistence. Past performance is noisy, with a low signal-to-noise ratio, and LPs need to observe an excessive number of past funds to identify PE firms with higher expected future returns with reasonable certainty. Even after observing 50 past funds, only 53% to 61% of the PE firms

that have generated top quartile past performance can be expected to continue to generate top quartile performance. In practice, to evaluate investments in PE funds, LPs need a substantial amount of information that goes well beyond just the performance of past funds.

We find that smaller funds have greater long-term persistence than larger funds. In particular, large VC funds have poor long-term and investable persistence. We find the least long-term persistence for PE firms located in the US, followed by Europe, and the greatest persistence for firms located in the rest of the world, although the ROW firms also have more volatile performance. Finally, we confirm the findings by Harris, Jenkinson, Kaplan, and Stucke (2014b) and Braun, Jenkinson, and Stoff (2016) that persistence has declined over our sample period. This decline is largest for VC firms, and BO and Other funds still show substantial long-term persistence, even post-2000.

Our results have practical implications. The low investable persistence can explain LPs' increasing focus on collecting detailed information about PE performance. For example, Ewens and Rhodes-Kropf (2015) study performance using deal- and partner-level information, and we find that such detailed information is necessary for LPs to evaluate PE investments. Our results provide a new explanation for why persistence is not competed away. When identifying top PE firms is as difficult as our results suggest, LPs with this ability can also be scarce, and these LPs should earn rents, resulting in persistence in net-of-fee returns.

Appendix A. Estimation procedure

We implement the model as a Bayesian multilevel hierarchical model, redefining the error terms to absorb the firm-specific random effects using hierarchical centering, as recommended by Gelfand, Sahu, and Carlin (1995). Matlab code for estimating the model is available on the authors' personal pages.

The performance of fund u of firm i is

$$y_{iu} = X_{iu}\beta + \sum_{\tau=t_{iu}}^{t_{iu}+9} \eta_{i\tau} + \varepsilon_{iu}. \quad (9)$$

The conditional distributions of the random effects are given as

$$\eta_{i\tau} | \gamma_i \sim \mathcal{N}(\gamma_i, \sigma_\eta^2) \quad (10)$$

and

$$\gamma_i \sim \mathcal{N}(0, \sigma_\gamma^2). \quad (11)$$

The fund-specific error term distribution is i.i.d. (we consider the mixture-of-normals specification below):

$$\varepsilon_{iu} \sim \mathcal{N}(0, \sigma_\varepsilon^2). \quad (12)$$

We need to estimate the parameter vector $\theta \equiv (\beta, \sigma_\gamma^2, \sigma_\eta^2, \sigma_\varepsilon^2)$, given a data set of fund returns, $\{y_{iu}\}$, the dates of inception and termination of each fund, and the set of observed fund-level covariates, X_{iu} . We augment the parameter vector with the latent γ s and η s, we and use a Bayesian estimation algorithm that produces a set of draws from the posterior distribution, $f(\theta, \{\gamma_i\}, \{\eta_{it}\} | data)$, using a Gibbs sampler [see Gelfand and Smith (1990) and Korteweg (2013) for a detailed description]. By the Hammersley-Clifford theorem, we can divide the posterior into five complete conditionals that are easy to sample from:

1. Latent firm-year random effects: $f(\{\eta_{it}\} | \{\gamma_i\}, \theta, data)$.
2. Variance of fund-specific error term and β coefficients:

$$f(\sigma_\varepsilon^2, \beta | \{\gamma_i\}, \{\eta_{it}\}, \sigma_\gamma^2, \sigma_\eta^2, data).$$

3. Latent firm random effects: $f(\{\gamma_i\} | \{\eta_{it}\}, \theta, data)$.
4. Variance of firm-year random effects: $f(\sigma_\eta^2 | \{\gamma_i\}, \{\eta_{it}\}, \beta, \sigma_\gamma^2, \sigma_\varepsilon^2, data)$.
5. Variance of firm random effects: $f(\sigma_\gamma^2 | \{\gamma_i\}, \{\eta_{it}\}, \beta, \sigma_\eta^2, \sigma_\varepsilon^2, data)$.

We sample from each distribution 1 through 5 in turn, conditional on the most recent draw of the other parameters. After Step 5, we return back to Step 1 and repeat. The resulting sequence of parameter draws forms a Markov chain, the stationary distribution of which is exactly the posterior distribution. Given a sample of draws of the posterior distribution, it is then straightforward to numerically integrate out the latent variables and obtain the marginal posterior of parameters, $f(\theta | data)$, or the distribution of the random effects, $f(\{\gamma_i\} | data)$ and $f(\{\eta_{it}\} | data)$. We now discuss how to draw from each conditional distribution.

A.1. Latent firm-year random effects

The firm-year random effects, η_{it} , are sampled using a Bayesian regression of the fund returns on a set of year indicator variables, with known variance. This is done on a firm-by-firm basis, as the random effects are assumed independent across firms (and time). For each firm, i , the regression model takes the form

$$y_i = X_i\beta + Z_i\eta_i + \varepsilon_i, \tag{13}$$

where y_i is a vector of stacked fund returns for the U_i funds of firm i , and X_i is the sub-matrix of the covariates $[X'_{i1} \dots X'_{iU_i}]'$ in which each row corresponds to a fund of firm i . The vector η_i contains the firm-year random effects for the years in which firm i has at least one active fund. The length of the vector η_i is denoted T_i and can vary by firm. The matrix Z_i is a $U_i \times T_i$ matrix of indicator variables. Each row represents a fund of firm i and contains ones in the columns that correspond to the years that the fund is active and zeros in all other columns.

Given the prior in Eq. (10), and using the standard Bayesian regression setup (e.g., Rossi, Allenby, and McCulloch (2005)), the posterior distribution is

$$\eta_i | \{\gamma_i\}, \theta, data \sim \mathcal{N}(\mu_\eta, \sigma_\varepsilon^2 \Omega^{-1}), \quad (14)$$

where

$$\Omega = \frac{\sigma_\varepsilon^2}{\sigma_\eta^2} \cdot \mathbb{I}_{T_i} + Z_i' Z_i \quad (15)$$

and

$$\mu_\eta = \Omega^{-1} \left(\gamma_i \cdot \frac{\sigma_\varepsilon^2}{\sigma_\eta^2} \cdot 1_{T_i} + Z_i' (y_i - X_i \beta) \right), \quad (16)$$

with $\mathbb{I}_{T_i}$ denoting the $T_i \times T_i$ identity matrix, and 1_{T_i} denoting a $T_i \times 1$ vector of ones.

A.2. Variance of fund-specific error term and β coefficients

Given the conditioning on the random effects, η_{it} , this step is a standard Bayesian regression. With the conjugate prior

$$\sigma_\varepsilon^2 \sim \mathcal{IG}(a_0, b_0) \quad (17)$$

and

$$\beta | \sigma_\varepsilon^2 \sim \mathcal{N}(\mu_0, \sigma_\varepsilon^2 \Sigma_0^{-1}), \quad (18)$$

the posterior distribution is

$$\sigma_\varepsilon^2 | \{\eta_i\}, data \sim \mathcal{IG}(a, b) \quad (19)$$

and

$$\beta | \sigma_\varepsilon^2, \{\eta_i\}, data \sim \mathcal{N}(\mu, \sigma_\varepsilon^2 \Sigma^{-1}), \quad (20)$$

where

$$a = a_0 + \sum_{i=1}^N U_i, \quad (21)$$

$$b = b_0 + e'e + (\mu - \mu_0) \Sigma_0 (\mu - \mu_0), \quad (22)$$

$$\Sigma = \Sigma_0 + X'X, \quad (23)$$

and

$$\mu = \Sigma^{-1} \cdot (\Sigma_0 \mu_0 + X'(y - Z\eta)). \quad (24)$$

The vector $y = [y'_1 \dots y'_N]'$ contains the fund returns stacked across the N firms, X is the matrix of stacked X_i , and Z is the stacked Z_i . The vector $e = y - Z\eta - X\mu$ contains the stacked error terms, using the stacked η vector from the current iteration.

