Decreasing Returns or Reversion to the Mean?
The Case of Private Equity Fund Growth

Andrea Rossi*

April 23, 2019

Abstract

When a private equity firm raises a larger fund, performance tends to decline relative to the previous funds it managed. This pattern is usually interpreted as evidence of decreasing returns to scale. I argue that this inference is unwarranted. In essence, private equity firms whose funds have grown the most were on average lucky in the past; as that luck reverts to zero, a spurious negative association between growth and returns is generated in the data. Controlling for this bias, the effect of growth on performance is about 90% and 80% smaller and statistically insignificant for buyout and venture capital funds, respectively. The silver lining is that, historically, private equity fund managers have been able to keep decreasing return at bay. Additional analyses suggest that doing so is harder for VCs than it is for buyout firms.

Keywords: Private Equity, Buyout, Venture Capital, Returns to Scale, Fund Size

JEL Classification: G11, G23, G24

*University of Arizona, Eller College of Management, email: rossi2@email.arizona.edu. I am deeply indebted to Michael Weisbach, Berk Sensoy, Itzhak Ben-David and Justin Birru for advice and guidance throughout my doctoral studies. I have also benefited from the comments of Gregory Brown, Scott Cederburg, Sergey Chernenko, Spencer Couts, Michael Ewens (discussant), Jana Fidrmuc, Shan Ge, Andrei Goncalves, Kewei Hou, Steven Kaplan, Arthur Korteweg, Juhani Linnainmaa (discussant), Pedro Matos, Bernadette Minton, Michael Schwert, René Stulz, Ingrid Werner and of seminar participants at The Ohio State University, UCLA, the 2017 European FMA Doctoral Consortium, the 2018 Southern California Private Equity Conference, PFMC 2018, MFA 2019, and from the feedback received during job market interviews and talks. I also thank PCRI and the NORC Data Enclave for providing data. First draft date: October 1, 2016. All remaining errors are my own.

I. Introduction

"There's been great smaller funds, but as those funds doubled, or tripled in size, their performance really suffered." - Kevin Kester, managing director at Siguler Guff & Co, a private equity investment firm.[1]

This article analyzes the interaction between the growth and the performance of the two most prominent groups of private equity funds: buyout funds and venture capital (VC) funds. These funds are usually structured as closed-end, limited-life vehicles. Therefore, in order to keep earning fees, private equity firms periodically attempt to raise a new fund, typically every 3 to 6 years. The success of a private equity firm's fundraising efforts critically depends on the performance of the previous funds it managed. In particular, the best-performing private equity firms are able to raise follow-on funds that are often significantly larger than their preceding funds. These dynamics motivate this study's main question: what is the impact of fund growth on performance?

The answer to this question is far from obvious. On the one hand, there are reasons to expect that private equity funds face decreasing returns to scale. Buyouts require monitoring by general partners (GPs) and unusually strong incentives for the portfolio companies' managers, while venture capitalists nurture young entrepreneurs by providing advice and experience that help them build successful businesses. When private equity firms raise a larger follow-on fund, they have to make more and/or larger investments than they did in their previous funds. Yet, general partners' time and energy are limited, and there are only so many potential portfolio companies for which returns are high enough to offset the funds fees and illiquidity.

On the other hand, private equity firms might be able to keep decreasing returns at bay. First, larger funds earn substantially higher management fees, which they can use to expand their operations, hire new associates to help source deals, and attract additional experienced general partners. Second, anecdotal evidence and academic studies suggests that a private equity firm's reputation (which is arguably correlated with fund growth and fund size) might help to generate higher returns. Specifically, studies have shown that more experienced buyout funds have better and cheaper access to credit (Ivashina and Kovner

[1] The full quote, from Inside the Mind of the Limited Partner II, Second Annual Conference, 2014 Duane Morris LP Institute Trans-Atlantic Simulcast: London-New York, is as follows: *"He [Kevin Kester, Managing Director at Siguler Guff & Co, a private equity investment firm] has seen plenty of situations "where there's been great smaller funds, but as those funds doubled, or tripled in size, their performance really suffered". Standbridge [Steven Standbridge, Managind Partner at Capstone Partners, a leading placement agent providing global fundraising and advisory services to private equity firms] cites the case where his group raised a fund that was more than two and a half times the prior fund. The response was "It's too big, it's too big, it's too big."*

(2011)). Moreover, young firms and startups might prefer to be financed by larger and more experienced VC firms because they are able to offer more resources, better networks, and higher certification value (e.g., Cong and Xiao (2017)). Indeed, empirical evidence shows that better-networked VCs experience better performance (Hochberg, Ljungqvist, and Lu (2007)) and that more reputable VCs are more likely to be able to access hot deals and obtain favorable contractual terms and pricing (e.g., Hsu (2004) and Nanda, Samila, and Sorensen (2017)).

Another important issue is whether the aggregate size of the private equity fund industry has an effect on average fund returns, for example, through competition for deals. In this research article, I focus exclusively on fund-level effects and abstract away from aggregate effects. Specifically, I compare returns across funds that operated in the same market environment by analyzing buyout and VC funds separately and controlling for vintage year-region fixed effects.

This article uses a commercially-available dataset of over 2,000 buyout and VC funds to study the effect of fund growth on performance. A first glance at the data seems to confirm previous findings: there is a strong negative correlation between changes in fund size and changes in returns *within* the funds managed by a given private equity firm. In other words, when a firm raises a larger follow-on fund, that fund tends to underperform the preceding ones. The seminal paper of Kaplan and Schoar (2005) was the first acedemic article to report that, when controlling for GP fixed effects, changes in fund size are negatively correlated with changes in performance. Although Kaplan and Schoar (2005) did not interpret their result causally, practitioners, academics and the financial media tend to implicitly or explicitly interpret a negative correlation between changes in size and changes in performance as evidence of decreasing returns. For example, institutional investors often report being disappointed by the returns of larger follow-on funds raised by private equity firms that were very successful in the past.

This study argues that this interpretation of the data is unwarranted because it fails to properly account for reversion to the mean in fund performance. Reversion to the mean is present is any setting where observed performance depends both on a fixed component (which is sometimes referred to as skill or persistence) and on a transitory component (which is sometimes referred to as idiosyncratic shocks, noise, or luck). This phenomenon has been observed, for instance, in the performance of corporate managers, fighter pilots, athletes and mutual fund managers.

I show that private equity fund data is subject to a selection mechanism that is likely to lead to incorrect inferences about returns to scale. Specifically, the data is biased towards finding spurious evidence of large decreasing returns. There are three factors contributing to this bias. First, higher past returns increase the probability that a private equity firm will be able to raise a follow-on fund (e.g., Kaplan and Schoar (2005) and Hochberg, Ljungqvist, and Vissing-Jorgensen (2014)). Second, a follow-on fund's growth is positively related to past returns (e.g., Chung, Sensoy, Stern and Weisbach 2012). Third, buyout and venture capital fund returns are extremely dispersed, and most (80% to 95%) of the variation is estimated to be driven by idiosyncratic shocks rather than by persistent differences in skill across fund managers (Korteweg and Sorensen (2017)). I show that these facts imply that (i) the returns of funds for which a follow-on was raised contained, on average, positive shocks; and that (ii) the higher the follow-ons growth relative to the prior fund, the higher the expected positive shock in prior fund performance. Therefore, since the mean expected idiosyncratic shock across follow-on funds is zero, we should expect high-growth follow-on funds to underperform their preceding funds even if there were no decreasing returns to scale.

Given that a decline in the returns of high-growth follow-on funds is to be expected, it is important to adjust one's estimates for this expectation when measuring the effect of fund growth on performance. Thus, the main objective of my analysis is to understand to which extent the negative relation between fund growth and performance within the funds managed by a given firm is attributable to actual decreasing returns (i.e., a causal effect) rather than to reversion to the mean (i.e., a spurious effect).

To address this issue, I use an approach that builds on the model of Korteweg and Sorensen (2017). Korteweg and Sorensen (2017) develop a variance decomposition model designed to identify long-term persistence in relative skill across private equity firms and separate it from idiosyncratic shocks. They propose an estimation procedure that uses Bayesian Markov Chain Monte Carlo methods. They find that the amount of the tremendous variation in private equity returns that can be attributed to persistent cross-sectional differences in managerial ability is only approximately 14% for buyout and 5% for venture capital. I estimate this model in my sample and confirm all their key findings. Based on these results, it is clear that, although some evidence of persistence can be statistically identified in the data, private equity fund returns are expected to exhibit a large degree of reversion to the mean.

I then use a simulation exercise to obtain estimates of the effect of fund growth on performance that account for reversion to the mean. Specifically, I generate panels of private equity data (separately for buyout and VC) similar to the sample data. Returns are drawn so that they have the same distribution and the same amount of persistence as in the data. In addition, I require that the probability of raising a follow-on fund and its growth are related to past returns in the same manner as in the actual data. Returns are drawn independently of fund size and growth, so that, by construction, there are no decreasing returns to scale in the simulated data.[2]

Despite the fact that the data were generated specifically not to contain decreasing returns, I find a negative and significant relation between fund growth and performance in these data. This suggests that the selection process that generates private equity data will lead one toward finding a negative association between fund growth and returns at the private equity firm level, even in the absence of true decreasing returns. This estimated spurious effect equals about 90& and 80% of the effect observed in the sample of buyout and venture capital data, respectively. These estimates are robust to reasonable variations in modeling choices. If I do not adjust for reversion to the mean, the estimates imply that a 100% increase in fund size is associated with a decline in internal rate of return (IRR) of 3.5 percentage points (p.p.) for buyout funds and of 6 p.p. for venture capital. Adjusting for the bias attributable to reversion to the mean, the same increase in size leads to a decrease in IRR of only 0.3 p.p. for buyout and of 1.1 p.p. for venture capital. Moreover, I can not reject the hypothesis that the adjusted estimates are different from zero, especially for buyout funds, and to a lesser extent for VC funds.

Thus, these results suggest that VC fund managers and especially buyout fund managers are able, on average, to keep decreasing returns at bay when they raise larger funds. To provide further insights into how they may be able to do so, I analyze a large sample of portfolio-level investments to determine how fund portfolios evolve across the sequence of funds managed by a given private equity firm. The analysis yields interesting results. When buyout firms raise a larger follow-on fund, they make only marginally more investments than in their smaller preceding funds. Rather, they simply tend to take over larger companies. Conversely, VC firms do not seem to be able to scale up only by making larger investments. When a VC firm raises a fund twice the size of its prior fund, it increases the number of portfolio companies by 30% and the average dollar investment made by 53%.

