

TECHNISCHE UNIVERSITÄT MÜNCHEN

Lehrstuhl für Entrepreneurial Finance, unterstützt durch die KfW Bankengruppe

Univ.-Prof. Dr. Dr. Ann-Kristin Achleitner

Essays on Challenges for Institutional Investors in the Maturing Private Equity Industry

Ingo Stoff

Vollständiger Abdruck der von der Fakultät für Wirtschaftswissenschaften der Technischen Universität München zur Erlangung des akademischen Grades eines Doktors der Wirtschaftswissenschaften (Dr. rer. pol.) genehmigten Dissertation.

Vorsitzender: Univ.-Prof. Dr. David Wozabal

Prüfer der Dissertation: 1. Univ.-Prof. Dr. Dr. Ann-Kristin Achleitner

2. Univ.-Prof. Dr. Christoph Kaserer

Die Dissertation wurde am 27.05.2014 bei der Technischen Universität München eingereicht und durch die Fakultät für Wirtschaftswissenschaften am 15.09.2014 angenommen.

“What I want is for them [General Partners (GPs)] to be able to pay for their administrative costs, their house, and a nice vacation with the management fee. [...] For that second and third home and that yacht, I want them to pay for that out of carry. I hope every one of our GPs gets rich, but after my members make money.”

Gary Bruebaker, Chief Investment Officer of the Washington State Investment Board, taken from “The People vs. Private Equity”, Bloomberg Businessweek, 23th, November 2011

Table of Contents - Overview

Table of Contents **IV**

List of Figures **VII**

List of Tables **VIII**

List of Abbreviations **IX**

1 Introduction **1**

 1.1 Motivation and Research Topic 1

 1.2 Development of Research Questions 8

2 Essays **20**

 2.1 Essay 1 - How Persistent is Private Equity Performance? Evidence from Deal-Level Data 20

 2.2 Essay 2 - The Cost of Private Equity 63

 2.3 Essay 3 - The Evolution of Private Equity Fund Terms Beyond 2 and 20 101

3 Conclusion **129**

References **136**

Table of Contents

- Table of Contents.....IV**
- List of Figures VII**
- List of Tables.....VIII**
- List of AbbreviationsIX**
- 1 Introduction..... 1**
 - 1.1 Motivation and Research Topic 1
 - 1.2 Development of Research Questions 8
- 2 Essays 20**
 - 2.1 Essay 1 - How Persistent is Private Equity Performance? Evidence from Deal-Level Data 20
 - 2.1.1 Introduction..... 21
 - 2.1.2 Data 31
 - 2.1.2.1 Funds and Fund Families 31
 - 2.1.2.2 Sample Selection Issues 33
 - 2.1.2.3 Deals..... 35
 - 2.1.2.4 Synthetic Funds 41
 - 2.1.3 Absolute Performance Persistence..... 44
 - 2.1.4 Relative Performance Persistence 47
 - 2.1.5 Top Quartile Performance Persistence..... 54
 - 2.1.6 Conclusion 60
 - 2.2 Essay 2 - The Cost of Private Equity 63
 - 2.2.1 Introduction..... 64
 - 2.2.2 Theoretical Background..... 69
 - 2.2.2.1 Literature Review..... 69

2.2.2.2	Fund Terms and Drivers of the Return Spread	71
2.2.2.3	GP-LP Relation and a Fund's Deployment Rate	75
2.2.3	Data Set and Descriptives	76
2.2.3.1	Data Sampling and Selection Bias	76
2.2.3.2	Summary Statistics.....	77
2.2.3.3	Return Spread Benchmark	80
2.2.3.4	Return Spreads over Time and Compared to Fund Characteristics	83
2.2.4	Analysis.....	86
2.2.4.1	How Have Return Spreads Developed over Time?	87
2.2.4.2	How does GP Investment Behavior Determine the Return Spread?	90
2.2.4.3	Is GP Investment Behavior Internally or Externally Driven?.....	94
2.2.5	Conclusion	97
2.3	Essay 3 - The Evolution of Private Equity Fund Terms Beyond 2 and 20.....	101
2.3.1	Introduction.....	102
2.3.2	Theoretical Background.....	105
2.3.3	Data Set and Descriptive Statistics	108
2.3.3.1	Data Set.....	108
2.3.3.2	Descriptive Statistics.....	109
2.3.4	Analysis.....	114
2.3.4.1	The Drivers of Fund Terms.....	115
2.3.4.2	Is There a New Normal of Fund Terms?	118
2.3.4.3	Alignment of Interests between LPs and GPs.....	120
2.3.4.4	How to Realign the Interests of LPs and GPs?	120
2.3.4.5	Simulation of Three Different Fund Terms	122
2.3.5	Conclusion	125

