Skill and Luck in Private Equity Performance

Arthur Korteweg

Morten Sorensen*

October 2014

Abstract

We evaluate the performance of private equity (PE) funds using a variance decomposition model that separates long-term, investable, and spurious persistence. We find a large amount of long-term persistence: the spread in the expected returns of top- and bottom-quartile PE firms is 7 to 8 percentage points annually. Performance is noisy, however, and top-quartile past performance does not imply top-quartile expected future returns, especially for VC firms. Based on past performance alone, an investor needs to observe an excessive number of funds to identify the PE firms with top-quartile expected returns, implying low investable persistence.

*The authors can be reached at the University of Southern California (korteweg@marshall.usc.edu) and Columbia Business School or Copenhagen Business School (ms3814@columbia.edu or mso.fi@cbs.dk). We are grateful to Ulf Axelson, Andrea Buraschi, Peter Cornelius, Wayne Ferson, Ulrich Hege, Chris Jones, Paul Pfleiderer, Matt Rhodes-Kropf, Per Strömberg, seminar participants at Arizona State University, Columbia Business School, Cornerstone Research, Hong Kong University, Syddansk Universitet, University of Alberta, UC-San Diego, University of Virginia Darden, the 2013 Spring JOIM conference on Private Equity, the 7th. Private Equity Findings Symposium at London Business School, and the 2014 Argentum Private Equity Symposium for helpful comments and feedback. Morten Sorensen gratefully acknowledges funding by the Danish Council for Independent Research ("Det Frie Forskningsråd") under the Sapere Aude program through grant number: DFF–4003-00095.

The persistence and predictability of returns is a central topic in finance. Studies of a wide range of asset classes, including individual stocks, mutual funds, and hedge funds, generally find that returns are unpredictable, and that investors cannot consistently outperform the market.[1] An important exception is private equity (PE), an asset class that includes venture capital (VC) and leveraged buyout (BO) firms, among others. A PE firm typically manages a sequence of funds, and Kaplan and Schoar [2005] find that the performance of fund number $N-1$ predicts the performance of the subsequent fund N by the same firm. Their interpretation is that PE firms differ in their skills and abilities, and that funds managed by skilled PE firms persistently outperform.

Kaplan and Schoar [2005] and subsequent studies[2] define persistence as a positive and statistically significant β coefficient in the regression:

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

where $y_{i,N}$ is the performance of fund number N managed by PE firm i. This regression is motivated by a cross-sectional intuition: some funds outperform and if these funds' predecessors also outperformed this is evidence of persistence. Formally, equation (1) is a time-series AR(1) model, however, and there is a tension between this cross-sectional intuition and its 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 of $E[y] = \frac{\alpha}{1-\beta}$. Hence, the AR(1) model is a model of performance persistence that, by construction, does not allow for any long-term performance

[1] Lack of aggregate persistence is documented, 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 there is substantial heterogeneity across managers. See, for example, Baks, Metrick and Wachter [2001], Kacperczyk, Sialm and Zheng [2005], Busse and Irvine [2006], and Koijen [2014] for mutual funds, and Titman and Tiu [2011] and Jagannathan, Malakhov and Novikov [2011] for hedge funds. For a comprehensive review of the literature on performance measurement for mutual and hedge funds, we refer to Ferson [2010] and Wermers [2011].

[2] Including Phalippou and Gottschalg [2009], Phalippou [2010], Chung [2012], Robinson and Sensoy [2013], Braun, Jenkinson, and Stoff [2013], Harris, Jenkinson, Kaplan, and Stucke [2014], Hochberg, Ljungqvist, and Vissing-Jorgensen [2014], and Li [2014].

differences, which seems undesirable.

We present a new variance-decomposition model of PE performance to better capture the cross-sectional intuition. Our model distinguishes three types of persistence. *Long-term* persistence arises when some PE firms have higher (or lower) expected returns and persistently outperform (or underperform), which is typically interpreted as reflecting PE firms' skills.[3] *Investable* persistence reflects the difficulty for investors in PE funds (limited partners or LPs) to learn which PE firms have higher expected returns. When performance is noisy, top-quartile performance may be due to luck, and it does not necessarily imply top-quartile expected future returns. The noisier is the performance, the more difficult these firms are to identify, and the lower is the *investable* persistence. *Spurious* persistence arises from the partial overlap of subsequent funds that are managed by the same PE firm. Partially overlapping funds are exposed to similar market conditions and management decisions by the PE manager. Even when these exposures are purely transitory, so that past performance does not predict future performance in any way, these common contemporaneous exposures introduce a positive correlation in the performance of subsequent funds, which introduces a spurious form of persistence and a positive beta coefficient in the AR(1) model. As far as we are aware, this is the first paper to formally distinguish between these types of persistence.

The distinction matters. We find a large amount of long-term persistence: PE firms with top-quartile expected performance have annual 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 due to luck, and LPs need to observe an excessive number of past funds (about 25 or more) to identify VC firms with top-quartile expected returns with reasonable certainty. To put this in perspective, only 21 out of the 831 PE firms in our sample have more than 8 funds in the data.

[3] We use the term "skill" informally. For our formal discussion we prefer the more precise term "expected returns," which differ from skill for two reasons: First, a recent mutual fund literature argues that net-of-fee performance is not equivalent to "manager skill," which is instead reflected by gross-of-fee performance adjusted for scale (e.g., Berk and Green [2004], Pastor and Stambaugh [2012], and Berk and van Binsbergen [2014]). Second, expected returns depend on risks. Although we control for common risk exposures, it is possible that expected returns contain some remaining premia for PE firm-specific risks.

Comparing subsamples, we find that smaller funds have greater long-term persistence and more investable persistence (higher signal-to-noise ratios) than larger funds, especially in VC. We find the least long-term persistence for PE firms in the U.S., followed by Europe, and the greatest persistence for PE firms located in the rest of the world (ROW), although the latter firms also have more volatile performance. We confirm the findings in Braun, Jenkinson, and Stoff [2013] and Harris, Jenkinson, Kaplan, and Stucke [2014] that long-term persistence has declined in the 2000s, relative to the 1990s. This decline is largest for VC firms, whereas BO and Other funds still show substantial persistence post 2000.

We measure performance and report results in terms of expected returns. LPs are sometimes thought to be risk-neutral, which means that they ignore risks and maximize expected return. For such LPs, our results are directly applicable. We also estimate specifications that control for common risk exposures, such as exposures to systematic market risks (CAPM beta). These specifications are quite general, and they accommodate time-varying exposures and risk premia, but they assume identical exposures for all PE funds with the same type and vintage year. This assumption is common in the PE literature (e.g., Driessen, Lin, Phalippou [2012], and Hochberg and Rauh [2013]). Under this assumption, our results show long-term persistence in risk-adjusted returns, which is the relevant metric for risk-averse investors, and is consistent with long-term persistence arising from differences in the skills of PE firms.

Our finding of large long-term persistence but low investable persistence has important implications. It explains LPs' increasing focus on obtaining more detailed information about PE firms and their past funds (such as firms' internal organization and culture, compensation structures and the alignment of incentives, processes, and deal sourcing) to help them attribute past performance (e.g., Ewens and Rhodes-Kropf [2014]) and Jenkinson, Jones, and Martinez [2014]). Such additional information is necessary for LPs to identify top PE firms, as past fund performance, by itself, is insufficient. Moreover, our results may explain why outperformance is not competed away, contrary to the prediction by Berk and Green [2004]. PE manager skills are scarce, but when performance is noisy, LPs with the ability to identify skilled PE firms may also be scarce, and those LPs earn rents. This empirical finding is consistent with the model by Hochberg, Ljungqvist, and Vissing-Jorgensen

[2014] where performance persistence arises from the asymmetric information problem between limited and general partners.[4] Finally, our findings confirm the economic realities behind the saying among VCs that "I'd rather be lucky than smart." For VC funds, the variation in performance that is due to luck is an order of magnitude greater than the variation that is due to skill.

We develop a new empirical variance decomposition model, or *hierarchical linear model*, that generalizes the classical *analysis of variance* (ANOVA) method to capture the particular features of PE fund performance. It has several advantages compared to the AR(1) regression. First, as mentioned, it distinguishes the different forms of persistence described above: *long-term, investable*, and *spurious* persistence. Second, it explicitly models the timing of funds, and it does not rely on their numbering, which is important for simultaneous funds, where it becomes arbitrary which one is labeled N and $N-1$. Third, the model accounts for overlap between funds and distinguishes situations where fund N follows fund $N-1$ by a few months from those where they are years apart. It is robust to missing data for intermediate funds. Fourth, unlike the AR(1) regressions, our empirical model accommodates firms with a single fund, which do not have a fund $N-1$. These firms are typically weaker firms, and omitting them may introduce systematic biases. Fifth, our model has a semi-parametric component, and it does not require fund returns to have a normal distribution, unlike standard ANOVA methods. Instead, we use a mixture-of-normals distribution and a Bayes factor test to determine the appropriate number of mixtures. This generality is especially important for VC funds, which have highly skewed returns. Sixth, our method captures estimation error in the model parameters and the effects of this parameter uncertainty on LPs' inference problem of identifying skilled firms. Seventh, 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. This is important since the parameters of interest are variances and ratios of variances (in the case of signal-to-noise ratios), which have non-standard asymptotic distributions.

[4] Lerner, Schoar, and Wongsunwai [2007], Hochberg and Rauh [2013], and Sensoy, Wang, and Weisbach [2014] show evidence that LPs have heterogeneous skills. Glode and Green [2011] present a similar model for hedge funds where the source of asymmetry is knowledge about the investment strategy of the manager. Marquez, Nanda, and Yavuz [2012] consider the information asymmetry between entrepreneurs and general partners.

Following Kaplan and Schoar [2005], several studies have investigated the persistence of PE performance using the AR(1) model.[5] 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 [2013] find persistence in a more recent sample than the original Kaplan and Schoar study, whereas Harris, Jenkinson, Kaplan, and Stucke [2014] and Braun, Jenkinson, and Stoff [2013] find that persistence has declined for BO firms post 2000. Li [2014] finds evidence of stronger persistence in BO compared to VC.

Our analysis focuses on cross-sectional heterogeneity in the performance within PE firms. We do not address questions related to aggregate skill and performance, or whether this asset class 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], and Harris, Jenkinson, and Kaplan [2014].

Methodologically, our model is related to the factor model approach to measuring performance persistence in the mutual fund and hedge fund literatures, though the application to PE is different due to the lower frequency of PE performance data and the overlapping fund issue.[6] Other contributions of our study are the introduction of the speed of learning as a measure of investable persistence, and the generalization to non-normal return distri-

[5] Persistence is also sometimes studied by estimating transition probabilities across fund quartiles. The assertion is that if fund returns are *i.i.d.*, then the probability that a top-quartile fund remains top quartile is 25%. More generally, $P[y_{i,N} \in Q | y_{i,N-1} \in Q] = 25\%$ where Q contains the performance for any quartile. Hence, the empirical finding that $P[y_{i,N} \in Q | y_{i,N-1} \in Q] \neq 25\%$ implies that performance cannot be *i.i.d.*, which is sometimes interpreted as evidence of persistence. However, this is neither a necessary nor sufficient condition. For example, suppose there are two types of PE firms, with an equal number of each type. The first type returns either +10% or -10% with equal probability, and the second type returns +25% or -25%, also with equal probability. Hence, the returns for the four quartiles are: +25%, +10%, -10%, and -25%. For each quartile, the transition probability is $P[y_{i,N} \in Q | y_{i,N-1} \in Q] = 50\%$, so returns are not *i.i.d.*, but there is no persistence in the conventional sense. Conversely, finding transition probabilities of $P[y_{i,N} \in Q | y_{i,N-1} \in Q] = 25\%$ does not imply an absence of persistence. Hence, the economic magnitudes and statistical significance of persistence are difficult to evaluate using transition probabilities across quartiles. Billingsley [1961] surveys the statistical issues that arise when estimating parameters and testing hypotheses in Markov chains.

[6] Within the vast literature on mutual and hedge fund persistence, our paper is most closely related to Barras, Scaillet and Wermers [2010] and Ferson and Chen [2014], who estimate the proportion of unskilled mutual fund managers in a frequentist setting. Bayesian methods to evaluate skill and persistence in mutual fund returns have been employed by Baks, Metrick, and Wachter [2001], Pastor and Stambaugh [2002a,b], Jones and Shanken [2005], Avramov and Wermers [2006], and Busse and Irvine [2006]. Finally, Kosowski, Timmermann, Wermers, and White [2006] develop a bootstrap method to allow for non-normal distributions of mutual fund alphas.

butions.

Our methodology of controlling for overlap of contemporaneous funds may also have application for mutual funds, where a manager may make similar investments across multiple funds under their control. For example, Wu, Wermers, and Zechner [2013] show that multi-fund managers are common and the average mutual fund manager in their sample manages 2.2 funds at a given time. More generally, separating skill from luck is a fundamental issue in economics and finance (and the broader social sciences), and our method may be useful for other applications, such as studies of performance persistence for serial entrepreneurs (e.g., Gompers, Kovner, Lerner, and Scharfstein [2010] and Bengtsson [2013]), corporate performance, CEOs, managerial performance, etc.

The paper proceeds as follows. In Section I we present the data. Section II presents our empirical model. Section III presents our results and discusses the evidence for long-term persistence in private equity performance. Section IV evaluates investable persistence. Section V analyzes various subsamples of the data, and section VI concludes.