These results are consistent with empirical evidence showing that, compared to buyouts, VC fund growth rates are lower on average and are less sensitive to past performance, and that VC firms are more likely to cap the size of their follow-on funds.

The analysis presented in this paper provides important insights to investors. Although my approach focuses on the econometrics of estimating returns to scale, it should be noted that the bias discussed in this paper is not just a detail that the econometrician should worry about when interpreting regression results. Rather, the bias is due to a very real selection mechanism to which private equity funds are subjected. This selection mechanism is likely

[2] This approach is similar in spirit and partially inspired by Fama and French (2010) and Linnainmaa (2013), who use simulation studies in order to study the cross-section of mutual fund skill while accounting for luck and survivorship bias, respectively.

to bias not only regression estimates, but also the beliefs of investors and practitioners. Consider, for instance, this article's opening quote: “there's been great smaller funds, but as those funds doubled, or tripled in size, their performance really suffered”. More generally, it is likely that many limiter partners (LPs) have experienced disappointing results when investing in high-growth follow-on funds of successful private equity firms, especially if their expectations were based on those firm's previous returns. My analysis suggests that most of the “disappointment” is due to luck in past winners reverting to zero rather than to negative effects of fund growth.

Returns to scale play a crucial role in our understanding of the asset management industry. Berk and Green (2004) show that a model where mutual fund managers face decreasing returns to scale can reconcile empirical facts such as a positive flow-performance relationship and the lack of performance persistence within a rational framework. In stark contrast to mutual funds, private equity firms do seem to exhibit a certain degree of net-of-fees performance persistence, as suggested by Kaplan and Schoar (2005) and by a number of other studies that are generally consistent with their initial evidence. Korteweg and Sorensen (2017) confirm those results using a Bayesian model that separates long-term persistence from spurious short-term persistence and noise. Finally, Harris, Jenkinson, Kaplan, and Stucke (2014b) confirm these results over the same sample periods, but find lack of statistically significant persistence in buyout funds in the subsample of funds raised after the year 2000.

The fact that there exists some evidence of persistence in private equity is at odds with the predictions of Berk and Green (2004) and appears particularly puzzling to academics. If differential skill exists, we would expect superior general partners to appropriate the benefits of their ability and boost their income by raising larger funds or increasing nominal fee rates. We know that the latter does not take place in most cases, and Hochberg et al. (2014) propose a model where information holdup on the part of the incumbent limited partners (LPs) explains why successful GPs don't increase the fees they charge. On the other hand, successful private equity GPs do tend to raise larger funds over time. The estimates presented in this paper help us to understand why fund growth is not sufficient to completely eliminate persistence, especially in venture capital funds, whose investment strategies are believed to be less scalable (e.g., Metrick and Yasuda (2010)).

There is an important nuance to the interpretation of the results presented here. The fact that a careful analysis of historical data fails to uncover evidence of statistically significant decreasing returns does not mean that they do not exist at all. On the contrary, it is plausible that capacity constraints exist, especially in venture capital, because by definition venture capital funds invest in small, young firms and take only minority equity stakes in them. Given their compensation structure, it is irrational for GPs to restrict capital raising unless they believe that they face significant decreasing returns to scale. However, many successful VC firms are known to restrict the size of the funds they raise, which often go oversubscribed, while buyout funds only rarely do so. Accordingly, in the data, VC fund growth is less sensitive to performance and lower on average than it is for buyout. In light of these facts and of the evidence presented, this article's findings are consistent with an *equilibrium* where the follow-on funds of successful private equity firms grow faster than other funds, yet stay below the point where decreasing returns become so strong that they would show up significantly in the data.

A corollary to this interpretation is that the results presented are valid for fund sizes and growth rates that have been explored as of the end of the sample studied. Buyout and venture capital firms may face larger decreasing returns if they grow past these levels.

It is important to notice that this paper's analysis applies to the cross-section of private equity funds, but it does not apply across asset classes. Specifically, my findings mean that, for a given cross-section of funds, fund growth does not significantly hinder the ability of a particular private equity firm (especially buyout firms) to outperform other firms whose funds grow less. However, they do not imply in any way that industry growth (in the form of an increase in the number of funds and/or in the average size of funds) is unrelated to average private equity performance. Harris, Jenkinson, and Kaplan (2014a) show that private equity funds raised in vintage years with particularly strong aggregate capital perform significantly worse than funds raised in different times. I present results suggesting that this phenomenon is unlikely to be driven by fund-level effects (e.g., dilution of senior partners value-creation ability). Instead, its causes are likely to be found in industry-wide or market-wide effects such as increased competition for deals (e.g., Gompers and Lerner (2000)), entry of less-skilled managers (e.g., Kaplan and Schoar (2005)), or credit market conditions (e.g., Axelson, Jenkinson, Stromberg, and Weisbach (2013)).

The findings presented in this paper complement the recent work of Braun, Jenkinson, and Stoff (2016) in fostering our understanding of the forces at play in the private equity industry. These authors study deal-level buyout data and are able to link increases in competition (proxied by aggregate fund raising) to decreases in persistence. At the same time, the present article finds that private equity managers have historically been able to avoid incurring significant fund-level decreasing returns. Overall, in private equity, the evidence seems consistent with the industry-level decreasing returns to scale hypothesis of Pástor and

Stambaugh (2012).[3]

The rest of the paper is organized as follows. Section II describes the data used in this study. Section III starts by documenting the existence of a negative association between fund growth and performance in the data. Then, it argues that the presence of reversion to the mean leads to misleading evidence about the nature of returns to scale in private equity. Finally, it details the methodology used to deal with the bias and presents bias-adjusted estimates. Section IV uses portfolio-level data in order to study how the composition of fund portfolios change across subsequent funds. Section V provides concluding remarks.

II. Data

A. Fund Performance Data

I use fund-level data for buyout and venture capital funds obtained from Preqin in October 2017. Brown, Harris, Jenkinson, Kaplan, and Robinson (2015) analyze four datasets from major commercial sources (including Preqin) and conclude that empirical inference about fund-level performance is very similar across the various sources.

The sample and variables used in the empirical analysis are constructed following the literature standards. In particular, I closely follow Chung, Sensoy, Stern, and Weisbach (2012) when defining "current" (or "preceding") and follow-on funds. Current funds are all funds with available performance (IRR or TVPI), with fund size (total capital committed to the fund) greater than \$5 million in 1990 dollars (as in Korteweg and Sorensen (2017)), and with vintage years ranging from 1969 to 2011. Funds with vintage years ranging from 2012 to 2017 are excluded from the sample of current funds for two reasons. First, total fund performance in the first few years of a funds activity can be noisy, and therefore I only analyze the performance of funds that have operated for at least six years. Second, I will analyze whether these funds had a follow-on fund, and since GPs usually raise a new fund every three to five years, I need to observe at least five years of fundraising activity after the last funds in the sample of current funds were raised.

The primary objective of this article is to study returns to scale in private equity. Different private equity investment strategies plausibly face different organizational challenges and capacity constraints. For this reason, I study buyout and venture capital separately. I exclude from the analysis all other private equity strategies, such as growth equity, balanced and turnaround. Some of the most prominent private equity firms have built on their success

[3] In the mutual fund literature, there is evidence of industry-level decreasing returns (e.g., Pástor, Stambaugh, and Taylor (2015)), while fund-level returns to scale results are mixed (e.g., Chen, Hong, Huang, and Kubik (2004), Pástor et al. (2015), Harvey and Liu (2017)).

and raised funds that operate in areas that are substantially different from their original scope. For example, The Blackstone Group started as a buyout firm but now manages funds that focus on real estate and infrastructure investments and even on hedge fund strategies. The initial screening eliminates all the non-core buyout and venture capital funds in the sample. Moreover, if a private equity firm manages both buyout and venture capital funds, or if it manages funds with different geographic focus, I treat these funds as belonging to different private equity firms. I further exclude all funds whose region of focus is outside North America or Europe.

The final sample of current buyout and venture capital funds includes, respectively, 1169 and 1034 funds managed by, respectively, 496 and 457 different private equity firms.

For each current fund, I identify the subsequent fund raised by the same private equity firm, if there is one. These funds make up the sample of follow-on funds. Any fund raised from the beginning of the sample period until 2017 is potentially classified as a follow-on fund, as long as its size is not missing. The final sample of follow-on funds comprises 996 buyout funds and 800 venture capital funds.

Table I presents summary statistics for the sample. Unless otherwise noted, in order to minimize the impact of outliers, all variables relating to fund performance, size, and growth are winsorized at the 1% and 99% levels. The sample only includes private equity firms whose first fund was raise before 2012. Panel A shows that buyout and venture capital firms have managed on average slightly less than 2.5 funds. The standard deviation of this figure is also relatively small, as over 90% of the GPs has managed at most 5 funds. Panel B and C present summary statistics for fund-level variables. The percentage of funds that are followed by another fund is 85% for buyout and 77% for VC. All fund-level variables present the typical patterns reported by prior research articles that study private equity fund data. In particular, the mean and median return is higher for buyout funds, while venture capital fund returns have greater dispersion. On average, follow-on funds tend to be larger than their preceding counterparts, and on average fund size grows faster for buyout funds than for VCs.

Preqin does not provide detailed LP-GP cash flow data for all funds, and therefore it is not possible to calculate the Public Market Equivalent (PME) performance measure for a significant portion of the sample. In this paper, I use net internal rates of returns (IRRs) as the measure of fund performance. All results are robust to using Total Value to Paid-In capital multiples (TVPI). Moreover, because all tests and inference in this paper control for vintage year fixed effects, all the results are most likely robust to using PMEs as an alternative measure of performance because IRRs and TVPIs have been shown to explain over 90% of the cross-sectional variation in PMEs (Harris et al. (2014a)).

Table I **Private equity fund sample.** This table reports summary statistics for the sample used in this paper. Panel A shows the number of private equity firms, or General Partners (GPs), and the distribution of the number of funds managed by each of them. Panel B and C present, respectively, the distribution of the buyout and venture capital fund-level variables used in the regressions shown in Table 2 and 3. Only funds that belong to the sample of "current" funds are included in this table. This sample spans the vintage years from 1969 to 2011. However, data on follow-on funds raised from 2012 to 2017 is also utilized to calculate the percentage of "current" funds that raise a follow-on and the size growth of those funds. See the data section for details on the construction of the sample of current and follow-on funds.