A.3. Latent firm random effects

Drawing the firm random effects, γ_i , is similar in spirit to simulating the firm-year random effects in Step 1. Write the estimation problem as a regression of the firm-year random effects on a set of indicator variables

$$\eta = W\gamma + \nu, \quad (25)$$

where $\eta = [\eta_1 \dots \eta_N]'$, $\gamma = [\gamma_1 \dots \gamma_N]'$, and $\nu \sim \mathcal{N}(0, \sigma_\eta^2 \cdot \mathbb{I}_N)$. The matrix W is a $\sum_{i=1}^N T_i \times N$ matrix of indicator variables. Each row of W represents a firm-year and contains a one in the column of the corresponding firm, and zeros in all other columns.

With the prior in Eq. (11), the posterior distribution is

$$\gamma | \{\eta_{it}\}, \theta, data \sim \mathcal{N}(\mu_\gamma, \sigma_\eta^2 A^{-1}), \quad (26)$$

where

$$A = \frac{\sigma_\eta^2}{\sigma_\gamma^2} \cdot \mathbb{I}_N + W'W \quad (27)$$

and

$$\mu_\gamma = A^{-1} (W'\eta). \quad (28)$$

A.4. Variance of firm-year random effects

The variance of the firm-year random effects, σ_η^2 , is the variance of the residuals $v = \eta - W\gamma$ from the regression in Step 3. Using the inverse gamma prior

$$\sigma_\eta^2 \sim \mathcal{IG}(c_0, d_0) \quad (29)$$

yields the posterior distribution

$$\sigma_\eta^2 | \{\gamma_i\}, \{\eta_i\}, data \sim \mathcal{IG} \left(c_0 + \sum_{i=1}^N T_i, d_0 + v'v \right). \quad (30)$$

A.5. Variance of firm random effects

The variance of the firm random effects, σ_γ^2 , using the inverse gamma prior

$$\sigma_\gamma^2 \sim \mathcal{IG}(f_0, g_0) \quad (31)$$

has posterior distribution

$$\sigma_\gamma^2 | \{\gamma_i\}, data \sim \mathcal{IG}(f_0 + N, g_0 + \gamma'\gamma). \quad (32)$$

A.6. Mixture-of-normals specification

For the mixture-of-normals specification we replace the distribution of the fund-specific error term in Eq. (12) with a mixture of K normal distributions,

$$\varepsilon_{iu} \sim \sum_{k=1}^K p_k \cdot \mathcal{N}(\mu_k, \sigma_{\varepsilon,k}^2). \quad (33)$$

Setting $K = 1$ reduces the model to the baseline normal specification in Eq. (12). We drop the intercept in X_{iu} because it is absorbed by the error term, which has unrestricted mean $E[\varepsilon_{iu}] = \sum_{k=1}^K p_k \mu_k$. This specification is equivalent to the specification with an intercept in X_{iu} and zero mean ε_{iu} , but it is easier to implement because it avoids enforcing cross-parameter restrictions on μ_k . To estimate the mixture model by Gibbs sampler, the procedure requires one additional latent variable, $\mathbb{I}_{iu}$, that indicates which of the K component distributions each observation is drawn from. This indicator is sampled independently for each observation from a multinomial distribution with the probability of drawing component k proportional to

$$p_k \phi\left(\frac{e_{iu} - \mu_k}{\sigma_{\varepsilon,k}}\right), \quad (34)$$

where $e_{iu} = y_{iu} - Z_{iu}\eta - X_{iu}\beta$, and $\phi(\cdot)$ is the standard normal probability density function. The mixture probabilities, $p = [p_1 \dots p_K]$, have a Dirichlet prior distribution with concentration hyperparameter δ . The posterior distribution is also of the Dirichlet family, with concentration equal to $\delta + N$, with $N = [N_1 \dots N_K]$ and N_k the number of observations for which $\mathbb{I}_{iu} = k$.

The mean and variance of each mixture component have conjugate normal and inverse gamma prior distributions that are identical for each component,

$$\sigma_k^2 \sim IG(a_k, b_k) \quad (35)$$

and

$$\mu_k \sim \mathcal{N}(0, \sigma_k^2 V^{-1}). \quad (36)$$

Their posterior distributions are

$$\sigma_k^2 | \{\mathbb{I}_{iu}\}, p, \beta, \{\eta_{it}\}, data \sim IG\left(a_k + N_k, b_k + \sum_{j \in \{\mathbb{I}_{iu}=k\}} (e_j - \bar{e}_k)^2\right) \quad (37)$$

and

$$\mu_k | \sigma_k^2, \{\mathbb{I}_{iu}\}, p, \beta, \{\eta_{it}\}, data \sim \mathcal{N}\left(\frac{N_k \bar{e}_k}{V + N_k}, \frac{\sigma_k^2}{V + N_k}\right), \quad (38)$$

where $\bar{e}_k$ is the sample mean of the residuals e_{iu} computed over the set of observations for which $\mathbb{I}_{iu} = k$.

Conditional on the indicators, $\mathbb{I}_{iu}$, the Gibbs steps used to estimate the other model parameters remain largely unchanged from the description above. See West (1992), Diebolt and Robert (1994), and Chen and Liu (2000) for more details on the estimation of mixture distributions.

A.7. Priors and starting values

Our Gibbs sampler uses 10,000 iterations for the initial burn-in phase, followed by 100,000 iterations to simulate the posterior distribution. We save every 10th draw of the simulation. During the burn-in phase, the parameters converge quickly. We use diffuse prior distributions, so that our results are driven by the data, not prior assumptions. First, we set $a_0 = 2.1$ and $b_0 = 1$. This implies that our prior belief is that $E[\sigma_\varepsilon] = 0.854$ and that σ_ε is between 0.362 and 2.874 with 99% probability (this is for ten-year fund returns, so the annualized volatility is about a factor three lower). Second, we set $c_0 = f_0 = 2.1$ and $d_0 = g_0 = 0.15^2$. Because both the γ s and η s are specified at the annual level, this implies that $E[\sigma_\gamma] = E[\sigma_\eta] = 0.128$ per year and that σ_γ and σ_η are between 0.054 and 0.431 (annually) with 99% probability. Conditional on X , the prior ten-year fund return variance, $100\sigma_\gamma^2 + 10\sigma_\eta^2 + \sigma_\varepsilon^2$, has an expected value of 1.658 and is between 0.861 and 4.666 with 99% probability. Finally, we set the prior mean for β equal to zero ($\mu_0 = 0$), implying a prior mean fund return of zero. We set Σ_0 equal to the identity matrix, so that the prior β s are between -3.1 and +3.1 with 99% probability.

For the mixture-of-normals specifications, we set the prior of each mixture component, $1, \dots, K$, equal to the prior of the error term ε in the normal model, i.e., mean zero and inverse gamma prior parameters equal to a_0 and b_0 . This

ensures that the prior distribution of y is the same across all K , so that the Bayes factor (see Subsection A.8) is a valid comparison across different mixtures. We set $V^{-1} = 1/10$. For the prior distribution of the mixture probabilities, we use $\delta = 1_K \cdot 10$, which implies that all distributions in the mixture have equal prior mean probability, $1/K$.

We start the algorithm with all γ s and β s equal to zero (their prior means). We initialize all variances (σ_γ^2 , σ_η^2 , and σ_ε^2) at their prior means. For the mixtures specification, we set the mixture probabilities to their prior mean, $1/K$, and randomly draw the initial indicators from this distribution. The starting means and variances of the component distributions are also set equal to their prior means. We do not need starting values for the η s, as they are the first variables we simulate.