I 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 voluntary reporting by some GPs and LPs. For each fund, Preqin reports aggregate fund performance, such as the internal rate of return (IRR) and the Total Value to Paid-In Capital multiple (TVPI). One limitation is that the data do not contain cash flows between LPs and the GP, and we cannot calculate the public market equivalent (PME) measure of fund performance, which has some advantages for evaluating performance (see Korteweg and Nagel [2014] and Sorensen and Jagannathan [2014]). We focus on the IRR, which is the annualized return to LPs net of performance fees ("carried interest" or "carry") and management fees. While the IRR has well-known limitations, it is the most widely available fund performance measure, and it is commonly used to analyze PE performance.[7] The IRR is an absolute performance measure, but our model controls for general market performance and systematic risk as discussed below.

Harris, Jenkinson, and Kaplan [2014] compare several datasets with PE fund performance. Most of these data are from commercial data providers (Preqin, Burgiss, and Cambridge Associates) and one dataset is from a large anonymous LP (studied by Robinson and Sensoy [2013]). For BO funds, they find that Preqin contains the largest total number of funds in the 1990s and 2000s (but not in the 1980s). For VC funds, Preqin has slightly weaker coverage in the 1980s and 1990s, but it is the most comprehensive dataset in the 2000s. 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 datasets from Burgiss and Cambridge Associates, though he only considers BO funds. Hence, when analyzing the performance and persistence of PE funds, the Preqin data are among the best data sets available.

The two main fund types are VC and BO funds, but Preqin also classifies funds as real-estate, fund-of-funds, infrastructure, turn-around, special situations, co-investment, and venture debt funds, which we collectively refer to as Other funds. The majority of these Other funds are real-estate and funds-of-funds, and while these two fund types are quite different, we find that they have (surprisingly) similar performance and persistence, and we combine all of these other fund types for most of our analysis.

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

[7] Other 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 [2013]; Robinson and Sensoy [2013]; Harris, Jenkinson, and Kaplan [2014]; Hochberg, Ljungqvist, and Vissing-Jorgensen [2014]; Li [2014]; and Sensoy, Wang and Weisbach [2014].

A Sample

We restrict our sample to funds with available performance information.[8] Our model does not require a fund to be preceded by another fund, and we include PE firms with only a single fund. To avoid concerns about funds' self-reported intermediate IRRs and NAVs (see Phalippou and Gottschalg [2009], Brown, Gredil and Kaplan [2014], Jenkinson, Sousa, and Stucke [2013], and Barber and Yasuda [2013]), we restrict our sample to fully liquidated funds. We eliminate funds with less than $5 million in committed capital (in 1990 U.S. dollars) to exclude smaller, idiosyncratic funds.

** TABLE I: SUMMARY STATISTICS **

** FIGURE 1: NUMBER OF FUNDS PER FIRM **

Table I shows summary 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.[9] 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, so including these firms may be important to assess the overall skill distribution.[10] The overlap of subsequent funds is important when assessing persistence. Table I shows that the average (median) overlap between two funds managed by the same PE firm is 5.8 years (6 years) for VC and BO, and 6.8 years (7 years) for Other.[11] Figure 1 shows histograms of the number of funds per PE firm, by fund type.

[8] Our estimates remain valid with randomly missing fund performance data. 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 that observed data is a sufficient statistic for survivorship. Similar (or stronger) assumptions are common in studies of performance persistence, and it is used, for example, in Baks, Metrick, and Wachter [2001] and Pastor and Stambaugh [2002a] when studying mutual funds. Formally, it always holds that $p(\gamma|data, survival) = [p(survival|data, \gamma)/p(survival|data)] \cdot p(\gamma|data)$, where γ are the model parameters. The γ estimates are valid when $p(\gamma|data, survival) = p(\gamma|data)$, which holds when the fraction in brackets equals 1.

[9] The 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.

[10] Our qualitative conclusions go through when using only PE firms with at least two funds in the data set.

[11] Note that we 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 will manage five partially overlapping funds (raised in years 0, 2, 4, 6, and 8). These five funds form ten fund pairs.

In terms of size, the average (median) VC fund has \$207 million (\$110 million) of committed capital, compared to \$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 I show that late-stage VC funds and distressed debt funds tend to be larger, whereas early-stage VC and natural resource funds tend to be smaller. There are no subclassifications of BO funds.

** TABLE II: FUND IRRs BY VINTAGE YEAR **

In terms of performance, Panel A of Table I 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 II shows IRRs by vintage year and fund type, and Figure 2 plots average IRRs. For VC funds we see very strong performance 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 still exposed to this bubble and show 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.

** FIGURE 2: IRRs BY VINTAGE YEAR **

Our empirical analysis uses total log-returns (i.e., continuously compounded returns) rather than 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:

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

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

II Variance Decomposition Model

Our empirical model is an *hierarchical linear model*, which generalizes the classical *analysis of variance* (ANOVA) decomposition. We exploit recently developed advances in numerical computing (Markov chain Monte Carlo, Gibbs sampling, and posterior augmentation) that make Bayesian methods particularly suited for estimating these models.

Hierarchical linear models were initially used for educational measurement, 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.[12] 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 also be extended to include data at additional levels, for example with data for individual deals, as in Braun, Jenkinson, and Stoff [2013] or with LPs' holdings of PE funds, as in Sensoy, Wang, and Weisbach [2014].

Hierarchical models address the *unit of analysis* problem (Burstein, Fischer, and Miller [1980]). When studying the persistence of PE performance, we are interested in differences in expected returns between PE firms, so the unit of analysis is the PE firm. However, the unit of observation is the underlying funds, which constitute a repeated measure of each PE firm's expected performance. Increasing the number of funds per firm improves the estimate of each firm's expected return, but not the number of firms that are compared. With few PE firms but many funds per firm, observing more funds per firm becomes uninformative, because the main sampling error arises from the sampling of the firms, not the funds. In contrast, observing more firms always improves the estimates. It is difficult for classical regression models, for example using PE firm fixed effects (FEs), to address this sampling problem, because these models only consider the sampling of funds for a given set of PE firms (i.e., a given set of PE firm FEs), not the sampling of the PE firms themselves (i.e., the sampling of these FEs from a larger population of potential FEs). Intuitively, our model can be thought of as consistently estimating the population variance of the FEs, which reflects the variance in the expected returns across PE firms.

[12] General overviews and discussions of hierarchical models are in Raudenbush and Bryk [2002] and De Leeuw and Meijer [2008].

A Economic Intuition

To illustrate the intuition behind our variance decomposition, consider a set of 60 PE firms. Suppose a firm makes two investments (or manages two funds), each of which either succeeds or fails. For the resulting 120 investments, if we observe one half failing and the other half succeeding, the unconditional success probability is 50%. If the investments were statistically independent, each firm would have 25% probability of zero successful investments, a 50% probability of a single success, and a 25% chance of two successes. We would then see 15 of the 60 PE firms with no successes, 30 with a single success, and the remaining 15 firms with two successful investments. 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 successes. In other words, the performance variation *between* PE firms exceeds the variation that would be implied by the variation *within* PE firms if they were statistically independent. In this case, investments cannot be independent, obviously, so some PE firms must have higher (or lower) success probabilities. In other words, some PE firms persistently show better (or worse) performance. 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 $E[p_i^2] = 33\%$. Based on this intuition, our model defines and measures persistence by comparing performance variation *within* PE firms to performance variation *between* firms. Excess variation between firms, as in this example, implies persistence.[13]

This intuition leads to a natural distinction between PE firms' past performance and their expected future performance. Continuing the example, with $p_i \sim U[0,1]$ and using Bayes rule, the posterior density of p_i conditional on observing two past successes is $f(p_i|SS) = 3p_i^2$. The probability that a firm with top-tercile past performance (two successes) has top-tercile expected future performance is just $\Pr(p_i \in [0.66,1]|SS) = 70\%$. The success probability of a new investment by firms with top-tercile past performance is

[13]Excess variation, and hence persistence, also implies a positive covariance in the outcomes of the investments made by each PE firm. In the example, let s_1 and $s_2$denote the outcomes of the two investments by a given PE firm, where each variable equals one if the corresponding investment is successful and is zero otherwise. When $p_i \sim U[0,1]$, the covariance between these outcomes is $\text{cov}(s_1,s_2) = E[s_1 s_2] - E[s_1]E[s_2] = 1/12$.

$E[p_i|SS] = 75\%$ whereas the success probability for firms with actual top-tercile expected performance is $E[p_i|p_i > 66\%] = 83\%$.

In this example, the distribution of the success probably is perfectly known (e.g., it is known that $p_i \sim U[0,1]$). In practice, when this distribution is estimated, estimation error introduces additional noise for identifying the top firms. Our formal model, which we discuss next, accounts for this parameter uncertainty as well.

B Formal Model

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 other characteristics of the fund and firm. The ten-year total log-return of fund u managed by firm i is specified as:

$$y_{iu} = X_{iu}'\beta + \sum_{\tau=t_{iu}}^{t_{iu}+9} (\gamma_i + \eta_{i\tau}) + \varepsilon_{iu}\,. \tag{3}$$

The sum runs over the fund's ten-year life, with year t_{iu} denoting the fund's first year of operation (vintage year). The three random effects that determine the covariance structure are γ_i, $\eta_{i\tau}$, and ε_{iu}. The fund-specific covariates, X_{iu}, contain either a single intercept term (Specification I) or vintage year fixed effects (Specification II). We estimate the model separately for each type of PE fund (VC, BO, and Other).

Random effects The γ_i term is constant for all funds managed by the same PE firm, and it captures long-term persistence. At birth, each PE firm receives an independent draw of γ_i, distributed $\gamma_i \sim \mathcal{N}\left(0, \sigma_\gamma^2\right)$, which remains constant throughout the firm's life. 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 above). The variation in γ_i across PE firms reflects differences in expected returns across PE firms. When there is little variation in γ_i, corresponding to a small σ_γ^2, then PE firms are similar, and there is little long-term persistent difference in their performances. 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 other random effects) is normalized to zero, and the industry level of expected returns is captured by the constant term in X_{iu}. The model is parameterized with γ_i inside the sum in equation (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.[14]

The covariance in the returns of partially overlapping PE funds is captured by the PE firm-time specific effect, distributed $\eta_{i\tau} \sim \mathcal{N}\left(0, \sigma_\eta^2\right)$ *i.i.d.*[15] 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 funds' performances. Economically, such correlations may arise due to common strategies, common risk exposures, or common investments in the same portfolio firms.[16] To illustrate, a PE firm that manages two funds with vintage years 1999 and 2001 may be focusing on investments in emerging markets, say, and 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 and significant coefficient, but this coefficient would not be evidence of persistence, as usually defined. It does not imply that past performance predicts future performance. When these spurious correlations are large, the estimated σ_η^2 is large.[17] Conversely, when they are small, the estimated σ_η^2 is small. In the limit, when there is no effect of shared exposures, the variance converges to zero.

The error term ε_{iu} captures fund-specific idiosyncratic performance shocks. It is *i.i.d.* across funds, across firms, and over time. Fund performance is skewed, and 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 Buyout and Other funds is captured by mixtures with just one or two normal distributions. More details are

[14]Equivalently, if γ_i were outside the sum, then it represent the total ten-year (not annualized) expected return of firm i relative to its peers.

[15]Although operating over much longer time scales, the autocorrelation due to the overlap of subsequent funds is closely related to the autocorrelation that arises from non-synchronous trading, which was introduced by Fisher [1966] and which has been extensively studied.

[16]Braun, Jenkinson, and Stoff [2013] find that an average BO fund with 15.6 investments contains 1.3 deals that are common with another fund that is managed by the same PE firm. They don't have data for VC or Other firms.

[17]We also estimate specifications that only account for overlap 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.

in the appendix.

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

$$\sigma_y^2 = 100\sigma_\gamma^2 + 10\sigma_\eta^2 + \sigma_\varepsilon^2\,, \tag{4}$$

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

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

This covariance relationship is plotted in Figure 3, and this figure also illustrates the identification of the model. The main parameters of interest are the variances of the three random effects, σ_γ^2, σ_η^2, and σ_ε^2. In Figure 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 equation (4).

We only observe each fund's total return, and the model cannot determine when this return is earned during each fund's life. The model *can* determine, however, how much of the variation in the funds' total performance is due to each of the three random effects.

** FIGURE 3: OVERLAP AND COVARIANCE **

In Specification I, X_{iu} contains just 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 captured by the $\eta_{i\tau}$ terms. To explicitly control for these risk exposures, Specification II adds vintage year FEs 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 CAPM beta or the same loadings on the Fama-French risk factors), then vintage-year FEs capture the common performance component that is due to these exposures that are shared by all firms and funds. Under these conditions, the γ_i is the PE firm's risk-adjusted annualized expected return relative to other PE firms. Note that γ_i is centered at zero, by assumption. It does not show performance relative to the market overall, and it cannot tell, for example, whether PE outperforms in the aggregate. It only measures variation in risk-adjusted expected returns among PE firms.

Apart from the assumption that PE firms have the same contemporaneous factor exposures, this risk-adjustment is quite general.[18] Each vintage year has an independent FE, so risk exposures and factor premia may vary over time, for example due to trends in leverage and credit market conditions. The exposures may also vary by fund type, since we estimate the model separately for each type (VC, BO, and Other). We also estimate the model separately for finer subsamples.