Panel A: Number of PE firms and funds, 1969-2011

	Number of PE firms	Number of funds	Distribution of funds by PE firm					
			Mean	Std Dev	25^{th}	50^{th}	75^{th}	90^{th}
Buyout	496	1169	2.4	1.7	1	2	3	5
Venture Capital	457	1034	2.3	1.8	1	2	3	5

Panel B: Distribution of fund-level variables: Buyout funds

	Obs	Mean	Std Dev	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
Fund size ($ Mn)	1169	691	1101	57	130	283	684	1935
Performance (IRR %)	1169	16.1	15.0	0.8	8.0	13.7	22.4	33.4
Fund has a follow-on fund	1169	85.2%						
Follow-on fund size growth	996	83%	138%	-37%	8%	53%	117%	216%

Panel C: Distribution of fund-level variables: VC funds

	Obs	Mean	Std Dev	10^{th}	25^{th}	50^{th}	75^{th}	90^{th}
Fund size ($ Mn)	1034	155	170	24	48	101	194	347
Performance (IRR %)	1034	12.5	25.7	-10.4	-1.0	7.9	18.0	38.1
Fund has a follow-on fund	1034	77.4%						
Follow-on fund size growth	800	65%	129%	-46%	-8%	33%	101%	192%

B. Portfolio Investments Data

In the last part of this article, I analyze the relationship between the size of a fund and the number of investments it makes. In order to do so, I use data on the portfolio investments made by a large sample of private equity funds provided by the Private Capital Research Institute (PCRI).[4] The data contain information on thousands of private capital deals, including an identifier for the fund involved in the deal and an identifier for private equity firms. I study a sub-sample of these data that includes funds whose type is classified as buyout or VC, whose size is greater than 10 million dollars, and that were raised before 2014.

After imposing these requirements on the data, I perform the following calculations. First, I sum all equity investments made by a given fund in a given company in order to obtain the total amount of equity invested in a given portfolio company. This is particularly relevant for VC funds that participate in multiple financing rounds of the same portfolio company. Then, I sum all investments made by a given fund to obtain the total capital invested by that fund.

In order to eliminate fund observations that would constitute outliers due to incomplete or misreported data, I eliminate the following cases: equity investments that are 100 time larger or 100 time smaller than the median investment made by the same fund; funds with less than 5 equity investments, because such level of concentration should not be possible since Limited Partnership contract usually requires a minimum level of diversification; funds for which the total amount of equity investments recorded in the database is less than 50% of fund size or over 120% of fund size. There are several reasons why the sum of the recorded amount of equity investments made by a fund would not equal 100% of fund size. First, we know that 10 to 30% of the capital committed to a fund usually goes towards management fees and other costs. Second, the partnership agreement usually gives the GP a certain degree of flexibility in deciding when and how much of the committed capital to invest. For instance, a GP might ask the LPs to provide more capital than initially agreed if a larger-than-expected investment opportunity becomes available after a substantial portion of the committed capital has already been called. In some cases, GPs are allowed to "recycle" some of the proceed from early exits to make additional investments instead of returning the capital to LPs immediately (Phalippou (2017)). After cleaning the sample, I have data on the portfolio composition of 279 Buyout funds and 1,411 VC funds.

I then calculate two main fund-level variables of interest: average equity investment in

[4]These data are stored at the NORC Data Enclave, and are available to researchers for academic purposes only. PCRI and NORC verify compliance of data usage with their access and confidentiality policies, but do not endorse the results of specific research projects.

each portfolio company and average number of portfolio companies. The first variable is simply the mean of the observed equity investments in each portfolio company, while the second variable is imputed using the formula: number of portfolio companies = fund size / average equity investment in portfolio company.

The imputed median number of investments across the sample is 16 for buyout funds and 24 for VC funds.[5] The variable is right-skewed, and first-time funds have a lower number of investments on average. The size of these portfolios appear in line with summary statistics reported in the literature, for instance, by Metrick and Yasuda (2010), who report a median of 12 and 25 investments for buyout and VC funds, respectively, and by Braun et al. (2016), who analyze a sample of buyout funds with a mean number of 15.6 deals per fund.

III. Returns to Scale in Private Equity

Fund size and fund flows play a central role in our understanding of skill, performance persistence, and managerial and investor behavior in the asset management industry. Arguably, returns to scale are crucial to estimate a fund's optimal size.

Private equity managers and investors certainly pay a lot of attention to this topic, and for good reason. The typical private equity fund is a closed-end, limited-life fund structured as a limited partnership. The size of a buyout or venture capital fund (i.e., the total amount of capital committed to it) is fixed at inception. Limited partners cannot withdraw capital from the fund before the underlying portfolio companies are liquidated, and the fund cannot accept more capital from incumbent or new investors. In other words, unlike in mutual funds and hedge funds, in private equity, capital cannot flow elastically in and out of a fund.

A. The Endogeneity Problem

Broadly speaking, when studying the asset management industry we are interested in two kinds of scale effects: at the manager level (returns to scale) and at the industry level (returns to aggregate scale). Although it is often assumed that both effects go in the same direction (i.e., the effect is negative), each effect is likely driven by different factors and has different implications. For instance, returns to scale should be taken into consideration when allocating across different managers, while knowledge of returns to aggregate scale can provide insight when allocating across investment styles or asset classes.

[5]Consistent with the confidentiality policies of the data provider, these figures do not refer to, or allow to identify, any specific investment or fund in the underlying data. These figures are the result of the estimation algorithm described here rather than information provided by individual contributors of data to PCRI.

The primary focus of the present article is on returns to scale at the private equity firm level. A natural way to study returns to scale in private equity is to analyze the relationship between fund performance and fund size in the data. An obstacle to the identification of returns to scale is the likely endogeneity of fund size and skill. This problem has recently been recognized in the mutual fund literature, with Pástor et al. (2015) arguing that cross-sectional regressions of performance on fund size are likely biased and that a fixed-effect specification is necessary to deal with the endogeneity.

It can be argued that cross-sectional regressions of performance on fund size in the context of private equity are even less likely to reveal the true effects of scale than in the context of mutual funds. The reason is that it is relatively easy to control for the relationship between investment opportunities and fund size in mutual funds, but in private equity it is not. For example, the average amount of assets under management is greater for large-cap mutual funds than it is for small-cap mutual funds; therefore, a regression of performance on fund size might pick up, for instance, differences in stock-picking opportunities across equities with different market capitalization rather than actual effects of scale. In mutual funds, this issue can be addressed quite effectively because researchers can control for the specific benchmark and the factor exposures of each funds.

Private equity firms tend to specialize in a particular investment style (e.g., there are venture capital firms that specialize in funding start-ups in the American Midwest), and each specific investment style is associated with a range of typical fund sizes. In the literature, the typical controls are limited to the general investment style (i.e., buyout, growth equity or venture capital), the broadly-defined region of focus (i.e., North America, Europe etc.), and the vintage year. As a consequence, cross-sectional regressions of performance on fund size are likely to be seriously confounded. In this context, the use of firm fixed effects appears necessary to correctly identify the effect of scale on performance, because fixed effects allow to control for both observable and unobservable time-invariant differences across private equity firms.

Interestingly, before the mutual fund literature recognized the necessity of fixed effects in the identification of returns to scale (Pástor et al. (2015)), the seminal article of Kaplan and Schoar (2005) demonstrated the empirical importance of fixed effects in the context of private equity (even though the paper did not explicitly discuss endogeneity). Kaplan and Schoar (2005) start by documenting a positive cross-sectional relationship between performance and fund size. Then, they reported that, once GP fixed effects are included, "the sign on the fund size and sequence number variables switch from positive to negative", and that the negative fund size coefficient is statistically significant. Finally, they concluded that "when a given GP raises a larger fund, fund returns decline for that GP". The purpose of the analysis carried out in the rest of this chapter is to understand whether the observed decline in the performance of larger follow-on funds may be interpreted as evidence that fund growth *causes* returns to decline.

B. Empirical Association between Fund Growth and Performance

I start the analysis by describing the empirical relation between changes in fund size and performance in the data. As discussed in the previous sections, regressions of performance on size are likely to be confounded. In order to alleviate the problem, I focus on changes in fund size as opposed to the level of fund size. This can simply be achieved by controlling for private equity firm fixed effects. Alternatively, in cross-sectional regressions, performance can be regressed on changes in fund size (i.e., fund growth) as opposed to fund size.

Specifically, for each fund in the sequence of funds managed by a given private equity firm, I define growth as the change in fund size of the focal fund relative to the first fund managed by the same firm. Using this definition of fund growth, the logarithm of fund size of the n^{th} fund managed by firm i can be written as $log(fund\, size_{i,n}) = log(1 + fund\, growth_{i,n}) + log(fund\ size_{i,1})$. Regressing performance on $log(1 + fund\ growth_{i,n})$ as opposed to on $log(fund\ size_{i,n})$ is helpful in identifying the causal effect of fund growth on performance because it removes the relation between the size of the first fund and performance from the estimated coefficient. This relation is linked to two of the confounding effects we should control for. Specifically, the size of the first fund managed by a given private equity firm is likely endogenous both to the investment strategy of the firm and to the perceived skill and reputation of the general partners of the firm at the time they raised their first fund. Notice that the above argument is relevant only for cross-sectional regressions: when firm fixed effects are included, there is no need to distinguish between fund size and fund growth, as the coefficient estimates for the two variables are by construction identical.

Table II illustrates the empirical results. Consistent the literature, standard errors are clustered at the private equity firm level[6] and vintage year fixed effects are included in all specifications in order to control for time patterns in average returns and to partial out potential effects of returns to aggregate scale. The sample includes North American and European private equity firms. Hence, for each fund, vintage year fixed effects are interacted with the region of focus of that fund. Panel A and Panel B report results for the sample of buyout funds and VC funds, respectively.

In the cross-sectional specification, there does not appear to be any significant relation between fund growth and performance for either type of fund. However, when firm fixed

[6] I verify that other clustering specification produce smaller standard errors.

effects are included, the coefficient on fund growth is negative and highly statistically significant. This result perfectly mirrors the findings of Kaplan and Schoar (2005): when a private equity firm raises a larger fund, the returns of that fund tend to be lower than those of the other funds managed by that firm. The effect is economically significant: a one percent increase in fund size is associated with a decline in IRR of 5.1 basis points for buyouts and of 8.7 basis points for VCs. Considering that buyout and VC follow-on funds respectively grow by 83% and 77% on average, the economic magnitude of these estimates appears extremely large, especially when compared to the cross-sectional estimates. In the next section, I argue that the observed negative association between fund growth and performance within funds managed by a given firm is primarily driven by reversion to the mean, and should therefore not be interpreted as a causal effect.