3	Conclusion	129
	References.....	136

List of Figures

Figure 1-1: PE Fundraising over Time	4
Figure 1-2: Absolute and Relative Dry Powder of Buyout Funds.....	6
Figure 2-1: Sample Distribution	37
Figure 2-2: Relation between Fund Gross Returns and Return Spreads	81
Figure 2-3: Gross Returns and Return Spreads over Time.....	83
Figure 2-4: Fund Sizes over Fund Vintage Years	109

List of Tables

Table 2-1: The Sample of Funds and Fund Families..... 33

Table 2-2: The Sample of Deals 39

Table 2-3: Distribution and Characteristics of Synthetic Funds..... 43

Table 2-4: Absolute Performance Persistence of PME 45

Table 2-5: Transition Matrices 49

Table 2-6: Fund Performance Quartiles Persistence 51

Table 2-7: Fund Top Performance Quartile Persistence..... 55

Table 2-8: Probability of Quartile Persistence (Current and Next Fund)..... 58

Table 2-9: Summary Statistics..... 78

Table 2-10: Return Spread Benchmarks..... 82

Table 2-11: Return Spread Descriptives..... 85

Table 2-12: Correlation Matrix..... 86

Table 2-13: Return Spreads over Time..... 88

Table 2-14: Influence of the Deployment Rate on Return Spreads..... 91

Table 2-15: Determinants of the Deployment Rate 96

Table 2-16: Fund Term Descriptives over Time and Fund Size Categories 110

Table 2-17: Management Fee by Time and Fund Size Categories..... 112

Table 2-18: Regressions of Fund Terms..... 116

Table 2-19: Simulation Results of Three Fund Term Cases..... 123

List of Abbreviations

CEFS	Center for Entrepreneurial and Financial Studies
e.g.	exempli gratia
EBITDA	Earnings before interest, taxes, depreciation, and amortization
et al.	et alii
etc.	et cetera
FoF	Fund-of-Fund
GP	General Partner
i.e.	id est
ICB	Industry Classification Benchmark
ILPA	Institutional Limited Partners Association
IRR	Internal Rate of Return
KKR	Kohlberg Kravis & Roberts
LBO	Leveraged Buyout
LN	Natural Logarithm
LP	Limited Partner
LPA	Limited Partnership Agreement
M&A	Merger and Acquisition
MM	Money Multiple
MSCI	Morgan Stanley Capital International
NAV	Net Asset Valuation
OLS	Ordinary Least Squares
p.	page
PE	Private Equity
PME	Public Market Equivalent
S&P	Standard & Poors

TVE	Thomson Venture Economics
U.S.	United States (of America)
USD	United States Dollar
VC	Venture Capital
vs.	versus

1 Introduction

1.1 Motivation and Research Topic

From the beginning of the 1980s until today, private equity (PE) has developed from a niche to a well-established asset class. Like hedge funds or infrastructure, PE is considered an alternative asset class, which tends to be less liquid, more risky, and more expensive than traditional stock or bond investments.¹ Yet, PE has historically promised above-market returns and enhanced portfolio diversification to its mainly institutional investors.² These benefits of the asset class have driven up capital flows into the PE industry by investors. Accordingly, yearly fundraising volumes have increased considerably from USD 7 billion in 1983, to USD 193 billion in 1998, and USD 496 billion in 2013 (Thomson ONE).³

These swelling capital flows have fuelled the evolution of the PE industry with typical side effects. Similar to the related hedge fund industry (Fung et al., 2008), the increasing capital commitments over time have shown negative correlations with fund returns in the PE industry (Harris et al., 2013a). As decreased fund returns might display the most obvious and for many institutional investors also the most relevant “new reality” of the PE asset class, the evolution of the industry also puts into question its industry mechanics and relation of market participants (Sensoy et al., 2013). For example, do investors have to reconsider their fund selection criteria? Are the existing governance structures still appropriate?

¹ Talmor and Vasvari (2011) provide an introduction into PE and point out the high costs (p. 35), low liquidity as well as high riskiness of the PE asset class (p. 39).

² A discussion on PE performance, especially in comparison to public benchmarks, is provided in section 1.2 of this introductory chapter. Talmor and Vasvari (2011, p. 27) elaborate on the diversification benefit. Franzoni et al. (2012) also point out the traditional view that PE provides enhanced diversification, but show in their study that similar liquidity risk factors of public and private alternative asset classes reduce this benefit.

³ These fundraising volumes include various PE segments according to the Thomson ONE definition: Balanced Stage, Buyouts, Core, Early Stage, Energy, Fund of Funds, Generalist, Later Stage, Mezzanine Stage, Opportunistic, Other Private Equity/Special Situations, Secondary Funds, Seed Stage, Turnaround/Distressed Debt, Value Add.

Related to these questions, this dissertation takes the perspective of institutional investors and analyzes whether these investors have to adjust their behavior to the new realities of the maturing PE industry and if so, in which way?

