

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

Andrea Rossi[†]

First Draft October 2016

Current Draft January 2018

Abstract

In private equity fund data, there exists an economically large negative association between fund growth and performance at the partnership level. This is usually interpreted as evidence of decreasing returns to scale. I argue that this inference is unwarranted. In essence, Bayesian-informed expectations reveal that the partnerships 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 80% smaller and statistically insignificant for both buyout and venture capital funds. Furthermore, I show that, historically, decreasing returns do not seem to have played a major role in the erosion of performance persistence in private equity. These results have implications for fund managers' and investors' decisions, and for our understanding of the private equity industry.

JEL classification: G11, G23, G24

Keywords: Private Equity, Buyout, Venture Capital, Skill and Luck, Returns to Scale, Fund Size, Asset Management

[†] Assistant Professor of Finance, The University of Arizona, Eller College of Management, McClelland Hall, Room 315J, Tucson AZ 85721-0108. E-mail: 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, Sergey Chernenko, Spencer Coutts, 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 seminar participants at The Ohio State University, UCLA, the 2017 European FMA Doctoral Consortium and the 2018 Southern California Private Equity Conference. All remaining errors are my own.

“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

1 Introduction

Private equity funds provide returns by making investments that could not be made by traditional sources of financing. Buyouts require monitoring by general partners (GPs) and unusually strong incentives for the portfolio firms’ managers, while venture capitalists nurture young entrepreneurs by providing advice and experience that help them build successful businesses. Yet, general partners’ time and energy are limited, and there are only so many potential portfolio firms for which private equity funds earn sufficient returns to offset their funds’ fees and illiquidity. These factors serve to limit the optimal fund size. In contrast, however, fund managers may be tempted to raise more capital than they can profitably invest, since the fees they receive increase with money under management.

It is not clear whether these limitations are sufficient to offset managers’ incentives to raise too much capital. For this reason, understanding whether returns decline with fund size is an important issue in the private equity industry. Returns to scale are likely to affect the optimal size of the funds that general partners decide to raise. Limited partners (LPs) have to choose between alternative funds of different sizes; presumably the extent to which fund performance is affected by fund size would affect their choices, especially when deciding to invest in a much larger follow-on fund of a successful partnership. Finally, academics want to understand the forces affecting the contractual relations in this industry, the effect of GP incentives on their fundraising activities, and the factors that drive fund returns.

The general consensus in the academic literature is that leveraged buyout and venture capital firms face large decreasing returns to scale at the partnership level. When a partnership raises a fund that is significantly larger than the funds it previously managed, the new fund will have a lower expected return than the partnership’s prior funds. This pattern is summarized in Panel A of Figure 1, which documents that when a follow-on fund is substantially larger than the previous fund managed by the same partnership, it tends to underperform the preceding fund. Panel regressions deliver similar inferences (e.g., Kaplan and Schoar 2005). Overall, the negative effect of fund growth on expected returns appears to be substantial.

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, Managing 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’”.*

This paper argues that conventional analyses of the data such as those mentioned above can be misleading. Private equity fund data is subject to a selection mechanism that is likely to lead to incorrect inferences about decreasing returns to scale at the general partner's (GP) level. There are three factors contributing to this bias. First, the higher a partnership's current fund return, the higher the probability that the partnership will be able to raise a follow-on fund (e.g., Kaplan and Schoar 2005 and Hochberg, Ljungqvist and Vissing-Jorgensen 2013). Second, the follow-on fund's growth is positively related to the current fund return (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 (or "noise") rather than by persistent differences in skill across GPs (Korteweg and Sorensen 2015). I show that these facts imply that (i) the returns of funds for which a follow-on was raised contained, on average, positive noise (i.e., luck); and that (ii) the higher the follow-on's growth relative to the current fund, the higher the expected positive noise in the current fund's return. 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 funds' returns to scale. How much should we expect returns to decline in the absence of decreasing returns to scale? How can we measure the magnitude of the positive noise in the returns of preceding funds whose follow-ons grow the most?

To address this issue, I use a Bayesian approach that builds on the model of Korteweg and Sorensen (2015). Korteweg and Sorensen develop a variance decomposition model designed to identify long-term persistence in relative skill across private equity partnerships and separate it from noise. 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 returns² 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.

Using a simple Bayesian learning model based on these estimates, I show that a substantial portion of the spread in realized returns between funds whose follow-ons grow the most and funds whose follow-ons grow the least is attributable to noise or, in other words, luck. The amount of noise is so large that,

² Private equity fund returns display by far the largest cross-sectional dispersion among the major asset management categories. Hochberg et al. (2013) report that the dispersion of after-fee venture capital fund returns is 2.5 times greater than that of mutual funds and 1.4 times greater than that of hedge funds. Buyout returns are just behind venture capital in terms of return dispersion (see Figure 1 in Hochberg et al. 2013).

combined with a positive fundraising-performance relationship, it can easily generate a spurious decreasing returns to scale effect that is of the same order of magnitude as the effect observed in the data.

Having determined that the bias exists and is likely to be large enough to distort inferences about the effects of size and growth on expected fund returns, I next obtain an estimate of the magnitude of the bias. To estimate the bias in the coefficient on fund growth in panel regressions, we need to take the expectation of that coefficient with respect to the full joint distribution of the variables involved (i.e., fund returns, growth rates and sequence numbers), after setting the true effect of fund growth to 0. To do so, I generate panels of private equity data (buyout and venture capital, separately) similar to the sample panel data. Returns are simulated so that they have the same distribution and the same amount of GP-level persistence as in the sample 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 sample. Returns are drawn independently of fund size, so that, by construction, there are no decreasing returns to scale in the simulated data.

I then estimate panel regressions on the simulated data. Despite the fact that the data were generated specifically *not* to contain size effects, I find a negative and significant relation between fund growth and performance. This suggests that the selection process that generates private equity data will lead one toward finding a negative association between fund size and returns at the GP level, even in the absence of true decreasing returns. This estimated bias equals about 80% of the effect calculated in both the sample buyout data and the sample venture capital data. These estimates are robust to reasonable variations in modeling choices. If I do not adjust for selection bias, the estimates imply that a 150% increase in fund size is associated with a decline in internal rate of return (IRR) of 2.7 percentage points (p.p.) for buyout funds and of 5.4 p.p. for venture capital. Adjusting for the bias, the same increase in size leads to a decrease in IRR of only 0.5 p.p. for buyout and of 0.9 p.p. for venture capital. Moreover, the bias-adjusted estimates are not statistically different from zero.

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 to bias not only ordinary least squares (OLS) estimators, but also the opinions 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 LPs have experienced disappointing results when investing in high-growth follow-on funds of successful partnerships, especially if their expectations were based on those partnerships' 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. Growth does not appear to be a variable of first-order importance when allocating capital across private equity funds³, especially buyout funds, and only extreme fund growth over many subsequent funds is likely to be detrimental.

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 partnerships 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 (2015) 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 (2014) 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 some evidence of persistence exists 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 and increasing fees. We know that the latter does not take place in most cases, and Hochberg et al. (2013) 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).

I propose a new way to assess the impact of fund growth on fund return persistence. In order to do so, I compare the cross-sectional distribution of partnership skill estimated in a variance-decomposition framework before and after controlling for decreasing returns to scale. The difference between the two distributions is minimal, suggesting that, historically, returns to scale have eroded only a small amount of performance persistence. This analysis can be interpreted as a direct test of the returns to scale mechanism predicted by the model of Berk and Green (2004). To the best of my knowledge, this is the first paper to perform this kind of test, which could be applied in other contexts, especially mutual funds and hedge funds.

This paper's findings have strong implications for our understanding of the private equity industry. Whenever partnership-level returns to scale are mentioned in academic research articles about private

³ However, it is likely important when allocating across asset classes.

equity, either they are found to be negative or they are assumed to be negative. For example, Kaplan and Lerner (2010) survey the venture capital literature and state that “fund size is the enemy of persistence”.⁴ This study is the first to show that GP-level diminishing returns to scale estimated from panel data are severely biased and that, accounting for the bias, the estimated effects are small and statistically insignificant.

There is an important nuance to the interpretation of these results. 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 diminishing returns to scale. However, many successful venture capital partnerships 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, venture capital fund growth is less sensitive to performance and lower on average than it is for buyout. In light of these facts and the evidence presented, this article’s findings are consistent with an equilibrium where the follow-on funds of successful GPs 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 used. Buyout and venture capital partnerships could potentially face larger diminishing 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 partnership (especially buyout partnerships) to outperform other partnerships 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 (2014) shows that private equity funds raised in vintage years with particularly strong capital raising at the industry level perform significantly worse than other funds. I present results suggesting that this phenomenon is unlikely to be driven by partnership-level effects (e.g., dilution of senior partners’ value-creation ability). Instead, its causes are likely to be found in industry-wide effects, for example increased competition for deals (e.g.,

⁴ The full citation is “Third and last, fund size is the enemy of persistence. GPs with funds that have produced high returns tend get larger (while GPs of funds with poor returns either get smaller or are unable to raise additional funds). Kaplan and Schoar find that, for funds raised by the same GP, a 50% increase in fund size is associated with roughly a 0.07 decline in PME, which translates into a 1.5% to 2% decline in a fund’s IRR”.

Gompers and Lerner 2000), entry of less-skilled partnerships, or credit market conditions (e.g., Axelson, Jenkinson, Stromberg and Weisbach 2013).

The findings presented in this paper complement the 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 partnership-level decreasing returns. Overall, in private equity, the evidence seems consistent with the industry-level decreasing returns to scale hypothesis of Pastor and Stambaugh (2012). In the mutual fund literature, there is evidence of industry-level decreasing returns (e.g., Pastor, Stambaugh and Taylor 2015), while fund-level returns to scale results are mixed (e.g., Chen, Hong, Huang and Kubik 2004, Pastor et al. 2015, Harvey and Liu 2017).

The rest of the paper is organized as follows. Section 2 describes the data used in this study, which is standard fund-level data provided by Preqin and used in several academic articles. Section 3 starts by documenting the existence of a negative association between fund growth and performance in the data. Then, it argues that the presence of selection bias 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 4 proposes a test of the effects of diminishing returns on performance persistence. Section 5 relates partnership-level returns to scale to broad patterns at the private equity industry level. Section 6 concludes by discussing some implications of my findings.

2 Data

I use fund-level data for buyout and venture capital funds provided by Preqin at the beginning of 2017. Brown et al. (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 et al. (2012) when defining “current” (or “preceding”) and follow-on funds. Current funds are all funds with available performance (IRR), with fund size (total capital committed to the fund) greater than \$5 million in 1990 dollars, 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 fund’s 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 fund-raising 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 partnerships have built on their success 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 a different geographic focus⁵, I treat these funds as belonging to different partnerships.

The final sample of current buyout and venture capital funds includes, respectively, 1333 and 1177 funds managed by, respectively, 582 and 536 different GPs.

For each current fund, I identify the subsequent fund raised by the same GP, if there is one. These funds make up the sample of follow-on funds. Any fund raised from the beginning of the sample period to 2017 can be a follow-on fund, as long as its size is not missing. The final sample of follow-on funds comprises 1069 buyout funds and 884 venture capital funds.

Table 1 presents summary statistics for the sample. 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. Panel A shows that buyout and venture capital GPs have managed on average slightly less than 3 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. 80% (75%) of current buyout (venture capital) funds in the sample have had a follow-on fund⁶. 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. Moreover, on average, follow-on buyout funds have grown more than follow-on venture capital funds.

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

⁵ Based on the region focus variable provided by Preqin, which includes US, Asia, Europe, etc.

⁶ These percentages are generally in line with those reported in the literature. However, they are slightly downward biased because of the definition of GP used here. For example, consider the case of a US-based partnership that, throughout the sample period, is successful in raising follow-on funds in the US, but attempts to expand into emerging markets by raising an Asia-focused fund and then subsequently fails to raise follow-on funds in Asia.

results are robust to using Total Value to Paid-In capital multiples (TVPI)⁷. Moreover, because all tests and inference in this paper control for vintage year fixed effects, all the results are most likely robust to using PME as performance measure because we know that IRRs and TVPIs multiples explain over 90% of the cross-sectional variation in PMEs (Harris, Jenkinson and Kaplan 2014).

In section 5 I focus on the average performance of the buyout and venture capital fund industry. In that setting, I also use dollar-weighted aggregate vintage year PMEs as a measure of average private equity performance. I use PMEs from 1984 to 2010 calculated using data provided by Burgiss and obtained from Table 8 of Brown et al. (2015). From the same paper (Brown et al. 2015, Table 3), I also obtain the total amount of capital committed to new buyout and venture capital funds in North America in each vintage year (this variable was provided by Private Equity Analyst).