The last four model specifications in Table II introduce fund sequence numbers. In particular, consistently with the literature (e.g., Kaplan and Schoar (2005)), I take the natural logarithm of the sequence number. On average, private equity firms tend to raise follow-on funds that are larger than their initial funds, therefore sequence number and fund growth are highly correlated. In the third and fourth columns of the table, fund returns are regressed on sequence numbers without controlling for fund growth or size. Results for buyout and VC firms differ substantially. In the cross-section, the performance of buyout funds is completely unrelated to their sequence number. By contrast, this relation is positive and significant for VC funds, suggesting that more experienced VC firms are able to deliver higher returns than less experienced ones. The latter result is generally consistent with mechanisms that have been discussed in the literature, specifically, that VC networks (e.g., Hochberg et al. (2007)) and access to deal-flow due to reputational effects (Cong and Xiao (2017) and Nanda et al. (2017)) are beneficial for fund returns. When controlling for private equity firm fixed effects, however, the coefficients on the sequence number turn sharply negative for both buyout and VC, and it is statistically significant at the 5% level for buyout funds.

The models estimated in the last two columns of the table include both fund growth and sequence number. In the cross-section (specification 5), buyout returns are unrelated to growth or sequence number, as the coefficient estimates are both economically small and statistically insignificant. The results for VC funds, on the other hand, are sharply different. The cross-sectional performance of VC funds continues to be positively related to their sequence number. However, controlling for the sequence number, the coefficient on fund growth goes from positive (0.81) and indistinguishable from 0 to negative (-2.59) and marginally statistically significant. This suggest that, in the cross-section, funds managed by more experience VC firms on average have higher returns, but at the same time funds

Table II **Observed relation between fund growth and performance.** This table reports results from regressions of fund performance (IRR) on the natural logarithm of one plus cumulated fund growth (i.e., fund growth relative to the first fund managed by the same private equity firm) and on the natural logarithm of each fund's sequence number. All specifications control for vintage year-region fixed effects. The second, fourth, and sixth specification also include PE firm fixed effects. The t-statistics presented in parenthesis are based on standard errors clustered at the PE firm level and robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Buyout funds

	(1)	(2)	(3)	(4)	(5)	(6)
Log(1+fund growth)	-0.25	-5.12***			-0.62	-4.75***
	(-0.77)	(-4.71)			(-0.72)	(-4.06)
Log(fund sequence #)			-0.00	-6.69**	0.81	-3.05
			(-0.01)	(-2.55)	(0.48)	(-1.12)
PE firm F.E.	No	Yes	No	Yes	No	Yes
Vintage-region F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.16	0.34	0.16	0.31	0.16	0.34
Observations	1,169	1,169	1,169	1,169	1,169	1,169

Panel B: VC funds

	(1)	(2)	(3)	(4)	(5)	(6)
Log(1+fund growth)	0.81	-8.68***			-2.59*	-8.84***
	(1.04)	(-4.28)			(-1.91)	(-4.30)
Log(fund sequence #)			3.07***	-5.34	5.81***	0.89
			(2.69)	(-1.46)	(2.83)	(0.25)
PE firm F.E.	No	Yes	No	Yes	No	Yes
Vintage-region F.E.	Yes	Yes	Yes	Yes	Yes	Yes
Adj. R^2	0.25	0.41	0.25	0.38	0.25	0.41
Observations	1,034	1,034	1,034	1,034	1,034	1,034

that have grown more than other funds with the same sequence number tend to have lower returns. When firm fixed effects are introduced (column (6)), fund growth is significantly negatively associated with performance for both buyout and VC funds, while the sequence number is not.

Given that in the data there exists such a strong negative relationship between fund growth and performance at the private equity firm level, it is not surprising that many practitioners believe that the performance of private equity funds are strongly impacted by decreasing returns. The next sections argue that this interpretation of the data might not be correct. Specifically, I find that at least a large share of the negative relation between fund growth and performance is a statistical artifact due to reversion to the mean.

C. Selection Mechanism, Reversion to the Mean, and Econometric Bias

There are reasons to believe that reversion to the mean is the primary driver of the negative association between fund growth and performance we observe in the data within the funds managed by the same private equity firm. Hence, this association should not be interpreted as definitive causal evidence of decreasing returns.

The combination of three strong empirical facts contribute to generating spurious evidence of decreasing returns. First, the probability that a private equity firm raises a follow-on fund is positively related to the return of preceding fund managed by the same firm. Second, the growth in the size of a follow-on fund is positively related to the return of the preceding fund. Third, fund returns are extremely dispersed and, although some evidence of persistence at the GP level exists, most of the variation is driven by noise (Korteweg and Sorensen 2015).

Taken together, these facts imply the following. First, the returns of funds managed by firms that are able to raise a subsequent fund must have contained, on average, positive noise. Second, the higher the growth in the follow-on funds, the higher, on average, the noise in the preceding funds. Since the average expected return shock across follow-on funds is zero, we should expect that high-growth follow-on funds perform poorly compared to the preceding funds managed by the same firm. In other words, firms that have been able to raise additional funds after their early funds must, on average, have had good luck, especially those firms whose follow-on funds have grown the most. Since there is no reason why that luck will continue, the follow-on funds raised by lucky firms should on average underperform their preceding funds. This will lead to finding spurious evidence of diminishing returns to scale in analyses such as those carried out in Table II.

C.1. How Past Performance Influences Future Fundraising

Buyout and venture capital investments are carried out though closed-end funds with limited life of, typically, ten years. To continue to earn revenue after the fund's life is over, it is crucial for a private equity firm to raise follow-on funds. They usually attempt to do so after three to six years following the inception of the current fund, and advertise the performance of the current and prior funds to facilitate raising a follow-on fund. This process will create a positive relation between past returns and future fundraising, similar to that observed in mutual funds and hedge funds.[7]

Consistent with this prediction, Kaplan and Schoar (2005) find that better performing partnerships are more likely to raise follow-on funds and larger funds. Chung et al. (2012) show that the relationship between current fund returns and future fundraising is particularly strong and, as a consequence, the "indirect pay for performance from future fundraising is of the same order of magnitude as direct pay for performance from carried interest".

The empirical relation between past performance and future fundraising is illustrated in Table III. The first four specifications estimate the probability that an existing private equity firm will raise another fund as a function of the performance of its prior funds. For each current (focal) fund n, the dependent variable is an indicator which takes the value of one if the firm raised a follow-on fund and zero otherwise, and the independent variables are the return on the current fund n, the sequence number of the current fund, and the interaction of these two variables. I also include the return of fund $n-1$, which seems to help predicting the probability of raising a follow-on fund even when controlling for the current fund's performance. Because for first-time funds do not have a prior fund, the performance of fund $n-1$ is set equal to zero when the focal fund is a first-time fund; hence, the model includes an indicator for first-time funds in order to capture the difference in the conditional mean due to setting the lagged fund return equal to zero. The model is estimated using both a linear probability model or a probit specification.

The estimates in columns 1 through 4 confirm prior literature findings. Past performance strongly predicts the likelihood that buyout and venture capital firms will raise another fund. Moreover, a funds sequence number increases the probability of a follow-on fund and it causes a decrease in the sensitivity of that probability to current fund returns. The inference is similar across OLS and the probit specifications. These findings are consistent with the idea that investors learn about a partnerships ability over time (see Chung et al. (2012)).

[7] In the case of mutual funds and hedge funds, poor performance increases the chance that the manager will be fired and that the fund will be shut down and merged with more successful funds or completely liquidated, while strong performance leads to the growth of the assets under management. In private equity, poor performance decreases the probability that a firm will be able to raise a follow-on fund, while strong performance leads to larger follow-on funds.

Table III **The relation between past performance and future fundraising.** This table describes the relationship between past fund returns and future fundraising. Columns 1 and 3 present estimates from linear probability models where the dependent variable is one if the focal (current) fund has a follow-on and zero otherwise, while columns 2 and 4 report coefficients from probit models. Columns 5 and 6 present estimates for regressions predicting the growth rate of the follow-on fund, conditional on the current fund having a follow-on fund. The dependent variable is the logarithm of one plus the growth of the follow-on fund relative to the focal (current) fund. For probit models, pseudo R^2s are reported. The t-statistics presented in parenthesis are based on standard errors clustered at the private equity firm level and robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Probability of raising fund $n+1$				Growth of fund $n+1$	
	Buyout		VC		Buyout	VC
	OLS	Probit	OLS	Probit	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Fund n performance	0.79***	4.24***	0.47***	2.04***	0.94***	0.61***
	(7.34)	(4.94)	(6.27)	(3.95)	(3.85)	(4.04)
Log(fund n sequence #)	0.17***	0.70***	0.09**	0.37**	-0.12*	-0.17***
	(5.14)	(3.28)	(2.56)	(2.49)	(-1.84)	(-2.75)
Fund n performance*	-0.29***	-0.97	-0.20***	-0.88*	-0.11	-0.32***
log(fund n sequence #)	(-2.75)	(-1.03)	(-3.35)	(-1.94)	(-0.35)	(-3.02)
Fund $n-1$ performance	0.37***	2.66***	0.12**	1.32***	0.47**	0.03
	(3.69)	(3.38)	(2.02)	(2.71)	(2.08)	(0.24)
I(fund n is first fund)	0.13***	0.83***	0.01	0.26	0.10	-0.07
	(2.65)	(3.14)	(0.15)	(1.32)	(1.06)	(-0.76)
Vintage-region F.E.	Yes	n.a.	Yes	n.a.	Yes	Yes
Adj. R^2	0.12	0.16	0.11	0.08	0.15	0.21
Observations	1,169	1,169	1,034	1,034	996	800

The right side of Table III presents evidence on the relationship between past fund performance and the size of the follow-on fund conditional on the follow-on fund being raised. The independent variables are the same as those used to predict the probability of raising a follow-on fund. The dependent variable is the natural logarithm of one plus the growth in size of fund $n + 1$ relative to fund n. The results indicate that fund growth is strongly related to past performance, with t-statistics around 4. Furthermore, average fund growth and growth sensitivity to past returns tend to decrease as firm matures. The latter result is statistically significant for venture capital funds, but not for buyout funds. In addition to the performance of fund n, the return on fund $n - 1$ helps to predict the future fund growth of buyout firms, but not that of VC firms. The coefficient estimates for the sensitivity of fundraising to past performance are larger for buyout firms than for VC firms across all specifications, suggesting that buyout fundraising is more sensitive to past returns than VC fundraising.