III Results

A IRR Regressions

We first confirm the findings by Kaplan and Schoar [2005] using our data. Table III (which follows the layout of Table VII in Kaplan and Schoar) shows OLS regressions of $IRR_{i,N}$ on $IRR_{i,N-1}$ with various controls, including further performance lags. All specifications have vintage year FEs 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 III show estimates of the AR(1) model for each fund type. 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 it is smaller and insignificant when including fund $N-2$, although this weaker statistical

[18]See Hochberg and Rauh [2013] for another use of vintage year fixed effects to control for systematic risk.

result may be due to the smaller sample size.

In these specifications, fund $N-2$ may also overlap with fund N, and the positive coefficients may still reflect the overlap rather than actual performance persistence. Panel B of Table II shows estimates using a sample of just the funds that are entirely non-overlapping. This restriction further reduces the sample size and leaves no remaining signs of persistence, but this weaker result may now be due to the lower statistical power of the smaller sample.

** TABLE III: IRR REGRESSIONS **

A natural interpretation of the AR(1) results in Table III is that BO funds have the most 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 (smallest coefficients and R^2, and insignificant with fund $N-2$). This ranking from BO, VC to Other funds, however, changes when we distinguish the different forms of persistence using our model.

B Long-Term Persistence

Table IV reports estimates of our model, estimated separately 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).[19] 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, we use two specifications. In Specification I, X_{iu} only contains an intercept. Specification II includes vintage year FEs. We first discuss Specification I.

** TABLE IV: PARAMETER ESTIMATES **

[19] We use a Bayesian estimator, but we discuss results using standard frequentist terminology: The "point estimate" is the mean of the posterior distribution, and the "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 non-linear transformations of the estimated parameters, and all reported inference is calculated this way and does not use any asymptotic approximations.

B.1 Venture Capital

For VC funds, Specification I in Table IV shows a total unconditional variance (σ_y^2) of 6.933. This variance can be 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).

The expected abnormal annual return that a PE firm earns, relative to the average return, is given by γ_i, which is distributed $\mathcal{N}\left(0, \sigma_\gamma^2\right)$. We find that σ_γ^2 is statistically significant, consistent with the findings from the AR(1) regressions.[20] 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 (i.e., worst) top-quartile VC firm has a γ_i of 3.30% annually. The spread in the expected returns of the marginal top- and bottom-quartile firms is then $q_\gamma(75\%) - q_\gamma(25\%) = 6.60\%$ annually. These percentiles are calculated from the point estimate of the standard deviation, assuming it is perfectly estimated. In Specification I in Table IV the estimate of σ_γ for VC funds has a standard error of 0.007. Our estimation procedure, however, generates the full posterior distribution of σ_γ, and using this distribution we generate the corresponding posterior distribution of the spread $q_\gamma(75\%) - q_\gamma(25\%)$. The estimate of the spread based on this posterior distribution now accounts for the parameter uncertainty. The resulting spread is 6.59%, as reported in Table IV. This estimate is very close to 6.60%, which was calculated from the point estimate of σ_γ, suggesting that parameter uncertainty is a minor problem for these spreads. Nevertheless, because the calculation is simple, all reported gamma spreads in Table IV 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.

Note that these gamma spreads cannot be calculated by simply subtracting the observed IRRs of the bottom quartile funds from those of the 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 where performance

[20]Testing 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.

is noisy, so σ^2_ε is large, but there is little difference between PE firms, so σ^2_γ is small. In this case, there is little long-term persistence and the spread $q_\gamma(75\%) - q_\gamma(25\%)$ is small. The noisy performance, however, still leads to a large difference in observed fund IRRs, so the empirically observed difference between top- and bottom-quartile funds may still be large, albeit entirely due to random noise. Conversely, the empirical difference may also understate long-term persistence. In periods when many particularly high-quality (or low-quality) PE firms are active, the empirical difference may be too small, because it is calculated from funds in a narrower range of the γ_i distribution. For this reason, it is important that our model accommodates PE firms that only manage a single fund. These firms are likely to be from the lower tail of the γ_i distribution, and excluding them could introduce a downward bias in σ^2_γ and underestimate the amount of long-term persistence.

B.2 Buyout

For BO funds, the variance that is due to long-term persistence ($100 \times \sigma^2_\gamma$) is 0.361, which is higher than for VC funds. Hence, the gamma spreads are also higher. 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 effect ($10 \times \sigma^2_\eta$) is 0.216, which is smaller than for VC funds, although this difference disappears with vintage-year FEs. Importantly, BO funds have smaller idiosyncratic risk (σ^2_ε) than VC funds, so the returns of BO funds are less noisy. The smaller noise in the returns of BO funds may also explain the stronger persistence results for BO in the AR(1) regressions, since less noisy returns give stronger statistical power in the AR(1) model.

B.3 Other

For Other funds the overlap and long-term persistence effects are largely similar to those for BO and VC 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 also similar. The idiosyncratic volatility of Other funds, however, is significantly lower than the volatilities of both VC and BO, resulting in a better signal-to-noise ratio for Other funds. We discuss the signal-to-noise ratio in more detail in the next section.

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. Moreover, these gamma spreads are calculated from the population distribution across PE firms, so 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 may earn by identifying PE firms with higher expected returns.

C Overlap Effects

The returns of overlapping funds are correlated. In Table IV Panel B, Specifications I and II for VC funds show overlap effects of 0.675 and 0.386. Without vintage-year FEs, 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. As discussed above, to control for these shared exposures, Specification II includes vintage-year FEs. 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. When evaluating two overlapping outperforming funds of the same PE firm, most likely, more of their outperformance is due to their overlap than to the firm generating persistently higher returns.

The estimates in Table IV allow us to quantify the effect of the 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 actually predict future performance. Table I shows an average overlap of subsequent funds of 5.8 to 6.8 years. Using equation (5), the estimates in Table IV imply that funds with such average overlaps have total covariances of 0.37 to 0.64. But 25.8% to 61.8% of this covariance is due to the overlap, suggesting that the AR(1) coefficients as a measure of performance predictability may be upward biased by 34% to 168%. The overlap effect for Other funds is smaller, implying a smaller upward bias in the AR(1) coefficients.

IV Learning and Investable Persistence

The estimates from the previous section show that there are substantial 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 gamma 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 intuition based on the updating of beliefs about gamma. The disadvantage of this ratio is that it assumes normal distributed returns. Consequently, as a second approach, we use the full model to estimate how many past funds an LP must observe to estimate a PE firm's expected return with reasonable certainty. Overall, 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. We find that an LP needs to observe an excessive number of past funds to evaluate a firm's expected return with reasonable certainty. In practice, LPs need additional information, such as detailed information about individual deals, individual partners associated with these deals (see Ewens and Rhodes-Kropf [2014]), or a firm's internal organization and culture.

A Signal-to-Noise Ratios

Our model contains two types of shocks: Transitory shocks are drawn independently each period, and are given by the $\eta_{i\tau}$ and ε_{iu} terms. The persistent shock, as given by γ_i, generates the long-term heterogeneity 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:[21]

$$s_\gamma = \frac{100\sigma_\gamma^2}{\sigma_y^2}. \tag{6}$$

This signal-to-noise ratio, which is bounded between zero and one, has a straightforward 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 $\mathcal{N}\left(\gamma_{i,N}, \sigma_{i,N}^2\right)$. After ob-

[21] For example, Cochrane [1988] uses a similar variance ratio to evaluate the persistence of GDP shocks.

serving the performance of one additional (non-overlapping) fund, the LP's updated beliefs are $\mathcal{N}\left(\gamma_{i,N+1}, \sigma^2_{i,N+1}\right)$ where:

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

and

$$\sigma^2_{i,N+1} = \left(1 - s_\gamma\right) \cdot \sigma^2_{i,N}\,. \tag{8}$$

Equation (7) gives the mean of the LP's updated beliefs. This updated mean is a combination of two terms, weighted by the signal-to-noise ratio. The first term (the fraction) is the new information, specifically the surprise performance of the new fund relative to 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 Equation (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, so the LP learns faster. When the signal-to-noise ratio is low, new performance is largely uninformative about the PE firm's expected return, and it is difficult for the LP to infer γ_i.

Point estimates of s_γ are reported in Table IV. Figure 4 plots the posterior distributions of s_γ for VC, BO, and Other firms, with and without vintage-year FEs. 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 the performance less informative. Other funds have less noise, and the best 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 **

B Identifying PE Firms

Figure 5 plots the probability that top-quartile performance implies top-quartile expected returns. Formally, this probability is $\mathrm{P}\left[\gamma_i \geq q_\gamma(75\%) \,|\, \frac{1}{N}\sum_{n=1}^{N} y_{i,n} \geq Q_N\right]$, where Q_N is the average performance of the marginal top-quartile firm with N past funds. To interpret Figure 5, consider the limit case where there is very little long-term persistence and σ^2_γ

converges to zero. In this uninformative case, 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 just 25%. As long-term persistence increases and performance becomes more informative, this probability increases. At the other limit, when σ_γ^2 becomes very large (relative to σ_y^2), persistence dominates. Now, top-quartile performance perfectly identifies the PE firms with top-quartile expected returns, and the probability converges to 100%. In practice, for a reasonable number of funds, the probability remains well below this limit. For example, Figure 5 shows that for VC firms with five past funds, having top-quartile past performance only implies a 37% probability of also having top-quartile expected future returns, which is only 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 improves 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 shows the probabilities for up to 50 past funds. As Figure 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 actually have top-quartile expected future returns. Additionally, these probabilities are calculated under the assumption that the funds are non-overlapping. With overlapping funds, learning will be slower.

** FIGURE 5: LEARNING SPEED **

C Investable Persistence

The previous analysis shows that it is easier to identify Other firms with high expected returns, but it does not account for the value of identifying these firms. As discussed above, Table IV shows that BO firms have more long-term persistence, as measured by σ_γ^2. Hence, even if it is more difficult, there is also more to gain from identifying top BO firms. Figure 6 shows the combination of these two effects. This figure plots the expected gamma for a PE firm with top-quartile performance, where top-quartile performance is calculated among all PE firms with a given number of past funds, N. Formally, Figure 6 plots $\mathrm{E}\left[\gamma_i \,\middle|\, \frac{1}{N}\sum_{n=1}^{N} y_{i,n} \geq Q_N\right]$. 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 4 to 5 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 5 or more funds, BO funds have the greatest investable persistence.

** FIGURE 6: INVESTABLE PERSISTENCE **

VC firms have poor investable persistence. Their signal-to-noise ratio is low, so it is difficult to identify better firms. Moreover, the long-term persistence and gamma spreads are also modest: they are similar to Other firms and well below the spread of BO firms. So the gains from identifying better VC firms are also lower. This overall poor investable persistence of VC firms clearly shows in Figures 5 and 6, where VC firms are dominated by Other and BO firms.

V Subsamples

We divide the sample into subsamples, and Table V reports estimates for each subsample independently. For example, when we compare the persistence of small and large funds in Panel A of Table V, the model is estimated twice, so 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_η^2, σ_ε^2, and risk factor loadings. The subsamples also provide a more nuanced picture of persistence. Unsurprisingly, our general finding is that PE firms in more developed and competitive markets, such as larger and more recent funds located in the U.S., show less performance persistence.

** TABLE V: SUBSAMPLES **

A Fund Size

Table V shows the long-term persistence of small and large funds. Across VC, BO, and Other firms, we find that smaller funds have more long-term persistence and greater gamma spreads than larger funds. This performance difference is not simply due to greater volatility of smaller funds. While smaller funds do have higher volatilities (except for Other funds), this volatility is captured by the σ_ε^2 term. Despite this greater volatility, the estimated σ_γ^2 parameters are also larger, so the signal-to-noise ratios are higher, and Panel B of Table V shows that the performance of smaller funds is more informative about the PE firm's expected returns.

B GP Location

Table V Panel A shows that PE firms located in the rest of the world (ROW) have more long-term persistence, followed by firms in Europe, while U.S.-based firms have the least persistence. Total volatility follows a different pattern, with ROW being most volatile, followed by the U.S., 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 U.S.-based funds is relatively less informative, which is consistent with the U.S. PE industry being more mature.

C Investment Style

VC and Other funds can be further categorized by their investment styles, as classified by Preqin. Table V 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 distinguish real estate from fund-of-funds. These fund types are very different, but they 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, performance of real estate is more informative than the performance of both VC and BO funds.

D Time Period

We confirm the findings by Braun, Jenkinson, and Stoff [2013] and Harris, Jenkinson, Kaplan, and Stucke [2014] that persistence has been declining. Table V shows estimates for the earlier and later half 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, although this decline is particularly pronounced for VC funds, while it is more marginal for BO and Other funds. This finding is consistent with the recent increase in LPs' focus on collecting additional information about the PE firms and their underlying funds.

VI Conclusion

We decompose the persistence of private equity (PE) performance into long-term, investable, and spurious persistence. Across all types of PE firms, we find a large amount of long-term persistence: the spread in expected returns between the top- and bottom-quartile PE firms is 7 to 8 percentage points, annually. In contrast, we find a small amount of 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. For example, after observing 50 past funds, only 53% to 61% of the PE firms that have generated top-quartile past performance also have top-quartile expected returns and can be expected to generate top-quartile performance in the future. In practice, to evaluate investments in PE funds, LPs need a substantial amount of information that goes beyond just the performance of past funds.