To further motivate this research topic, the following section starts with providing an overview of the PE industry and its development, with a special focus on PE fundraising. Thereby, the dissertation differentiates between early-stage venture capital (VC) investing and later-stage buyout investing.⁴ While this work focusses on buyout funds only, the related VC segment is presented as a deterrent example of how over-allocation of institutional money can harm an asset class. Based on this example, section 1.1 continues by pointing out the current challenges and the necessities to reconsider the industry mechanics for the buyout segment. Section 1.2 provides the research gap and develops three detailed research questions.

The starting point of the PE industry can be traced back to the 1930s and 1940s when wealthy individuals began to act as angel investors to provide start-up financing. Until the end of the 1970s, the PE industry developed slowly while keeping its focus on VC investments. In 1966, Warburg Pincus, the oldest of today's large PE firms, started its business by raising money individually for each deal (Talmor and Vasvari, 2011, p. 5). However, during the 1960s other PE firms already started to raise funds and to collect capital commitments from investors (Metrick, 2007, p. 11), which became the industry standard in the 1970s. For example, the PE firms Thomas Lee Partners in 1974 and Kohlberg Kravis & Roberts (KKR) in 1978 raised their first funds from institutional investors (Talmor and Vasvari, 2011, p. 5). The breakthrough to the modern era of the PE industry happened in 1979 when the investment rules for U.S. pension funds changed to allow these large institutional investors allocating capital to PE funds (Metrick, 2007, p. 11). In the following decades until today, the

⁴ Generally, the PE asset class is sub-divided into buyout and VC funds (e.g., see Metrick and Yasuda, 2011).

PE industry has mainly been dominated by the rising funds that used – partly excessive amounts of – leverage to finance the acquisitions of large and established companies (Kaplan and Strömberg, 2009). As these later-stage investments tend to be financed with a considerable portion of debt, these buyout deals are often called leveraged buyouts (LBOs).

The PE industry consists of three main players: *investors*, *fund sponsors*, and *portfolio companies*. Institutional investors such as banks, insurances, pension funds, but also wealthy individuals take on the role of limited partners (LPs). These LPs commit capital to funds of the general partners (GPs), the well-known PE fund sponsors such as Bain Capital or KKR, that invest the committed capital on their own discretion into portfolio companies (Kaplan and Strömberg, 2009; Metrick and Yasuda, 2011).

The typical fund life cycle of an individual fund consists of two phases. During the investment stage, ranging between 3 to 5 years, the main task of GPs is to employ the committed capital to acquire portfolio companies. In the following divestment phase, lasting 5 to 7 years, GPs focus on selling the existing portfolio companies with the aim to liquidate the fund thereafter (Litvak, 2009; Metrick and Yasuda, 2010).

As this fund life cycle allows GPs to deploy the funds' committed capital only during the investment phase, they tend to raise new funds every 3 to 5 years (Metrick and Yasuda, 2010), that is each time the investment phase of the current fund ends. Accordingly, established GPs, such as Bain Capital, have by now established a series of funds.⁵ Typically, fundraising takes around 18 months before a fund's final closing (Talmor and Vasvari, 2011, p. 30). Consequently, fundraising is an often overlooked immanent task of GPs that ensures

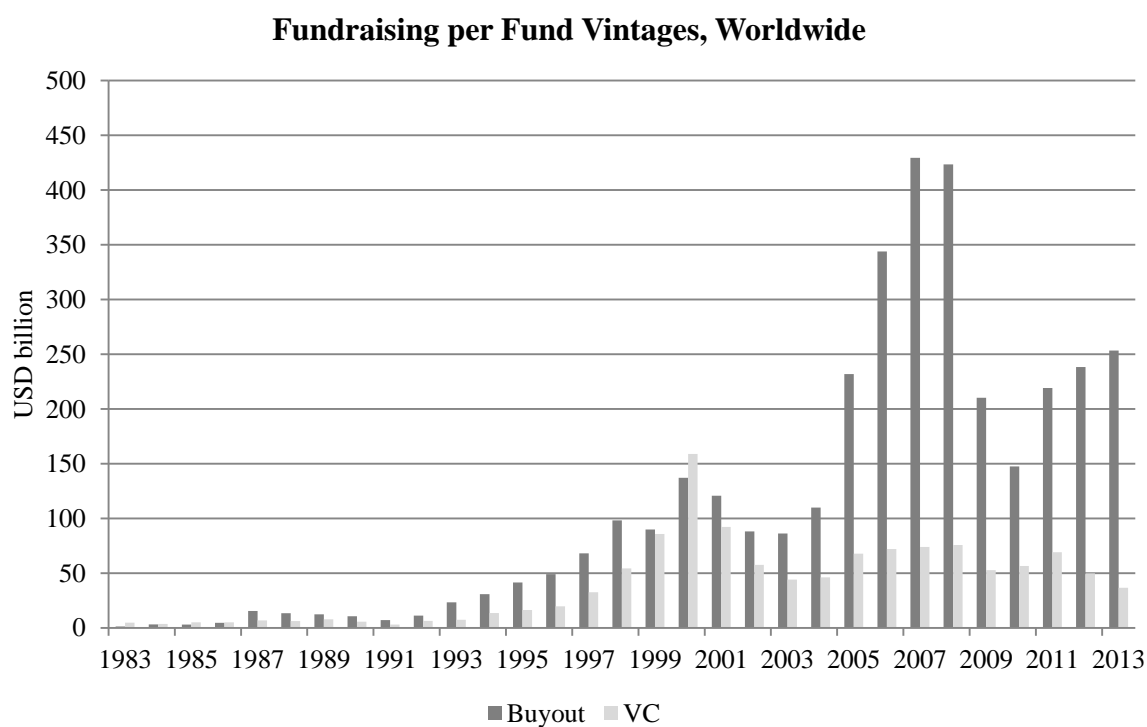
⁵ For example, Bain Capital currently raises the fund Bain Capital XI (e.g., see Chicago Tribune, Apr. 15th, 2013, "Bain nears USD 3 billion fundraising close.").