C.2. Performance Persistence

The extent to which variation in performance across fund managers is attributable to skill rather than luck is one of the most important questions in asset management.[8] This issue is crucial to understanding why the empirical analysis of private equity fund returns is biased towards finding evidence of decreasing returns to scale.

Korteweg and Sorensen (2017) develop a variance decomposition model to disentangle skill from luck in private equity performance. An objective of their model is to understand how much of the extraordinary amount of variation in fund returns is due to noise and how much is due to persistent differences in skill across private equity firms. To this end, they define the total log-return of the n^{th} of firm i, y_{in}, as $10 * log(1 + IRR_{in})$ and model it as:

$$y_{in} = X'_{in}\beta + \sum_{\tau=t_{in}}^{t_{in}+9} (\gamma_i + \eta_{i\tau}) + \epsilon_{in}, \qquad (1)$$

where X_{in} contains vintage year fixed effects and potentially other controls, γ_i is a private equity firm-level random effect which is constant across all funds managed by the same firm, $\eta_{i\tau}$ is a firm-time random effect that is designed to capture the covariance in returns of overlapping funds of the same firm, and ϵ_{in} is a fund-specific random shock. In this model, the mean γ_i across all the private equity firms in the sample is 0. Hence, γ_i is a measure of relative skill and the model focusses on the cross-section of skill rather than on the average skill of private equity fund managers. The model is estimated as a hierarchical

[8] For evidence on the cross-section of equity mutual fund returns see, among others, Kosowski, Timmermann, Wermers, and White (2006), Fama and French (2010) and Linnainmaa (2013).

Table IV **Parameter Estimates for the Variance Decomposition Model.** This table reports the estimated posterior means of the parameters of the variance decomposition model of Korteweg and Sorensen (2017) estimated on my sample by Markov Chain Monte Carlo (MCMC) method with 10,000 burn-ins and 100,000 simulation runs. The error term is modeled as a mixture of 2 normal distributions for buyout funds and 3 normal distribution for venture capital funds. The model includes vintage year fixed effects. Bayesian standard errors for each posterior mean estimate are in parenthesis. See the text (section III.C.2) and Korteweg and Sorensen (2017) for further details.

	Parameter estimates			Signal-to-noise
	σ_γ	σ_η	σ_ϵ	S_γ
	(1)	(2)	(3)	(4)
Buyout	0.038	0.141	1.280	0.073
	(0.004)	(0.022)	(0.076)	(0.026)
Venture capital	0.055	0.142	2.023	0.066
	(0.007)	(0.039)	(0.105)	(0.024)

model by fitting it to the data via Bayesian Markov Chain Monte Carlo (MCMC). The return components γ_i and $\eta_{i\tau}$ are assumed to be normally distributed with mean 0 and variance σ_γ^2 and σ_η^2 respectively, while the error term ϵ_{in} is modeled as a mixture of normal distributions with overall variance σ_ϵ^2. This procedure allows us to decompose the variance in total returns σ_y^2 into three components:

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

where $100\sigma_\gamma^2$ is the amount of variation due to real, long-term differences in skill across private equity firms, while $10\sigma_\eta^2$ is variation due to time-varying random noise shared by partially overlapping funds managed by the same firm and σ_ϵ^2 is variation coming from pure fund-level random shocks.

In this setting, the signal-to-noise ratio is defined as the proportion of the total variation that is due to variation in managerial skill:

$$S_\gamma = 100\sigma_\gamma^2/\sigma_y^2. \tag{3}$$

I use this model to decompose the total variation in fund returns in my sample. Following Korteweg and Sorensen (2017), the model is estimated via MCMC simulations with 10,000 burn-ins and 100,000 simulations. Table IV reports the results. The first column shows that the estimated amount of variation in skill across funds, σ_γ, is 3.8% for buyout and 5.4% for venture capital. These estimates are large, as they imply that the spread in expected returns between top and bottom quartile partnerships is over 5 and over 7 percentage points annually for buyout and venture capital, respectively. However, the amount of return variation that is due to noise is much larger, and so the signal-to-noise ratio is only 7.3% for buyout and 6.7% for venture capital. This result implies that, even though some long-term persistence in returns exists, it is extremely hard to identify it in real time. The amount of persistence estimated in my sample is almost identical to that estimated by Korteweg and Sorensen (2017) for venture capital, but it is lower for buyout. Consistent with Harris et al. (2014b), which shows that performance persistence of buyout funds has decline in recent years, in unreported tests I find that the lower estimate of buyout persistence in my sample relative to the estimates of Korteweg and Sorensen (2017) is driven by funds raised in the most recent decade.

In this setting, the investor's learning process about a firm's skill level γ_i can be modeled using a Gaussian learning model. The prior for a firm's skill γ_i is distributed as $\mathcal{N}(0, \sigma_\gamma^2)$, and after observing N+1 fund returns the belief is updated to $\mathcal{N}(\gamma_{i,N+1}, \sigma^2_{\gamma_{i,N+1}})$, where

$$\gamma_{i,N+1} = S_\gamma \cdot \frac{y_{i,N+1} - X'_{i,N+1}\beta}{10} + (1 - S_\gamma) \cdot \gamma_{i,N} \tag{4}$$

and

$$\sigma^2_{\gamma_{i,N+1}} = (1 - S_\gamma) \cdot \sigma^2_{\gamma_{i,N}}. \tag{5}$$

As discussed above, the models estimates reveal that there exists a significant amount of persistence in skill, but, at the same time, it is difficult for investors to identify it, because the amount of noise in fund returns is particularly large. A numerical example illustrates the issue. Suppose we just observed the return of the first-time fund of a buyout partnership. Recall that, for buyout partnerships, the standard deviation of the distribution of differential skill γ_i, i.e., σ_γ, is estimated to be 3.8% and the signal-to-noise ratio S_γ is 7.3%. Suppose the return we observed is 25%, and we also know that the sample mean buyout return is approximately 16%. Hence, the observed fund has outperformed the average fund by 9%. Using the Bayesian updating method illustrated above, after observing this fund return the posterior for the GPs skill is distributed as $\mathcal{N}(0.0066, 0.035^2)$. Hence, in expectation, only a small part of the large outperformance of 9 percentage points can be attributed to skill (i.e., 0.66 percentage points), while the majority should be attributed to luck.

The example provided in the previous paragraph is instrumental to understanding why a simple analysis of the data is biased towards finding strong evidence of decreasing returns to scale at the private equity firm level. On average, even if a certain amount of skill can be statistically detected in the data, most of the past outperformance generated by firms whose follow-on funds have grown the most must have been driven by random shocks. Since the average expected shock across follow-on funds is zero, we should expect that the funds that grow the most will underperform their preceding funds by a large amount, regardless of whether actual decreasing returns to scale exist in the management of private equity firms.

D. Estimating the True Effect of Growth on Returns

The estimates presented in Table II, if taken at face value, seem to suggest that private equity firms face large decreasing returns to scale when they raise larger funds. However, we have argued that these estimates are likely biased. The next step is to obtain an unbiased assessment of the effect of fund growth on performance. Is the measured empirical relationship between growth and returns entirely generated by the selection mechanism, or is there any evidence of decreasing returns left once we control for it?

To answer this question, I start from the private equity firm-level returns to scale regression specification,

$$fund\,return_{i,n} = a_0 + a_1 \cdot log(1 + fund\,growth_{i,n}) + \mu_v + \nu_i + \epsilon_{i,n}, \tag{6}$$

where $fund return_{i,n}$ is the return of the n^{th} fund managed by firm i, μ_v is the vintage year-region fixed effect, ν_i is the private equity firm fixed effect, and $fund growth_{i,n}$ is size growth of the n^{th} fund relative to the first fund managed by the same firm. As argued above, this model specification has the advantage of controlling for the endogeneity of fund size, skill, and investment strategy. However, the coefficient of interest, a_1, is likely downward-biased because of reversion to the mean.

In order to obtain an unbiased estimate of decreasing returns in this context, we need to know what value coefficient a_1 would take if the empirical relations creating the bias were in place, but no true decreasing returns to scale existed. To do so, we need to take the expectation of a_1 with respect to the full joint distribution of the variables involved (i.e., fund returns, fund growth rates, and sequence numbers) after setting decreasing returns to scale to zero.

I use a simulation technique to analyze this problem. As with previous analyses, this exercise is carried out separately for buyout and venture capital funds. I start by initializing one private equity firm for each firm in the data (496 buyout firms and 457 venture capital firms). In order to model the return-generating process, I use the model proposed by Korteweg and Sorensen (2017), specifically, I generate returns by drawing from equation (1).

The underlying parameters, i.e, σ_γ, σ_η, and σ_ϵ are estimated from the data and have been presented in Table IV. In particular, in order to match the non-normal distribution of fund returns, the fund-specific shock component of fund returns, ϵ_{in}, is modeled using Gaussian-mixtures of 2 normal distributions for buyout and 3 normal distributions for venture capital. Using parameters estimated from the sample ensures that the simulated returns have the same distribution as well as the same ratio of skill to noise and the same amount of spurious short-term persistence as in the data.[9]

In the simulation exercise, each private equity firm is assigned a skill level γ_i by drawing randomly from the estimated distribution of differential skill $\mathcal{N}(0, \sigma_\gamma^2)$. I want to generate a panel of data where no decreasing returns to scale exist. To do so, I draw γ_i independently of fund size or fund growth, and I keep it constant across all the funds managed by the same firm i. Each simulated firm's first-time fund has the same vintage year as the corresponding fund in the sample data. The return of each fund is generated from equation (1). Finally, $X'\beta$ is a constant that is added to the sum of the fund return components, $\gamma_i's$, $\eta_i's$, and $\epsilon_i's$ in order to ensure that the mean generated return is equal to the mean return observed in the sample data.

After drawing the returns of each first-time fund, I obtain the probability that they would have a follow-on fund using the probit model estimated in columns (2) and (4) of Table III. I then determine which funds will actually have a follow-on fund by drawing randomly accordingly to the probability predicted by the probit model. For each firm that raises a follow-on fund, the growth in the size of the follow-on is obtained using the growth-performance equations estimated in columns (5) and (6) of Table III. The error term is drawn directly from a probability density function fitted to the empirical error distribution from these equation. This way, the simulated data displays the same relationship between past performance and fundraising as the actual data. Each follow-on funds vintage year is 3 to 5 years after that of the preceding fund, and the number of years between funds is chosen according to probabilities based on the sample data.