A.8. Hypothesis tests

We consider two sets of hypothesis tests for our model. The first set of tests determines the number of mixtures of normal distributions in ε_{iu} . This is a Bayes factor test and relies on the marginal log-likelihood, which integrates out all parameters from the likelihood function (Kass and Raftery (1995)). We use the method of Chib (1995) to compute the marginal log-likelihood from the MCMC output and the algorithm proposed by Berkhof, Van Mechelen, and Gelman (2003) and Marin and Robert (2008) to deal with the well-known label switching problem. Fig. A1 plots the marginal log-likelihood as a function of the number of mixture distributions, by type and model specification (with and without vintage year fixed effects). The optimal number of mixture distributions used in the main results of the paper (as reported in Table 4) are those with the highest marginal likelihood.

** INSERT FIGURE A1 NEAR HERE **

To test for the presence of the long-term persistence and overlap random effects, we use Bayes factors to test $H_0 : \sigma_\gamma^2 = 0$ against $H_A : \sigma_\gamma^2 > 0$ and $H_0 : \sigma_\eta^2 = 0$ against $H_A : \sigma_\eta^2 > 0$. Table A1 reports the Bayes factors for

each type and model specification. The long-term persistence random effect is significant for all types and models, except Other funds when vintage year fixed effects are included. The random effect overlap term is not significant in many cases, a perhaps surprising result given the posterior standard deviations on the σ_η estimates. Though it does not affect the conclusions that we draw from the paper, it does underscore the importance of performing proper hypothesis tests in the presence of small samples and non-normal distributions.

** INSERT TABLE A1 NEAR HERE **

A.9. IRR and TVPI parameter estimates in common sample

To determine whether differences in IRR and TVPI model estimates are driven by sample selection, we estimate the model on a sample of funds for which both return measures are observed (the joint sample). Table A2 shows that the parameter estimates for the joint sample are close to the full sample estimates of Tables 4 and 6 (for IRR and TVPI, respectively), across specifications and fund types. The most notable differences are in the σ_η estimates, which tend to be slightly lower in the joint sample, resulting in somewhat higher signal-to-noise estimates. This does not affect the conclusions of the paper. Moreover, the relative ranking of fund types remains unchanged.

** INSERT TABLE A2 NEAR HERE **

References

- Acharya, V. V., Gottschalg, O. F., Hahn, M., Kehoe, C., 2013. Corporate governance and value creation: Evidence from private equity. *Review of Financial Studies* 26 (2), 368–402.
- Ang, A., Chen, B., Goetzmann, W. N., Phalippou, L., 2014. Estimating private equity returns from limited partner cash flows. Unpublished working paper. Columbia University, Yale University, and University of Oxford.
- Avramov, D., Wermers, R., 2006. Investing in mutual funds when returns are predictable. *Journal of Financial Economics* 81 (2), 339–377.
- Baks, K. P., Metrick, A., Wachter, J., 2001. Should investors avoid all actively managed mutual funds? A study in Bayesian performance evaluation. *Journal of Finance* 56 (1), 45–85.
- Barber, B. M., Odean, T., 2000. Trading is hazardous to your wealth: The common stock investment performance of individual investors. *Journal of Finance* 55 (2), 773–806.
- Barber, B. M., Yasuda, A., 2016. Interim fund performance and fundraising in private equity. *Journal of Financial Economics*, forthcoming.
- Barras, L., Scaillet, O., Wermers, R., 2010. False discoveries in mutual fund performance: Measuring luck in estimated alphas. *Journal of Finance* 65 (1), 179–216.
- Bengtsson, O., 2013. Relational venture capital financing of serial founders. *Journal of Financial Intermediation* 22 (3), 308–334.
- Berk, J. B., Green, R. C., 2004. Mutual fund flows and performance in rational markets. *Journal of Political Economy* 112 (6), 1269–1295.
- Berk, J. B., Van Binsbergen, J. H., 2015. Measuring skill in the mutual fund industry. *Journal of Financial Economics* 118 (1), 1–20.

- Berkhof, J., Van Mechelen, I., Gelman, A., 2003. A Bayesian approach to the selection and testing of mixture models. *Statistica Sinica* 13 (2), 423–442.
- Bertrand, M., Schoar, A., 2003. Managing with style: The effect of managers on firm policies. *The Quarterly Journal of Economics* 118 (4), 1169–1208.
- Braun, R., Jenkinson, T., Stoff, I., 2016. How persistent is private equity performance? Evidence from deal-level data. *Journal of Financial Economics*, forthcoming.
- Brown, G. W., Gredil, O., Kaplan, S. N., 2016. Do private equity funds game returns? Fama-Miller working paper. University of North Carolina, Tulane University, and University of Chicago.
- Brown, S. J., Goetzmann, W. N., Ibbotson, R. G., 1999. Offshore hedge funds: Survival and performance, 1989-95. *Journal of Business* 72 (1), 91–117.
- Busse, J. A., Goyal, A., Wahal, S., 2010. Performance and persistence in institutional investment management. *Journal of Finance* 65 (2), 765–790.
- Busse, J. A., Irvine, P. J., 2006. Bayesian alphas and mutual fund persistence. *Journal of Finance* 61 (5), 2251–2288.
- Carhart, M. M., 1997. On persistence in mutual fund performance. *Journal of Finance* 52 (1), 57–82.
- Caselli, S., Garcia-Appendini, E., Ippolito, F., 2013. Contracts and returns in private equity investments. *Journal of Financial Intermediation* 22 (2), 201–217.
- Chen, R., Liu, J. S., 2000. Mixture Kalman filters. *Journal of the Royal Statistical Society: Series B* 62 (3), 493–508.
- Chib, S., 1995. Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* 90 (432), 1313–1321.

- Chung, J.-W., 2012. Performance persistence in private equity funds. Unpublished working paper. Chinese University of Hong Kong.
- Chung, J.-W., Sensoy, B. A., Stern, L., Weisbach, M. S., 2012. Pay for performance from future fund flows: The case of private equity. *Review of Financial Studies* 25 (11), 3259–3304.
- Cochrane, J. H., 1988. How big is the random walk in GNP? *Journal of Political Economy* 96 (5), 893–920.
- De Leeuw, J., Meijer, E., Goldstein, H., 2008. *Handbook of multilevel analysis*. Springer, New York, NY.
- Diebolt, J., Robert, C. P., 1994. Estimation of finite mixture distributions through Bayesian sampling. *Journal of the Royal Statistical Society: Series B* 56 (2), 363–375.
- Driessen, J., Lin, T.-C., Phalippou, L., 2012. A new method to estimate risk and return of nontraded assets from cash flows: The case of private equity funds. *Journal of Financial and Quantitative Analysis* 47 (3), 511–535.
- Efron, B., Morris, C., 1975. Data analysis using Stein’s estimator and its generalizations. *Journal of the American Statistical Association* 70 (350), 311–319.
- Ewens, M., Rhodes-Kropf, M., 2015. Is a VC partnership greater than the sum of its partners? *Journal of Finance* 70 (3), 1081–1113.
- Fama, E. F., French, K. R., 2010. Luck versus skill in the cross-section of mutual fund returns. *Journal of Finance* 65 (5), 1915–1947.
- Ferson, W. E., 2010. Investment performance evaluation. *Annual Review of Financial Economics* 2 (1), 207–234.
- Ferson, W. E., Chen, Y., 2015. How many good and bad fund managers are there, really? Unpublished working paper. University of Southern California and Texas A&M University.