3 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 industry⁸. 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 topic⁹, 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; LPs cannot withdraw capital from the fund before the underlying portfolio companies are liquidated¹⁰; and the GP 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.

3.1 Empirical Association between Fund Growth and Performance

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

⁷ These results are available upon request to the author.

⁸ Papers discussing the role of fund size and flows include, for example, Chevalier and Ellison (1997), Chen, Hong, Huang and Kubik (2004), Berk and Green (2004), Pollet and Wilson (2008) for mutual funds and Teo (2009), Ramadorai (2013) and Yin (2016) for hedge funds.

⁹ See for example "Inside the Mind of the Limited Partner II, Second Annual Conference, 2014 Duane Morris LP Institute Trans-Atlantic Simulcast: London-New York", which contains a section dedicated to the topic of fund growth titled "Tension Around Fund Size: Beware of Bias Toward Big".

¹⁰ In recent years, the private equity industry has witnessed the birth of a secondary market for LP stakes in private equity partners, which allows incumbent investors to sell their stakes to new investors, thus partially alleviating the illiquidity of these investments. However, these transactions are costly for incumbents; see Nadauld et al. (2016) for more evidence on secondary market transactions in private equity.

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.

The primary focus of the present article is on returns to scale at the partnership 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 Pastor 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 endogeneity¹¹.

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 of each fund and for factor exposures.

Private equity firms tend to specialize in a particular investment style (e.g., there are venture capital partnerships 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 partnership fixed

¹¹ Interestingly, Pastor et al. (2015) argue that, in the context of mutual funds, returns to scale estimates are downward biased because of the positive contemporaneous correlation between changes in fund size and unexpected fund returns, leading to a small-sample bias (Stambaugh 1999). This contemporaneous correlation does not exist for private equity funds, because fund size (total capital committed) is fixed at the inception of the fund and does not change. They do not consider the two features of the flow-performance relation that are studied in this paper, namely, (i) fund flows are related past returns, and (ii) the probability of a fund remaining in the sample is related to past returns. Pastor et al. (2015) propose solving the problem by running “a panel regression of forward-demeaned returns on forward-demeaned fund size, while instrumenting for the latter quantity by its backward-demeaned counterpart”. This methodology is clearly not feasible in private equity data, because the mean number of observations for each GP is around 3.

effects appears necessary to correctly identify the effect of scale on performance, because it allows to control for differences – both observable and unobservable - across partnerships¹².

Interestingly, before the mutual fund literature recognized the importance of fixed effects in the identification of returns to scale (Pastor et al. 2015), the seminal private equity article of Kaplan and Schoar (2005) demonstrated the empirical importance of fixed effects (even though the paper did not explicitly discuss endogeneity). First, the authors documented 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”.

I start my analysis by confirming this result. In Table 2, total fund returns (measured by the fund’s IRR) are regressed on the natural logarithm of fund size:

$$fund\ ret_{it} = a_0 + a_1 \cdot \log(fund\ size_{it}) + \mu_t + \nu_i + \varepsilon_{it} \quad (1)$$

where i is the GP index and t is the vintage year index. In accordance with the argument made above, partnership fixed effects are included in all specifications. Moreover, consistent the literature, vintage year fixed effects are included in order to control for time patterns in average returns and to partial out potential effects of returns to aggregate scale. Column 1 and 5 report results for buyout and venture capital, respectively. The results confirm the literature findings (e.g., Kaplan and Schoar (2005) and Kaplan and Lerner (2010)) that larger funds earn lower returns than smaller funds managed by the same partnership. The effect is economically significant: a one standard deviation increase in the independent variable is associated with a decline in returns (IRR) of approximately 5.5 p.p. and 8 p.p. for buyout and venture capital, respectively. The fund size coefficients are also strongly statistically significant, with t -statistics around -5. Following the literature, standard errors are robust to heteroskedasticity and clustered at the private equity firm level. Other clustering specifications produce smaller standard errors.

In column 3 and 7, the natural logarithm of one plus cumulated fund growth is used as the covariate of interest. This variable is defined as the size of the focal fund divided by the size of the first fund managed by the same partnership, minus one. The regression specification is as follows:

$$fund\ ret_{it} = b_0 + b_1 \cdot \log(1 + cumulated\ fund\ growth_{it}) + \mu_t + \nu_i + \varepsilon_{it} \quad (2)$$

¹² As discussed in the data section, if a partnership raises a fund with a different investment style or geographic focus than its other funds, I consider that fund as part of a separate partnership.

Notice that, when GP fixed effects are included, specification (1), which has fund size as the independent variable, is mathematically identical to specification (2). However, the coefficients reported in the table are not exactly the same because the independent variables are winsorized at the 1% and 99% levels. In the next sections of this article, I will analyze a potential bias in this kind of returns to scale regression. Because modelling fund growth is technically more straightforward than modelling fund size directly, and given that specifications (1) and (2) are identical, I will use (2) as the main returns to scale specification in the rest of the paper.

The odd-numbered columns of Table 2 report results that control for the sequence number of each fund. The size and cumulated fund growth coefficients are essentially unchanged, and the sequence number variable is insignificant. Therefore, we can rule out the hypothesis that the negative association between changes in fund size and performance may be mechanically driven by follow-on funds always delivering lower returns than their preceding funds, regardless of growth. The results presented are robust to winsorizing or trimming the most extreme observations at any reasonable level, and to using the TVPI multiple as a measure of performance.

Panel A of Figure 1 shows the relationship between two consecutive funds managed by the same partnership conditional on fund growth. Within each vintage year, all follow-on funds are sorted into three groups depending on their growth relative to the preceding fund. The figure shows that, for both buyout and venture capital, follow-on funds that grow the most tend to drastically underperform their preceding funds, while funds that grow less do not. This kind of analysis is simple and intuitive, and it is likely that many institutional advisors and investors have produced similar graphs in their analysis of historical private equity fund data.

Panel B of Figure 1 graphically illustrate the relationship between returns and fund growth within a GP using a more sophisticated analysis based on panel regression specification (2). Compared to sorting funds into groups, panel regressions allow us to exploit the full history of funds managed by each partnership and are less influenced by outliers. The figure reports both linear and non-parametric estimates and suggests that the functional form in specification (1) and (2) (i.e., regressing returns on the natural logarithm of size or of cumulated growth) is appropriate for both buyout and venture capital funds. Using the test proposed by Härdle and Mammen (1993), the null hypothesis that the linear and the non-parametric fits are not different cannot be rejected. Overall, panel regression results are consistent with the simple analysis based on sorting funds into different groups based on their growth in size as presented in Panel A of Figure 1.

Given that in the data there exists such a strong relationship between fund growth and performance, it is not surprising that many academics and practitioners believe that private equity partnership face large

decreasing returns to scale. In the next section, I explain why interpreting the data in this way may be incorrect.

3.2 Selection Mechanism and Econometric Bias

There are reasons to believe that selection bias is the primary driver of the negative association between fund growth and performance that we observe in the data. Ergo, this association should not be interpreted as definitive causal evidence of decreasing returns.

The combination of three strong empirical facts regarding private equity funds contribute to generating spurious evidence of diminishing returns. First, the probability that a partnership raises a follow-on fund is positively related to the return of preceding fund managed by the same partnership. 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 the funds of partnerships 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 noise 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 GP. In other words, the funds that have had follow-ons will on average have had positive luck, especially those whose follow-ons have grown the most. Since there is no reason why that luck will continue, the follow-on funds raised by lucky GPs should on average underperform their previous funds. In analyses such as those carried out in Figure 1 and Table 2, this effect will lead to finding spurious evidence of diminishing returns to scale.

3.3 How Current Fund Performance Influences Future Fund-Raising

Buyout and venture capital investments are carried out through 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 partnership to raise follow-on funds. They usually attempt to do so after three to five years following the inception of the current fund, and advertise the performance of the current fund 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 funds¹³. Consistent with this prediction, Kaplan and Schoar (2005)

¹³ 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 of the strong-performing funds. In Private Equity, poor performance in the current fund decreases the probability that a partnership will raise a follow-on fund, while strong performance leads to larger follow-on funds.

study a large dataset of PE funds and 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 fund-raising is of the same order of magnitude as direct pay for performance from carried interest”.

The left half of Table 3 presents regression results from probability models for the probability of raising a follow-on fund. In the simulations presented in section 3.5, this relationship is modeled using a probit specification. Here I present results from a linear probability model because the inference is identical, but the interpretation of the interaction coefficients is more intuitive. For each current (focal) fund the dependent variable is an indicator which takes the value of one if the partnership raised a follow-on fund and zero otherwise, and the independent variables are the return on the current fund, the sequence number of the current fund, and the interaction of these two variables. I also include the return on the previous fund, which seems to help predicting the probability of raising a follow-on fund even when controlling for the current fund return. In this specification, the lagged return for sequence-one funds does not exist and thus it is set to zero, and so I add an indicator for sequence-one funds to capture the difference in the conditional mean for these funds. The estimated equation is:

$$I_{has\ followon,it} = c_0 + c_1 \cdot fund\ ret_{it} + c_2 \cdot \log(seq\#_{it}) + c_3 \cdot \log(seq\#_{it}) \cdot fund\ ret_{it} + c_4 \cdot fund\ ret_{it-1} + c_5 \cdot I(seq\# = 1) + \mu_t + \varepsilon_{it}. \quad (3)$$

The estimates in columns 1 through 4 confirm prior literature findings. Current fund returns strongly predict the likelihood that buyout and venture capital partnerships will raise another fund. Moreover, a fund’s 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. These findings are consistent with the idea that investors learn about a partnership’s ability over time (see Chung et al. 2012).

The right half of Table 3 presents evidence on the relationship between the performance of the current fund and the size of the follow-on fund conditional on the follow-on fund being raised. The model uses the same independent variables as in model (3), specifically:

$$\log(1 + followon\ growth_{it}) = d_0 + d_1 \cdot fund\ ret_{it} + d_2 \cdot \log(seq\#_{it}) + d_3 \cdot \log(seq\#_{it}) \cdot fund\ ret_{it} + d_4 \cdot fund\ ret_{it-1} + d_5 \cdot I(seq\# = 1) + \mu_t + \varepsilon_{it}. \quad (4)$$

The results indicate that future fund growth is strongly related to current returns. A partnership’s average fund growth and its sensitivity to returns tends to decrease as the partnership matures; the latter result is statistically significant for venture capital funds, but not for buyout funds. In addition to the current fund return, the preceding fund return seems to help to predict the size growth of the follow-on fund.

Based on the previous literature and on the evidence presented above, it is clear that a statistically robust fundraising-performance relationship exists in private equity. However, what does this relationship look like in practice? Table 4 is designed to help us visualize it. In this table, all the first-time¹⁴ funds raised by the buyout (Panel A) and venture capital (Panel B) partnerships in the sample are first sorted into funds that had no follow-on (column 1) and funds that did have a follow-on (column 2). Out of 582 buyout partnerships in the sample, 124 (i.e., 21%) failed to raise a follow-on fund after their first-time fund. The reason why they failed to do so is clear when we compare the mean IRR in columns 1 and 2 of Panel A: while the funds of partnership that managed to raise a follow-on fund delivered average returns of 19.7%, the funds that had no follow-on returned only 5.4% on average. Evidently, the limited partners of the latter funds, as well as other prospective investors, found these fund returns to be too low and declined to reinvest with those partnerships.

In columns 3 to 5 of Panel A, the funds that did have a follow-on are sorted into three groups based on the growth of their follow-on funds. Column 2 shows that the mean growth rate was 102%; however, the sort into growth terciles reveals substantial heterogeneity in the growth rate. On average, the follow-ons of funds in the low-growth tercile managed about the same amount of money as their preceding funds (i.e., the growth rate was -4%), while the follow-ons of funds in the high-growth tercile managed 3.5 times more (i.e., the growth rate was 251%). Moving our attention to returns, preceding funds in the high-growth tercile returned 9.3% more than those in the low-growth tercile. A glance at Panel B reveals that the exact same patterns emerge for venture capital funds, the only noticeable difference being that both returns and growth rates are slightly lower here than they are for buyout funds.