We find that the subsample of 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 U.S., followed by Europe, and the greatest persistence for firms located in the rest of the world ("ROW"), although the latter firms also have more volatile performance. Finally, we confirm the findings by Harris, Jenkinson, Kaplan, and Stucke [2014] and Braun, Jenkinson, and Stoff [2013] that persistence has declined over our sample period. This decline is largest for VC firms, though, and we find that BO and Other funds still show substantial remaining long-term persistence, even post 2000.

Our results have three practical implications: First, the low investable persistence may explain LPs' increasing focus on collecting detailed information about PE performance. For example, Ewens and Rhodes-Kropf [2014] study performance using deal- and partner-level information. We show that such detailed information is necessary for LPs to evaluate PE investments. Second, our results provide a new explanation for why persistence is not competed away. When identifying top PE firms is difficult, LPs with this ability may be as scarce as good PE firms, and these LPs should earn rents. Third, our results prove the economic realities behind the common saying among VCs that "I'd rather be lucky than smart."

Appendix: Estimation Procedure

We implement the model as a Bayesian multi-level hierarchical model, redefining the error terms to absorb the firm-specific random effects using hierarchical centering, as recommended by Gelfand, Sahu, and Carlin [1995]. The performance of fund u of firm i is

$$y_{iu} = X_{iu}\beta + \sum_{\tau=t_{iu}}^{t_{iu}+9} \eta_{i\tau} + \varepsilon_{iu}, \tag{A.1}$$

The conditional distributions of the random effects are given as

$$\eta_{i\tau}|\gamma_i \sim \mathcal{N}\left(\gamma_i, \sigma_\eta^2\right), \tag{A.2}$$

$$\gamma_i \sim \mathcal{N}\left(0, \sigma_\gamma^2\right). \tag{A.3}$$

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

$$\varepsilon_{iu} \sim \mathcal{N}\left(0, \sigma_\varepsilon^2\right). \tag{A.4}$$

We are interested in estimating the parameter vector $\theta \equiv \left(\beta, \sigma_\gamma^2, \sigma_\eta^2, \sigma_\varepsilon^2\right)$, given a dataset 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 $\eta' s$, and use a Bayesian estimation algorithm that produces a set of draws from the posterior distribution, $f\left(\theta, \{\gamma_i\}, \{\eta_{it}\} | data\right)$, using a Gibbs sampler (Gelfand and Smith [1990]. See also 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: $f\left(\{\eta_{it}\} | \{\gamma_i\}, \theta, data\right)$
2. Variance of fund-specific error term and β-coefficients: $f\left(\sigma_\varepsilon^2, \beta | \{\gamma_i\}, \{\eta_{it}\}, \sigma_\gamma^2, \sigma_\eta^2, data\right)$
3. Latent firm random effects: $f\left(\{\gamma_i\} | \{\eta_{it}\}, \theta, data\right)$
4. Variance of firm-year random effects: $f\left(\sigma_\eta^2 | \{\gamma_i\}, \{\eta_{it}\}, \beta, \sigma_\gamma^2, \sigma_\varepsilon^2, data\right)$
5. Variance of firm random effects: $f\left(\sigma_\gamma^2 | \{\gamma_i\}, \{\eta_{it}\}, \beta, \sigma_\eta^2, \sigma_\varepsilon^2, data\right)$

We sample from each distribution 1 through 5 in turn, after which 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 $f(\{\eta_{it}\}|data)$, for example. We now discuss how to draw from each conditional distribution.

A1 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

$$y_i = X_i\beta + Z_i\eta_i + \varepsilon_i, \tag{A.5}$$

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 $\left[X_{i1}' \ldots X_{iU_i}'\right]'$ for which each row correspond 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 may 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 equation (A.2), and using the standard Bayesian regression setup (e.g., Rossi, Allenby, and McCulloch [2005]), the posterior distribution is

$$\eta_i|\{\gamma_i\},\theta,data \sim \mathcal{N}\left(\mu_\eta,\sigma_\varepsilon^2\Omega^{-1}\right), \tag{A.6}$$

where

$$
\begin{aligned}
\Omega &= \frac{\sigma_\varepsilon^2}{\sigma_\eta^2} \cdot \mathbb{I}_{T_i} + Z_i' Z_i, && \text{(A.7)} \\
\mu_\eta &= \Omega^{-1} \left(\gamma_i \cdot \frac{\sigma_\varepsilon^2}{\sigma_\eta^2} \cdot 1_{T_i} + Z_i' (y_i - X_i \beta) \right), && \text{(A.8)}
\end{aligned}
$$

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

A2 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

$$
\begin{aligned}
\sigma_\varepsilon^2 &\sim I\mathcal{G}(a_0, b_0), && \text{(A.9)} \\
\beta | \sigma_\varepsilon^2 &\sim \mathcal{N}\left(\mu_0, \sigma_\varepsilon^2 \Sigma_0^{-1} \right), && \text{(A.10)}
\end{aligned}
$$

the posterior distribution is

$$
\begin{aligned}
\sigma_\varepsilon^2 | \{\eta_i\}, data &\sim I\mathcal{G}(a, b), && \text{(A.11)} \\
\beta | \sigma_\varepsilon^2, \{\eta_i\}, data &\sim \mathcal{N}\left(\mu, \sigma_\varepsilon^2 \Sigma^{-1} \right), && \text{(A.12)}
\end{aligned}
$$

where

$$
\begin{aligned}
a &= a_0 + \sum_{i=1}^{N} U_i, && \text{(A.13)} \\
b &= b_0 + e'e + (\mu - \mu_0) \Sigma_0 (\mu - \mu_0), && \text{(A.14)} \\
\Sigma &= \Sigma_0 + X'X, && \text{(A.15)} \\
\mu &= \Sigma^{-1} \cdot \left(\Sigma_0 \mu_0 + X' (y - Z\eta) \right). && \text{(A.16)}
\end{aligned}
$$

The vector $y = \left[y_i' \ldots y_N' \right]'$ contains the fund returns stacked across the N firms, X is the matrix of stacked X_i and Z the stacked Z_i. The vector $e = y - Z\eta - X\mu$ contains the stacked error terms.

A3 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

$$\eta = W\gamma + \nu, \tag{A.17}$$

where $\eta = [\eta_1 \ldots \eta_N]'$, and $\gamma = [\gamma_1 \ldots \gamma_N]'$, and $\nu \sim \mathcal{N}\left(0, \sigma_\eta^2 \cdot \mathbb{I}_N\right)$. 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 equation (A.3), the posterior distribution is

$$\gamma | \{\eta_{it}\}, \theta, data \quad \sim \quad \mathcal{N}\left(\mu_\gamma, \sigma_\eta^2 A^{-1}\right), \tag{A.18}$$

where

$$A = \frac{\sigma_\eta^2}{\sigma_\gamma^2} \cdot \mathbb{I}_N + W'W, \tag{A.19}$$

$$\mu_\gamma = A^{-1}\left(W'\eta\right). \tag{A.20}$$

A4 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

$$\sigma_\eta^2 \sim I\mathcal{G}\left(c_0, d_0\right), \tag{A.21}$$

yields the posterior distribution

$$\sigma_\eta^2 | \{\gamma_i\}, \{\eta_i\}, data \sim I\mathcal{G}\left(c, d\right), \tag{A.22}$$

where

$$
\begin{aligned}
c &= c_0 + \sum_{i=1}^{N} T_i, &\text{(A.23)}\\
d &= d_0 + v'v. &\text{(A.24)}
\end{aligned}
$$

A5 Variance of firm random effects

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

$$\sigma_\gamma^2 \sim \mathcal{IG}\left(f_0, g_0\right), \tag{A.25}$$

has posterior distribution

$$\sigma_\gamma^2 | \left\{\gamma_i\right\}, data \sim \mathcal{IG}\left(f, g\right), \tag{A.26}$$

with parameters

$$
\begin{aligned}
f &= f_0 + N, &\text{(A.27)}\\
g &= g_0 + \gamma'\gamma. &\text{(A.28)}
\end{aligned}
$$

A6 Mixture of Normals Specification

For the mixture-of-normals specification we replace the distribution of the fund-specific error term in equation (A.4), with a mixture of K normal distributions,

$$\varepsilon_{iu} \sim \sum_{k=1}^{K} p_k \cdot \mathcal{N}\left(\mu_k, \sigma_{\varepsilon,k}^2\right). \tag{A.29}$$

Setting K=1 reduces the model to the baseline normal specification in (A.4). We drop the intercept in X_{iu} because it is absorbed by the error term, which has mean $E\left[\varepsilon_{iu}\right] = \sum_{k=1}^{K} p_k \mu_k$. This specification is equivalent to the specification with an intercept in X_{iu} and zero mean error ε_{iu}, but it is easier to implement because it avoids enforcing cross-parameter restrictions on the μ_k. To estimate the mixture model by Gibbs sampler, the procedure requires one more latent variable that indicates which of the K normal distributions each observation is drawn from. Conditional on this indicator, the Gibbs steps described above remain largely unchanged. For details on estimation of the latent indicator and the parameters of the mixture components we refer to West [1992], Diebolt and Robert [1994], and Chen and Liu [2000].

A7 Priors and Starting Values

Our Gibbs sampler uses 10,000 iterations for the initial burn-in, followed by 100,000 iterations to simulate the posterior distribution. We save every 10th draw of the simulation. During the burn-in phase, the simulations converge quickly. We use diffuse prior distributions for the parameters, so that our results are driven by the data rather than 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 (note that this is for ten-year fund returns, so the annualized volatility is about a factor 3 lower). Second, we set $c_0 = f_0 = 2.1$, and $d_0 = g_0 = 0.15^2$. Since both the $\gamma's$ and η's are specified at the annual level, this implies that $E[\sigma_\gamma] = E[\sigma_\eta] = 0.128$ per year, and 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, $1,\ldots,K$, 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 below) is a valid comparison across different mixtures. The prior distribution of the mixture probabilities, p, is the conjugate Dirichlet distribution, $Dir(K,\delta)$, with $\delta = 1_K \cdot 10$. This 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, $1/K$. We do not need starting values for the η's, since they are the first variables we simulate.

A8 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 Chib's [1995] method 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. Figure 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 FEs). The optimal number of mixture distributions used in the main results of the paper (as reported in Table IV) are those with the highest marginal likelihood.

Second, we test for the presence of the long-term persistence and overlap random effects. To be precise, 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.

References

Acharya, Viral, Oliver Gottschalg, Moritz Hahn, and Conor Kehoe, 2013, Corporate Governance and Value Creation: Evidence from Private Equity, *Review of Financial Studies* 26, 368–402.

Ang, Andrew, Bingxu Chen, William N. Goetzmann, and Ludovic Phalippou, 2013, Estimating Private Equity Returns from Limited Partner Cash Flows, *Working Paper*, Columbia, Yale and Oxford Universities.

Avramov, Doron, and Russ Wermers, 2006, Investing in Mutual Funds when Returns are Predictable, *Journal of Financial Economics* 81, 339–377.

Baks, Klaas, Andrew Metrick, and Jessica Wachter, 2001, Should Investors Avoid All Actively Managed Mutual Funds? A Study in Bayesian Performance Evaluation, *Journal of Finance* 56, 45–85.

Barber, Brad, and Terrance Odean, 2000, Trading Is Hazardous to Your Wealth: The Common Stock Investment Performance of Individual Investors, *Journal of Finance* 55, 773–806.

Barber, Brad, and Ayako Yasuda, 2013, Interim Fund Performance and Fundraising in Private Equity, *Working Paper*, UC Davis.

Barras, Laurent, Olivier Scaillet, and Russ Wermers, 2010, False Discoveries in Mutual Fund Performance: Measuring Luck in Estimated Alphas, *Journal of Finance* 65, 179–216.

Bengtsson, Ola, 2013, Relational Venture Capital Financing of Serial Founders, *Journal of Financial Intermediation* 3, 308–344.

Berk, Jonathan, and Richard Green, 2004, Mutual Fund Flows and Performance in Rational Markets, *Journal of Political Economy* 112, 1269–1295.

Berkhof, Johannes, Iven van Mechelen, and Andrew Gelman, 2003, A Bayesian Approach to the Selection and Testing of Mixture Models, *Statistica Sinica* 13, 423–442.

Billingsley, Patrick, 1961, Statistical Methods in Markov Chains, *The Annals of Mathematical Statistics* 32, 12–40.

Braun, Reiner, Tim Jenkinson, and Ingo Stoff, 2013, How Persistent is Private Equity Performance? Evidence from Deal-Level Data, *Working Paper*, Oxford University.

Brown, Gregory, Oleg Gredil, and Steven Kaplan, 2014, Do Private Equity Funds Game Returns?, *Working Paper*, University of North Caroline and University of Chicago.

Brown, Stephen, William Goetzmann, and Roger Ibbotson, 1999, Offshore Hedge Funds: Survival and Performance 1989-1995, *Journal of Business* 72, 91–117.

Burstein, Leigh, Kathleen Fischer, and M. David Miller, 1980, The Multilevel Effects of Background On Science Achievement: A Cross National Comparison, *Journal of Sociology* 53, 215–225.

Busse, Jeffrey, Amit Goyal, and Sunil Wahal, 2010, Performance and Persistence in Institutional Investment Management , *Journal of Finance* 65, 765–790.

Busse, Jeffrey, and Paul Irvine, 2006, Bayesian Alphas and Mutual Fund Persistence, *Journal of Finance* 61, 2251–2288.