- Franzoni, F., Nowak, E., Phalippou, L., 2012. Private equity performance and liquidity risk. *Journal of Finance* 67 (6), 2341–2373.
- Garleanu, N., Pedersen, L. H., 2016. Efficiently inefficient markets for assets and asset management. Unpublished working paper. Copenhagen Business School and UC Berkeley.
- Gelfand, A. E., Sahu, S. K., Carlin, B. P., 1995. Efficient parametrisations for normal linear mixed models. *Biometrika* 82 (3), 479–488.
- Gelfand, A. E., Smith, A. F. M., 1990. Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association* 85 (410), 398–409.
- Globe, V., Green, R. C., 2011. Information spillovers and performance persistence for hedge funds. *Journal of Financial Economics* 101 (1), 1–17.
- Gompers, P., Kovner, A., Lerner, J., Scharfstein, D., 2010. Performance persistence in entrepreneurship. *Journal of Financial Economics* 96 (1), 18–32.
- Graham, J. R., Harvey, C. R., 1996. Market timing ability and volatility implied in investment newsletters’ asset allocation recommendations. *Journal of Financial Economics* 42 (3), 397–421.
- Griffin, J. M., Xu, J., 2009. How smart are the smart guys? A unique view from hedge fund stock holdings. *Review of Financial Studies* 22 (7), 2531–2570.
- Gruber, M. J., 1996. Another puzzle: The growth in actively managed mutual funds. *Journal of Finance* 51 (3), 783–810.
- Harris, R. S., Jenkinson, T., Kaplan, S. N., 2014a. Private equity performance: What do we know? *Journal of Finance* 69 (5), 1851–1882.
- Harris, R. S., Jenkinson, T., Kaplan, S. N., Stucke, R., 2014b. Has persistence persisted in private equity? Evidence from buyout and venture capital funds. Unpublished working paper. University of Virginia, University of Chicago, and University of Oxford.

- Higson, C., Stucke, R., 2012. The performance of private equity. Unpublished working paper. London Business School and University of Oxford.
- Hochberg, Y. V., Ljungqvist, A., Vissing-Jorgensen, A., 2014. Informational holdup and performance persistence in venture capital. *Review of Financial Studies* 27 (1), 102–152.
- Hochberg, Y. V., Rauh, J. D., 2013. Local overweighting and underperformance: Evidence from limited partner private equity investments. *Review of Financial Studies* 26 (2), 403–451.
- Jagannathan, R., Malakhov, A., Novikov, D., 2010. Do hot hands exist among hedge fund managers? An empirical evaluation. *Journal of Finance* 65 (1), 217–255.
- Jenkinson, T., Jones, H., Martinez, J. V., 2016. Picking winners? Investment consultants’ recommendations of fund managers. *Journal of Finance* 71, 2333–2370.
- Jenkinson, T., Sousa, M., Stucke, R., 2013. How fair are the valuations of private equity funds? Unpublished working paper. University of Oxford and University of Porto.
- Jensen, M. C., 1968. The performance of mutual funds in the period 1945–1964. *Journal of Finance* 23 (2), 389–416.
- Jones, C. S., Shanken, J., 2005. Mutual fund performance with learning across funds. *Journal of Financial Economics* 78 (3), 507–552.
- Kacperczyk, M., Sialm, C., Zheng, L., 2005. On the industry concentration of actively managed equity mutual funds. *Journal of Finance* 60 (4), 1983–2011.
- Kaplan, S. N., Schoar, A., 2005. Private equity performance: Returns, persistence, and capital flows. *Journal of Finance* 60 (4), 1791–1823.
- Kaplan, S. N., Stromberg, P., 2009. Leveraged buyouts and private equity. *Journal of Economic Perspectives* 23 (1), 121–46.

- Kass, R. E., Raftery, A. E., 1995. Bayes factors. *Journal of the American Statistical Association* 90 (430), 773–795.
- Koijen, R. S., 2014. The cross-section of managerial ability, incentives, and risk preferences. *Journal of Finance* 69 (3), 1051–1098.
- Korteweg, A. G., 2013. Markov chain Monte Carlo methods in corporate finance. In: Damien, P., Dellaportas, P., Polson, N. G., Stephens, D. A. (Eds.), *Bayesian Theory and Applications*. Oxford University Press, Oxford, UK, pp. 516–545.
- Korteweg, A. G., Nagel, S., 2016. Risk-adjusting the returns to venture capital. *Journal of Finance* 71 (3), 1437–1470.
- Korteweg, A. G., Sorensen, M., 2010. Risk and return characteristics of venture capital-backed entrepreneurial companies. *Review of Financial Studies* 23 (10), 3738–3772.
- Kosowski, R., Timmermann, A., Wermers, R., White, H., 2006. Can mutual fund “stars” really pick stocks? New evidence from a bootstrap analysis. *Journal of Finance* 61 (6), 2551–2595.
- Kuckertz, A., Kollmann, T., Rohm, P., Middelberg, N., 2015. The interplay of track record and trustworthiness in venture capital fundraising. *Journal of Business Venturing Insights* 4, 6–13.
- Lerner, J., Schoar, A., Wongsunwai, W., 2007. Smart institutions, foolish choices: The limited partner performance puzzle. *Journal of Finance* 62 (2), 731–764.
- Li, Y., 2014. Reputation, volatility and performance persistence of private equity. Unpublished working paper. Federal Reserve Board of Governors.
- Lindley, D. V., 1962. Discussion on Professor Stein’s paper. *Journal of the Royal Statistical Society: Series B* 24 (2), 285–287.

- Ljungqvist, A., Richardson, M., 2003. The cash flow, return and risk characteristics of private equity. Unpublished working paper. NYU Stern.
- Ljungqvist, A., Richardson, M., Wolfenzon, D., 2008. The investment behavior of buyout funds: Theory and evidence. Unpublished working paper. NYU Stern.
- Lopez-de Silanes, F., Phalippou, L., Gottschalg, O., 2015. Giants at the gate: Investment returns and diseconomies of scale in private equity. *Journal of Financial and Quantitative Analysis* 50 (3), 377–411.
- Malkiel, B. G., 1995. Returns from investing in equity mutual funds 1971 to 1991. *Journal of Finance* 50 (2), 549–572.
- Marin, J.-M., Robert, C. P., 2008. Approximating the marginal likelihood in mixture models. *Indian Bayesian Society Newsletter* 5 (1), 2–7.
- Marquez, R., Nanda, V., Yavuz, M. D., 2015. Private equity fund returns and performance persistence. *Review of Finance* 19 (5), 1783–1823.
- Merton, R. C., 1980. On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics* 8 (4), 323–361.
- Pástor, L., Stambaugh, R. F., 2002. Mutual fund performance and seemingly unrelated assets. *Journal of Financial Economics* 63 (3), 315–349.
- Phalippou, L., 2010. Venture capital funds: Flow-performance relationship and performance persistence. *Journal of Banking & Finance* 34 (3), 568–577.
- Phalippou, L., 2014. Performance of buyout funds revisited? *Review of Finance* 18 (1), 189–218.
- Phalippou, L., Gottschalg, O., 2009. The performance of private equity funds. *Review of Financial Studies* 22 (4), 1747–1776.
- Rao, C. R., 1972. Estimation of variance and covariance components in linear models. *Journal of the American Statistical Association* 67 (337), 112–115.