3.4 Skill, Luck and Fundraising in Private Equity

An important question in asset management is how much of the variation in performance across fund managers is due to actual skill rather than luck.¹⁵ 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 (2015) 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 private equity fund returns is due to noise and how much is due to

¹⁴ This table focuses on the first fund of each partnership in the sample for simplicity and because funds with sequence number equal to one are, by construction, the largest group. Consistent with the regressions in Table 3, the same table made using funds with sequence number greater than one looks similar to Table 4 except for the fact that, as explained when interpreting the results of Table 3, the relationship between current returns and future fundraising becomes slightly weaker as the sequence number of the current fund increases.

¹⁵ For evidence on the cross-section of equity mutual fund returns see, among others, Kosowski, Timmerman, Wermers and White (2006), Fama and French (2010) and Linnainmaa (2013).

persistent differences in skill across private equity firms. To this end, they define the total log-return of fund u of firm I , y_{iu} , as $10 \cdot \log(1 + \text{IRR}_{iu})$ and model it as:

$$y_{iu} = X_{iu}'\beta + \sum_{\tau=t_{iu}}^{t_{iu}+9} (\gamma_i + \eta_{i\tau}) + \varepsilon_{iu} \quad (5)$$

where X_{iu} contains vintage year fixed effects and potentially other controls, γ_i is a GP-level random effect which is constant across all funds of the same partnership, $\eta_{i\tau}$ is a GP-time random effect that is designed to capture the covariance in returns of overlapping funds of the same partnership, and ε_{iu} is a fund-specific random shock. In this model, the mean γ_i across all the GPs 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 model by fitting it to the data via Bayesian Markov Chain Monte Carlo.¹⁶ 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_{iu} 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:

$$\sigma_y^2 = 100\sigma_\gamma^2 + 10\sigma_\eta^2 + \sigma_\varepsilon^2 \quad (6)$$

where $100\sigma_\gamma^2$ is the amount of variation due to real, long-term differences in skill across partnerships, while $10\sigma_\eta^2$ is variation due to time-varying random noise shared by overlapping funds of the same GP 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:

$$s_\gamma = \frac{100\sigma_\gamma^2}{\sigma_y^2}. \quad (7)$$

I use this model to decompose the total variation in fund returns in my sample. The model is estimated via MCMC simulations with 10,000 burn-ins and 100,000 simulations. Panel A of Table 5 reports the results. The first row 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

¹⁶ See Korteweg and Sorensen (2015) for additional information, in particular see their appendix for a description of the estimation method. See Korteweg (2013) for a discussion of the use of Markov Chain Monte Carlo methods in corporate finance. Sample code for MCMC methods is available on Arthur Korteweg's website.

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 (2015) for venture capital, but it is lower for buyout. Consistent with Harris, Jenkinson, Kaplan and Stucke (2014), in unreported tests I find that the lower estimate of buyout persistence is driven by funds raised in the most recent decade.

In this setting, the investor's learning process about a partnership's skill level γ_i can be modeled using a Gaussian learning model. The prior for a partnership'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 $\mathcal{N}(\gamma_{i,N+1}, \sigma_{\gamma_{i,N+1}}^2)$ where

$$\gamma_{i,N+1} = s_\gamma \cdot \frac{y_{i,N+1} - X'_{i,N+1}\beta}{10} + (1 - s_\gamma) \cdot \gamma_{i,N} \quad (8)$$

and

$$\sigma_{\gamma_{i,N+1}}^2 = (1 - s_\gamma) \cdot \sigma_{\gamma_{i,N}}^2 \quad (9)$$

As discussed above, the model's estimates reveal that there exists a significant amount of persistence in skill, but, at the same time, it is difficult for market participants to identify it. 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%. The return we observed is 24.5%, and we also know that the sample mean buyout return is 15.7%. Hence, the observed fund has outperformed the average fund by 8.8%. Using the Bayesian updating method illustrated above, after observing this fund return the posterior for the GP's skill is distributed as $\mathcal{N}(0.0064, 0.035^2)$. In expectation, only a small part of the large outperformance of 8.8 p.p. can be attributed to skill (0.64 p.p.) while the majority is due to luck (8.16 p.p.). Suppose instead that the return on the first fund is only 15.2%; then the updated assessment on the skill of that partnership is $\mathcal{N}(-0.0004, 0.035^2)$, that is, skill is essentially 0 and the small underperformance relative to average buyout returns is due to noise.

The two hypothetical fund returns in the above example (i.e., 24.5% and 15.2%) have been chosen to match the mean returns in the high- and low-growth terciles in Panel A of Table 4. The difference in mean returns across the two groups is 9.3 p.p. Using the calculations from the previous examples we can infer that, on average, the spread in actual skill is only 0.7 p.p. while the remaining part (8.62 p.p.) is due to noise in past realized returns. Very similar patterns can be found in venture capital funds, that is, most of the spread in mean return between the high- and low- growth terciles is driven by noise – even more so than for buyout since the signal-to-noise ratio for venture capital is smaller than that for buyout.

The argument put forth in the previous paragraph is instrumental to understanding why a simple analysis of the data is biased towards finding strong evidence of partnership-level decreasing returns to scale. On average, most of the outperformance of funds whose follow-ons grow the most is driven by noise. Since the average expected noise across follow-on funds is zero, we should expect that the funds that grow the most underperform their preceding funds by a large amount, regardless of whether actual decreasing returns to scale exist in the management of private equity investments.

3.5 Bias-adjusted Returns to Scale

The estimates in Table 2 suggest that private equity firms face large decreasing returns to scale at the partnership level. 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 diminishing returns left once we control for it?

To answer this question, we must study the magnitude of the bias. In particular, in the context of regression (2), we need to know what value coefficient b_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 b_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.

The size of the bias in panel regressions is estimated via a simulation technique. As with the previous analysis, this exercise is carried out separately for buyout and for venture capital. I start by initializing one partnership for each partnership in the data (582 buyout GPs and 536 venture capital GPs). Returns are generated accordingly to equation (5). The underlying parameters σ_γ^2 , σ_η^2 and σ_ϵ^2 are those presented in Panel A of Table 5. In particular, in order to match the non-normal distribution of fund returns, the pure noise component σ_ϵ^2 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 GP-persistence to noise and the same amount of spurious short-term persistence as in the data. A simpler, less precise version of the simulation in which fund returns are simply bootstrapped from the empirical distribution delivers similar results.

Each GP is assigned a skill level γ_i by drawing randomly from the estimated distribution of persistent differential skill $\mathcal{N}(0, \sigma_\gamma^2)$. I want to generate a panel of data where no decreasing returns to scale exist. For this reason, I draw γ_i independently of fund size or fund growth and I keep it constant across different funds of the same GP. Each simulated partnership's first-time fund has the same vintage year as

the corresponding fund in the sample data. The return of each of these funds is generated from equation (5), where the ten GP-time random effects η_i 's and the fund-specific noise term ε_i are randomly drawn from the appropriate distributions as illustrated above. Finally, $X'\beta$ is a constant that is added to the sum of γ_i , η_i and ε_i to ensure that the mean generated IRR equals the mean IRR in the sample data. For each fund return, I obtain the probability of raising a follow-on fund by fitting equation (3) using a probit model (the coefficients used are reported in the Appendix). I then determine which funds will actually have a follow-on fund by drawing randomly accordingly to the probability given by the probit model. For each GP that raises a follow-on fund, the growth in the size of the follow-on is obtained by fitting equation (4). The error term is drawn directly from a probability density function fitted to the empirical error distribution from this equation. This way, the simulated data contains the same relationship between fundraising and performance and the same amount of noise as the actual sample data. Each follow-on fund's vintage year is 2 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, I simulate a return following a procedure similar to that used for the returns of first-time funds. The only difference is that not all of the ten GP-time random effects are drawn from scratch; instead, the first m η_i are the same as the last m η_i of the preceding fund of the same GP, 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 fund's probability of raising a sequence-three follow-on fund (again using equation (3)). Using those probabilities I randomly decide which funds will have a follow-on fund and then I determine the growth of those funds using equation (4). 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 our analysis, it is important that the distribution of the simulated variables closely resembles the distributions of the variables in the sample data. Figure 2 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. This is confirmed in Panel A of Table A1 in the Appendix, which presents summary statistics for both variables in the sample data and in the simulated data. Panel B of Table A1 shows that the joint distribution of these variables is similar to that found in the data. The characteristics of the simulated data is discussed in greater detail in the Appendix. 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.

An important detail of the simulation procedure is that the fundraising-performance relationship is modeled using the final total fund return of the current 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, etc.), 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 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 diminishing returns) and is therefore a prudent research design choice.

In Table 6 I present the main results relating to the estimation of the magnitude of the bias. I run regressions of fund returns on the natural logarithm of one plus cumulated fund growth with GP fixed effects (i.e., specification (2)) using the sample data and the simulated data and then I compare the estimates of the slope b_1 . Panel A and B show results for buyout and venture capital, respectively. Column 1 reports b_1 from the actual sample data; this is the same estimate already presented in Table 2.

I also run four 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 6, the estimates presented are the average coefficient b_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 fund-raising 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, the average slope b_1 in this case is virtually 0 and statistically insignificant in 100% of the simulations 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 in the data. The slope b_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 buyout and venture capital simulations, respectively. Moreover, its magnitude is not comparable to that of the coefficient found in the real 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 b_1 .

In the third shadow simulation, the probability of raising a follow-on is random, but its growth rate depends on past performance. In this case, the slope b_1 is always negative and its mean is similar in magnitude to the coefficient estimated in the data.

Column 5 shows the results for the main simulation. Here, both the probability of raising a follow-on fund and its growth depend on past performance. As discussed, the fundraising-performance relationship is modeled using specifications (3) and (4). Despite the fact that fund returns are generated independently of fund size (i.e., true b_1 is 0), b_1 is negative in 100% of the simulations, and statistically significant at the 1% level in 99% of the cases 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. The mean value of b_1 in the main simulation is the estimate of the bias in the empirical estimator of b_1 .

The analysis of the main simulation results continues in Panel C of Table 6 and Figure 3. The figure displays the probability density function (PDF) of the b_1 estimator across the 10,000 simulation and its 1st, 5th, 10th, 90th, 95th and 99th percentiles. In other words, this is the PDF of the bias. A feature of private equity data is that there is a relatively low number of funds (1333 buyout funds and 1177 venture capital funds). This feature, together with the fact that in the simulations the error terms in the fundraising relationship and the shocks to fund returns are drawn randomly from their empirical distributions, generates the variation across simulations depicted in the figure.

The main takeaways from the simulation study are illustrated in Figure 3. First, the PDF of the simulated slope b_1 (i.e., the bias) lies entirely to the left of 0. This strongly supports the hypothesis that the OLS fixed effect estimator of b_1 is negatively biased. Second, the bias is larger in magnitude for venture capital than for buyout. This is due to the fact that venture capital returns have a greater variance, are more skewed, and contain more noise (i.e., the signal-to-noise ratio is lower) than buyout funds. Finally, the biased OLS fixed effect estimator of b_1 (represented by the solid vertical line) lies to the left of the mean of the distribution of the bias, but within it. Specifically, the empirical p -value is 0.16 for buyout and 0.18 for venture capital.

Panel C of Table 6 provides a summary of the analysis carried out. For buyout, the OLS estimate of the returns to scale slope is -5.71, the bias is -4.66, and the bias-adjusted estimate is -1.05, which is 82% smaller than the unadjusted figure. For venture capital, the OLS estimate of b_1 is -8.49, the bias is -7.08, and the bias-adjusted estimate is -1.41, which is 83% less than the unadjusted figure. Moreover, for both private equity investment types, the slope is highly statistically significant before the adjustment, but becomes insignificant after the bias has been taken into account. Hence, this analysis indicates that,

although in the data fund growth and performance are negatively related, the actual effect of fund growth is small and statistically insignificant.

After the main simulation, I run a fourth and final shadow simulation. 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 b_1 is set equal to the estimated bias-adjusted figure reported in Panel C. Specifically, true b_1 is set equal to -1.05 for buyout and -1.41 for venture capital. Given these parameters, the regression estimate of b_1 in the simulated data is extremely similar to its counterpart in the actual sample data. Specifically, for buyout the coefficient is -5.71 using the sample data and -5.81 using the data generated in the fourth shadow simulation. For venture capital, these figures are -8.49 and -8.64, respectively. Hence, the simulated coefficients are extremely similar to those found in the data; both are approximately within 1.7% of the actual coefficient 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 negative association between fund growth and performance observed in the sample buyout and venture capital data likely reflects a much weaker true effect (about 80% smaller) coupled with a complex selection bias.