Carhart, Mark, 1997, On Persistence in Mutual Fund Performance, *Journal of Finance* 52, 57–82.

Caselli, Stefano, Emilia Garcia-Appendini, and Filippo Ippolito, 2013, Contracts and Returns in Private Equity, *Journal of Financial Intermediation* 22, 201–217.

Chen, Rong, and Jun S. Liu, 2000, Mixture Kalman Filters, *Journal of the Royal Statistical Society, Series B* 62, 493–508.

Chib, Siddhartha, 1995, Marginal Likelihood from the Gibbs Output, *Journal of the American Statistical Association* 90, 1313–1321.

Chung, Ji-Woong, 2012, Performance Persistence in Private Equity Funds, *Working Paper*, Chinese University of Hong Kong.

Chung, Ji-Woong, Berk Sensoy, Lea Stern, and Michael Weisbach, 2012, Pay for Performance from Future Fund Flows: The Case of Private Equity, *Review of Financial Studies* 25, 3259–3304.

Cochrane, John, 1988, How Big is the Random Walk Component in GDP?, *Journal of Political Economy* 96, 893–920.

De Leeuw, Jan, and Erik Meijer, 2008, Handbook of Multilevel Analysis, Springer, NY.

Diebolt, Jean, and Christian P. Robert, 1994, Estimation of Finite Mixture Distributions Through Bayesian Sampling, *Journal of the Royal Statistical Society, Series B,* 56, 363–375.

Ewens, Michael, and Matthew Rhodes-Kropf, 2014, Is a VC Partnership Greater Than the Sum of its Partners, forthcoming, *Journal of Finance*.

Ferson, Wayne, 2010, Investment Performance Evaluation, *Annual Review of Financial Economics* 2, 207–234.

Ferson, Wayne, and Yong Chen, 2014, How Many Good and Bad Fund Managers are there, Really?, *Working Paper*, University of Southern California and Texas A&M University.

Fisher, Lawrence, 1966, Some New Stock Market Indices, *Journal of Business* 39, 191–225.

Franzoni, Francesco, Eric Nowak, and Ludovic Phalippou, 2012, Private Equity Performance and Liquidity Risk, *Journal of Finance* 67, 2341–2373.

Gelfand, Alan E., Sujit K. Sahu, and Bradley P. Carlin, 1995, Efficient Parametrisations for Normal Linear Mixed Models, *Biometrika* 82, 479–488.

Gelfand, Alan E., and Adrian F.M. Smith, 1990, Sampling-based Approaches to Calculating Marginal Densities, *Journal of the American Statistical Association* 85, 398–409.

Glode, Vincent, and Richard C. Green, 2011, Information Spillovers and Performance Persistence for Hedge Funds, *Journal of Financial Economics* 101, 1–17.

Gompers, Paul, Anna Kovner, Josh Lerner, and David Scharfstein, 2010, Performance Persistence in Entrepreneurship, *Journal of Financial Economics* 96, 18–32.

Graham, John, and Campbell Harvey, 1996, Market Timing Ability and Volatility Implied in Investment Newsletters' Asset Allocation Recommendations, *Journal of Financial Economics* 42, 397–421.

Griffin, John, and Jin Xu, 2009, How Smart are the Smart Guys? A Unique View from Hedge Fund Stock Holdings, *Review of Financial Studies* 22, 2531–2570.

Gruber, Martin, 1996, Another Puzzle: The Growth in Actively Managed Mutual Funds, *Journal of Finance* 51, 783–810.

Harris, Robert, Tim Jenkinson, and Steven Kaplan, 2014, Private Equity Performance: What Do We Know?, forthcoming, *Journal of Finance*.

Harris, Robert, Tim Jenkinson, Steven Kaplan, and Ruediger Stucke, 2014, Has Persistence Persisted in Private Equity? Evidence From Buyout and Venture Capital Funds, *Working Paper*, University of Virginia, University of Chicago and Oxford University.

Higson, Chris, and Ruediger Stucke, 2012, The Performance of Private Equity, *Working Paper*, London Business School and Oxford University.

Hochberg, Yael, Alexander Ljungqvist, and Annette Vissing-Jorgensen, 2014, Informational Hold-Up and Performance Persistence in Venture Capital, *Review of Financial Studies* 27, 102–152.

Hochberg, Yael, and Joshua Rauh, 2013, Local Overweighting and Underperformance: Evidence from Limited Partner Private Equity Investments, *Review of Financial Studies* forthcoming, .

Jagannathan, Ravi, Alexey Malakhov, and Dmitry Novikov, 2010, Do Hot Hands Exist among Hedge Fund Managers? An Empirical Evaluation, *Journal of Finance* 65, 217–255.

Jenkinson, Tim, Howard Jones, and Jose Vicente Martinez, 2014, Picking Winners? Investment Consultants' Recommendations of Fund Managers, *Working Paper*, Oxford University.

Jenkinson, Tim, Miguel Sousa, and Ruediger Stucke, 2013, How Fair are the Valuations of Private Equity Funds?, *Working Paper*, Oxford University and University of Porto.

Jensen, Michael, 1968, The Performance of Mutual Funds in the Period 1945-1964, *Journal of Finance* 23, 389–416.

Jones, Chris, and Jay Shanken, 2005, Mutual Fund Performance with Learning across Funds , *Journal of Financial Economics* 78, 507–552.

Kacperczyk, Marcin, Clemens, Sialm, and Lu Zheng, 2005, On the Industry Concentration of Actively Managed Equity Mutual Funds, *Journal of Finance* 60, 1983–2011.

Kaplan, Steven, and Antoinette Schoar, 2005, Private Equity Performance: Returns, Persistence, and Capital Flows, *Journal of Finance* 60, 1791–1823.

Kaplan, Steven, and Per Stromberg, 2009, Leveraged Buyouts and Private Equity, *Journal of Economic Perspectives* 23, 121–146.

Kass, Robert E., and Adrian E. Raftery, 1995, Bayes Factors, *Journal of the American Statistical Association* 90, 773–795.

Koijen, Ralph, 2014, The Cross-Section of Managerial Ability, Incentives, and Risk Preferences, *Journal of Finance* 69, 1051–1098.

Korteweg, Arthur, 2013, Markov Chain Monte Carlo Methods in Corporate Finance, in P. Damien, P. Dellaportas, N. Polson, and D. Stephens, eds.: Bayesian Theory and Applications.

Korteweg, Arthur, and Stefan Nagel, 2014, Risk-adjusting the Returns to Venture Capital, *Working Paper*, NBER working paper #1934.

Korteweg, Arthur, and Morten Sorensen, 2010, Risk and Return Characteristics of Venture Capital-backed Entrepreneurial Companies, *Review of Financial Studies* 23, 3738–3772.

Kosowski, Robert, Allan Timmermann, Russ Wermers, and Hal White, 2006, Can Mutual Fund "Stars" Really Pick Stocks? New Evidence from a Bootstrap Analysis, *Journal of Finance* 61, 2551–2595.

Lerner, Josh, Antoinette Schoar, and Wan Wongsunwai, 2007, Smart Institutions, Foolish Choices? The Limited Partner Performance Puzzle, *Journal of Finance* 62, 731–764.

Li, Yi, 2014, Reputation, Volatility and Performance Persistence of Private Equity, *Working Paper*, Federal Reserve Board of Governers.

Ljungqvist, Alexander, and Matthew Richardson, 2003, The Cash Flow, Return and Risk Characteristics of Private Equity, *Working Paper*, New York University.

Ljungqvist, Alexander, Matthew Richardson, and Daniel Wolfenzon, 2008, The Investment Behavior of Buyout Funds: Theory and Evidence, *Working Paper*, NBER working paper #14180.

Lopez de Silanes, Florencio, Ludovic Phalippou, and Oliver Gottschalg, 2013, Giants at the Gate: Investment Returns and Diseconomies of Scale in Private Equity, forthcoming, *Journal of Financial and Quantitative Analysis*.

Malkiel, Burton, 1995, Returns from Investing in Mutual Funds, *Journal of Finance* 50, 549–572.

Marin, Jean-Michel, and Christian P. Robert, 2008, Approximating the Marginal Likelihood in Mixture Models, *Working Paper*, INRIA and Université Paris-Dauphine.

Marquez, Robert, Vikram Nanda, and Deniz Yavuz, 2012, Private Equity Fund Returns and Performance Persistence, *Working Paper*, UC Davis, Georgia Tech, and Purdue University.

Metrick, Andrew, and Ayako Yasuda, 2010, The Economics of Private Equity Funds, *Review of Financial Studies* 23, 2303–2341.

Pastor, Lubos, and Robert Stambaugh, 2002a, Mutual Fund Performance and Seemingly Unrelated Assets, *Journal of Financial Economics* 63, 315–349.

Pastor, Lubos, and Robert Stambaugh, 2002b, Investing in Equity Mutual Funds, *Journal of Financial Economics* 63, 351–380.

Phalippou, Ludovic, 2010, Venture Capital Funds: Flow-Performance Relationship and Performance Persistence, *Journal of Banking and Finance* 34, 568–577.

Phalippou, Ludovic, 2014, Performance of Buyout Funds Revisited?, *Review of Finance* 18, 189–218.

Phalippou, Ludovic, and Oliver Gottschalg, 2009, The Performance of Private Equity Funds, *Review of Financial Studies* 22, 1747–1776.

Raudenbush, Stephen, and Anthony Bryk, 2002, Hierarchical Linear Models: Applications and Data Analysis Methods, 2nd ed., Thousand Oaks: Sage Publication.

Robinson, David, and Berk Sensoy, 2013, Cyclicality, Performance Measurement, and Cash Flow Liquidity in Private Equity, *Working Paper*, Duke and Ohio State University.

Rossi, Peter E., Greg M. Allenby, and Robert McCullough, 2005, Bayesian Statistics and Marketing, John Wiley and Sons, Chichester, UK.

Sensoy, Berk, Yingdi Wang, and Michael Weisbach, 2014, Limited Partner Performance and the Maturing of the Private Equity Industry, *Journal of Financial Economics* 112, 320–343.

Sorensen, Morten, and Ravi Jagannathan, 2014, The Public Market Equivalent and Private Equity Performance, *Working Paper*, Columbia and Northwestern Universities.

Timmermann, Allan, and David Blake, 2005, International Asset Allocation with Time-varying Investment Opportunities, *Journal of Business* 78, 71–98.

Titman, Sheridan, and Cristian Tiu, 2011, Do the Best Hedge Funds Hedge?, *Review of Financial Studies* 24, 123–168.

Wermers, Russ, 2011, Performance Measurement of Mutual Funds, Hedge Funds, and Institutional Accounts, *Annual Review of Financial Economics* 3, 537–574.

Wu, Youchang, Russ Wermers, and Josef Zechner, 2013, Managerial Rents vs. Shareholder Value in Delegated Portfolio Management: The Case of Closed-End Funds, *Working Paper*, University of Wisconcin, University of Maryland, and Vienna University.

West, Mike, 1992, Modeling With Mixtures, in Bayesian Statistics 4, eds. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, Oxford, U.K.: Oxford University Press, 503–524.

Figure 1: Distribution of funds per firm. Histograms of the number of funds per firm, by type (VC, Buyout, and Other). For firms that manage multiple types of funds, each histogram only counts the number of funds of the particular type.