- Raudenbush, S. W., Bryk, A. S., 2002. Hierarchical linear models: Applications and data analysis methods, 2nd Edition. Sage Publications, Thousand Oaks.
- Robinson, D. T., Sensoy, B. A., 2016. Cyclicalities, performance measurement, and cash flow liquidity in private equity. *Journal of Financial Economics*, forthcoming.
- Rossi, P. E., Allenby, G. M., McCulloch, R., 2005. Bayesian statistics and marketing. John Wiley & Sons, Chichester, UK.
- Sensoy, B. A., Wang, Y., Weisbach, M. S., 2014. Limited partner performance and the maturing of the private equity industry. *Journal of Financial Economics* 112 (3), 320–343.
- Sorensen, M., Jagannathan, R., 2015. The public market equivalent and private equity performance. *Financial Analysts Journal* 71 (4), 43–50.
- Stein, C., 1956. Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. In: Neyman, J. (Ed.), *Proceedings of the Third Berkeley symposium on mathematical statistics and probability*. Vol. 1. University of California Press, Berkeley, CA, pp. 197–206.
- Timmermann, A., Blake, D., 2005. International asset allocation with time-varying investment opportunities. *Journal of Business* 78 (1), 71–98.
- Titman, S., Tiu, C., 2011. Do the best hedge funds hedge? *Review of Financial Studies* 24 (1), 123–168.
- Wermers, R., 2011. Performance measurement of mutual funds, hedge funds, and institutional accounts. *Annual Review of Financial Economics* 3, 537–574.
- West, M., 1992. Modeling with mixtures. In: Bernardo, J., Berger, J. O., Dawid, A. P., Smith, A. F. M. (Eds.), *Bayesian Statistics*. Vol. 4. Oxford University Press, Oxford, UK, pp. 503–524.

Wu, Y., Wermers, R., Zechner, J., 2016. Managerial rents vs. shareholder value in delegated portfolio management: The case of closed-end funds. *Review of Financial Studies*, forthcoming.

Table 1

Summary statistics

Descriptive statistics of the sample of private equity funds, by fund type [VC (venture capital), BO (buyout), and Other]. The sample contains 1,924 fully liquidated funds raised between 1969 and 2001, with at least \$5 million in committed capital (in 1990 dollars) and with non-missing returns data. The funds are raised by 831 individual private equity firms (some firms manage funds of more than one type). Fund size is the committed capital in millions of dollars. IRR is the fund's internal rate of return, net of fees. The ten-year log return is computed as $10 \cdot \ln(1 + IRR)$. Overlap is the number of years of overlap for funds of the same firm and type that overlap. Source: Preqin.

<i>Panel A: Broad fund categories</i>			
	VC	BO	Other
Number of funds	842	562	518
Number of firms	409	285	197
Number of funds / firm			
Mean	2.1	2.0	2.6
Median	1	1	2
Standard deviation	1.9	1.6	2.6
10th percentile	1	1	1
90th percentile	4	4	5
Fund size (millions of dollars)			
Mean	206.9	694.1	373.3
Median	110.0	300.0	206.8
Standard deviation	276.1	1,035.6	517.1
10th percentile	27.0	52.6	33.0
90th percentile	500.0	1,823.6	863.0
IRR (percent)			
Mean	17.7	16.9	13.9
Median	8.6	14.9	11.9
Standard deviation	54.8	18.6	12.9
10th percentile	-10.4	-1.7	0.4
90th percentile	46.0	37.9	28.9
Ten-year log return			
Mean	1.173	1.438	1.245
Median	0.825	1.385	1.124
Standard deviation	2.623	1.552	1.075
10th percentile	-1.101	-0.170	0.040
90th percentile	3.786	3.216	2.542
Overlap (years)			
Number of fund pairs	891	512	968
Mean	5.8	5.8	6.8
Median	6	6	7
Standard deviation	2.5	2.3	2.4
10th percentile	2	2	3
90th percentile	9	9	9

Table 1 (*continued*)
Panel B: Fund subcategories

Subcategory	Number of funds	Number of firms	Fund size (millions of dollars)		IRR (percent)	
			Mean	Median	Mean	Standard deviation
VC						
Early stage	177	105	186.2	92.0	23.7	6.9
Late stage	153	89	289.0	154.0	12.3	10.5
Generalist	512	239	187.5	104.0	17.2	8.5
Other						
Real estate	202	86	389.3	263.0	13.8	12.9
Fund-of-funds	144	48	349.3	138.3	11.2	8.3
Distressed debt	36	14	662.1	438.0	14.6	14.0
Natural resources	58	25	265.4	154.8	17.0	15.2
Secondaries	34	12	336.8	263.5	16.8	15.0
Infrastructure, turnaround, special situations, co-investments, venture debt	44	26	293.0	150.0	16.9	13.0
						18.6

Table 2

Fund internal rates of return by vintage year

This table shows the number of private equity funds in the sample, and their internal rates of return (IRRs), by vintage year and fund type [VC (venture capital), BO (buyout), and Other]. The sample contains 1,924 fully liquidated funds over the period 1969 to 2001, with at least \$5 million in committed capital (in 1990 dollars) and non-missing returns data. Weighted mean internal rates of return are weighted by fund size, computed using the funds for which size data are available. Source: Preqin.

Vintage year	Funds	VC			BO			Other		
		IRR (percent)			IRR (percent)			IRR (percent)		
		Mean	Median	Weighted mean	Mean	Median	Weighted mean	Mean	Median	Weighted mean
1969	1	8.7	8.7	8.7	0	—	—	—	—	—
1970	0	—	—	—	0	—	—	0	—	—
1971	0	—	—	—	0	—	—	0	—	—
1972	1	21.5	21.5	21.5	0	—	—	0	—	—
1973	0	—	—	—	0	—	—	0	—	—
1974	0	—	—	—	0	—	—	0	—	—
1975	0	—	—	—	0	—	—	0	—	—
1976	0	—	—	—	0	—	—	0	—	—
1977	0	—	—	—	0	—	—	0	—	—
1978	2	48.6	48.6	51.0	1	35.5	35.5	0	—	—
1979	1	18.5	18.5	18.5	1	19.4	19.4	1	18.0	18.0
1980	6	16.5	14.0	25.9	3	23.7	25.8	1	12.0	12.0
1981	7	18.9	11.3	19.1	0	—	—	2	11.1	11.1
1982	11	13.8	9.3	19.3	1	39.2	39.2	1	10.0	10.0
1983	12	9.8	9.3	9.4	2	21.9	21.9	3	14.7	5.9
1984	20	12.3	12.0	12.0	7	30.4	18.4	1	7.1	7.1
1985	21	13.1	13.0	14.9	4	13.4	10.7	4	4.4	2.7
1986	20	9.5	8.5	9.3	12	27.6	18.8	3	4.0	4.0
1987	22	13.4	14.8	12.6	9	16.8	18.9	8	8.9	8.3
1988	26	20.8	21.4	27.0	12	19.2	14.2	6	10.9	11.6
1989	35	22.0	16.4	30.1	12	29.8	27.5	8	13.6	11.1
1990	24	15.1	16.5	17.6	23	19.2	15.4	10	16.8	16.0
1991	18	23.1	21.1	27.2	10	28.0	25.2	10	15.5	12.2
1992	30	21.8	16.0	24.3	21	18.3	21.2	15	18.0	16.3
1993	36	32.6	29.5	35.9	23	19.2	16.9	16	23.2	19.8
1994	30	30.0	25.5	38.2	36	25.9	21.5	18	15.9	14.0
1995	41	45.2	17.5	42.5	33	17.1	17.6	30	16.8	17.0
1996	43	29.3	10.3	21.3	35	17.4	11.4	45	17.3	12.7
1997	62	42.8	20.6	37.6	54	11.6	9.0	45	11.5	8.4
1998	73	36.6	7.0	25.0	73	6.7	8.3	65	11.2	8.3
1999	95	-2.5	-3.5	-3.4	57	8.6	8.5	58	10.1	9.3
2000	120	0.1	-0.9	-0.2	90	16.8	17.5	80	12.1	12.4
2001	85	2.4	1.0	2.9	43	27.1	28.0	88	17.0	15.2
										22.3

Table 3

AR(1) regressions

The table shows AR(1) regressions using fund internal rates of return (IRRs) in percent, by fund type. The dependent variable is $IRR_{i,N}$, the IRR of fund N of private equity firm i . In Panel A, $IRR_{i,N-1}$ is the IRR of the most recent fund of the same type [VC (venture capital), BO (buyout), or Other] and the same firm, and $IRR_{i,N-2}$ is the IRR of the second previous fund. $VC=1$ is a dummy variable that equals one for VC funds and zero otherwise. The variables $\ln(size)$ and $\ln(sequence)$ are the natural logarithm of fund size and the sequence number of the fund within its class (VC, BO, Other) for a given PE firm. In Panel B, $IRR_{i,N-1}$ is the net IRR on the most recent, nonoverlapping fund of the same type and the same firm. All regressions include vintage year fixed effects. Standard errors are clustered by firm and are shown in brackets. ***, **, and * indicate statistical significance at the 1%, 5%, and 10% level, respectively.