As discussed, Panel C of Table 6 shows that the bias-adjusted estimate of b_1 is -1.05 for buyout and -1.41 for venture capital, and the null hypothesis that these estimates are equal to zero cannot be rejected. Lack of power does not seem to be the reason why I fail to reject. In fact, using the 10% level as a threshold, the simulation test can reject that “true” coefficients below -1.35 and -1.94 are equal to zero. Both these figures are less than 25% of the (biased) slope found in the sample data. In other words, if true decreasing returns were at least one fourth as large as the biased empirical estimate suggests, I would be able to reject the null hypothesis that they are equal to zero.

3.6 Interpreting the Results

The last two rows of Panel C show the estimated change in return (IRR) for a given change in fund size. In this discussion of the results, I will mostly focus on buyout funds (column 1 to 3); the inference is very similar for venture capital. 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 panel regressions indicate that a 150% size increase is associated with a 5.14 p.p. decrease in returns. However, after accounting for econometric bias, the actual effect of fund growth is only a statistically insignificant -0.96 p.p.

Living aside the technical details discussed in the previous section, the simple intuition behind this result is as follows. Partnerships that raise a fund that is 2.5 times larger than their previous fund are able to do so because the previous fund had an exceptional performance¹⁷. I estimate that part of that performance was due to a random return shock (in other words, luck) equal to, on average, +4.28% ($=5.24 - 0.96$). Thus, as luck reverts to zero, the return of the follow-on fund is expected to decrease by 4.28 p.p. Only the additional change in return (a decrease of 0.96 p.p.) can be attributed to the impact of scale on the ability of GPs to generate returns.

The analysis presented in this paper is carried out in the context of panel regressions with vintage year fixed effects. Hence, the estimated effect of decreasing returns (-0.96 p.p. for a size increase of 150%) is intended with respect to funds raised in the same vintage year and whose size growth is 0. Average fund growth, however, is positive in all vintage years. The median fund growth rate is 56% for buyout and 32% for venture capital. Hence, compared to a median-growth follow-on fund, the effect of a 150% size increase is only -0.50 p.p for a buyout fund ($=(\log(2.5)-\log(1.56))*-1.05$) and -0.89 p.p. for a venture capital fund ($=(\log(2.5)-\log(1.32))*-1.41$). To put these figures into perspective, the cross-sectional standard deviation of fund returns (winsorized at the 1% and 99% level) is 14.8% for buyout and 25.5% for venture capital. Based on these considerations, and on the fact that the estimated slope b_1 is not statistically significantly different from 0, we can conclude that, historically, even relatively large fund growth rates (e.g., +150%) have had only a marginal impact on fund returns.

4 Returns to Scale and Performance Persistence

The findings presented in the previous section have important implications for our understanding of the private equity industry. In particular, they might help us to understand why there seems to be an economically significant amount of performance persistence in buyout and venture capital fund returns.

4.1 How Much Persistence Has Historically Been Eroded by Decreasing Returns?

In an influential article, Berk and Green (2004) model an economy where asset managers possess differential skills and investors rationally learn about skill by observing past fund returns and try to invest with the managers that have positive alpha. In equilibrium, they argue, every manager will have zero net-of-fees alpha. This equilibrium assumes that funds face diminishing returns to scale and obtains because skilled fund managers can appropriate the value of their abilities by increasing fees and/or the size of their

¹⁷ For simplicity, I explain the intuition behind the result as if each partnership managed only 2 funds; in reality, the estimates are based on panel regressions and therefore account for the full history of funds managed by each partnership.

funds. The lack of net-of-fees alpha and of performance persistence in mutual funds appear consistent with the model's prediction.

However, the evidence of performance persistence in private equity appears at odds with the predictions of Berk and Green (2004). Since the work of Kaplan and Schoar (2005), this fact has puzzled financial economists. Hochberg et al. (2013) explore the fees channel. In particular, since GPs almost never increase the fees they charge, they ask "why don't VCs eliminate excess demand for follow-on funds by raising fees?". They propose a model where investors in the early fund of a partnership have holdup power when the latter attempts to raise a follow-on fund. As a result, skilled venture capitalists refrain from increasing fees to the point where performance persistence is cancelled out, and fund managers and early investors share the economic rents generated by skill.

The second channel through which skilled managers can appropriate the value of their skill is fund growth. Empirical evidence confirms that past winners in the cross-section of fund returns do in fact tend to increase the size of their follow-on funds more than other managers. Assuming that private equity managers face economically large decreasing returns to scale, this fundraising-performance relationship could potentially lead the private equity industry towards an equilibrium where performance persistence disappears quickly.

I propose a way to estimate the impact of decreasing returns on the long-term persistence of private equity returns. The following example helps to understand the reasoning behind it. The example assumes that GPs and investors learn about managerial ability in the spirit of Holmstrom (1999); the setup is essentially that of the Berk and Green (2004) model, simplified and adapted to private equity. Suppose that the minimum expected return that investors consider acceptable is 15%, that partnerships face decreasing returns, and that there are two private equity partnerships, Y and Z. These GPs are endowed with an endogenous level of skill, based on which they can initially generate total fund returns (IRRs) of 20% and 15%, respectively. Investors and GPs cannot observe the skill level directly, and learn about it by observing fund returns.

Suppose that the two GPs raise a first-time fund of the same size. In the absence of random shocks, Y and Z will deliver returns of 20% and 15%, respectively. Following the logic of Berk and Green (2004), in competitive rational markets the following should take place. GP Z raises a follow-on fund with the same size as the preceding one. GP Y is instead able to raise a larger fund; investors commit capital to it until the expected net present value of the investment is 0, that is, until the expected return has decreased to 15% and committing capital to the fund has become zero net-present value investment. Assuming, for example, that each time a fund doubles in size the expected return of that fund decreases by 5%, and that this is

known by GPs and investors, the second fund raised by partnership Z would be two times larger than the preceding one, and therefore its expected return would be 15%.

In this example, the private equity industry quickly reaches an equilibrium where there is no performance persistence, even though differential skill exists before decreasing returns. Then, it would be impossible for the econometrician to detect any performance persistence across funds managed by the same GP. Contrary to this prediction, the literature has reported empirical evidence suggesting that persistence exists across subsequent funds (e.g., Kaplan and Schoar 2005) and across all the funds managed by a given partnership (Korteweg and Sorensen 2015).

In the example above, the realized return of the second fund raised by GP Y (the skilled one) was 15%, but it would have been 20% before the effects of scale. If we used the latter return instead of the former, we would find evidence of persistence in the private equity economy described in the example. Arguably, the fund returns we observe in private equity data (e.g., in databases such as Preqin or Burgiss) are the returns that have been realized after scale effects. This raises an interesting counterfactual question. How much persistence would we be able to find in the data, if fund growth had no effect on returns? I propose a test to answer this question.

Recall the Bayesian variance decomposition model proposed by Korteweg and Sorensen (2015) and illustrated in section 3.4. Using that model, we estimated the extent to which private equity fund returns reflect long-term differential skill across partnerships, after controlling for luck and short-term persistence. In the sample used in this paper (funds with vintage years from 1969 up to 2011), the standard deviation of the cross-sectional distribution of differential skill, σ_γ , was found to be 3.8% for buyout and 5.4% for venture capital. These estimates are based on returns realized after scale effects. How much more persistence would there have been, absent decreasing returns to scale?

I answer this question by re-estimating σ_γ after adding back the effects of fund growth on performance. In order to do so, I control for the natural logarithm of one plus cumulated fund growth within the model itself:

$$y_{iu} = X_{iu}'\beta + \sum_{\tau=t_{iu}}^{t_{iu}+9} (\gamma_i + \eta_{i\tau}) + \varepsilon_{iu} \quad (5)$$

The growth variable is added to the model as a covariate in the X matrix. Because this is a hierarchical, mixed effects model (i.e., it contains vintage year fixed effects and GP random effects), in order to ensure consistency with the preceding analyses I demean the fund growth variable at the partnership level. Consistently with the choices made for the other parameters, the prior for the growth coefficient is diffuse, so that the results reflect the data rather than prior beliefs. Specifically, similarly to the other β s, the prior

is normally distributed and lies between -3.1 and 3.1 with 99% probability. The model is again estimated via Markov Chain Monte Carlo method, with 10,000 burn-ins and 100,000 simulations.

Panel B of Table 5 presents the estimates from this model. As expected, the coefficient for the cumulated fund growth variable has a negative sign for both buyout and venture capital, and it is about 10 times and 4 times larger than the corresponding Bayesian standard error, respectively. Because the dependent variable is the total log return of the fund, $10 \cdot \log(1 + \text{IRR}_{it})$, and because the error terms are explicitly modeled as random shocks from a Gaussian-mixture distribution, the coefficient is not directly comparable to that obtained from OLS regressions. In unreported tests, I verify that this coefficient is negatively biased, for the same reasons discussed in section 3.2.

In order to understand the impact that decreasing returns to scale have had on performance persistence over the sample, we should focus on the estimate of the standard deviation of differential skill σ_γ . Compared to Panel B, in Panel A σ_γ is 13% smaller for buyout and 4% smaller for venture capital. Notice that these percentages should be considered as an upper bound, because the effect of fund growth on returns that is added back in Panel B is overestimated due to the bias described in section 3.2. Therefore, this test suggests that, historically, returns to scale at the partnership level have had little impact on performance persistence in private equity. These findings are also consistent with the results presented in section 3.5, i.e., that the true effect of fund growth on performance in the data is likely negative, but small and statistically insignificant.

The analysis carried out above can only identify the effect of GP-level fund growth on persistence. Clearly, other factors are likely to influence the evolution of performance in the private equity industry. For example, using deal-level data, Braun, Jenkinson and Stoff (2015) present evidence suggesting that industry-level competition destroys persistence. I focus on returns to scale at the partnership level, and I control for vintage year effects in all specification. Therefore, the evidence presented in this paper is orthogonal to Braun et al. (2015)'s analysis, and complementary to it.

Notice that the analysis described above can be considered as a test of the model of Berk and Green (2004). To the best of my knowledge, this paper is the first to test the decreasing return mechanism predicted by that model in a unified framework. Several papers in the asset management literature, especially those studying mutual funds and hedge funds, assert that they find evidence consistent with the model of Berk and Green (2004), in particular when they find that fund flows respond to past performance, that fund size is negatively related to returns, or that there is lack of performance persistence. These pieces of evidence are suggestive, but they do not represent a real test of the mechanisms and equilibrium predicted by the model. A proper test of the model should be carried out in a unified setting. In particular, after confirming

that there exists a positive flow-performance relationship, the following two conditions should be verified¹⁸: (i) the distribution of manager-level long-term differential skill is narrow and tends to zero when using net realized returns, and (2) the same distribution is significantly wider once the manager-level effects of fund growth are added back to returns¹⁹. According to the model, performance persistence and differential skill across managers should exist at least ex-ante, before fund flows and decreasing returns have eroded them. Unless this condition is verified the data, it is not possible to claim that the investors’ return-chasing behavior can be reconciled with rationality, which, in my view, is the key insight of the Berk and Green (2004) model. However, this point is usually ignored in the literature.

Based on this argument, the evidence presented in Table 5 is actually consistent with the ‘rationality’ hypothesis of the Berk and Green (2004) model in the sense that, in private equity, the investors’ return-chasing behavior does not appear irrational, because long-term differences in expected returns across partnerships do seem to exist in private equity. However, as noted in the literature, capital flows do not appear to have completely destroyed persistence in private equity returns, especially for venture capital funds. The next section investigates plausible reasons why this is the case.

4.2 Fundraising, speed of learning, and the erosion of performance persistence

In the previous section, we estimated the impact of returns to scale on performance persistence by comparing the distribution of partnership random effects (in long-term returns) in a variance decomposition model before and after controlling for cumulated fund growth. The results presented suggest that, despite the fact that private equity investors tend to commit more capital with past winners, decreasing returns have eroded only a small fraction of performance persistence that initially existed.

I perform a simple test to check whether this result is plausible. The test is carried out via a simulation. I initialize one million buyout GPs and one million venture capital GPs. The size of the first fund managed by each GP is random. Similarly to the simulation study presented in section 3.5, I assign to each partnership a level of differential skill γ_i randomly drawn from its estimated empirical distribution $\mathcal{N}(0, \sigma_\gamma^2)$, where σ_γ^2 is 0.038 for buyout and 0.055 for venture capital (Table 5, Panel A). Then, using equation (5), I draw returns for the first-time fund of each GP. Based on those returns, I use equation (3) to estimate the probability that each GP will raise a follow-on fund, and randomly decide whether the follow-on fund will exist using that probability. The size of the follow-on funds is determined using equation (4).

¹⁸ I am focusing on the returns to scale mechanism, and intentionally ignoring the fees channel, because fee changes are rare and usually economically inconsequential. Clearly, researchers can also control for fee changes when carrying out this test.