Panel A: VC

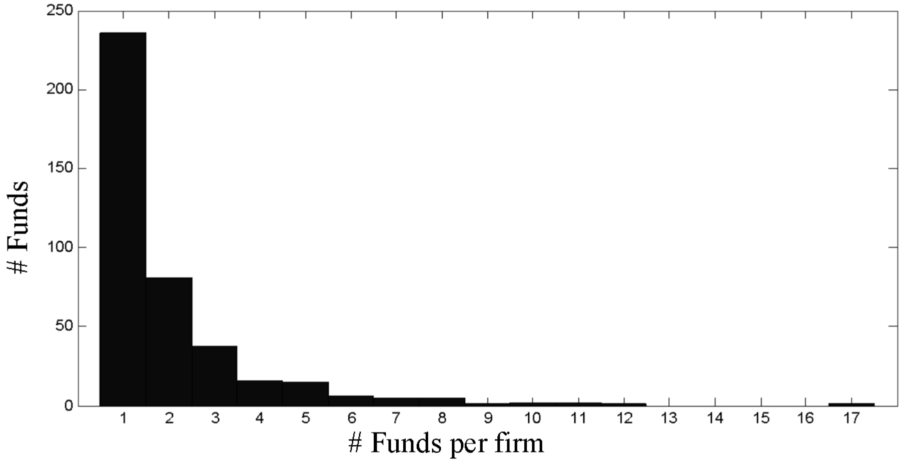

Panel B: Buyout

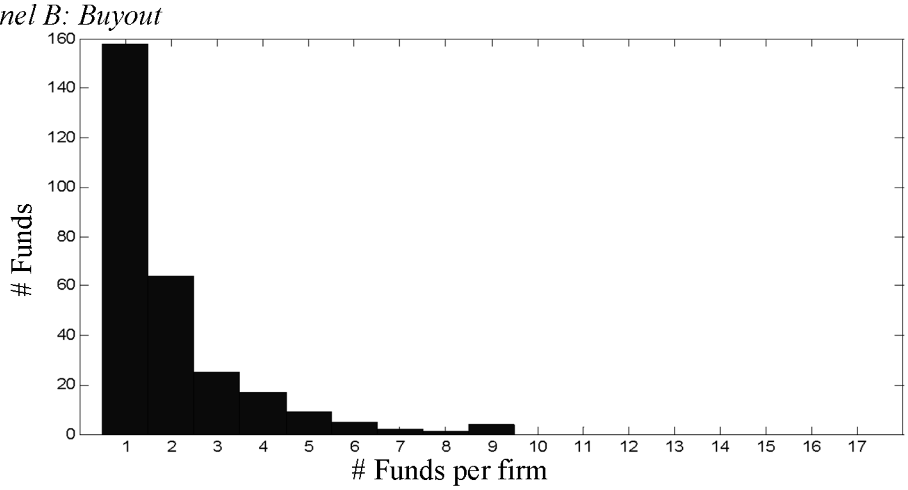

Panel C: Other

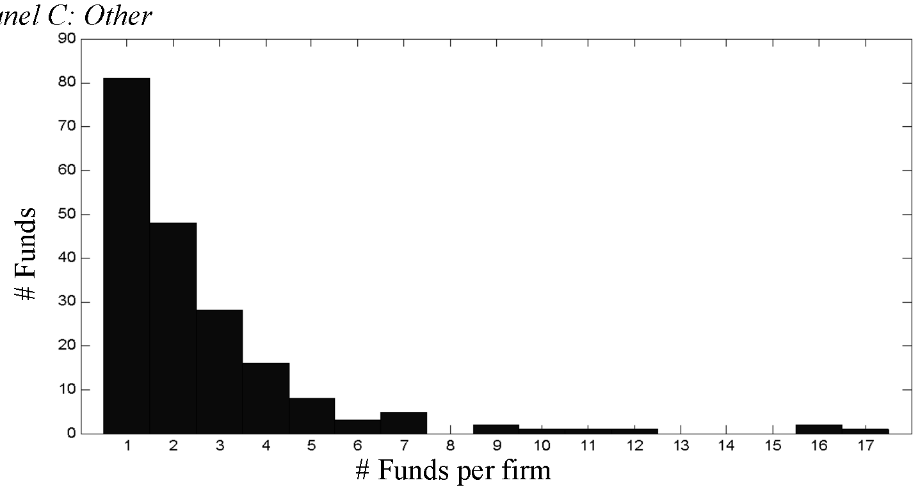

Figure 2: IRRs and Vintage Year Fixed Effects. The striped lines are the average fund IRR (per annum) in each vintage year, by fund type. The solid lines are the posterior means of the vintage year fixed effects from specification II in Table IV. The vintage year fixed effects are transformed to IRR equivalents (annualized, and in percent). The shaded bands represents the (1%, 99%) Bayesian credible interval (confidence bounds).

Panel A: VC

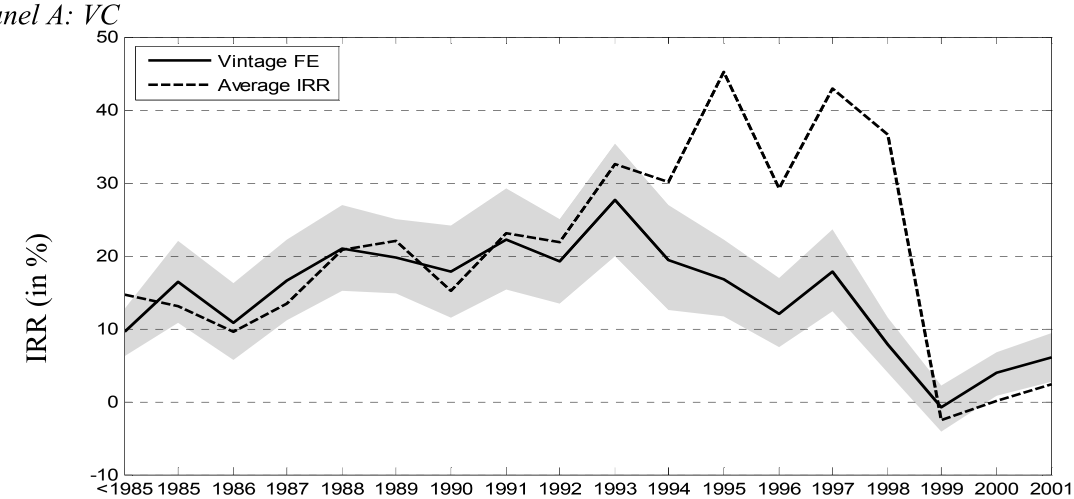

Panel B: Buyout

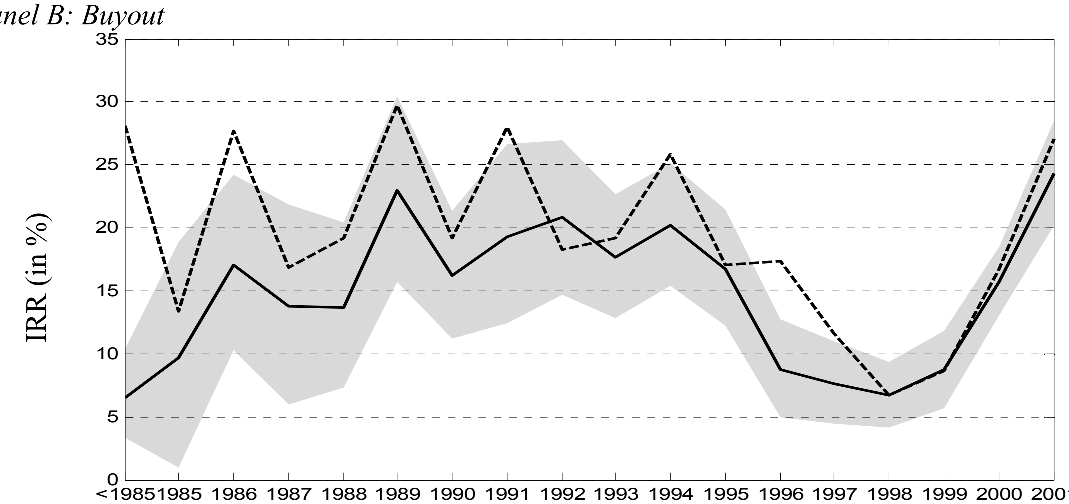

Panel C: Other

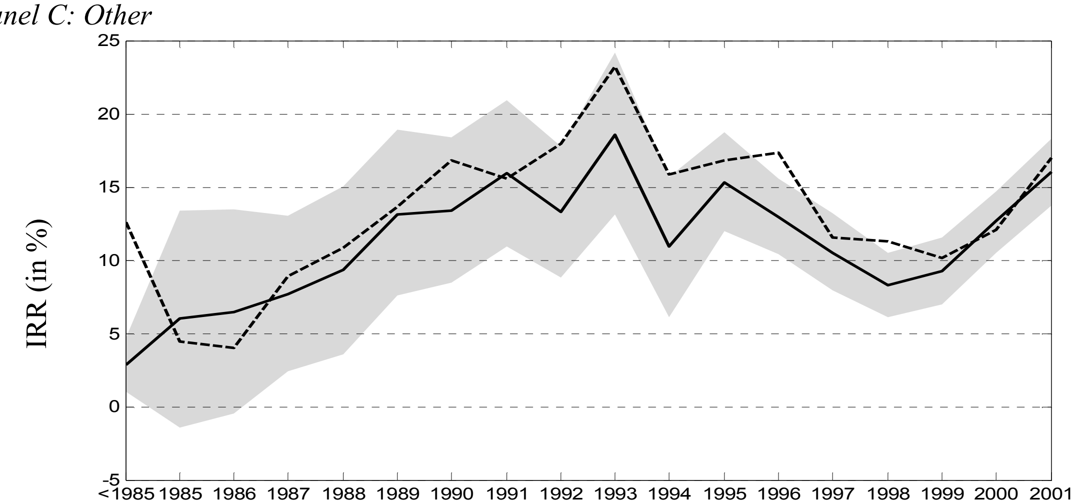

Figure 3: Fund Overlap and Covariance. The figure shows the covariance between total fund returns as a function of the overlap (in years) between two funds managed by the same firm, using the variance estimates for specification II in Table IV.

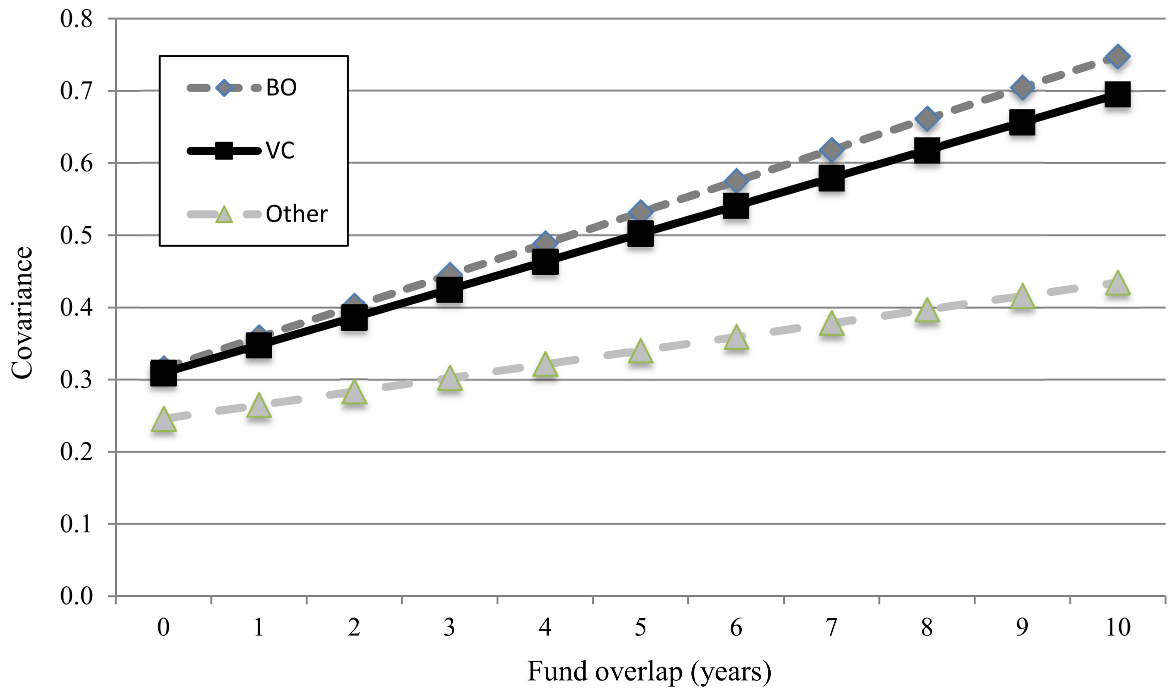

Figure 4: Estimates of Signal-to-Noise Ratio. Posterior distribution of the signal-to-noise ratio, s_γ, by fund type, from the specifications reported in Table IV. The solid line is the kernel plot for specification I (without vintage year fixed effects), and the striped line is the kernel plot for specification II in Table IV (with vintage year fixed effects).