Panel A: All funds										
Variable	VC and BO			VC only			BO only			Other only
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$IRR_{i,N-1}$	0.162* (0.071)	0.304*** (0.053)			0.297*** (0.056)	0.129 (0.079)	0.270*** (0.059)	0.314*** (0.067)	0.299*** (0.072)	0.255*** (0.091)
$IRR_{i,N-2}$		-0.063 (0.053)	0.052 (0.053)		-0.064 (0.055)		-0.045 (0.063)		0.152 (0.102)	0.148 (0.097)
$IRR_{i,N-3}$				-0.044 (0.026)						
$VC=1$	1.260 (2.696)	1.988 (4.024)	3.268 (4.714)	4.867 (6.237)	2.032 (4.835)					
$\ln(size)$					0.375 (1.649)					
$\ln(sequence)$					1.769 (4.506)					
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	673	375	398	244	365	411	240	262	135	167
Adj. R^2	0.136	0.173	0.107	0.082	0.162	0.195	0.209	0.308	0.398	0.104
Panel B: Non-overlapping funds										
Variable	VC and BO		VC		BO		Other			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$IRR_{i,N-1}$	0.065 (0.090)	0.067 (0.067)	0.195 (0.181)		-0.296 (0.163)					
$VC=1$	8.567 (7.901)									
Year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	207	147	60	38	38	38	38	38	38	38
Adj. R^2	0.064	0.087	-0.006		0.063					

Table 4

Parameter estimates

This table reports posterior means of parameters of the model described in the text. Panel A shows the parameter estimates, and Panel B reports the variance decomposition estimates for the same parameters. Panel C shows the spread in gammas across percentiles of the posterior distribution. The model includes either a single intercept (Specification I) or vintage year fixed effects, grouping the pre-1985 vintages into one bucket (Specification II). The error term ε_{iu} is a mixture of K normal distributions, where K is chosen as the best fit according to models' marginal log-likelihood. The model is estimated separately for each fund type [VC (venture capital), BO (buyout), and Other], by Markov chain Monte Carlo (MCMC) using ten thousand burn-in cycles followed by 100,000 samples, saving every tenth draw. Posterior standard deviations (Bayesian standard errors) are in brackets.

	VC		BO		Other	
	I	II	I	II	I	II
<i>Panel A: Parameter estimates</i>						
σ_γ	0.049 (0.007)	0.055 (0.008)	0.060 (0.008)	0.056 (0.008)	0.049 (0.006)	0.049 (0.006)
σ_η	0.258 (0.031)	0.193 (0.037)	0.142 (0.039)	0.203 (0.043)	0.202 (0.028)	0.135 (0.028)
σ_ε	2.449 (0.123)	2.326 (0.121)	1.359 (0.058)	1.225 (0.073)	0.807 (0.039)	0.865 (0.050)
Vintage fixed effects	No	Yes	No	Yes	No	Yes
K	3	3	2	2	1	3
N	842	842	562	562	518	518
<i>Panel B: Variance decomposition</i>						
$100 \cdot \sigma_\gamma^2$	0.243 (0.067)	0.309 (0.087)	0.361 (0.094)	0.316 (0.089)	0.244 (0.065)	0.246 (0.061)
$10 \cdot \sigma_\eta^2$	0.675 (0.158)	0.386 (0.141)	0.216 (0.113)	0.432 (0.160)	0.416 (0.111)	0.189 (0.076)
σ_ε^2	6.015 (0.604)	5.426 (0.567)	1.852 (0.159)	1.505 (0.180)	0.654 (0.064)	0.751 (0.087)
σ_y^2	6.933 (0.596)	6.120 (0.561)	2.428 (0.152)	2.253 (0.168)	1.314 (0.084)	1.186 (0.090)
Signal-to-noise	0.035 (0.010)	0.051 (0.015)	0.148 (0.037)	0.141 (0.040)	0.185 (0.045)	0.208 (0.048)
<i>Panel C: Gamma spread</i>						
$q_\gamma(75\%) - q_\gamma(25\%)$	6.59% (0.90%)	7.42% (1.05%)	8.03% (1.05%)	7.51% (1.06%)	6.61% (0.87%)	6.64% (0.81%)
$q_\gamma(87.5\%) - q_\gamma(12.5\%)$	11.24% (1.54%)	12.66% (1.78%)	13.70% (1.79%)	12.81% (1.81%)	11.27% (1.49%)	11.34% (1.40%)

Table 5

Subsamples

Panel A shows point estimates (posterior means) of the variance decomposition parameters for different subsamples and the spread in long-term persistence across firms as described in Table 4. Panel B shows the signal-to-noise ratio, s_γ , for the same subsamples. The columns labeled “Mean” and “SD” report the posterior mean and standard deviation of s_γ , respectively. IQR is the interquartile range of the posterior distribution of s_γ , i.e., the spread in between the 75th and 25th percentile of the posterior distribution. N is the number of funds in the subsample. All estimates are based on model specification II in Table 4 (which includes vintage year fixed effects). The top row in each panel corresponds to the estimates of the full sample, as reported in Table 4 and Fig. 4. The number of funds in the fund size subsamples do not add up to the number of funds in the full sample, because size is not observed for some funds. VC = venture capital; BO = buyout.

<i>Panel A: Variance decomposition</i>												
Sample	VC				BO				Other			
	$100\sigma_\gamma^2$	$10\sigma_\gamma^2$	σ_ε^2	$q_\gamma(75\%) - q_\gamma(25\%)$	$100\sigma_\gamma^2$	$10\sigma_\gamma^2$	σ_ε^2	$q_\gamma(75\%) - q_\gamma(25\%)$	$100\sigma_\gamma^2$	$10\sigma_\gamma^2$	σ_ε^2	$q_\gamma(75\%) - q_\gamma(25\%)$
Full sample	0.309	0.386	6.120	7.42%	0.316	0.432	2.253	7.51%	0.246	0.189	1.186	6.64%
Fund size												
Small (< median)	0.823	0.522	7.537	12.00%	0.489	0.207	2.667	9.25%	0.326	0.245	1.110	7.54%
Large ($\geq$ median)	0.239	0.328	4.808	6.46%	0.234	0.452	1.107	6.40%	0.203	0.086	1.217	5.98%
General partner location												
US	0.300	0.339	5.897	7.26%	0.286	0.438	2.046	6.99%	0.089	0.319	1.130	3.89%
Europe (including UK)	0.491	0.151	2.743	9.19%	0.427	0.146	1.578	8.68%	0.327	0.121	1.031	7.49%
Rest of the world	0.562	0.258	10.465	9.72%	0.444	0.152	5.643	8.56%	1.434	0.180	2.306	15.45%
Style												
Early stage	0.217	0.106	9.866	6.13%	—	—	—	—	—	—	—	—
Late stage	0.272	0.272	1.949	6.87%	—	—	—	—	—	—	—	—
Generalist	0.377	0.386	5.674	8.10%	—	—	—	—	—	—	—	—
Real estate	—	—	—	—	—	—	—	—	0.149	0.079	0.674	5.16%
Fund-of-funds	—	—	—	—	—	—	—	—	0.182	0.074	1.056	5.65%
Sample period												
Early (< 1997)	0.529	0.429	4.947	9.68%	0.611	0.372	2.879	10.26%	0.338	0.123	1.239	7.71%
Late ($\geq$ 1997)	0.288	0.117	7.383	7.11%	0.303	0.156	1.709	7.33%	0.279	0.097	1.143	7.04%
<i>Panel B: Signal-to-noise ratio</i>												
Sample	VC				BO				Other			
	Mean	SD	IQR	N	Mean	SD	IQR	N	Mean	SD	IQR	N
Full sample	0.051	0.015	0.020	842	0.141	0.040	0.053	562	0.208	0.048	0.065	518
Fund size												
Small (< median)	0.111	0.043	0.057	373	0.183	0.068	0.096	261	0.291	0.103	0.146	210
Large ($\geq$ median)	0.051	0.022	0.028	374	0.123	0.049	0.064	270	0.168	0.056	0.077	210
GP location												
US	0.051	0.020	0.027	675	0.140	0.065	0.092	416	0.079	0.042	0.050	439
Europe (including UK)	0.180	0.079	0.110	93	0.269	0.081	0.112	113	0.310	0.116	0.159	59
Rest of the world	0.057	0.036	0.042	74	0.083	0.058	0.061	33	0.584	0.187	0.292	20
Style												
Early stage	0.023	0.011	0.013	177	—	—	—	—	—	—	—	—
Late stage	0.141	0.063	0.080	153	—	—	—	—	—	—	—	—
Generalist	0.067	0.029	0.040	512	—	—	—	—	—	—	—	—
Real estate	—	—	—	—	—	—	—	—	0.222	0.061	0.084	202
Fund-of-funds	—	—	—	—	—	—	—	—	0.173	0.061	0.080	144
Sample period												
Early (< 1997)	0.108	0.035	0.047	407	0.211	0.091	0.129	245	0.271	0.088	0.125	227
Late ($\geq$ 1997)	0.040	0.016	0.021	435	0.178	0.058	0.080	317	0.243	0.066	0.091	291