For each sequence-two fund, the return is simulated following a procedure similar to that used for the returns of first-time funds. The only difference is that not all of the ten firm-time random effects are drawn randomly again; instead, the first m $\eta_i's$ are the same as the last m $\eta_i's$ of the preceding fund of the same firm, where m is the number of overlapping years shared by the preceding fund and follow-on fund. Then, using the returns of sequence-two funds, I calculate each funds probability of raising a sequence-three follow-on fund (again,

[9] A less sophisticated, less precise version of the simulation in which fund returns are simply bootstrapped from the empirical distribution (as opposed to being drawn based on the estimates of the Korteweg and Sorensen (2017) model) delivers similar results.

Figure 1: **Distribution of Sample Variables and Simulated Variables.** This figure shows the distribution of IRR and fund growth rates for the private equity panel data simulated following the procedure described in section III.D (main simulation). The actual distribution of the same variables in the sample obtained from Preqin is also presented. The four panels show, in clockwise order from the top left, the distribution of buyout IRRs, that of venture capital IRRs, that of venture capital fund growth rates, and that of buyout fund growth rates.

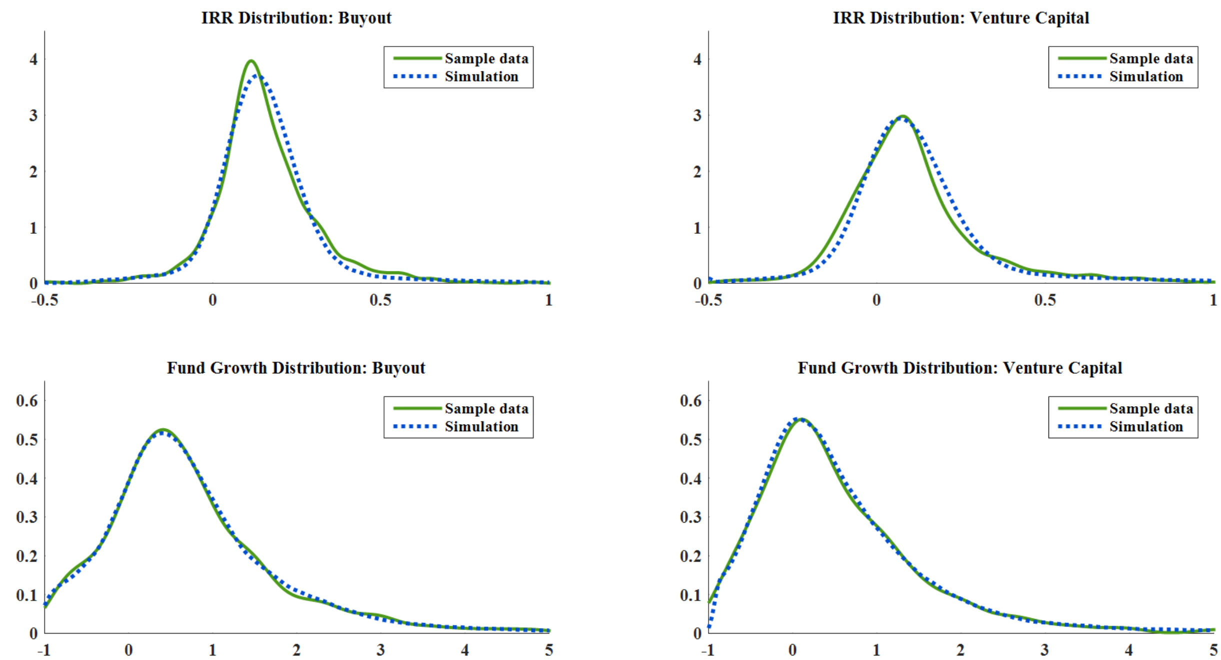

using the probit model estimated in the data). Using those probabilities I randomly decide which funds will have a follow-on fund and then I determine the growth of those funds using the growth-performance equations estimated in the data. This process continues until the last vintage year, i.e., 2011, is reached. I run this algorithm 10,000 times in order to obtain 10,000 independent panels of simulated data.

For the purposes of this analysis, it is important that the distribution of the simulated variables closely matches the distributions of the variables in the sample data. Figure 1 shows the simulated distribution of the two key continuous variables, i.e., fund returns and growth rates, and compares them with the distributions of the same variables as found in the data. The simulated and the sample distributions are very similar for both buyout and venture capital. A Kolmogorov-Smirnov test comparing the distributions of these variables in the sample and in the simulated data indicates that the empirical distribution and the simulated distribution do not differ significantly. The test and the characteristics of the simulated data is discussed in greater detail in the Internet Appendix.[10] We can therefore conclude that we have successfully generated panels of simulated data that have the desired characteristics in terms of distribution and joint distribution and where, notably, there are no decreasing returns to scale by construction.

An important detail of the simulation procedure is that the fundraising-performance relationship is modeled using the final total fund return of the current (focal) fund as opposed to the interim return up to the moment of fundraising. This specification, although used extensively in the literature (e.g., Kaplan and Schoar (2005), Chung et al. (2012), Hochberg et al. (2014)), might lead to a sort of error-in-variable problem, because the final return is not known with certainty at the moment of fundraising. As a result, the sensitivity of fundraising to current fund performance may be underestimated due to attenuation bias. Nonetheless, for the purpose of the simulation, using the final fund return is preferable. The reason is that the objective of the simulation is to estimate the bias in the OLS estimator of the effect of fund growth on performance in a panel regression setting where performance is measured as the *final* fund return. Therefore, modelling the fundraising-performance relation using interim returns would lead to overestimating the magnitude of the bias. It should also be noted that using the final fund return not only is conceptually correct, but it also makes it more difficult to find evidence consistent with the thesis of this paper (i.e., that in private equity data there exists a bias towards finding evidence of decreasing returns) and is therefore a prudent research design choice.

Table V presents the main results of the simulation exercise described above. I run regressions of fund returns on the natural logarithm of one plus cumulated fund growth with private equity firm fixed effects (i.e., specification (6)) using the sample data and the simulated data and then I compare the estimates of the returns to scale slope a_1. Panel A and B show results for buyout and venture capital, respectively. Column 1 reports a_1 estimated in the sample data; this is the same estimate already presented in the second column of Table II.

I also run three different shadow simulations in parallel to the main simulation. These auxiliary simulations are designed to guide us through the simulation exercise while also verifying that the methodology used produces sensible results. In columns 2 through 5 of Panels A and B, the estimates presented are the average coefficient a_1 and t-statistic across the 10,000 simulations.

In the first shadow simulation, returns to scale are set to 0 and there is no relationship between performance and fundraising in the simulated data. That is, both the probability of raising a follow-on fund and its size are independent of past performance. As expected,

[10]The Internet Appendix is available from the author upon request.

Table V **Returns to Scale Estimates Adjusted for Reversion to the Mean.** This table reports the main results of the simulation study designed to estimate and control for the econometric bias in returns to scale regressions in private equity data. Panel A and B report returns to scale slope a_1 from regression specification (6) estimated in the sample data and in the simulated data. Section III.D provides details regarding the simulation exercise. Panel C summarizes the results presented in the first two panels. *** indicates statistical significance at the 1% level.

Panel A: Buyout funds

	Data	Shadow simul. 1	Shadow simul. 2	Main simul.	Shadow simul. 3
	(1)	(2)	(3)	(4)	(5)
True a_1 set to	n.a.	0	0	0	-0.46
$Prob(fund_{n+1})=f(past\ returns)$	n.a.	No	Yes	Yes	Yes
$Growth(fund_{n+1})=f(past\ returns)$	n.a.	No	No	Yes	Yes
Estimate of a_1	-5.12***	-0.01	-0.10	-4.66***	-5.17***
t-statistic	(-4.71)	(-0.01)	(-0.11)	(-5.31)	(-6.12)

Panel B: Venture capital funds

	Data	Shadow simul. 1	Shadow simul. 2	Main simul.	Shadow simul. 3
	(1)	(2)	(3)	(4)	(5)
True a_1 set to	n.a.	0	0	0	-1.60
$Prob(fund_{n+1})=f(past\ returns)$	n.a.	No	Yes	Yes	Yes
$Growth(fund_{n+1})=f(past\ returns)$	n.a.	No	No	Yes	Yes
Estimate of a_1	-8.68***	0.00	-0.24	-7.08***	-8.81***
t-statistic	(-4.28)	(0.00)	(-0.16)	(-5.46)	(-6.71)

Panel C: Adjusted estimates

	Buyout funds		Venture capital funds	
	Observed (biased)	Adjusted estimate	Observed (biased)	Adjusted estimate
Returns to scale slope a_1	-5.12***	-0.46	-8.68***	-1.60
p-value	(0.00)	(0.31)	(0.00)	(0.12)
Δ(fund return) for +100% size	-3.55	-0.32	-6.02	-1.11

the average returns to scale slope a_1 in this case is virtually 0 and statistically insignificant in all of the 10,000 simulation runs for both buyout and venture capital.

In the second shadow simulation, the probability of raising a follow-on is no longer random; instead, it is a function of past fund returns, just as it is in the data. In this case, the slope a_1 is on average slightly negative for both private equity fund types, but it is statistically significantly negative at the 5% level in only 3% and 3.8% of the 10,000 buyout and venture capital simulation runs, respectively. Moreover, its magnitude is not comparable to that of the coefficient found in the actual data. Therefore, in this context, controlling for whether a follow-on fund exists (i.e., using a selection-correction model, e.g., Heckman (1979)) is not sufficient to deal with the bias in the estimator of the slope a_1.

Column 4 shows the results for the main simulation. Here, both the probability of raising a follow-on fund and its growth depend on past performance. Despite the fact that fund returns are generated independently of fund size (i.e., true a_1 is 0), the estimated slope a_1 is negative in 100% of the 10,000 simulation runs, and statistically significant at the 1% level in 99% of the runs for buyout and 99.4% for venture capital. The average t-statistic is below -5 for both private equity fund types. This strongly supports the thesis of this paper. That is, in private equity data there exists a strong bias towards finding spurious evidence of decreasing returns to scale at the firm level. The mean value of a_1 in the main simulation is the estimate of the bias in the empirical estimator of a_1.