¹⁹ In the light of the econometric bias described in section 3.2, which likely applies to mutual funds and hedge funds just as much as to private equity, researchers attempting to verify the two conditions listed above should keep in mind that when they add back the effect of fund growth to realized returns, part of what they are adding back is luck.

The return of these funds will depend on the level of skill γ_i initially assigned, but also on the cumulated growth of the fund. Performance is adjusted for fund growth using the bias-adjusted coefficient b_1 from Panel C of Table 6, that is, -1.05 for buyout and -1.41 for venture capital. More than 90% of the partnerships in the data managed 6 or less funds, so this procedure is repeated to generate up to 6 funds for each simulated GP.

The results of the simulation test are presented in Table 7. All the GPs are sorted into four groups based on their initial level of skill γ_i . In order to assess the impact of scale on performance persistence over time, I focus on the evolution of the average fund return and size of top and bottom skill quartiles. By construction, the first-time funds managed by top quartile GPs earn, on average, higher returns, and so, on average, their follow-on funds grow more. Column 2 and 3 show the ratio of the size of top quartile GPs to the size of bottom quartile GPs for buyout and venture capital, respectively. As expected, the ratio increases with the fund sequence number. Empirically, future fundraising is positively related to current and past fund performance, and that this relationship is stronger for buyout funds (see Table 3). For this reason, the ratio should grow faster for buyout funds. However, the standard deviation of skill is greater for venture capital than it is for buyout, and so the spread between top and bottom GP performance is greater in venture capital. The two effects tend to offset each other, and overall the ratio grows at approximately the same speed for the two private equity investment styles.

Next, let us focus on the evolution of performance. In this simulation exercise, I define performance persistence as the difference between the average return generated by the funds of top and of bottom quartile GPs. Importantly, GPs are sorted on the initial level of skill assigned, not on realized returns. By construction, performance persistence is at its maximum across first-time funds. Over time, more successful partnerships tend to raise larger funds, and so, under the assumption that there are decreasing returns to scale, performance persistence will decline. Column 4 and 5 of Table 7 display the portion of the initial amount of persistence that is eroded by diminishing returns. As expected, persistence declines as the sequence number increases. However, only a small amount of the initial persistence is destroyed every time the GPs raise a new cross-section of follow-on funds. For buyout funds, every time a new cohort of funds is raised, persistence declines by about 1.4%. The pattern is slightly weaker for venture capital, despite the fact that the decreasing returns to scale slope b_1 is steeper for these funds than for buyout funds (-1.41 and -1.05, respectively). The reason is that top-quartile buyout funds grow faster than venture capital funds.

Overall, the simulation results presented in Table 7 are consistent with the test presented in Panel B of Table 5 and discussed in section 4.1. The results of the latter test suggest that, historically, decreasing returns have eroded somewhere between 0% and 13% of buyout persistence and between 0% and 4% of venture capital persistence. These estimates seem in accordance with the figures presented in column 4 and

5 of Table 7, especially when considering that the vast majority of GPs have managed at most 6 funds and that the mean number of funds managed is slightly below 3. Moreover, this simulation test seem consistent with the idea that decreasing returns have had a greater impact on buyout than on venture capital.

The analysis carried out above suggests that the estimates obtained in section 4.1 regarding the impact of returns to scale on persistence are plausible. A key question remains. Given that I have simulated a Berk and Green (2004)-like economy where skilled partnerships deliver higher returns, raise larger funds, and face decreasing returns, why is it the case that the zero-persistence equilibrium is not reached? Notice that, in the simulation, I have assumed that the relation between fund growth and past performance is the same as observed in the data. The fundraising-performance regressions estimated in Table 3 indicate that, overall, investors seem to be learning about GP skill in a way that is not inconsistent with Bayesian learning. The simulation results are actually consistent with the predictions of the Berk and Green (2004) model, except for the fact that the decreasing returns channel seems very slow in destroying persistence.

There are two reasons why this is the case. First, fund returns are very noisy, and so top-skilled GPs takes a long time, on average, before they are able to raise significantly larger funds. According to the empirical relation between fund growth and performance, it takes 6 fundraising cycles before the average top skill-quartile GPs can raise a fund that is twice as large as that of a bottom-quartile fund. To see why that is the case, consider that, because fund returns contain so much noise, the chance that a fund managed by a top skill-quartile buyout GP will deliver top-quartile realized returns is only about 40%, and there is a 13% chance that a bottom skill-quartile GP will manage a fund that delivers a top-quartile return entirely by luck. Therefore, consistent with arguments made in Korteweg and Sorensen (2015), in private equity, learning about skill by observing fund returns is a very slow and uncertain process.

Second, as demonstrated in this paper, decreasing returns to scale are actually relatively small. Therefore, it is plausible that the decreasing returns to scale mechanism has historically been weaker than we might have expected, and therefore it has failed to destroy a significant amount of performance persistence in private equity. The point can be illustrated with a numerical example. Suppose that a buyout partnership initially has a level of differential skill that allows it to generate an IRR that is 1.00 p.p. higher than that of the average fund. According to the empirical relation between fund growth and returns (Table 3, column 5), the follow-on fund of this partnership will grow 1.23% more than the follow-ons of average-performing partnerships. This excess growth will lead to an expected decrease in IRR of approximately 0.01 p.p. ($=1.23\% \times 1.05$), which only reduces the initial level of skill by 1% (from 1.00 p.p. to 0.99 p.p.). Decreasing returns would need to be several times stronger in order to quickly erode differences in expected returns across partnerships.

5 How do partnership-level returns to scale impact the performance of the private equity fund industry?

This article argues that the growth in size of successive funds managed by the same partnership has only a marginal effect on the ability of that partnership to outperform the funds of partnerships that experience slower fund growth. However, this does not imply that the size of the buyout or venture capital fund industry is unrelated to its average performance. In fact, empirical evidence suggests that funds raised in vintage years in which aggregate capital commitments to private equity are particularly high earn significantly lower returns (e.g., Harris, Jenkinson, and Kaplan, 2014).

In Panel A of Table 8, I present evidence consistent with the early findings in the literature. This section focusses on aggregate private equity industry size and performance; therefore, only funds that operate in North America are included in the analysis²⁰. I use three measures of vintage year performance: equal-weighted IRRs, cap-weighted IRRs (i.e., the IRR of each fund is weighted by the amount of capital committed to that fund), and cap-weighted PME. Following the literature (e.g., Kaplan and Stromberg, 2009), vintage year industry size is defined as the total amount of capital commitments to all funds raised in that year, scaled by the total value of the US stock market at the beginning of that year. Controlling for a time trend, all measures of average performance are negatively associated with capital commitments for both the buyout and venture capital fund industry. The relationship is statistically significant in all specifications, except for buyout funds when performance is measured using the equal-weighted IRR. The relation appear economically large. For example, a one standard deviation increase in buyout commitments is associated with a decline of 2.6 p.p. in cap-weighted IRR. This relationship is even stronger for venture capital funds.

These patterns are clearly important and potentially troubling for the large institutions that invest in these funds. In particular, the fact that cap-weighted PMEs are so strongly negatively associated with capital commitments is especially relevant for investors' allocation choices across asset classes. This result mirrors, to some extent, the findings of Dichev and Yu (2011), who show that the dollar-weighted returns earned by investors in hedge funds are significantly lower than the buy-and-hold returns generated by those same funds.

What are the sources of these strong time-series patterns? Several factors may contribute to generating them. First, a simple explanation might be that investors (knowingly, or, more likely, unintentionally) commit more capital to private equity funds at times when expected equity returns are low. However, PMEs control for realized equity returns and therefore this cannot be the sole explanation.

²⁰ For example, a Europe-focusses buyout fund is unlikely to compete for deals with US-focusses funds. In general, North American funds represent the majority of funds in Preqin.

Second, exogenous industry-wide conditions might play a role. In particular, credit-market conditions are likely an important determinant of buyout fund returns. Consistent with this narrative, Axelson et al. (2013) find that the performance of buyout funds is positively related to the high-yield bond spread over the LIBOR. Finally, capital commitments might be causally related to average performance through industry- or partnership-level channels. Potential causal industry-wide effects might arise if increased competition for deals (i) pushes investment prices up at the time of entry and (ii) forces funds to deploy capital in relatively less attractive deals. Finally, there may be causal effects at the partnership level such as those that have been discussed in the previous sections of this article (e.g., limited scalability of the human capital of private equity firms).

Disentangling and testing these stories in a unified framework is arguably a difficult task²¹. Panel B and Panel C of Table 8 show that some of the vintage-year variables (a time trend, that is, the vintage year itself; aggregate capital commitments; the number of funds raised; and the average fund growth at the partnership level) that might reasonably be used to test or proxy for various effect are all negatively associated with average performance. Moreover, these variables are also correlated with each other.

Due to these issues, I focus solely on the effects of partnership-level returns to scale as defined in this article. In particular, I aim to contribute to the literature on aggregate private equity performance by isolating and quantifying the effect that fund growth over the life of a partnership has on aggregate performance.

For each follow-on fund in each vintage year, I estimate the expected effect of growth on IRR performance by multiplying the natural logarithm of one plus the growth of the fund relative to the first fund managed by the same partnership (i.e., $\log(1 + \text{cumulated fund growth})$) by the appropriate bias-adjusted coefficient b_1 reported in Panel C Table 6. This effect is also translated into a PME effect. To do so, I follow the suggestion of Harris, Jenkinson and Kaplan (2014) and use the coefficient from a univariate regression of buyout and venture capital fund PMEs on fund IRRs with vintage year controls (the coefficients are obtained from Table IA.IV of Harris, Jenkinson and Kaplan 2014). Then I take the equal-weighted and the cap-weighted average of these effect at the vintage year level. I call this the vintage year average effect of decreasing returns to scale on performance; this is labelled “effect of DRS on performance” in Panel D of Table 8.

As discussed above, because many variables that could be used to test different drivers of the negative relation between industry size and performance are highly correlated with each other (see Panel B

²¹ For example, Harris, Jenkinson and Kaplan (2014) do not attempt a causal interpretation of their results showing that performance is negatively related to capital commitments, and Kaplan and Stromberg (2009) define their regression evidence “illustrative”.

and C of Table 8), regression methods are unlikely to help us identify any causal relationship. However, I show that a simple analysis of the time series variation in vintage year performance and average effects of decreasing returns to scale can be useful to understand the underlying patterns.

This analysis is presented in Panel D of Table 8. This panel shows summary statistics for vintage year performance and vintage year average effects of decreasing returns to scale on performance for both buyout and venture capital funds. Because funds on average grow over time, the mean effect of DRS on performance is negative across all performance measures. All performance measures are positively correlated with DRS effects. This is not surprising, since we have seen that performance is inversely related to aggregate capital commitments. Moreover, the correlation is statistically significant in three out of six cases.

At a first glance, this might suggest that DRS play a big role in the broad patterns of vintage year performance. However, the ratio of the standard deviation of DRS effects to the standard deviation of performance is small (4% to 14%). Moreover, the higher the ratio, the lower the correlation. Hence, DRS effects are unlikely to be a major driver of time patterns in aggregate performance. The same conclusion can be drawn from Figure 4. This figure plots vintage year performance (IRRs and PME) and estimated DRS effects for buyout and venture capital funds over the sample period. The patterns displayed in Figure 5 strongly suggest that, when compared to the overall variation in performance over time, partnership-level returns to scale are just a side-show.

There are two reasons why partnership-level diminishing returns have only a small impact on vintage year returns. First, in a typical year, the average fund raised is not much larger than the first fund managed by the same partnership. This happens because, over the sample analyzed, new private equity firms keep entering the market with first-time funds; as a result, buyout and venture capital GPs have managed, respectively, 2.89 and 2.73 funds on average (see Table 1, Panel A). Second, as argued in section 4 of this article, diminishing returns are much smaller and statistically weaker than previously thought.

6 Concluding Remarks

Since the seminal work of Kaplan and Schoar (2005), the academic literature has either found or assumed that private equity partnerships face particularly large decreasing returns to scale. Several investors and buy-side practitioners seem to share this belief and have expressed concern regarding the growth of certain private equity funds²². This conjecture appears to be supported by a negative association between fund growth and performance found in the data.

²² See, for example, Inside the Mind of the Limited Partner II, Second Annual Conference, 2014 Duane Morris LP Institute Trans-Atlantic Simulcast: London-New York, and reports by the Kauffman Foundation and CALPERS.