Panel A: VC

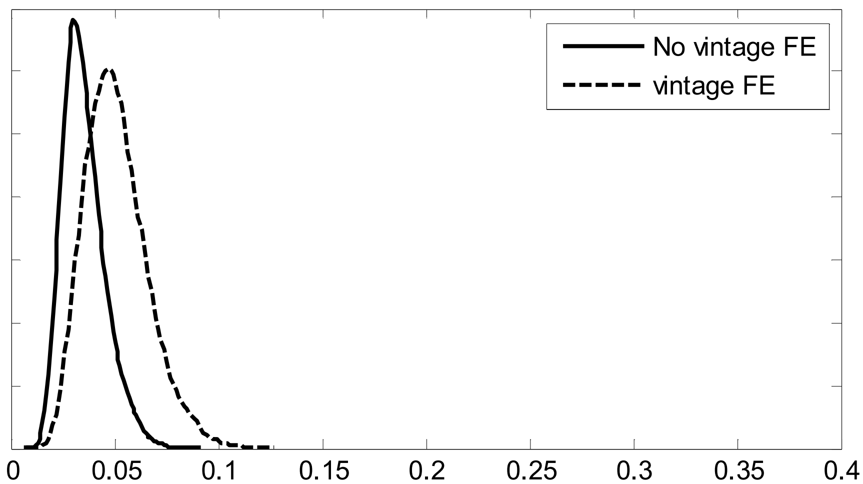

Panel B: Buyout

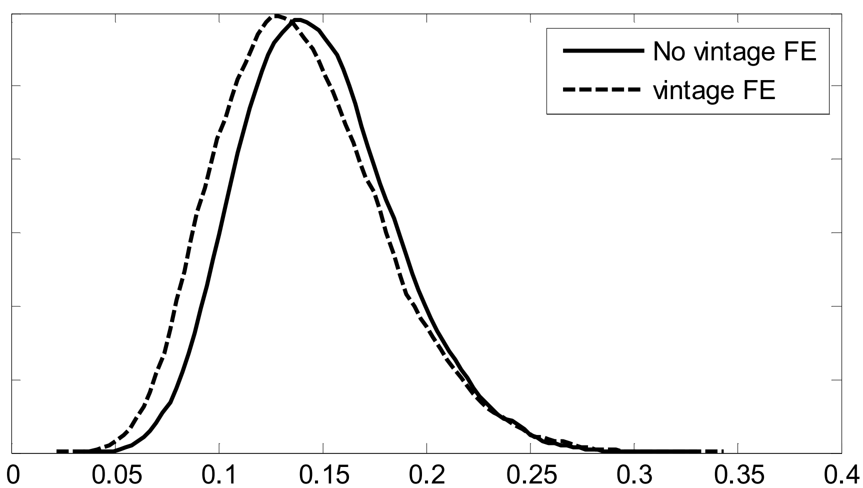

Panel C: Other

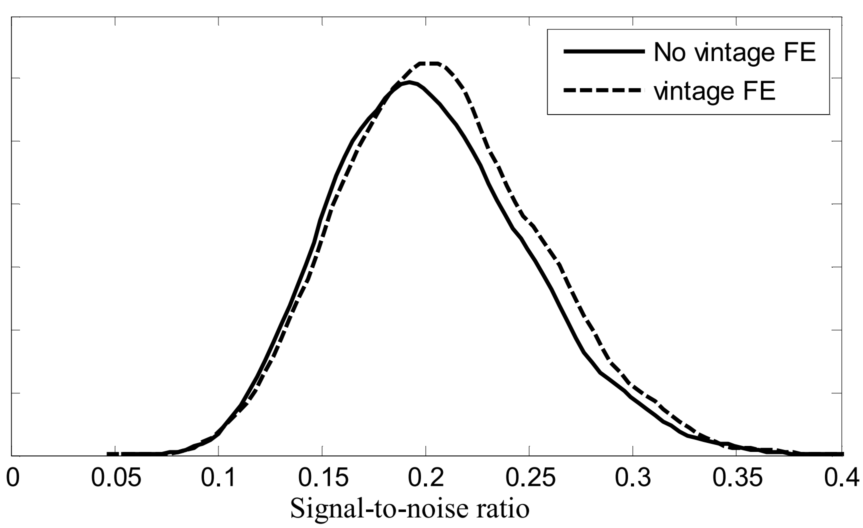

Figure 5: Speed of Learning. Graph of the posterior probability that a fund is in the top quartile of funds. Probabilities are calculated from 100,000 simulations of a panel of 100 firms, each with a different gamma that is drawn from the top 25% of the distribution. Each firm produces a sequence of 50 non-overlapping 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*, on the horizontal axis). The figure uses the parameter estimates from Table IV specification II, with vintage year fixed effects.

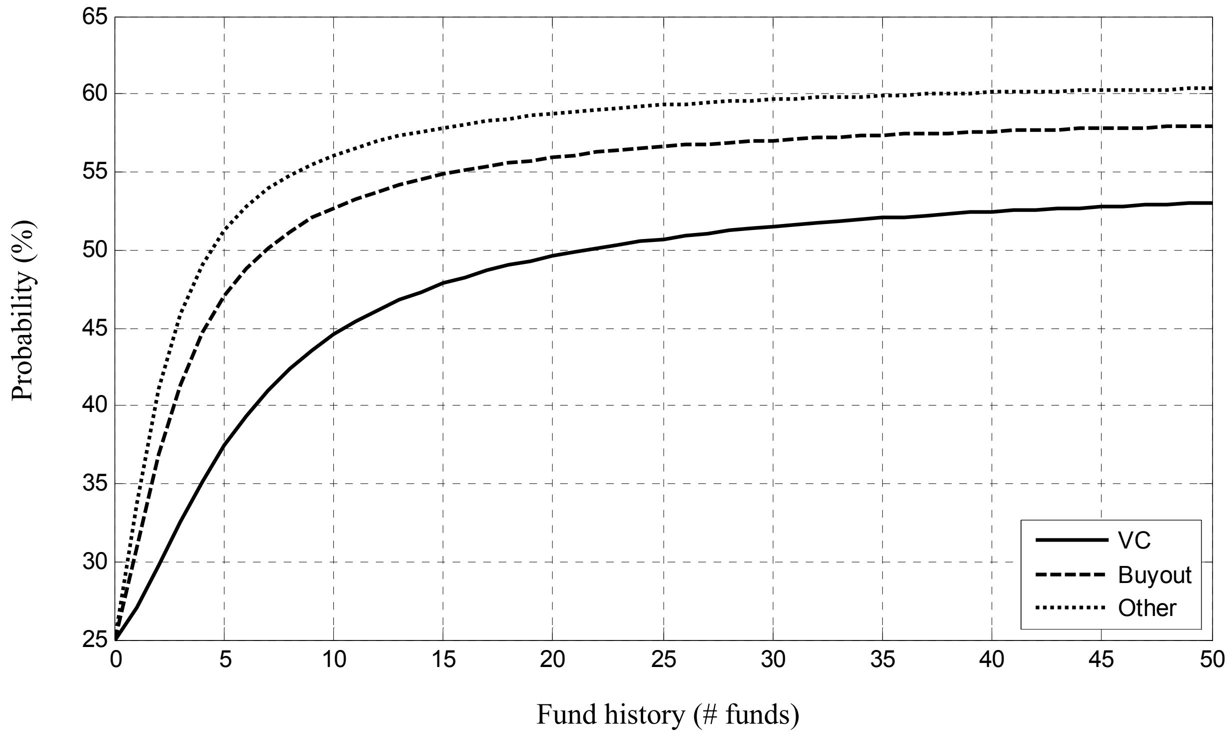

Figure 6: Investable Persistence. This figure shows the expected (true) gamma of investing in funds raised by PE firms with top-quartile performance as observed after a given number of realized fund returns for each firm (*Fund history*). Calculations are based on 100,000 simulations of a panel of fund histories for 100 firms, using the parameter estimates from Table IV specification II (with vintage year fixed effects).

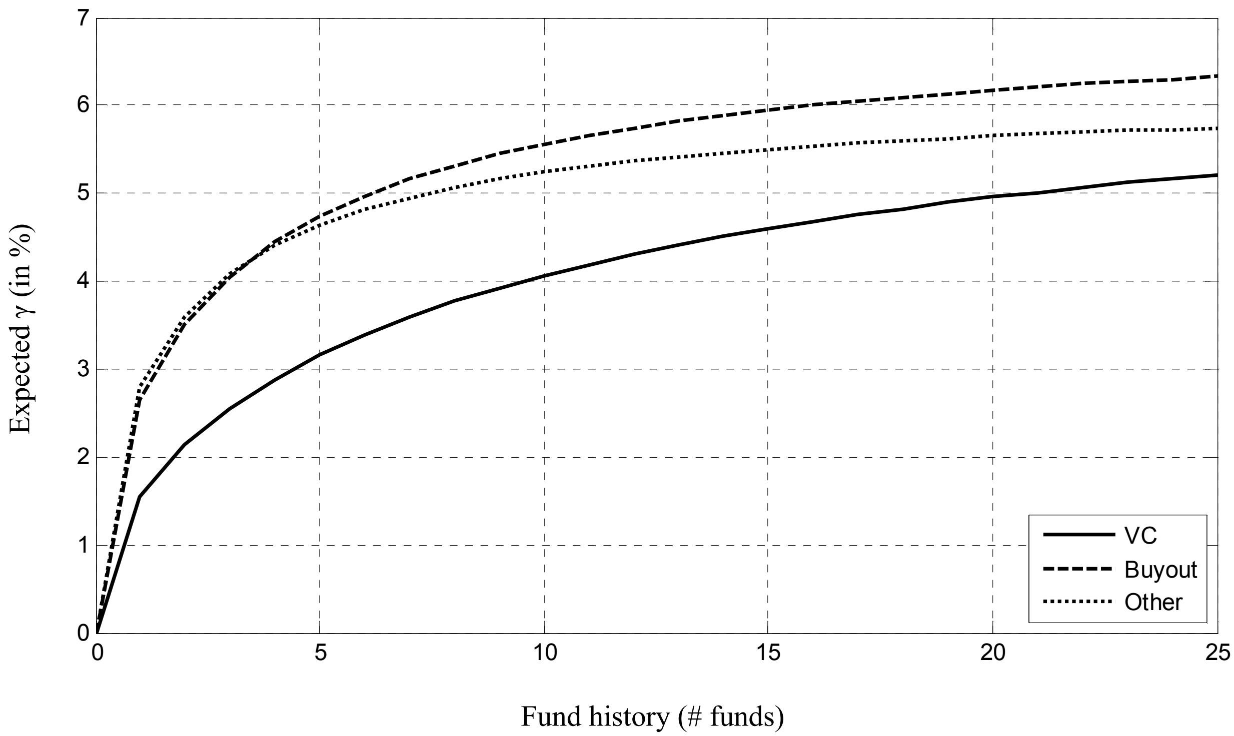

Figure A1: Marginal Log-likelihood. Plots of the marginal log-likelihood as a function of the number of Normal mixtures in the error term distribution, by fund type. The solid line represents specification I of Table IV (which has no vintage year fixed effects), and the striped line represents specification II (with vintage year fixed effects).

Panel A: VC

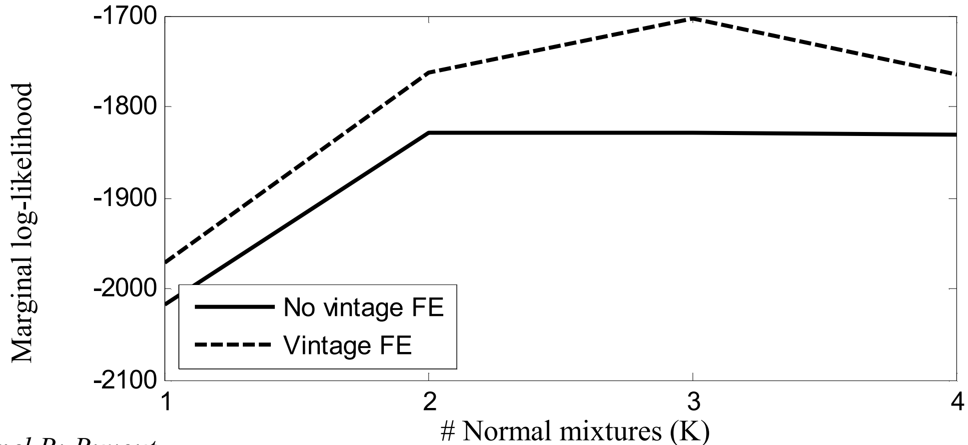

Panel B: Buyout

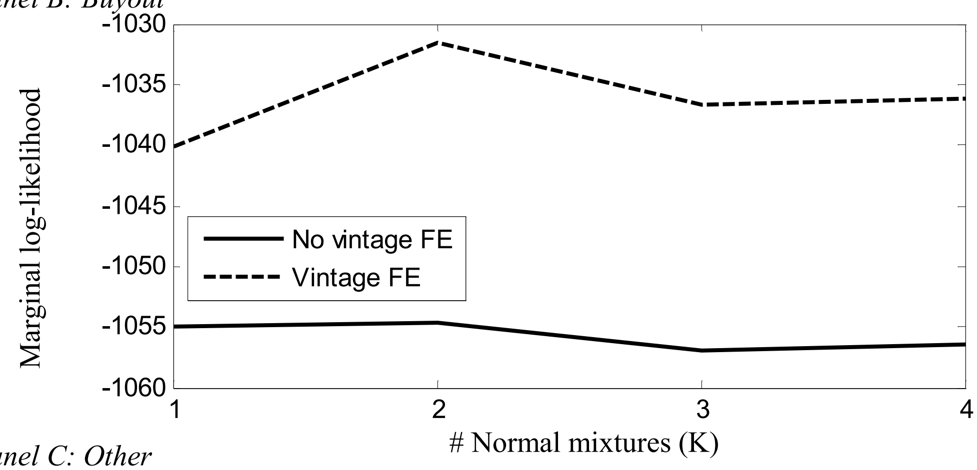

Panel C: Other

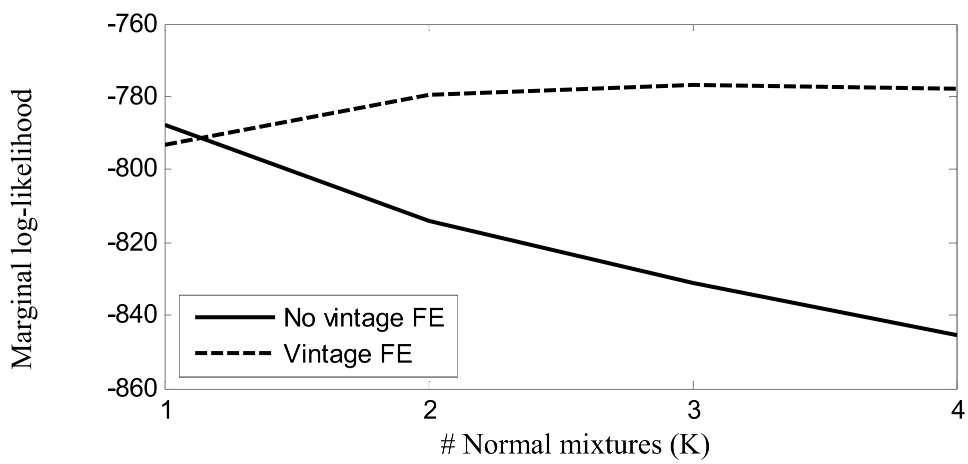

Table I: Summary Statistics. Descriptive statistics of the sample of private equity funds, by fund type (VC, Buyout, or Other). The sample contains 1,924 fully-liquidated funds raised between 1969 and 2001, with at least \$5 million in committed capital (in 1990 dollars) and with non-missing returns data. The funds are raised by 831 individual PE firms (some firms manage funds of more than one fund type). Fund size is the committed capital in millions of dollars. *IRR* is the fund's internal rate of return, net of fees (where 0.1 represents a 10% return). 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.