Panel C of Table V provides a summary of the analysis carried out. Specifically, it reports the observed slope coefficient as well as the bias-adjusted estimate. The latter is calculated as the difference between the coefficient estimated in the data before any adjustment and the slope obtained from the main simulation, which represent the estimated bias. For buyout, the OLS estimate of the returns to scale slope is -5.12, while the adjusted estimate is -0.46, which is 91% smaller. For venture capital, the OLS estimate of the slope is -8.68, while the adjusted estimate is -1.60, which is 81.6% smaller. p-values for the adjusted estimates are simulation-based empirical p-values, i.e., they measure the proportion of simulation runs in which the adjusted estimate is positive. For both private equity investment types, the slope is highly statistically significant before the adjustment, but becomes insignificant after the adjustment. The empirical p-value for buyout and venture capital funds is 0.31 and 0.12, respectively. Overall, this analysis indicates that, although in the data fund growth and performance are negatively related, most of this relationship is attributable to reversion to the mean, and the actual effect of fund growth is small and statistically insignificant.

After the main simulation, I run an additional shadow simulation whose results are reported in the fifth column of Panels A and B. Similarly to the main simulation, the data is generated so as to contain the same fundraising-performance relation as in the sample data. However, in this case the true slope a_1 is set equal to the estimated bias-adjusted figure reported in Panel C. Specifically, true a_1 is set equal to -0.46 for buyout and -1.60 for venture capital. Given these parameters, the regression estimate of a_1 in the simulated data is extremely similar to its counterpart in the actual sample data. Specifically, for buyout the coefficient is -5.12 using the sample data and -5.17 using the data generated in the fourth shadow simulation. For venture capital, these figures are -8.68 and -8.81, respectively. Hence, the simulated coefficients are approximately within 1.5% of the coefficient observed in the data for both private equity investment styles. This result serves as a robustness test for the methodology employed and confirms the interpretation given to the results: the large negative association between fund growth and performance observed in private equity data likely reflects a much weaker true effect (about 80% to 90% smaller in magnitude) coupled with a complex selection bias.

The last two rows of Panel C show the estimated change in return (IRR) for a 100% increase in fund size. As highlighted in the introduction, many investors are wary of committing capital to a fund that is two or three times larger than the preceding fund managed by the same partnership. A first glance at the data might suggest that investors concerns are not baseless, because the unadjusted panel regression estimates suggest that a 100% size increase is associated with a decrease in returns of 3.55 p.p. for buyout and of 6.02 p.p. for venture capital. However, after adjusting these estimates for econometric bias, the expected change in IRR is only a statistically insignificant decline of 0.32 p.p. for buyout and of 1.11 p.p. for venture capital.

IV. What Do Private Equity Firms Do When They Raise a Larger Fund?

The analysis conducted in the previous section indicates that the negative relation between fund growth and performance we observe in the data is primarily a statistical artefact due to reversion to the mean. Estimates that adjust for this effect indicate that the causal effect of decreasing returns on performance is economically small and statistically insignificant. The silver lining of this result is that, historically, private equity firms that have raised large follow-on funds have been able to avoid incurring strong decreasing returns, on average.

To provide some insights into how private equity firms change their fund-level activities when they raise a larger fund, in this section I analyze the relationship between the size of a fund, defined as total capital committed to it, and the number of portfolio companies it invests in. To do so, I use deal-level data from a large sample of buyout and VC funds provided by PCRI. These data is described in section II.B.

The objective of this analysis can be illustrated with an example. Suppose a private equity firm's first-time fund has $100 million in capital commitments and has delivered good returns to investors. As a result, after a few years, the firm is able to raise a second fund with a size of $200 million. Suppose the first fund had 10 portfolio companies, in which it invested on average $10 million. How will the firm invest the capital committed to its second fund? For instance, it could make 10 investments of $20 million each. Alternatively, it could decide to invest an average of $10 million in 20 portfolio companies. More likely, compared to its prior fund, the firm could increase both the number of investments and the average size of the investments it makes. The optimal decision will depend both on the investment opportunities the firms faces as well as on the resources the firm has to identify and analyze those opportunities and then manage the investments made.

Prior research has address the issue of scalability of buyout and VC investments from different angles. Metrick and Yasuda (2010) study how contract terms vary across a sample of buyout and VC funds and how these terms interact with fund size in determining the expected revenue of private equity firms. They estimate that revenue per partner in follow-on funds increases significantly faster for buyout funds than for VCs, and conclude that "the buyout business is more scalable than the VC business". Consistent with the results presented in Table III of the current article, other authors report that buyout follow-on funds tend to grow more than VC follow-ons, and that the former are more sensitive to past performance than the latter (e.g., Kaplan and Schoar (2005), Chung et al. (2012)). However, to the best of my knowledge, the literature has not provided a formal answer to the question of how the *composition* of the portfolios of private equity funds evolves as a function of their size.

Here, I provide empirical evidence to answer this question. Specifically, I analyze how the number of portfolio companies and the average investment size vary across funds with different sizes. The total size of a fund's portfolio can be defined as fund size = number of portfolio companies*average investment size. Taking the natural logarithm, we obtain log(fund size) = log(number of portfolio companies) + log(average investment size). Then, we can separately regress log(number of portfolio companies) and log(average investment size) on log(fund size) in order to determine the percentage change in the number of portfolio companies and in the average investment size associated with a change in fund size.

The results of this analysis are presented in Table VI. The regressions are estimated separately for buyout and VC funds. The specifications without PE firm fixed effects describe the cross-section of fund portfolios, while the specifications with PE firm fixed effects describe how the portfolios of funds managed by a given PE firm change as a function of fund size.

Table VI **Large Funds: More Investments or Larger Investments?** This table describes how fund portfolios vary across private equity funds of different size. Fund size is defined as the total capital committed to a fund. Regressions are used to decompose differences in fund size into differences in the number of portfolio companies and in the average investment size in each portfolio company. The specifications without PE firm fixed effects describe the cross-section of fund portfolios, while the specifications with PE firm fixed effects describe how the portfolios of funds managed by a given PE firm change as a function of fund size. By construction, the coefficient estimates in each specification sum to 1. The t-statistics presented are based on standard errors clustered at the private equity firm level and robust to heteroskedasticity. * and *** indicate statistical significance at the 10% and 1% levels, respectively. See section II.B for a detailed description of the data

	% change in number of portfolio companies for a 1% increase in fund size	% change in average investment size for a 1% increase in fund size	PE firm F.E.
Buyout	0.171*** [t-stat=4.69]	0.829*** [t-stat=22.70]	No
Buyout	0.208* [t-stat=1.68]	0.792*** [t-stat=6.39]	Yes
VC	0.350*** [t-stat=26.33]	0.650*** [t-stat=48.81]	No
VC	0.384*** [t-stat=10.75]	0.616*** [t-stat=17.26]	Yes
H_0: Buyout=VC	p-value=0.000	p-value=0.000	No
H_0: Buyout=VC	p-value=0.069	p-value=0.069	Yes

By construction, the coefficient estimates in each specification sum to 1.

The cross-sectional results indicate that, compared to small buyout funds, larger buyout funds predominantly make larger investments rather than more investments. Specifically, 83% of an increase in fund size can be attributed to increases in the size of the investments made, while only 17% can be attributed to increases in the number of investments made. The results are almost identical when private equity firm fixed effects are included. Hence, when a buyout firm raises a larger follow-on fund, it primarily takes over larger companies rather than making more investments.

The results are rather different for VC funds. In the cross-section, changes in the number of portfolio companies and in the average investment size account for 35% and 65% of a 1% change in fund size, respectively. When private equity firms fixed effects are included, these estimates become approximately 38% and 62%, respectively. Hence, compared to buyout funds, larger VC funds do not seem to simply make larger investments; on the contrary, they also make a significantly larger number of investments. The last two rows of Table VI test whether the measured relationship between portfolio composition and fund size are statistically different for buyout funds relative to VC funds. The results indicate that they are indeed statistically different. The rejection of the null hypothesis tested is less strong when private equity firms fixed effects are included because the test uses standard errors clustered at the firm level. That is, to be conservative, I report the p-value based on clustered and robust standard errors, i.e., 0.069; without clustering, the p-value is 0.000.

Figure 2: **How Do Portfolios Change as Funds Grow?** This figure shows the hypothetical evolution of the portfolio composition of a sequence of four funds managed by a buyout firm and by a VC firm. The size of the first-time funds is set equal to the 25^{th} percentile of the distribution of fund size in the Preqin sample (i.e., \$130 million for buyout funds and \$48 million for VC funds). Each subsequent fund is assumed to grow at the median growth rate for funds with the same sequence number. Estimates from Table VI are used to calculate the number of portfolio companies and the average investment size of each fund as a function of its fund size.