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. In particular, any analysis of the effect of fund growth at the partnership level is biased towards finding strong evidence of decreasing returns. I show that, controlling for this bias, there is little evidence that buyout and venture capital fund returns have been impacted by fund growth.

These findings have several important implications for academics and practitioners alike. First, the fact that partnership-level decreasing returns to scale are much weaker than previously thought provides us with the last piece to solve the puzzle of why a certain amount of long-term performance persistence exists in private equity, especially in venture capital. A plausible explanation is that, historically, diminishing returns have not been large enough to push the industry towards a Berk and Green (2004) equilibrium. I present tests that support this hypothesis and suggest that the industry is in fact moving towards such equilibrium, but this is happening slowly because, due to the fact that private equity fund returns are infrequent and noisy, learning about GP skill is a slow and arduous process.

Second, my results have implications for GPs' fundraising choices. In fact, they already seem to be acting in a manner that is inconsistent with (biased) estimates that suggest the existence of extremely large decreasing returns. On the contrary, their behavior appears to be broadly consistent with my findings. Specifically, successful partnerships usually raise larger follow-on funds. Only the most successful partnerships, which could grow their funds a great deal more than the average GP does (especially in venture capital), put a cap to the size of their follow-on funds and turn down further capital commitments.

Third, understanding the extent to which scale impacts expected returns is critical for investors' decisions. Average buyout and venture capital vintage year performance (both absolute and relative to equity markets) is negatively related to the amount of capital raised at the industry level. 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 diminishing returns at the partnership level. I provide some estimates that suggest that, when compared to the overall variation in private equity performance over time, partnership-level returns to scale are just a side-show.

The final remark is addressed to the institutions that invest in private equity funds and to the professionals who advise them. The findings of this paper imply the following. Although the average realized performance of the follow-on funds raised by the most successful partnership has certainly been disappointing for limited partners, it should be understood that most of the disappointment is due to positive idiosyncratic shocks, or "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 (e.g., using Bayesian updating), one-standard-deviation cross-sectional differences in fund growth no longer appear to be a first-order concern, especially when allocating across buyout funds.

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, J. B., and R. 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.
- Chevalier, J., and G. Ellison, 1997, Risk Taking by Mutual Funds as a Response to Incentives, *Journal of Political Economy* 105, 1167-1200.
- Chen, J., H. Hong, M. Huang, and J. D. Kubik, 2004, Does Fund Size Erode Mutual Fund Performance? The Role of Liquidity and Organization, *American Economic Review* 94, 1276-1302.
- 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.
- Dichev, I. D., and G. Yu, 2011, Higher Risk, Lower Returns: What Hedge Fund Investors Really Earn, *Journal of Financial Economics* 100, 248-263.
- Fama, Eugene F., and Kenneth 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
- Härdle, W., and E. Mammen, 1993, Comparing Nonparametric versus Parametric Regression Fits. *Annals of Statistics* 21: 1926-1947.
- Harris, R. S., Jenkinson, T. and Kaplan, S. N., 2014, Private Equity Performance: What Do We Know?, *Journal of Finance* 69, 1851-1882.
- Harris, R. S., Jenkinson, T., Kaplan, S. N., and Stucke, R., 2014, Has Persistence Persisted in Private Equity? Evidence From Buyout and Venture Capital Funds, Working Paper, University of Virginia, University of Chicago and Oxford University.
- Harvey, C. R., and Y. Liu, 2017, Decreasing Returns to Scale, Fund Flows, and Performance, Duke I&E Research Paper No. 2017-13.
- 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.
- Holmstrom, B., 1999, Managerial Incentive Problems: A Dynamic Perspective, *Review of Economic Studies* 66, 169-182.
- Kaplan, S. N., and J. Lerner, 2010, It Ain't Broke: The Past, Present, and Future of Venture Capital, *Journal of Applied Corporate Finance* 22, 36-47.

- Kaplan, S. N., and A. Schoar, 2005, Private Equity Performance: Returns, Persistence, and Capital Flows, *Journal of Finance* 60, 1791-1823.
- Kaplan, Steven, and Per Stromberg, 2009, Leveraged Buyouts and Private Equity, *Journal of Economic Perspectives* 23, 121-146.
- Korteweg, Arthur, 2013, Markov Chain Monte Carlo Methods in Corporate Finance, in P. Damien, P. Dellaportas, N. Polson, and D. Stephens, eds.: Bayesian Theory and Applications.
- Korteweg, S., and M. Sorensen, 2015, Skill and Luck in Private Equity Performance, forthcoming, *Journal of Financial Economics*.
- Kosowski, R., A. Timmermann, R. Wermers, and H. 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.
- Mulcahy, D., B. Weeks and H. Bradley, 2012, “WE HAVE MET THE ENEMY... AND HE IS US” Lessons from Twenty Years of the Kauffman Foundation’s Investments in Venture Capital Funds and The Triumph of Hope over Experience, Working paper, Kauffman Foundation.
- Nadauld, T., B. Sensoy, K. Vorkink and M. S. Weisbach, 2016, The Liquidity Cost of Private Equity Investments, Evidence from Secondary Market Transactions, *Fisher College of Business Working Paper*.
- Pastor, L., and R. F. Stambaugh, 2012, On the Size of the Active Management Industry, *Journal of Political Economy* 120, 740–781.
- Pastor, L., R. F. Stambaugh, and L. A. Taylor, 2015, Scale and Skill in Active Management, *Journal of Financial Economics* 116, 23-45.
- Pollet, J., and M. Wilson, 2008, How Does Size Affect Mutual Fund Behavior?, *Journal of Finance* 63, 2941-2969.
- Ramadorai, T., 2013, Capacity Constraints, Investor Information, and Hedge Fund Returns, *Journal of Financial Economics* 107, 401-416
- Robinson, P. M., 1988, Root-N-consistent Semiparametric Regression, *Econometrica*, 56, 931-954.
- Stambaugh, R. F., 1999, Predictive Regressions, *Journal of Financial Economics* 54, 375–421.
- Teo, M., 2009, Does Size Matter in the Hedge Fund Industry?, Working Paper, BNP Paribas Hedge Fund Centre, Singapore Management University.
- Yin, C., 2016, The Optimal Size of Hedge Funds: Conflict between Investors and Fund Managers, *Journal of Finance* 71, 1857-1894.

Figure 1. Empirical Association Between Fund Growth and Performance at the Partnership Level

This figure displays the empirical association between changes in fund size and changes in fund growth within a partnership. In Panel A, within each vintage year, all buyout and venture capital follow-on funds are sorted into three groups based on each fund's growth rate over the preceding fund managed by the same partnership. For each of the three groups, I plot the mean change in returns (current fund IRR minus preceding fund IRR) against the mean fund growth within each growth tercile. Panel B plots returns to scale slope b_1 from regression specification (2) estimated using OLS (dashed line) and the semiparametric regression of Robinson (1988) (solid line). Härdle and Mammen's (1993) specification test is used to test the null hypothesis that the linear and the non-parametric fits are not different. The null cannot be rejected for both buyout (p -value = 0.16) and venture capital (p -value = 0.68) funds.

Figure 1.A: Funds Sorted into Growth Terciles

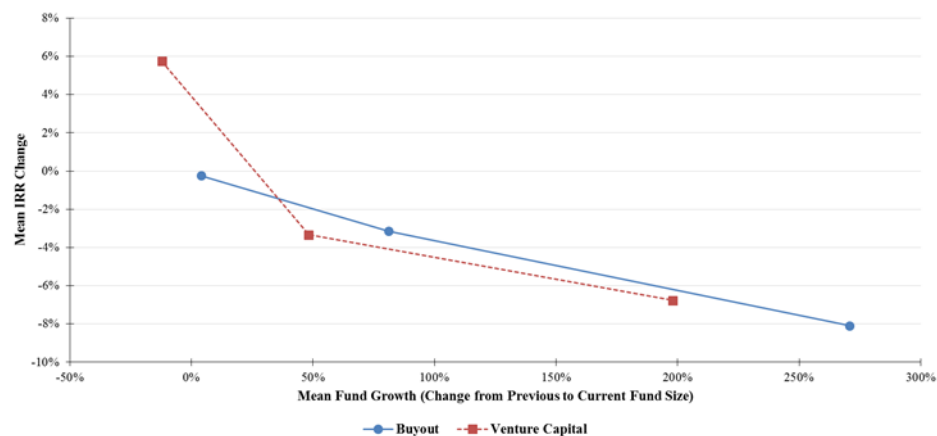


Figure 1.B: Panel Regressions

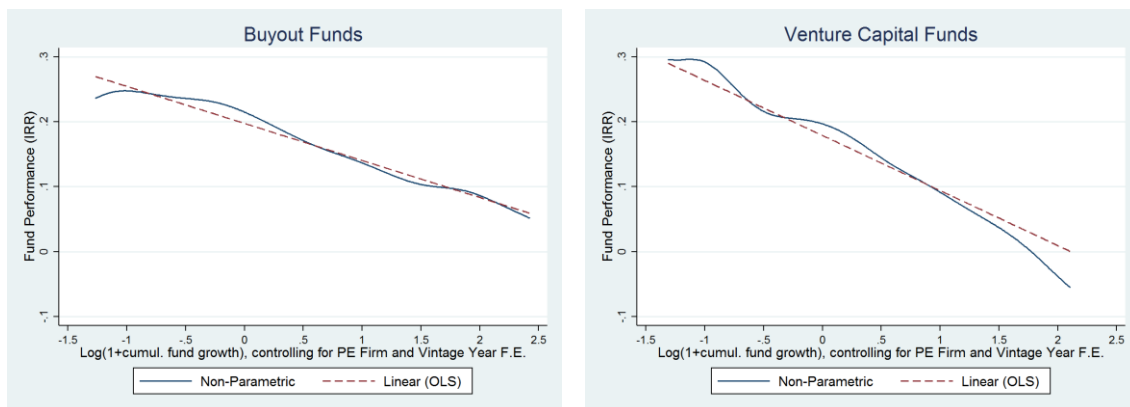


Figure 2. Distribution of Sample Variables and Simulated Variables

This figure shows the distribution of IRR and fund growth rates for the private equity panel data simulated following the procedure described in section 3.5 (main simulation). Summary statistics for these variables are presented in Panel A of Table A1. 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.

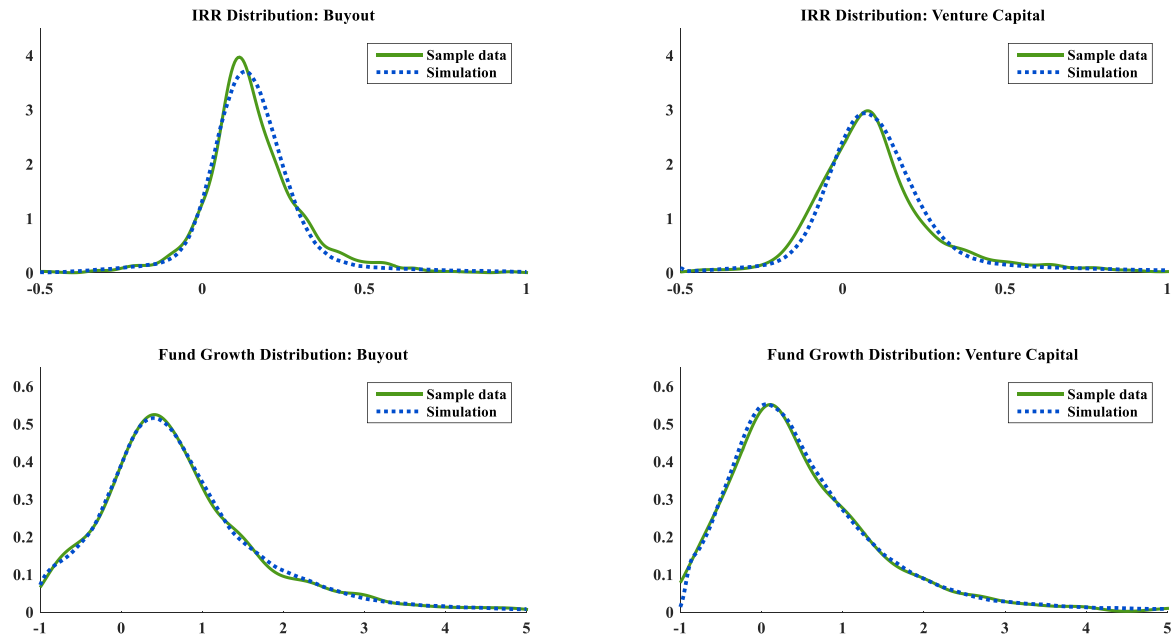


Figure 3. Bias in Returns to Scale Regressions: Simulation Results

This figure displays the results of the simulation study presented in section 3.5 (main simulation) for buyout funds (top) and venture capital funds (bottom). The dashed line is the probability density function (PDF) of the returns to scale slope b_1 from regression specification (2) across 10,000 panels of simulated data in which the true value of b_1 is set to zero. This represents an estimate of the PDF of the bias in returns to scale regression caused by the selection mechanism that generates private equity fund data. The vertical dotted lines indicate the 1st, 5th, 10th, 90th, 95th and 99th percentiles of the distribution. The solid vertical line represents the value of coefficient b_1 estimated in the sample data before controlling for econometric bias. The point-estimate of the bias and the bias-adjusted slope are presented in Table 6.