their going concern. In turn, if GPs are not able to raise a follow-on fund, they will disappear from the market after the liquidation of their final fund.

The probability of raising a successor fund, its fund size (i.e., the amount of capital commitments by LPs), and its timing depend on internal GP-related and external market factors. A key GP-related driver of a successor's fund size is the return of previous funds (Chung et al., 2012). The main external driver is the stage of the business cycle (Talmor and Vasvari, 2011, p. 30). The typical boom and bust cycles can be seen in Figure 1-1. This figure presents the worldwide fundraising volumes by fund vintage years, which are defined as the year when a fund starts investing, differentiated into buyout and VC for a 30 year period until 2013.

Figure 1-1: PE Fundraising over Time

This figure presents the global yearly fundraising volumes in USD billion per fund vintage years between 1983 and 2013, differentiated by buyout and VC funds. The figure includes 5,273 buyout and 12,119 VC funds. The data is provided by Thomson ONE.



The figure displays, for example, a strong increase in fundraising around the economic boom years 1999 and 2000 as well as 2005 to 2007, and sharp drops thereafter.⁶ Before the vintage year 2000, buyout and VC funds showed similar yearly fundraising volumes, except that buyout funds already attracted more institutional money between 1993 and 1998 and VC funds caught up in the booming dot-com years 1999 and 2000. Yet, from 2001 onwards, capital flows into buyout and VC funds exhibited a fundamentally different pattern. Worldwide VC fund commitments stabilized at around USD 50 billion per year, whereas buyout funds attracted new record commitments of USD 400 billion in 2007 and 2008 and still collected USD 150 billion to USD 250 billion per year after the financial crisis.

The stagnation of VC in comparison to buyout commitments after 2001 can be explained by disappointed LPs. As the majority of VC funds have generated negative net returns for most vintage years since 1998 (Harris et al., 2013a; Talmor and Vasvari, 2011, p. 9), the VC model was increasingly questioned and described as “broken”.⁷ Thus, large institutional investors reduced their allocations to VC funds.⁸ Lerner (2011) points out the severe consequences of such a market withdrawal: innovative growth companies might suffer from insufficient access to VC financing. In turn, this shortage can prevent companies from developing new products and growing their business. In that sense, the “broken VC model” is a deterrent example of how capital over-allocation and the following lack of investor interest can harm an asset class leading to negative consequences for the overall economy.

⁶ Fundraising usually follows the economic activity with a time lag as raising a fund takes around 18 months. Thus, large commitments still flew to funds with a vintage year 2008, although the global economy already strongly contracted in this year.

⁷ See The Wall Street Journal, Jun. 29th, 2009, “Majority of VCs in survey call industry ‘broken’.”


⁸ See Kauffman Foundation, May 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.”

What can the buyout segment, specifically the LPs investing in buyout funds, learn from the over-allocation in the VC segment? Is there an over-allocation of capital commitments to buyout funds? And if so, what are the resulting consequences and challenges for the buyout industry, specifically for LPs?

To analyze a potential over-allocation in the buyout segment, Figure 1-2 provides the worldwide absolute and relative level of unused capital commitments (“dry powder”).

Figure 1-2: Absolute and Relative Dry Powder of Buyout Funds

This figure presents the absolute global yearly unused commitment capital (“dry powder”) of buyout funds in USD billions between 2000 and 2013 (left axis). Furthermore, the figure shows the duration of dry powder for each year (right axis). This “relative” dry powder is defined as the absolute dry powder divided by the actual (for years 2000 to 2007) or projected (for years 2008 to 2013) value of transactions to be completed in future years. The figure is based on the industry statistics presented in the Global Private Equity Report 2014 of Bain & Company (p. 33).