Table 6

Parameter estimates using total value to paid-in capital (TVPI)

This table reports posterior means of parameters of the model described in the text, using the natural logarithm of TVPI as the dependent variable. Panel A shows the parameter estimates, and Panel B shows the variance decomposition estimates for the same parameters. Panel C shows the spread in gammas across percentiles of the posterior distribution. The model and parameters are described in Table 4. VC = venture capital; BO = buyout.

	VC		BO		Other	
	I	II	I	II	I	II
<i>Panel A: Parameter estimates</i>						
σ_γ	0.032 (0.003)	0.032 (0.003)	0.028 (0.002)	0.029 (0.002)	0.030 (0.002)	0.029 (0.002)
σ_η	0.123 (0.013)	0.073 (0.012)	0.080 (0.010)	0.071 (0.010)	0.066 (0.008)	0.055 (0.008)
σ_ε	0.747 (0.030)	0.727 (0.028)	0.499 (0.030)	0.481 (0.031)	0.346 (0.020)	0.353 (0.020)
Vintage fixed effects	No	Yes	No	Yes	No	Yes
K	2	2	2	2	2	3
N	970	970	624	624	563	563
<i>Panel B: Variance decomposition</i>						
$100 \cdot \sigma_\gamma^2$	0.104 (0.018)	0.106 (0.018)	0.077 (0.011)	0.086 (0.012)	0.088 (0.014)	0.083 (0.013)
$10 \cdot \sigma_\eta^2$	0.154 (0.032)	0.055 (0.017)	0.066 (0.015)	0.051 (0.014)	0.044 (0.011)	0.030 (0.009)
σ_ε^2	0.559 (0.044)	0.529 (0.041)	0.250 (0.031)	0.233 (0.031)	0.120 (0.014)	0.125 (0.014)
σ_y^2	0.816 (0.041)	0.690 (0.039)	0.393 (0.030)	0.369 (0.030)	0.252 (0.018)	0.125 (0.014)
Signal-to-noise	0.127 (0.022)	0.154 (0.026)	0.196 (0.029)	0.233 (0.034)	0.349 (0.046)	0.348 (0.045)
<i>Panel C: $\ln(TVPI)$ spread due to firm effect</i>						
$q_\gamma(75\%) - q_\gamma(25\%)$	0.432 (0.038)	0.438 (0.038)	0.373 (0.027)	0.394 (0.028)	0.399 (0.032)	0.388 (0.030)
$q_\gamma(87.5\%) - q_\gamma(12.5\%)$	0.738 (0.064)	0.747 (0.064)	0.636 (0.047)	0.671 (0.048)	0.680 (0.055)	0.662 (0.051)

Table A1

Model specification tests

This table shows tests of the model specification. Column 1 reproduces Specification I of Table 4, which includes an intercept but no vintage year fixed effects. Column 2 drops the transient firm effect, η , from the model, and Column 3 drops the long-run firm-specific effect, γ . Columns 4 to 6 show the same for Specification II of Table 4, which includes vintage year fixed effects. The Bayes factor represents the ratio of marginal likelihoods, indicating the weight of evidence of each model, relative to the full model specification in Column 4. A Bayes factor of one indicates that the two models have equal support in the data. For each model, the number of distributions in the error term (K) is chosen to find the best model fit by marginal log-likelihood. Posterior standard deviations (Bayesian standard errors) are in brackets. VC = venture capital; BO = buyout.

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: VC</i>						
σ_γ	0.049 (0.007)	0.042 (0.004)		0.055 (0.008)	0.040 (0.004)	
σ_η	0.258 (0.031)		0.313 (0.023)	0.193 (0.037)		0.268 (0.025)
σ_ε	2.449 (0.123)	2.604 (0.107)	2.431 (0.114)	2.326 (0.121)	2.465 (0.115)	2.316 (0.123)
σ_y^2	6.933 (0.596)	6.971 (0.562)	6.906 (0.553)	6.120 (0.561)	6.255 (0.572)	6.106 (0.567)
Vintage fixed effects	No	No	No	Yes	Yes	Yes
K	3	3	2	3	3	3
N	842	842	842	842	842	842
Marginal log-L	-1,829.5	-1,755.4	-1,823.6	-1,703.8	-1,669.2	-1,727.4
Bayes factor	0.000	0.000	0.000	N/A	1.0E+15	0.000
<i>Panel B: BO</i>						
σ_γ	0.060 (0.008)	0.043 (0.005)		0.056 (0.008)	0.044 (0.005)	
σ_η	0.142 (0.039)		0.277 (0.034)	0.203 (0.043)		0.279 (0.020)
σ_ε	1.359 (0.058)	1.542 (0.051)	1.261 (0.065)	1.225 (0.073)	1.469 (0.054)	1.196 (0.071)
σ_y^2	2.428 (0.152)	2.567 (0.163)	2.373 (0.121)	2.253 (0.168)	2.359 (0.163)	2.217 (0.167)
Vintage fixed effects	No	No	No	Yes	Yes	Yes
K	2	2	1	2	2	2
N	562	562	562	562	562	562
Marginal log-L	-1,054.6	-1,039.7	-1,054.6	-1,031.6	-1,019.7	-1,035.0
Bayes factor	0.000	0.000	0.000	N/A	1.4E+05	0.033
<i>Panel C: Other</i>						
σ_γ	0.049 (0.006)	0.039 (0.004)		0.049 (0.006)	0.037 (0.004)	
σ_η	0.202 (0.028)		0.245 (0.022)	0.135 (0.028)		0.199 (0.017)
σ_ε	0.807 (0.039)	1.059 (0.033)	0.793 (0.041)	0.865 (0.050)	1.038 (0.037)	0.848 (0.050)
σ_y^2	1.314 (0.084)	1.272 (0.075)	1.234 (0.079)	1.186 (0.090)	1.217 (0.081)	1.119 (0.087)
Vintage fixed effects	No	No	No	Yes	Yes	Yes
K	1	2	1	3	3	3
N	518	518	518	518	518	518
Marginal log-L	-787.8	-777.6	-781.7	-776.9	-771.3	-774.5
Bayes factor	0.000	0.491	0.008	N/A	256.185	10.677