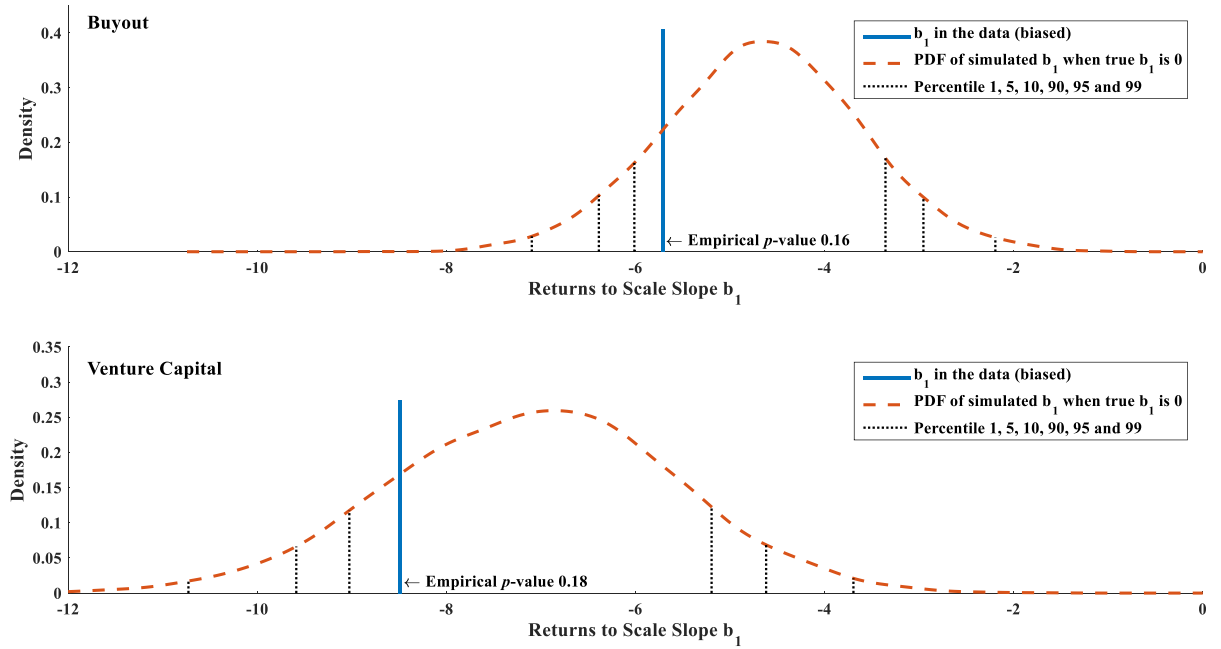


Figure 4. Vintage Year Performance and Partnership-level Returns to Scale

This figure shows average buyout and venture capital performance by vintage year and the estimated effect of partnership-level decreasing returns to scale (DRS) on vintage year performance. See Table 8 and section 5 for variable definitions.

Figure 4a: Equal-weighted Buyout IRRs




Figure 4b: Equal-weighted Venture Capital IRRs




Figure 4. Continued.

Figure 4c: Cap-weighted Buyout PME




Figure 4d: Cap-weighted Venture Capital PME




Table 1. Descriptive Statistics

This table reports summary statistics for the sample used in this paper. Panel A shows the number of private equity partnerships, 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 Private Equity Partnerships and Fund Sequence

	Number of Partnerships	Number of Funds per Partnership						
		Mean	Std. Dev.	25 th perc	Median	75 th perc	90 th perc	Max
Buyout funds	582	2.89	1.86	2	2	4	5	13
Venture Capital funds	536	2.73	2.09	1	2	3	5	16

Panel B - Variables for Fundraising-Performance and Returns to Scale Regressions: Buyout

	Observations	Mean	Std. Dev.	5 th perc	25 th perc	Median	75 th perc	95 th perc
Fund sequence number	1333	2.29	1.67	1	1	2	3	6
Fund performance (IRR)	1333	15.7	14.8	-5.3	7.7	13.7	22.4	43.9
Percentage of funds that raise a follow-on	80.2%							
Growth in follow-on fund size	1069	89%	142%	-66%	10%	56%	125%	321%
Log(fund size)	1333	6.01	1.37	3.71	5.14	5.99	6.91	8.30
Log(1+cumulated fund growth)	1333	0.71	1.03	0.00	0.00	0.13	1.22	2.92

Panel C - Variables for Fundraising-Performance and Returns to Scale Regressions: Venture Capital

	Observations	Mean	Std. Dev.	5 th perc	25 th perc	Median	75 th perc	95 th perc
Fund sequence number	1177	2.46	2.09	1	1	2	3	7
Fund performance (IRR)	1177	12.8	25.5	-14.1	-0.9	8.1	18.2	63.3
Percentage of funds that raise a follow-on	75.1%							
Growth in follow-on fund size	884	62%	122%	-68%	-9%	32%	101%	275%
Log(fund size)	1177	4.77	1.13	2.72	4.01	4.83	5.56	6.51
Log(1+cumulated fund growth)	1177	0.60	0.94	-0.11	0.00	0.00	1.10	2.56

Table 2. Panel Data Estimates of Decreasing Returns to Scale at the Partnership Level

This table reports results from regressions of fund performance (IRR) on the natural logarithm of fund size and 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 partnership). The models presented in even-numbered columns also control for the fund's sequence number. Within- R^2 statistics are reported. The t -statistics presented in parenthesis are based on standard errors clustered at the private equity partnership (GP) level and robust to heteroskedasticity. *** indicates statistical significance at the 1% level.

Dep Var: Fund IRR	Buyout				Venture Capital			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(fund size)	-5.54*** (-5.52)	-5.15*** (-4.75)			-8.07*** (-4.67)	-7.98*** (-4.50)		
Log(fund sequence number)		-3.17 (-1.25)		-2.59 (-1.01)		-0.67 (-0.22)		0.33 (0.10)
Log(1+cumulated fund growth)			-5.71*** (-5.97)	-5.36*** (-5.01)			-8.49*** (-4.63)	-8.55*** (-4.51)
Vintage Year F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
PE Firm F.E.	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	GP	GP	GP	GP	GP	GP	GP	GP
R^2	0.08	0.08	0.08	0.08	0.04	0.04	0.04	0.04
Number of Observations	1333	1333	1333	1333	1177	1177	1177	1177

Table 3. The Fundraising-Performance Relation in Private Equity: Regression Analysis

This table presents regression results that describe the relationship between current fund returns and future fund-raising. Columns 1 to 4 present estimates from linear probability models where the dependent variable is one if the focal (current) fund has a follow-on and zero otherwise. Columns 5 to 8 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. Within- R^2 statistics are reported. The t -statistics presented in parenthesis are based on standard errors clustered at the private equity partnership (GP) level and robust to heteroskedasticity. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Probability of Raising a Follow-on Fund				Growth of Follow-on Fund Size			
	Buyout		Venture Capital		Buyout		Venture Capital	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Current fund IRR	0.96*** (9.97)	0.89*** (8.86)	0.44*** (7.15)	0.46*** (6.67)	1.23*** (5.04)	1.13*** (4.42)	0.88*** (7.73)	0.69*** (5.60)
Log(fund sequence number)	0.14*** (3.64)	0.17*** (4.93)	0.12*** (3.96)	0.12*** (3.63)	-0.18** (-2.34)	-0.13* (-1.83)	-0.16*** (-3.28)	-0.14*** (-2.61)
Current fund IRR*log(fund sequence number)	-0.43*** (-3.96)	-0.33*** (-3.05)	-0.25*** (-4.56)	-0.22*** (-3.95)	-0.21 (-0.66)	-0.27 (-0.86)	-0.29*** (-2.79)	-0.34*** (-3.23)
Lag IRR	0.57*** (5.92)	0.41*** (4.71)	0.27*** (4.23)	0.15** (2.11)	0.65*** (3.05)	0.59*** (2.77)	0.49*** (5.02)	0.19* (1.72)
I(fund sequence number = 1)	0.14*** (2.92)	0.12** (2.53)	0.09* (1.90)	0.05 (1.00)	0.08 (0.86)	0.10 (1.09)	0.03 (0.36)	-0.02 (-0.24)
Vintage Year F.E.	No	Yes	No	Yes	No	Yes	No	Yes
Clustering	GP	GP	GP	GP	GP	GP	GP	GP
R^2	0.12	0.12	0.06	0.06	0.10	0.07	0.13	0.07
Number of Observations	1303	1303	1133	1133	1048	1048	847	847

Table 4. The Fundraising-Performance Relation in Private Equity: A Closer Look

This table illustrates the relationship between current fund returns and future fund-raising by sorting sequence-one funds (i.e., the first fund raised by each partnership) into groups based on whether they have a follow-on fund and how large the latter is compared to the current fund. Panel A and B present statistics calculated for buyout and venture capital, respectively. In column 1 and 2 all sequence-one funds are sorted into two groups depending on whether they have a follow-on. Funds that have a follow-on are then sorted into terciles based on the growth of the follow-on fund relative to the current fund (column 3 to 5).

Panel A: Sequence-one Buyout Funds Sorted on Their Follow-On Fund's Growth

	No Follow-On	Has Follow-On			
			Follow-on Growth Tercile		
		All	Tercile 1	Tercile 2	Tercile 3
	(1)	(2)	(3)	(4)	(5)
Number of funds	124	458	152	155	151
Mean growth of follow-on	n.a.	108%	-4%	79%	251%
Mean IRR	5.4	19.7	15.2	19.3	24.5

Panel B: Sequence-one Venture Capital Funds Sorted on their Follow-On Fund's Growth

	No Follow-On	Has Follow-On			
			Follow-on Growth Tercile		
		All	Tercile 1	Tercile 2	Tercile 3
	(1)	(2)	(3)	(4)	(5)
Number of funds	151	385	128	133	124
Mean growth of follow-on	n.a.	83%	-24%	52%	228%
Mean IRR	1.9	16.8	8.3	15.0	27.5

Table 5. Parameter Estimates for the Variance Decomposition Model

Panel A reports posterior means of parameters for the variance decomposition model of Korteweg and Sorensen (2015) 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. Bayesian standard errors for each estimate are in parenthesis. See the text and Korteweg and Sorensen (2015) for further details. In Panel B, I report estimates for the same model, but after controlling for cumulated fund growth (c.f.g.), that is, fund growth relative to the first fund managed by the same private equity partnership. See section 4.1 for more details.

Panel A: Parameter Estimates Before Controlling for Fund Growth

	σ_{γ}	σ_{η}	σ_{ϵ}	Signal-to-Noise
	(1)	(2)	(3)	(4)
Buyout	0.038 (0.004)	0.141 (0.022)	1.280 (0.076)	0.073 (0.026)
Venture Capital	0.055 (0.007)	0.142 (0.039)	2.023 (0.105)	0.066 (0.024)

Panel B: Parameter Estimates After Controlling for Fund Growth

	σ_{γ}	σ_{η}	σ_{ϵ}	Signal-to-Noise	$\log(1+c.f.g.)$
	(1)	(2)	(3)	(4)	(5)
Buyout	0.044 (0.004)	0.107 (0.018)	1.260 (0.075)	0.098 (0.033)	-0.437 (0.045)
Venture Capital	0.057 (0.007)	0.140 (0.036)	1.953 (0.102)	0.071 (0.025)	-0.284 (0.067)

Table 6. Adjusting Returns to Scale Estimates for Econometric Bias

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 panel data. The first two panels report returns to scale slope b_1 from regression specification (2) estimated in the sample data and in simulated data. Panel A and Panel B report results for buyout and venture capital funds, respectively. Section 3.5 provides details regarding the simulation study. The coefficients reported in column 1 is the same as the coefficient reported in Table 2. Columns 2 through 6 report the mean coefficient and mean t -statistic obtained from 10,000 regressions using 10,000 panels of independently simulated data. Panel C summarizes the results presented in the first two panels. In this panel, p -values are presented in parenthesis. The p -value reporter for the bias-adjusted estimate of b_1 (column 3 and 6) is an empirical p -value based on the comparison of the distribution of the simulated bias and the biased empirical coefficient. *** indicates statistical significance at the 1% level. See Figure 3 for a graphical representation of these results.