Absolute dry powder amounts to USD 399 billion per end of 2013, having increased dramatically since 2004 from USD 176 billion to record levels of USD 479 billion in 2008. The dry powder duration (“relative dry powder”), measured in the number of years it takes to completely consume the existing dry powder, shows a similar development with record levels of about 4.9 years in 2007 and 2008 staying on a high level until end of 2013 with 3.4 years.

Thus, Figure 1-2 provides evidence for a capital over-allocation in the buyout segment and raises the question of related consequences and specific challenges for LPs.

On the one hand, the increased commitments to the buyout industry show that the asset class has developed from a niche to a well-accepted asset class among institutional investors. On the other hand, the evolution of the buyout industry with increased capital flows also creates side effects, like in the example of the VC segment. The most obvious side effect in the buyout industry is a decrease in returns reported by Harris et al. (2013a). Besides lower returns, the evolution of the buyout industry might also affect the existing industry economics (Sensoy et al., 2013), and raises fundamental questions concerning the behavior of the market participants: For example, does the governance system (terms and conditions) still provide a sufficient alignment of interests between GPs and LPs? Or can LPs still rely on old archetypes of fund selection?

This dissertation builds on the current challenges of changing industry dynamics and provides relevant implications for LPs on how to adapt to the new realities. Addressing the challenges will not only benefit LPs, but will also help the overall buyout industry to establish a sustainable way of doing business, so that market participants will not have to talk about “a broken buyout model” in some years from now.

1.2 Development of Research Questions

This section, first, starts by providing a short systematic overview of the existing PE literature, specifically about buyouts, and classifies this dissertation within this field. Second, the principal-agent theory as a theoretical framework of this dissertation is described and applied in the PE context. Third, the three research questions are developed and the related research gaps are highlighted. Lastly, this section outlines the structure of this dissertation.

The academic PE literature generally differentiates between buyout and VC studies, whereas some academic studies include both market segments (for survey papers see Cumming et al., 2007; Kaplan and Strömberg, 2009; Metrick and Yasuda, 2011). In contrast, a clear differentiation of existing research is possible between *deal-level* studies, *fund-level* studies, and emerging research that analyzes the *PE asset class* at a comprehensive level.

The buyout deal-level⁹ research predominantly analyzes the performance of individual deals and especially the underlying value-creation drivers (at a portfolio company level). The buyout fund-level research mainly focusses on the controversial discussion of fund return measures¹⁰ as well as on fund returns¹¹, their drivers and their persistence. Furthermore,

⁹ Lopez-de-Silanes et al. (2013) report the first large-scale study on deal returns with more than 7,000 observations. Various studies with smaller data sets have explored how deal returns are generated in PE (see Achleitner et al., 2010; Kaserer, 2011). More focused studies, such as Acharya et al. (2013), Guo et al. (2011), and Kaplan (1989) advance the literature specifically by focusing on operational value creation, whereas Achleitner et al. (2011), Axelson et al. (2013), and Demiroglu and James (2010) explicitly investigate financial engineering or pricing of LBOs.

¹⁰ The discussion of fund returns is complex due to a variety of fund return measures, namely the internal rate of return (IRR), money multiple (MM), and public market equivalent (PME), and due to a variety of various data sets (e.g., see Phalippou (2008) for a critical discussion of the IRR measure and Stucke (2011) for a skeptical perspective on different databases).

¹¹ Most of the controversial discussion about PE returns is related to the PME return introduced by Kaplan and Schoar (2005). This PME measure benchmarks PE fund returns to a public equity index such as the S&P 500. Some studies show a relative underperformance of buyout net returns (Phalippou and Gottschalg, 2009) or at least that buyout funds do not significantly outperform public benchmarks (Kaplan and Schoar, 2005; Kaserer and Diller, 2005). In contrast, more recent academic work shows superior fund net returns in comparison to public benchmarks (see Harris et al., 2013a; Robinson and Sensoy, 2013a; Stucke, 2011).

existing fund-level research investigates the structure and level of GP compensation.¹² An emerging strand of PE literature takes a more holistic view and employs perspectives and approaches of other (related) asset classes, to model the inherent illiquidity, risk, and alpha returns of the PE asset class.¹³

This dissertation takes a fund-level perspective throughout its analysis. After this introductory overview of the academic literature, the focus now turns more specifically to the theoretical background followed by the three research questions.

The PE asset class, with its major players GPs and LPs, is a suitable example of the principal-agent theory, which provides the theoretical framework of this dissertation.¹⁴ Jensen and Meckling (1976, p. 308) define the agency relation in their seminal paper as “a contract under which one or more persons (the principal(s)) engage another person (the agent) to perform some service on their behalf which involves delegating some decision making authority to the agent.” In the PE context, LPs exhibit the role of the principal by committing money to GPs. In turn, GPs act as the agents and invest the committed capital on their own discretion, but in the name of their LPs. This GP-LP relation is prone to opportunistic behavior by the agent. Hence, the rights and duties of both parties as well as certain incentive structures – to align the interests of both parties – are formulated in the limited partnership agreement (LPA).