Table A2

Parameter estimates in sample with both internal rate of return (IRR) and total value to paid-in capital (TVPI)

This table reports posterior means of parameters of the model described in the text. Panels A and C use the ten-year log return as the performance metric (computed from the IRR), and Panels B and D use the natural logarithm of TVPI. Panels A and B include vintage year fixed effects, whereas Panels C and D do not. The full sample is the original sample with funds that report IRRs (Panels A and C) or TVPI (Panels B and D). For example, the columns labeled “Full sample” in Panel A replicate the results in Table 4 Specification II, and in Panel B they replicate the result in Table 6 Specification II. The joint sample is the subsample of funds that report both the IRR and TVPI metrics. The error term ε_{iu} is a mixture of K normal distributions, where K is chosen as the best fit according to models’ marginal log-likelihood. Posterior standard deviations (Bayesian standard errors) are in brackets. VC = venture capital; BO = buyout.

	VC		BO		Other	
	Full sample	Joint sample	Full sample	Joint sample	Full sample	Joint sample
<i>Panel A: Ten-year log return, with vintage year fixed effects</i>						
σ_γ	0.055 (0.008)	0.055 (0.008)	0.056 (0.008)	0.057 (0.008)	0.049 (0.006)	0.053 (0.007)
σ_η	0.193 (0.037)	0.179 (0.044)	0.203 (0.043)	0.197 (0.040)	0.135 (0.028)	0.128 (0.029)
σ_ε	2.326 (0.121)	2.346 (0.129)	1.225 (0.073)	1.237 (0.073)	0.865 (0.050)	0.835 (0.051)
Signal-to-noise	0.051 (0.015)	0.051 (0.015)	0.141 (0.040)	0.144 (0.039)	0.208 (0.048)	0.244 (0.057)
Vintage fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
K	3	3	2	2	3	3
N	842	793	562	548	518	459
<i>Panel B: Ln(TVPI), with vintage year fixed effects</i>						
σ_γ	0.032 (0.003)	0.033 (0.003)	0.029 (0.002)	0.030 (0.002)	0.029 (0.002)	0.030 (0.002)
σ_η	0.073 (0.012)	0.074 (0.013)	0.071 (0.010)	0.064 (0.009)	0.055 (0.008)	0.056 (0.008)
σ_ε	0.727 (0.028)	0.657 (0.028)	0.481 (0.031)	0.461 (0.029)	0.353 (0.020)	0.307 (0.020)
Signal-to-noise	0.154 (0.026)	0.188 (0.031)	0.233 (0.034)	0.260 (0.037)	0.348 (0.045)	0.420 (0.052)
Vintage fixed effects	Yes	Yes	Yes	Yes	Yes	Yes
K	2	2	2	2	3	3
N	970	793	624	548	563	459
<i>Panel C: Ten-year log return, no vintage year fixed effects</i>						
σ_γ	0.049 (0.007)	0.051 (0.007)	0.060 (0.008)	0.062 (0.008)	0.049 (0.006)	0.059 (0.008)
σ_η	0.258 (0.031)	0.256 (0.032)	0.142 (0.039)	0.125 (0.035)	0.202 (0.028)	0.157 (0.031)
σ_ε	2.449 (0.123)	2.468 (0.129)	1.359 (0.058)	1.373 (0.058)	0.807 (0.039)	0.855 (0.045)
Signal-to-noise	0.035 (0.010)	0.038 (0.011)	0.148 (0.037)	0.159 (0.038)	0.185 (0.045)	0.264 (0.061)
Vintage fixed effects	No	No	No	No	No	No
K	3	3	2	2	3	3
N	842	793	562	548	518	459

Table A2 (*continued*)*Panel D: $\ln(TVPI)$, no vintage year fixed effects*

σ_γ	0.032 (0.003)	0.034 (0.003)	0.028 (0.002)	0.028 (0.002)	0.030 (0.002)	0.032 (0.003)
σ_η	0.123 (0.013)	0.119 (0.014)	0.080 (0.010)	0.071 (0.009)	0.066 (0.008)	0.065 (0.009)
σ_ε	0.747 (0.030)	0.694 (0.031)	0.499 (0.030)	0.484 (0.029)	0.346 (0.020)	0.297 (0.018)
Signal-to-noise	0.127 (0.022)	0.154 (0.027)	0.196 (0.029)	0.220 (0.031)	0.349 (0.046)	0.432 (0.051)
Vintage fixed effects	No	No	No	No	No	No
K	2	2	2	2	2	2
N	970	793	624	548	563	459

Panel A: VC

Panel B: BO

Panel C: Other

Fig. 1. This figure depicts histograms of the number of funds per firm, by fund type [VC (venture capital), BO (buyout), and Other]. For firms that manage different types of funds, the histograms count only the number of funds of the particular type indicated.

Panel A: VC

Panel B: BO

Panel C: Other

Fig. 2. Internal rates of return (IRRs) and vintage year fixed effects. The striped lines are the average fund IRR (per annum) in each vintage year, by fund type. The solid lines are the posterior means of the vintage year fixed effects from specification II in Table 4. The vintage year fixed effects are transformed to IRR equivalents (annualized and in percent). The shaded bands represents the (1%, 99%) Bayesian credible interval (confidence bounds). VC = venture capital; BO = buyout.

Fig. 3. Fund overlap and covariance. The figure shows the covariance between total fund returns as a function of the overlap (in years) between two funds managed by the same firm, using the variance estimates for Specification II in Table 4. VC = venture capital; BO = buyout.

Panel A: VC

Panel B: BO

Panel C: Other

Fig. 4. Estimates of signal-to-noise ratio. The figure depicts posterior distribution of the signal-to-noise ratio, s_γ , by fund type, from the specifications reported in Table 4. The solid line is the kernel plot for Specification I (without vintage year fixed effects), and the striped line is the kernel plot for Specification II in Table 4 (with vintage year fixed effects). VC = venture capital; BO = buyout.

Fig. 5. Speed of learning. The figure shows the posterior probability that a fund is in the top quartile of funds. Probabilities are calculated from 100,000 simulations of a panel of one hundred firms, each with a different gamma that is drawn from the top 25% of the distribution. Each firm produces a sequence of 50 nonoverlapping fund returns. Reported probabilities are averages of the posterior mean probability across the simulated firms after observing a given number of realized fund returns for each firm (“fund history”). The figure uses the parameter estimates from Table 4 Specification II (with vintage year fixed effects). VC = venture capital; BO = buyout.

Fig. 6. Investable persistence. This figure shows the expected (true) gamma of investing in funds raised by private equity firms with top quartile performance as observed after a given number of realized fund returns for each firm (“fund history”). Calculations are based on 100,000 simulations of a panel of fund histories for one hundred firms, using the parameter estimates from Table 4 Specification II (with vintage year fixed effects). VC = venture capital; BO = buyout.

Panel A: Speed of learning

Panel B: Investable persistence

Fig. 7. Speed of learning and investable persistence using total value to paid-in capital (TVPI). Panels A and B replicate the speed of learning and investable persistence graphs in Figs. 5 and 6, respectively, using TVPI to measure fund performance. VC = venture capital; BO = buyout.

Panel A: VC

Panel B: BO

Panel C: Other

Fig. A1. This figure plots the marginal log-likelihood as a function of the number of normal mixtures in the error term distribution, by fund type [VC (venture capital), BO (buyout), and Other]. The solid line represents Specification I of Table 4 (which has no vintage year fixed effects), and the striped line represents Specification II (with vintage year fixed effects).