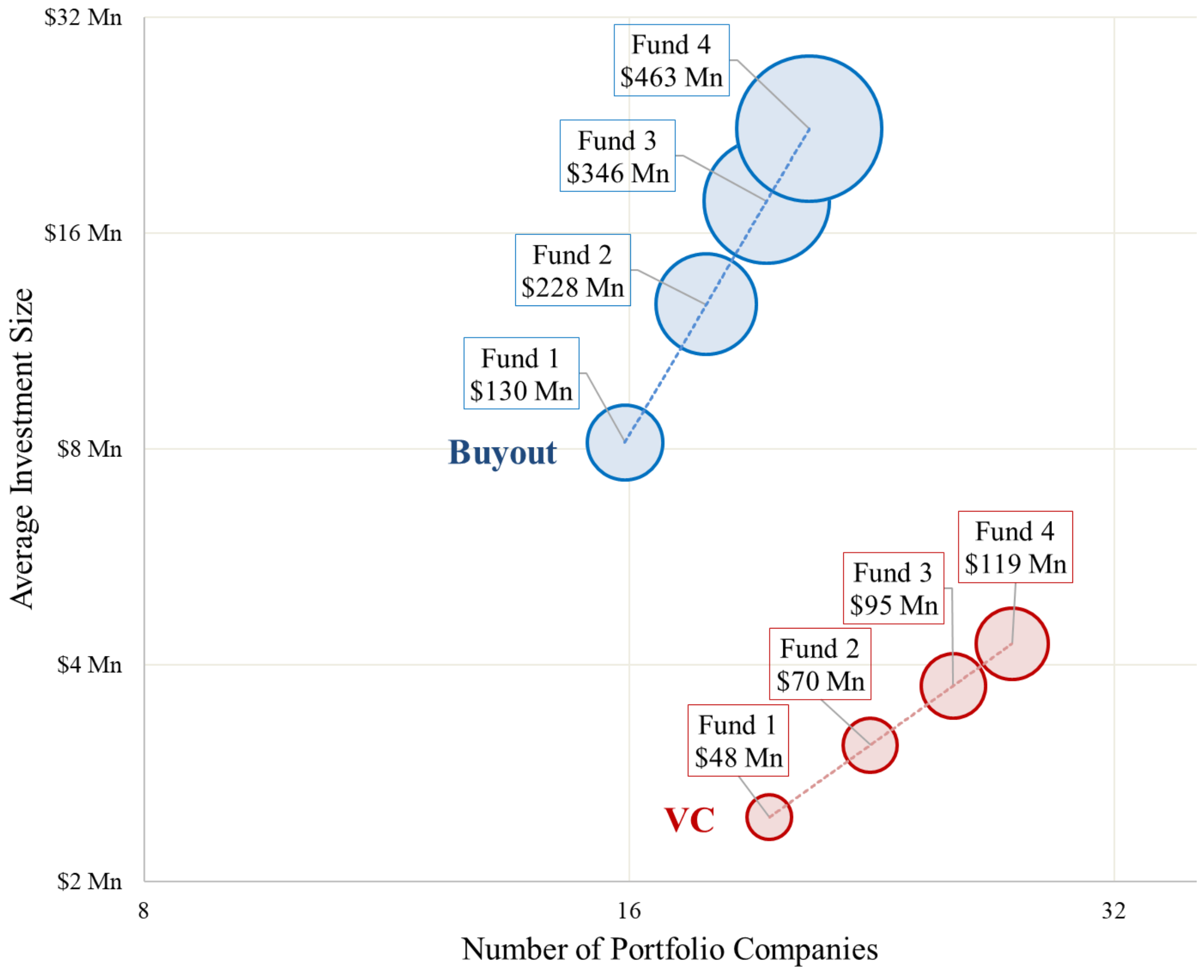

To provide economic context to these results, I use the estimates obtained from the regression analysis to provide an example that illustrates how the portfolios of subsequent funds managed by a given firm would change with fund size. The examples discussed here is illustrated in Figure 2.

Suppose the first-time funds of a buyout firm has a portfolio of equity investments of total size equal to $130 million dollars, which is the size of a 25^{th} percentile fund in Preqin. Using the regression results discussed above, we can estimate that a buyout fund of this size would have approximately 16 portfolio companies with an average equity investment of $8.1 million. Conditional on its existence, the median size growth rate of a sequence-two buyout fund is 75%. Hence, suppose the second fund raised by the buyout firm in this example has a size of $228 million (228=130*175%). Using the estimates reported in Table VI, a 75% increase in fund size will on average result in a 12.3% increase in the number of portfolio companies ($0.123 = e^{0.208*ln(1+0.75)} - 1$) and in a 55.8% increase in the average investment size ($0.558 = e^{0.792*log(1+0.75)} - 1$). Hence, the second fund will have approximately 18 investments of $12.7 million each, on average. Using actual fund growth rates measured in the data and the same logic applied above, the third fund raised by the hypothetical buyout firm would have a portfolio with an average of 19.5 investments totalling $346 million, and the fourth fund would have a portfolio with an average of 20.7 investments for a total of $463 million. Hence, the fourth fund is about 3.5 times larger than the first fund managed by the same buyout firm, and yet the number of portfolio companies in the fourth fund is only about 30% larger (i.e., the number of companies went from 16 to 20.7).

Figure 2 also illustrates the portfolio characteristics of a series of hypothetical funds managed by a venture capital firm. The logic used to generate the funds and their portfolio characteristics is the same as that used in the buyout example above. The first-time fund has a portfolio size of $48 million, which is equal to the size of a 25^{th} percentile VC fund in Preqin. Using the regression estimates, this fund would have an average of 19.5 investments of around $2.5 million each. The figure shows how the size of the portfolio would change for three funds raised after the initial one. Compared to buyout funds, follow-on VC funds tend to grow slower, and yet they have to increase the number of investments they make faster than buyout funds.

Overall, the analysis carried out in this section provides us with insights that complement the main finding of this paper, i.e., that estimates of the true effect of fund growth on performance are economically small and statistically insignificant, especially for buyout funds, and to a lesser extent for VC funds. This result implies that, historically, the managers of successful private equity firms–those that have been able to raise larger follow-on funds–have managed to scale up effectively and have avoided incurring severe decreasing returns to scale. The portfolio-level analysis suggests that private equity firms, especially buyout firms, are able to grow their funds primarily by making larger investments rather than more investments, thus avoiding to spread their skill too thin. Moreover, while the number of investments in a portfolio appears to grow less than proportionally with fund size, management fees would normally grow proportionally with fund size. Hence, large follow-on funds have significantly more resources to help them make and manage their investments, which can help to explain how fund managers are able, on average, to keep decreasing returns at bay.

V. Conclusion

When a private equity firm raises a larger fund, this fund's performance tends to be substantially lower than the performance of the funds the firm previously managed. If taken at face value, this would imply that fund-level decreasing returns are large and strongly binding in equilibrium. This would be symptomatic of a serious conflict of interest between fund managers and investors.

In this paper, however, I challenge this view. I show that the selection mechanism that governs private equity fundraising causes a bias which can lead to incorrect inference about returns to scale. Specifically, an analysis of the correlation between fund growth and performance is likely to misinterpret simple reversion to the mean as being *caused* by fund growth. I show that, controlling for reversion to the mean, the estimated impact of fund growth on performance becomes approximately 90% and 80% smaller for buyout and venture capital funds, respectively, and is not statistically different from zero, especially so for buyout funds.

These findings have several important implications for academics and practitioners alike. First, the fact that fund-level decreasing returns to scale are much weaker than previously thought provides us with an important piece to solve the puzzle of why a certain amount of long-term performance persistence has been observed in private equity. A plausible explanation is that, historically, decreasing returns have not been large enough to push the industry towards a Berk and Green (2004)-type equilibrium.

Moreover, understanding the extent to which scale impacts expected returns is critical for investors's allocation decisions. Average buyout and venture capital vintage year performance (both absolute and relative to equity markets) is negatively related to the aggregate amount of capital committed to new funds in that year. This fact is particularly troubling for limited partners because it implies that, in aggregate, the dollar-weighted returns they earn on their private equity investments are lower than the average fund return. The findings presented in this paper suggest that this negative relationship is unlikely to be driven by decreasing returns at the individual fund level. On the contrary, the negative relationship between aggregate fundraising and future performance is likely attributable to market-wide or industry-wide factors.

In light of this paper's results, a remark should be addressed to the institutions that invest in private equity funds and to the professionals who advise them. Although the average realized performance of the follow-on funds raised by the most successful private equity firms has certainly been disappointing for limited partners, most of the "disappointment" should be attributed to positive idiosyncratic shocks (that is, "luck") in past fund returns reverting to zero, rather than to a negative impact of fund growth on performance. When the noisy nature of private equity returns is taken into account and we form appropriate expectations that account for expected reversion to the mean, reasonable cross-sectional differences in fund growth rates no longer appear to be a first-order concern, especially when allocating across buyout funds.

Finally, although the findings presented imply that, historically, private equity managers have been able to scale up effectively, we should bear in mind that these result are, of course, valid only in-sample. Private equity firms might encounter stronger decreasing returns if they grow at faster rates and past the size levels currently observed in the data. In fact, the estimates suggest that the fastest-growing VC funds might almost have reached the boundary where fund growth starts to hinder performance significantly.

REFERENCES

Axelson, U., T. Jenkinson, P. Stromberg, and M. Weisbach, 2013, Borrow cheap, buy high? The determinants of leverage and pricing in buyouts, *Journal of Finance* 68, 2223–2267.

Berk, Jonathan B., and Richard C. Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.

Braun, R., T. Jenkinson, and I. Stoff, 2016, How persistent is private equity performance? Evidence from deal-level data, *Journal of Financial Economics* 123, 273–291.

Brown, G. W., R. S. Harris, T. Jenkinson, S. N. Kaplan, and D. Robinson, 2015, What do different commercial data sets tell us about private equity performance?, Working paper, NBER.

Chen, Joseph, Harrison Hong, Ming Huang, and Jeffrey D. Kubik, 2004, Does fund size erode mutual fund performance? The role of liquidity and organization, *American Economic Review* 94, 2941–2969.

Chung, J. W., B. A. Sensoy, L. Stern, and M. S. Weisbach, 2012, Pay for performance from future fund flows: The case of private equity, *Review of Financial Studies* 25, 3259–3304.

Cong, L. W., and Y. Xiao, 2017, Persistent blessings of luck, Working paper, University of Chicago.

Fama, Eugene F, and Kenneth R French, 2010, Luck versus skill in the cross-section of mutual fund returns, *Journal of Finance* 65, 1915–1947.

Gompers, P., and J. Lerner, 2000, Money chasing deals? The impact of fund inflows on private equity valuations, *Journal of Financial Economics* 55, 281–325.

Harris, R. S., T. Jenkinson, and S. N. Kaplan, 2014a, Private equity performance: What do we know?, *Journal of Finance* 69, 1851–1882.

Harris, R. S., T. Jenkinson, S. N. Kaplan, and R. Stucke, 2014b, Has persistence persisted in private equity? Evidence from buyout and venture capital funds, Working paper, University of Virginia.

Harvey, Campbell R., and Yan Liu, 2017, Decreasing returns to scale, fund flows, and performance, Working paper, Duke University.

Heckman, J. J., 1979, Sample selection bias as a specification error, *Econometrica* 47, 153–161.

Hochberg, Y., A. Ljungqvist, and A. Vissing-Jorgensen, 2014, Informational hold-up and performance persistence in venture capital, *Review of Financial Studies* 27, 102–152.

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

Hsu, D, 2004, What do entrepreneurs pay for venture capital affiliation?, *Journal of Finance* 59, 1805–1844.

Ivashina, V., and A. Kovner, 2011, The private equity advantage: Leveraged buyout firms and relationship banking, *Review of Financial Studies* 24, 2462–2498.

Kaplan, S. N., and A. Schoar, 2005, Private equity performance: Returns, persistence, and capital flows, *Journal of Finance* 60, 1791–1823.

Korteweg, A., and M. Sorensen, 2017, Skill and luck in private equity performance, *Journal of Financial Economics* 124, 535–562.

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.

Linnainmaa, J., 2013, Reverse survivorship bias, *Journal of Finance* 68, 789–813.

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

Nanda, R., S. Samila, and O. Sorensen, 2017, The persistent effect of initial success: Evidence from venture capital, Working paper 17-065, Harvard Business School.

Pástor, Luboš, and Robert F Stambaugh, 2012, On the size of the active management industry, *Journal of Political Economy* 120, 740–781.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2015, Scale and skill in active management, *Journal of Financial Economics* 116, 23–45.

Phalippou, L., 2017, *Private Equity Laid Bare*, first edition (CreateSpace).