¹² A related stream of fund-level literature bridges the gap between fund gross returns – before GP income is deducted – and net returns by analyzing the fund terms defined in the limited partnership agreement. Chung et al. (2012) and Metrick and Yasuda (2010), for example, investigate the proportion of fixed and variable GP compensation. Robinson and Sensoy (2013b) show that GPs with more costly fund terms generate similar net returns in comparison to less costly peers.

¹³ Ang et al. (2012) analyze the influence of illiquidity on choices in a multi-asset portfolio, whereas Ang et al. (2013) determine a perceived alpha of holdings of endowments for PE and hedge funds. Driessen et al. (2011) and Gredil et al. (2014) develop a new methodology to calculate the alpha and beta of PE funds.

¹⁴ See McCahery and Vermeulen (2013) and Phalippou (2009).

From a theoretical point of view this dissertation builds on the principal-agent theory and analyzes whether the mechanisms to minimize the agency costs are still suitable in the buyout industry. The underlying assumption of this question is that LPs might not have sufficiently adjusted their view on the maturing buyout asset class, and thus their relation to GPs, in answer to the swelling capital flows, lower returns, and changing industry dynamics.

The principal-agency theory arises due to information asymmetries and conflicting goals between the principal and the agent. Jensen and Meckling (1976) describe the asymmetric information of the principal-agent theory, as the agent usually has an informational advantage in comparison to the principal. In this context, Leland and Pyle (1977) point out that information asymmetries are especially pronounced in financial markets. These informational asymmetries generally favor the delegation of decision rights to the agent, however, also feature the risk of opportunistic behavior of the agent, when his actions are not completely observable. Ross (1973) formulates that the agency problem is based on the conflicting goals of the principal and the agent to maximize their own expected utility. Thus, there is a considerable probability that the agent will deviate from the decision that will maximize the utility of the principal, while maximizing his own interests. The PE industry also features information asymmetries and conflicting goals between LPs and GPs. GPs, which take investment decisions in the name of their LPs, might act in a way that maximizes their fees, but not necessarily the returns for their investors.

The principal-agent theory also includes three specific information-related problems (see Picot et al., 2008, p. 74): (i) adverse selection, (ii) moral hazard, and (iii) hold-up. Yet, most academic literature focusses mainly on adverse selection and moral hazard, as these problems are more relevant in practice and often occur simultaneously (see for example Picard (1987)). As only adverse selection and moral hazard are relevant for the PE industry, the following

paragraphs focus on these two problems. Picard (1987) briefly explains adverse selection by the inability of the principal to determine the agent's private information, whereas moral hazard is based on the missing ability to monitor the agent's behavior.

Akerlof (1970) describes the problem of adverse selection in more detail with an example related to the market of used cars, in which information asymmetries between the principal and the agent are present ex-ante, thus, before the principal decides to engage in business with the agent. In the PE context, LPs frequently face the problem of adverse selection in their immanent task to select funds to which they allocate their money. For LPs it is hard to evaluate the true capabilities and characteristics of GPs ex-ante to generate above-market returns with their next fund.

The academic literature provides three ways to mitigate the problem of adverse selection (Picot et al., 2008, p. 76). First, GPs can signal ("signaling") their quality (see also Spence (1973) for an example of signaling related to job markets). In the PE industry, GPs tend to market the success of their last fund to attract LPs for the next fund (Harris et al., 2012). Furthermore, the fund managers usually commit their own private capital into a fund to signal their confidence and commitment as they have "skin in the game".¹⁵ Second, principals can screen ("screening") prospective agents, which is done intensively by LPs during the due diligence process when a GP is seeking capital for a successor fund (Talmor and Vasvari, 2011, p. 93). Third, adverse selection can also be reduced when principals introduce different offers for agents. Agents will choose offers ("self selection") according to their characteristics and will thus reveal information to the principal. Currently, this self selection is hardly observed in the PE industry. In addition to these three ways to mitigate adverse selection,

¹⁵ Usually GP fund managers inject 1% (of the fund's committed capital) own capital into the fund (Metrick and Yasuda, 2010).

reputational concerns of agents (“reputation”) can also reduce adverse selection in repeated interactions. If a high reputation – in the long-run – pays off for the agent, he tends to behave in the best interest of the principals. In the PE industry, LPs tend to stick to established brand names with a series of funds, enhancing reputation incentives.

Similar to adverse selection, the problem of moral hazard is based on information asymmetries. However, the asymmetries in moral hazard do not exist *ex-ante* before the principal’s decision to engage in business with the agent, but *ex-post* when both have started working together (Picot et al., 2008, p. 75). If the principal is not able to observe and evaluate the agent’s actions, the agent might deviate from the best interest of the principal to maximize his own utility.