Panel A: Buyout Simulations

Regression specification: $IRR_{it} = b_0 + b_1 \cdot \log(1 + \text{cumulated fund growth}) + \mu_i + v_i + \varepsilon_{it}$ (eq. 2)						
	Data	Shadow Sim. 1	Shadow Sim. 2	Shadow Sim. 3	Main Simulation	Shadow sim. 4
	(1)	(2)	(3)	(4)	(5)	(6)
True b_1 in simulation set to:	n.a.	0	0	0	0	-1.05
Follow-on prob. depends on performance	n.a.	No	Yes	No	Yes	Yes
Follow-on growth depends on performance	n.a.	No	No	Yes	Yes	Yes
Regression estimate of b_1	-5.71*** (-5.97)	-0.01 (-0.01)	-0.10 (-0.11)	-4.31*** (-5.20)	-4.66*** (-5.31)	-5.81*** (-6.62)

Panel B: Venture Capital Simulations

Regression specification: $IRR_{it} = b_0 + b_1 \cdot \log(1 + \text{cumulated fund growth}) + \mu_i + v_i + \varepsilon_{it}$ (eq. 2)						
	Data	Shadow Sim. 1	Shadow Sim. 2	Shadow Sim. 3	Main Simulation	Shadow sim. 4
	(1)	(2)	(3)	(4)	(5)	(6)
True b_1 in simulation set to:	n.a.	0	0	0	0	-1.41
Follow-on prob. depends on performance	n.a.	No	Yes	No	Yes	Yes
Follow-on growth depends on performance	n.a.	No	No	Yes	Yes	Yes
Regression estimate of b_1	-8.49*** (-4.63)	0.00 (0.00)	-0.24 (-0.16)	-6.06*** (-4.68)	-7.08*** (-5.46)	-8.64*** (-6.66)

Panel C: Bias-adjusted Returns to Scale Estimates

	Buyout Funds			Venture Capital Funds		
	OLS Estimate (Biased)	Estimated Bias	Bias-adjusted Estimate	OLS Estimate (Biased)	Estimated Bias	Bias-adjusted Estimate
	(1)	(2)	(3)	(4)	(5)	(6)
Returns to scale slope b_1	-5.71***	-4.66***	-1.05	-8.49***	-7.08***	-1.41
p -value	(0.00)	(0.00)	(0.16)	(0.00)	(0.00)	(0.18)
Effect for a median size increase (Buyout +56%, Venture Capital +32%)	-2.53	n.a.	-0.47	-2.39	n.a.	-0.39
Effect for a 150% size increase	-5.24	n.a.	-0.96	-7.78	n.a.	-1.29

Table 7. Fundraising-Performance Relation and the Erosion of Performance Persistence

This table shows the results of a simulation in the spirit of Berk and Green (2004) model. The intuition is that investors learn about managerial skill from past returns and re-allocate capital accordingly; however, managers face decreasing returns and therefore capital flows towards high-performing funds destroy performance persistence. Consistent with the rest of the paper, buyout funds and venture capital funds are considered separately. 1,000,000 partnerships (GPs) are initialized with an initial level of differential skill drawn from its empirical distribution estimated via the variance-decomposition model proposed by Korteweg and Sorensen (2015) using the sample data and presented in Panel A of Table 5. The table focusses on the differences between GPs in the top and bottom skill quartile across six consecutive funds. Importantly, funds are sorted into quartiles based on the level of skill initially assigned as opposed to realized returns. Column 2 and 3 show how the funds raised by top quartile GPs grow over time compared to bottom quartile GPs according to the empirical fundraising-performance relation estimated using model (3) and (4) presented in Table 3. The negative impact of fund growth is calculated using the bias-adjusted estimates of returns to scale presented in Panel C of Table 6. Column 4 and 5 show the fraction of the initial performance persistence that has been eliminated by decreasing returns. See section 4.2 for more details.

Fund Sequence Number	Ratio of Size of Top Skill-Quartile GPs to Size of Bottom Skill-Quartile GPs		% of Performance Persistence (Top minus Bottom Skill-Quartile) Eroded by Decreasing Returns to Scale	
	Buyout	Venture Capital	Buyout	Venture Capital
(1)	(2)	(3)	(4)	(5)
1	1.00	1.00	0.0%	0.0%
2	1.14	1.13	2.1%	1.4%
3	1.35	1.32	3.2%	2.2%
4	1.57	1.53	5.7%	3.3%
5	1.83	1.78	6.0%	4.5%
6	2.10	2.08	7.0%	5.5%

Table 8. Industry Size, Vintage Year Performance, and Partnership-level Returns to Scale

Panel A reports results for regressions of average (equal-weighted and cap-weighted) vintage-year buyout and venture capital industry performance on the size of the industry. Industry size in a particular vintage year is measured as the total amount of capital commitments to North American buyout or venture capital funds raised in a that vintage year scaled by the value of the U.S. stock market at the beginning of that vintage year. In the regressions, this variable is standardized to a z-score. The time trend is the vintage year. Panel B (C) shows correlation coefficient for buyout (venture capital) fund variables at the vintage year level. $\log(1 + \text{cumul. fund growth})$ is the natural logarithm of one plus the growth of the fund relative to the first fund managed by the same partnership. Panel D presents summary statistics for vintage year performance and for the estimated effects of partnership-level decreasing returns to scale. The latter variable is estimated as follows. For each follow-on fund in each vintage year, the expected effect of growth on IRR performance is obtained by multiplying the natural logarithm of one plus the growth of the fund relative to the first fund managed by the same partnership (i.e., $\log(1 + \text{cumul. fund growth})$) by the appropriate bias-adjusted coefficient b_1 reported in Panel C of Table 6. This effect is translated into a PME effect using the coefficients provided by Harris, Jenkinson and Kaplan (2014). See section 5 for more details. The t -statistics presented in parenthesis in Panel A and the p -values presented in parenthesis in Panel B, C, and D are based on standard errors robust to heteroskedasticity and serial correlation. *, **, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

Panel A: Capital Commitments and Average Fund Performance

Dependent Variable	Equal-weighted IRR		Cap-weighted IRR		Cap-weighted PME	
	BO	VC	BO	VC	BO	VC
	(1)	(2)	(3)	(4)	(5)	(6)
Capital Commitments	-0.74 (-0.65)	-6.04*** (-3.88)	-2.62* (-1.73)	-7.14*** (-3.71)	-0.10** (-2.45)	-0.21*** (-2.60)
Time Trend	-0.57*** (-2.84)	-0.52** (-2.34)	-0.26 (-1.24)	-0.61** (-2.19)	0.00 (0.18)	0.00 (-0.29)
R ²	0.36	0.34	0.21	0.30	0.21	0.06
Number of Observations	27	27	27	27	27	27

Panel B: Correlation Coefficients for Aggregate Vintage Year Variables (Buyout)

	Eq.-w IRR		Cap-w IRR		Cap-w PME		Eq.-w Log(1+gr ₁)		Cap-w Log(1+gr ₁)	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Time Trend	-0.60	(0.00)	-0.40	(0.03)	-0.24	(0.19)	0.35	(0.09)	0.18	(0.45)
Capital Commitments	-0.42	(0.02)	-0.43	(0.02)	-0.46	(0.00)	0.31	(0.09)	0.40	(0.02)
Number of Funds Raised	-0.63	(0.00)	-0.52	(0.00)	-0.24	(0.18)	0.37	(0.04)	0.28	(0.16)
Log(1+cumul. fund growth)	-0.49	(0.00)	-0.28	(0.19)	-0.25	(0.22)				

Table 8. Continued.

Panel C: Correlation Coefficients for Aggregate Vintage Year Variables (Venture Capital)

	Eq.-w IRR		Cap-w IRR		Cap-w PME		Eq.-w Log(1+gr ₁)		Cap-w Log(1+gr ₁)	
	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value	Coeff	p-value
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Time Trend	-0.41	(0.00)	-0.39	(0.01)	-0.10	(0.43)	0.56	(0.00)	0.68	(0.00)
Capital Commitments	-0.51	(0.00)	-0.48	(0.00)	-0.24	(0.01)	0.78	(0.00)	0.53	(0.00)
Number of Funds Raised	-0.47	(0.00)	-0.42	(0.00)	-0.16	(0.18)	0.83	(0.00)	0.60	(0.00)
Log(1+cumul. fund growth)	-0.57	(0.00)	-0.47	(0.00)	-0.24	(0.12)				

Panel D: Estimated Effect of Partnership-level Decreasing Returns to Scale (DRS) and Vintage Year Performance

Aggregate performance measure:	Equal-w Vintage Year IRR		Cap-w Vintage Year IRR		Cap-w Vintage Year PME	
	BO	VC	BO	VC	BO	VC
	(1)	(2)	(3)	(4)	(5)	(6)
Number of Vintage Years	27	27	27	27	27	27
Mean performance	18.56	16.60	17.84	18.83	1.25	1.46
STD of performance	8.31	14.12	9.11	17.69	0.20	0.92
Mean effect of DRS on performance	-0.82	-1.16	-1.77	-2.00	-0.05	-0.07
STD of effect of DRS	0.38	0.61	1.06	1.20	0.03	0.04
Corr (performance, effect of DRS)	0.49	0.57	0.28	0.47	0.25	0.24
p-value	(0.00)	(0.00)	(0.19)	(0.00)	(0.22)	(0.12)
Ratio of the STD of effect of DRS to the STD of performance	0.05	0.04	0.12	0.07	0.14	0.05

Appendix

Table A. Distribution and Joint Distribution of Simulated Variables

Panel A shows the distribution of fund performance and follow-on fund growth in the simulated data generated for the simulation study discussed in section 3.5 (main simulation) and compares it to the sample data. The p -value for a Kolmogorov-Smirnov test comparing the distributions of these variables in the sample and in the simulated data is reported in the last row. In all cases, the test suggests that the empirical distribution and the simulated distribution do not differ significantly. Panel B shows the fundraising-performance regression coefficients estimated in the sample data and in the simulated data. The specification for the probit models is the one presented in equation (3). The specification for the fund growth models is the one presented in equation (4). The simulation coefficients are the mean across 10,000 independent simulations of private equity panel data. See section 3.5 for more details.

Panel A: Distribution of Simulated Variables

	Fund Performance (IRR)				Growth in Follow-on Fund Size			
	Buyout		Venture Capital		Buyout		Venture Capital	
	Data	Simulation	Data	Simulation	Data	Simulation	Data	Simulation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mean	15.7	15.3	12.8	12.6	89%	84%	62%	63%
Std. Dev.	14.8	16.0	25.5	22.8	142%	127%	122%	116%
5 th perc	-5.3	-5.0	-14.1	-15.5	-66%	-63%	-68%	-67%
25 th perc	7.7	7.1	-0.9	0.3	10%	10%	-9%	-12%
Median	13.7	14.0	8.1	8.8	56%	58%	32%	38%
75 th perc	22.4	21.9	18.2	20.0	125%	127%	101%	106%
95 th perc	43.8	38.8	63.2	54.6	321%	298%	275%	274%
K-S test p -value	0.41		0.25		1.00		0.32	

Panel B: Joint Distribution of Simulated Variables (Fundraising-Performance Relationship)

	Probit Model for the Probability Of Raising a Follow-on Fund				OLS Regression for $\log(1+\text{Growth of Fund Size})$ Conditional on Raising a Follow-on Fund			
	Buyout		Venture Capital		Buyout		Venture Capital	
	Data	Simulation	Data	Simulation	Data	Simulation	Data	Simulation
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Current fund IRR	4.31	4.34	1.89	1.91	1.23	1.20	0.88	0.86
Log(fund sequence number)	0.50	0.50	0.46	0.46	-0.18	-0.18	-0.16	-0.17
Current fund IRR*log(fund sequence number)	-1.71	-1.71	-1.06	-1.07	-0.21	-0.20	-0.29	-0.29
Lag IRR	0.63	0.63	0.38	0.38	0.08	0.08	0.03	0.03
I(fund sequence number = 1)	2.77	2.78	1.13	1.16	0.65	0.65	0.49	0.49
Number of Observations	1303	1303	1133	1133	1048	1048	887	887