Panel A: Broad fund categories

	VC	Buyout	Other
# Funds	842	562	518
# Firms	409	285	197
# Funds / firm			
Mean	2.1	2.0	2.6
Median	1	1	2
Std. dev.	1.9	1.6	2.6
10th percentile	1	1	1
90th percentile	4	4	5
Fund size (\$m)			
Mean	206.9	694.1	373.3
Median	110.0	300.0	206.8
Std. dev	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 (in %)			
Mean	17.7	16.9	13.9
Median	8.6	14.9	11.9
Std. dev.	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
Std. dev.	2.623	1.552	1.075
10th percentile	-1.101	-0.170	0.040
90th percentile	3.786	3.216	2.542
Overlap (years)			
# fund pairs	891	512	968
Mean	5.8	5.8	6.8
Median	6	6	7
Std. dev.	2.5	2.3	2.4
10th percentile	2	2	3
90th percentile	9	9	9

Panel B: Fund sub-categories

	# funds	# firms	Fund size ($m)		IRR (in %)		
			mean	median	mean	median	std. dev.
VC							
Early-stage	177	105	186.2	92.0	23.7	6.9	90.8
Late-stage	153	89	289.0	154.0	12.3	10.5	17.0
Generalist	512	239	187.5	104.0	17.2	8.5	44.7
Other							
Real estate	202	86	389.3	263.0	13.8	12.9	9.7
Fund-of-funds	144	48	349.3	138.3	11.2	8.3	13.7
Distressed Debt	36	14	662.1	438.0	14.6	14.0	10.8
Natural Resources	58	25	265.4	154.8	17.0	15.2	15.7
Secondaries	34	12	336.8	263.5	16.8	15.0	12.2
Infrastructure, Turnaround, Special Situations, Co-Investments, Venture Debt	44	26	293.0	150.0	16.9	13.0	18.6

Table II: 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, by vintage year and fund type (VC, Buyout, and Other). The sample comprises 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 average IRRs are based on funds for which size data is available. Source: Preqin.

	VC				Buyout				Other			
		IRR (in %)				IRR (in %)				IRR (in %)		
Vintage year	Funds	Avg.	Median	Weighted Avg.	Funds	Avg.	Median	Weighted Avg.	Funds	Avg.	Median	Weighted Avg.
1969	1	8.7	8.7	8.7	0	-	-	-	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	-	-	-	1	35.5	35.5	35.5	0	-	-	-
1978	2	48.6	48.6	51.0	0	-	-	-	0	-	-	-
1979	1	18.5	18.5	18.5	1	19.4	19.4	19.4	1	18.0	18.0	18.0
1980	6	16.5	14.0	25.9	3	23.7	25.8	25.7	1	12.0	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	39.2	1	10.0	10.0	10.0
1983	12	9.8	9.3	9.4	2	21.9	21.9	21.4	3	14.7	5.9	36.5
1984	20	12.3	12.0	12.0	7	30.4	18.4	30.4	1	7.1	7.1	7.1
1985	21	13.1	13.0	14.9	4	13.4	10.7	17.9	4	4.4	2.7	15.3
1986	20	9.5	8.5	9.3	12	27.6	18.8	25.6	3	4.0	4.0	8.5
1987	22	13.4	14.8	12.6	9	16.8	18.9	9.7	8	8.9	8.3	10.7
1988	26	20.8	21.4	27.0	12	19.2	14.2	14.6	6	10.9	11.6	9.0
1989	35	22.0	16.4	30.1	12	29.8	27.5	28.6	8	13.6	11.1	14.7
1990	24	15.1	16.5	17.6	23	19.2	15.4	15.5	10	16.8	16.0	27.6
1991	18	23.1	21.1	27.2	10	28.0	25.2	28.4	10	15.5	12.2	18.3

1992	30	21.8	16.0	24.3	21	18.3	21.2	31.5	15	18.0	16.3	21.1
1993	36	32.6	29.5	35.9	23	19.2	16.9	19.7	16	23.2	19.8	24.6
1994	30	30.0	25.5	38.2	36	25.9	21.5	35.5	18	15.9	14.0	10.1
1995	41	45.2	17.5	42.5	33	17.1	17.6	15.8	30	16.8	17.0	17.6
1996	43	29.3	10.3	21.3	35	17.4	10.4	11.4	45	17.3	12.7	14.9
1997	62	42.8	20.6	37.6	54	11.6	9.0	9.3	45	11.5	8.4	10.8
1998	73	36.6	7.0	25.0	73	6.7	8.3	5.7	65	11.2	8.3	11.7
1999	95	-2.5	-3.5	-3.4	57	8.6	8.5	7.1	58	10.1	9.3	8.1
2000	120	0.1	-0.9	-0.2	90	16.8	17.5	16.3	80	12.1	12.4	13.4
2001	85	2.4	1.0	2.9	43	27.1	28.0	28.0	88	17.0	15.2	22.3

Table III: AR(1) Regressions. The table shows AR(1) regressions using fund IRRs (in %), by fund type. The dependent variable is $IRR_{i,N}$, the IRR of fund *N* of PE firm *i*. In panel A, $IRR_{i,N-1}$ is the IRR of the most recent fund of the same type (VC, Buyout, 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, Buyout, Other) for a given PE firm. In Panel B, $IRR_{i,N-1}$ is the net IRR on the most recent, non-overlapping 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% levels, respectively.

Panel A: All funds

	VC and Buyout					VC only		Buyout only		Other only	
$IRR_{i,N-1}$	0.162*	0.304***			0.297***	0.129	0.270***	0.314***	0.299***	0.255**	0.102
	(0.071)	(0.053)			(0.056)	(0.079)	(0.059)	(0.067)	(0.072)	(0.091)	(0.074)
$IRR_{i,N-2}$		-0.063	0.052		-0.064		-0.045		0.152		0.148
		(0.053)	(0.053)		(0.055)		(0.063)		(0.102)		(0.097)
$IRR_{i,N-3}$				-0.044							
				(0.026)							
VC=1	1.260	1.988	3.268	4.867	2.032						
	(2.696)	(4.024)	(4.714)	(6.237)	(4.835)						
ln(Size)					0.375						
					(1.649)						
ln(Sequence)					1.769						
					(4.506)						
Year FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
N	673	375	398	244	365	411	240	262	135	295	167
Adj R^2	0.136	0.173	0.107	0.082	0.162	0.195	0.209	0.308	0.398	0.104	0.069

Panel B: Non-overlapping funds

	VC and Buyout	VC	Buyout	Other
$IRR_{i,N-1}$	0.065	0.067	0.195	-0.296
	(0.090)	(0.067)	(0.181)	(0.163)
VC = 1	8.567			
	(7.901)			
Year FE	Y	Y	Y	Y
N	207	147	60	38
Adj R^2	0.064	0.087	-0.006	0.063

Table IV: 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 shows 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, Buyout, Other), by Markov chain Monte Carlo (MCMC) using 10,000 burn-in cycles followed by 100,000 samples, saving every 10th draw. Posterior standard deviations (Bayesian standard errors) are in brackets.

Panel A: Parameter estimates

	VC		Buyout		Other	
	I	II	I	II	I	II
σ_γ	0.049	0.055	0.060	0.056	0.049	0.049
	(0.007)	(0.008)	(0.008)	(0.008)	(0.006)	(0.006)
σ_η	0.258	0.193	0.142	0.203	0.202	0.135
	(0.031)	(0.037)	(0.039)	(0.043)	(0.028)	(0.028)
σ_ε	2.449	2.326	1.359	1.225	0.807	0.865
	(0.123)	(0.121)	(0.058)	(0.073)	(0.039)	(0.050)
Vintage FE	N	Y	N	Y	N	Y
K	3	3	2	2	1	3
N	842	842	562	562	518	518

Panel B: Variance decomposition

	VC		Buyout		Other	
	I	II	I	II	I	II
$100\cdot\sigma_\gamma^2$	0.243	0.309	0.361	0.316	0.244	0.246
	(0.067)	(0.087)	(0.094)	(0.089)	(0.065)	(0.061)
$10\cdot\sigma_\eta^2$	0.675	0.386	0.216	0.432	0.416	0.189
	(0.158)	(0.141)	(0.113)	(0.160)	(0.111)	(0.076)
σ_ε^2	6.015	5.426	1.852	1.505	0.654	0.751
	(0.604)	(0.567)	(0.159)	(0.180)	(0.064)	(0.087)
σ_y^2	6.933	6.120	2.428	2.253	1.314	1.186
	(0.596)	(0.561)	(0.152)	(0.168)	(0.084)	(0.090)
Signal-to-noise	0.035	0.051	0.148	0.141	0.185	0.208
	(0.010)	(0.015)	(0.037)	(0.040)	(0.045)	(0.048)

Panel C: Gamma spread

	VC		Buyout		Other	
	I	II	I	II	I	II
$q_\gamma(75\%)$-$q_\gamma(25\%)$	6.59%	7.42%	8.03%	7.51%	6.61%	6.64%
	(0.90%)	(1.05%)	(1.05%)	(1.06%)	(0.87%)	(0.81%)

$q_{\gamma}(87.5\%)-q_{\gamma}(12.5\%)$	11.24%	12.66%	13.70%	12.81%	11.27%	11.34%
	(1.54%)	(1.78%)	(1.79%)	(1.81%)	(1.49%)	(1.40%)

Table V: 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 IV. Panel B shows the signal-to-noise ratio, s_γ, for the same subsamples. The columns labeled *mean* and *std. dev.* 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 75[th] and 25[th] percentile of the posterior distribution. *N* is the number of funds in the subsample. All estimates are based on model specification II in Table IV (which includes vintage year fixed effects). The top row in each panel corresponds to the estimates of the full sample, as reported in Table IV and Figure 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.

Panel A: Variance decomposition

	VC				Buyout				Other			
	$100\sigma_\gamma^2$	$10\sigma_\eta^2$	σ_y^2	$q_\gamma(75\%)$-$q_\gamma(25\%)$	$100\sigma_\gamma^2$	$10\sigma_\eta^2$	σ_y^2	$q_\gamma(75\%)$-$q_\gamma(25\%)$	$100\sigma_\gamma^2$	$10\sigma_\eta^2$	σ_y^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 (>= 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%
GP 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 (incl. 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%
ROW	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 (>= 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%

Panel B: Signal-to-noise ratio

	VC				Buyout				Other			
	mean	std. dev.	IQR	N	mean	std. dev.	IQR	N	mean	std. dev.	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 (>= 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 (incl. UK)	0.180	0.079	0.110	93	0.269	0.081	0.112	113	0.310	0.116	0.159	59
ROW	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 (>= 1997)	0.040	0.016	0.021	435	0.178	0.058	0.080	317	0.243	0.066	0.091	291

Table A1: Model Specification Tests. This table shows tests of the model specification. Column I reproduces specification I of Table IV, which includes an intercept but no vintage year fixed effects. Column II drops the transient firm effect, η, from the model, and Column III drops the long-run firm-specific effect, γ. Columns IV to VI show the same for specification II of Table IV, 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 IV, where 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.

Panel A: VC

	I	II	III	IV	V	VI
σ_γ	0.049	0.042		0.055	0.040	
	(0.007)	(0.004)		(0.008)	(0.004)	
σ_η	0.258		0.313	0.193		0.268
	(0.031)		(0.023)	(0.037)		(0.025)
σ_ε	2.449	2.604	2.431	2.326	2.465	2.316
	(0.123)	(0.107)	(0.114)	(0.121)	(0.115)	(0.123)
σ_y^2	6.933	6.971	6.906	6.120	6.255	6.106
	(0.596)	(0.562)	(0.553)	(0.561)	(0.572)	(0.567)
Vintage FE	N	N	N	Y	Y	Y
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

Panel B: Buyout

	I	II	III	IV	V	VI
σ_γ	0.060	0.043		0.056	0.044	
	(0.008)	(0.005)		(0.008)	(0.005)	
σ_η	0.142		0.277	0.203		0.279
	(0.039)		(0.034)	(0.043)		(0.020)
σ_ε	1.359	1.542	1.261	1.225	1.469	1.196
	(0.058)	(0.051)	(0.065)	(0.073)	(0.054)	(0.071)
σ_y^2	2.428	2.567	2.373	2.253	2.359	2.217
	(0.152)	(0.163)	(0.121)	(0.168)	(0.163)	(0.167)
Vintage FE	N	N	N	Y	Y	Y
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

Panel C: Other

	I	II	III	IV	V	VI
σ_γ	0.049	0.039		0.049	0.037	
	(0.006)	(0.004)		(0.006)	(0.004)	
σ_η	0.202		0.245	0.135		0.199
	(0.028)		(0.022)	(0.028)		(0.017)
σ_ε	0.807	1.059	0.793	0.865	1.038	0.848
	(0.039)	(0.033)	(0.041)	(0.050)	(0.037)	(0.050)
σ_y^2	1.314	1.272	1.234	1.186	1.217	1.119
	(0.084)	(0.075)	(0.079)	(0.090)	(0.081)	(0.087)
Vintage FE	N	N	N	Y	Y	Y
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