Moral hazard is also observed in the PE industry. LPs commit capital to GPs for a long time horizon, but are hardly able to evaluate the GPs’ efforts within the fund life time as the final fund performance becomes obvious only after the closing of the fund. The academic literature displays two levers to reduce moral hazard (e.g., see Eisenhardt, 1989; Picot et al., 2008, p. 79). First, principals can reduce the information asymmetries by monitoring activities, as explained by Fama and Jensen (1983) who point out the monitoring activities of boards. LPs monitor GPs, for example, during annual meetings of the respective fund and, more intensively, with the help of quarterly reports provided by GPs. Second and most importantly, principals should establish an adequate incentive system, which aligns the interest of principals and their agents (Berhold, 1971). The PE industry has implemented this idea by introducing a GP compensation that consists mainly of a fixed management fee and a success-based compensation (Metrick and Yasuda, 2010).

Lastly, the principal-agent theory includes the so-called hold-up problem, also described by Picot et al. (2008, p. 75). This problem arises due to information asymmetries between

both, the principals and the agents, to another third party, mainly a juristic institution. Thus, even if principals have the full knowledge of the agent's actions, it might be difficult to prove potential deviations in court. However, the hold-up problem only occurs when the principal depends in some way on the agent and therefore cannot act on the agent's misbehavior. In practice, the hold-up problem is far less important than adverse selection and moral hazard, which is also valid for the PE industry. LPs can easily walk away from one GP after the current fund is fully realized or might even sell their current share of the fund in a secondary transaction.

Overall, the principal-agent theory sheds light on the relation between LPs and GPs. In line with the theory, the PE industry also features various measures to mitigate the omnipresent risk that GPs might engage in opportunistic behavior. As explained above, GPs try to signal their quality, LPs conduct a thorough GP due diligence, and an incentive system is in place to align the interests of LPs and GPs.

Now, the important question in the light of the maturing buyout industry is whether the measures are still effective in mitigating the principal-agent problem between LPs and GPs. For example, is the incentive system in the buyout industry still appropriate to minimize the agency costs? How reliable are signals by GPs in a more competitive and more efficient buyout industry? The overarching question of this dissertation can be summarized as follows: *Do LPs have to adjust their behavior to the new realities of the maturing buyout industry and if so, in which way?*

Based on this overarching question, this dissertation develops three specific research questions in the following paragraphs, which are based upon the following questions LPs are facing in their daily business: Should I invest in the new fund of a GP? If yes, how much money should I invest? To which fund terms (fees and conditions) should I invest?

The first question, whether a LP should invest in a new fund of a GP, addresses the agency problem of adverse selection. At the inception of a new fund, LPs face an information disadvantage in comparison to a GP, concerning the GP's inherent capabilities to generate above-average returns with the next fund. Consequently, LPs usually intensively screen GPs to reduce asymmetric information. They conduct a GP due diligence and analyze, for example, the investment team, the investment strategy, and the performance as well as the value-creation strategy of past deals or past funds (Talmor and Vasvari, 2011, p. 93). In turn, GPs usually proactively signal the performance of past fund returns. For example, Harris et al. (2012) show that surprisingly more than 25% of all funds in a given vintage year claim to be a top quartile fund. GPs most likely try hard to sell the top quartile quality label, justified or not, as LPs heavily depend their fund selection on a GP's signal of their past fund performance. This behavior can be considered as rational based on the seminal work by Kaplan and Schoar (2005). They document a distinctive feature of the PE asset class that the successor fund of a GP has a higher probability to achieve a top ranking when the previous fund has already been successful in comparison to its peers. As the study by Kaplan and Schoar (2005) only includes fund vintages until about 1996, the question arises as to whether this performance persistence is still existent in the asset class. Potentially, the performance persistence might simply be an artefact of the first two decades in the buyout industry, but might have disappeared with the growing capital commitments. Consequently, chapter 2.1 addresses the following research question:

(1) Is the past fund performance of a GP still a good predictor for the performance of the successor fund?

As performance persistence targets the most elementary decision LPs have to make, namely whether to invest in a new fund or not, various researchers have followed Kaplan und

Schoar (2005) in analyzing this topic. For example, Chung (2012) shows that excessive fund growth, driven by superior returns of a previous fund, decreases performance persistence. In recent years, Harris et al. (2013b) also document lower persistence, while Sensoy et al. (2013) find disappeared persistence. The essay in section 2.1 takes a more detailed look on performance persistence. In contrast to the existing fund-level studies, this essay employs a rich deal-level data set of over 10,000 LBOs and a new approach of sequencing deals and repacking them into synthetic funds.¹⁶ This approach circumvents various technical issues of using fund-level data, for example, that consecutive funds frequently include the same deals.¹⁷

The second research question concerns a GP's fund size. If a LP decided to invest into a new fund of a GP, the LP would have to determine how much money to invest. This decision includes the adverse selection and moral hazard agency problems, which in this case focus both on the amount of committed capital for a specific fund.

First, the adverse selection problem is relevant for LPs when analyzing in their due diligence whether the target fund size of a new fund is appropriate. Thus, LPs have to evaluate whether the respective GP is able to employ all the committed capital based on a sufficient deal flow of high quality transactions. The appropriateness of a realistic target fund size is particularly important as the capital flows into the buyout industry have shown a long-term upward trend. For example, some funds reached sizes above USD 10 billion in 2013, and the unused committed capital for buyout funds in 2013 was close to the 2009 record levels of USD 399 billion.¹⁸

¹⁶ This study draws on a data set from the Center for Entrepreneurial and Financial Studies (CEFS), which includes deal-level data from three large international LPs.

¹⁷ GPs quite often finance the last deal of the current fund with money from the successor fund when the capital commitments of the current funds are already fully drawn.

¹⁸ See Bain & Company, Global Private Equity Report 2014, p. 6.

Second, the moral hazard problem is present after the inception of a fund. An extreme and peculiar example of the moral hazard problem is the issue of the so-called zombie funds.¹⁹ The GPs of these funds simply stop investing – usually after it has turned out that the few investments so far have been write-offs so that the variable compensation will be zero anyway – and continue to extract the fixed management fee from the fund. Yet, this extreme form of moral hazard does not entirely fit the definition within the principal-agent framework, as the refusal to conduct deals can easily be observed by LPs.

A more intuitive moral hazard problem arises when a GP recognizes that the current fund with a specific size faces too few attractive deals. On the one hand, a GP could act rigorously and invest only, for example, half of the committed capital in the few attractive deals so that net returns might still be satisfying for the LPs. On the other hand, GPs might simply invest in less promising deals to reduce the unused committed capital (i.e., dry powder) while accepting lower returns for their LPs, and not to be confronted with critical questions related to the target fund size of a potential successor fund.

As the dry powder has reached record levels in recent years, it has become increasingly important for LPs to understand the costs and drivers of this unused capital. Thus, LPs need – more than ever before – to consider the disadvantages of capital over-allocation to a specific GP and the buyout industry as a whole in their investment strategy. Consequently, the second research question is:

(2) What are the costs and drivers of dry powder in buyout funds?

The cost of investing into buyout funds has been addressed by several academic studies, for example, by Chung et al. (2012) and Metrick and Yasuda (2010) who analyze the mix of

¹⁹ See The Economist, March 23th, 2013, “Zombies at the gates.”

fixed and variable GP compensation. These studies base their cost analyses mainly on assumptions and simulations. In contrast, the essay in section 2.2 employs actual data of 358 nearly realized buyout funds including their difference between gross and net returns (“return spread”), which makes it possible to expand the existing literature in several ways.

First, in contrast to the existing studies this essay provides a cost benchmark for different levels of gross returns. Second, the essay displays the return spread over time and investigates the drivers of the return spread changes. Third, it identifies the GP investment behavior, namely the increasing amounts of unused capital over time as an important driver of the increased costs of PE investing, which has been so far mainly disregarded in academic research. Fourth, the essay investigates drivers of dry powder and differentiates between internal GP-related and external market-related drivers.

The third research question again focusses on the moral hazard problem and is related to the fund terms that are negotiated between LPs and GPs and documented in the LPA. Historically, the terms for buyout funds (so-called 2 and 20 structure) include a fixed management fee of about 2% of the committed capital per year (during the investment phase) and a variable compensation of 20% of all fund profits (“carried interest”) if a specific return threshold (“hurdle rate”) is reached (Metrick and Yasuda, 2010). The fixed management fee has initially been intended to cover the ongoing operating costs of the GP for managing the fund, whereas the purpose of the carried interest has been to align the interests of LPs and GPs. In other words, GPs should only see substantial compensation when they deliver decent returns for their LPs. Yet, the exact reason for the 20% carried interest in the PE industry is unknown. Metrick and Yasuda (2010, p. 2311) summarize the speculation about the origin of the 20% carried interest as being related to “Venetian merchants in the Middle Ages, speculative sea voyages in the age of exploration, and even the book of Genesis as sources.”

Still, it is not the origin of the compensation scheme, but its effectiveness that is relevant for LPs. For example, the first fund of KKR with the vintage year 1976 contained only USD 31 million, leading to USD 0.6 million yearly management fees (assuming a 2% management fee). In contrast, the largest KKR fund so far from 2006 with a size of USD 17.6 billion would yield an astronomical yearly fee of USD 352 million.²⁰ This amount certainly does not reflect the initial purpose of management fees to simply cover the operating expenses of a fund, nor does it allow the carried interest to be the major source of GP income. Clearly, the massive increase in fund sizes over the last decades (with strong cyclical movements) has weakened the effectiveness of the compensation scheme in the PE industry to reduce the moral hazard problem. In this respect, industry participants have raised concerns that GPs might focus on maximizing their assets under management, but not necessarily the returns for their LPs (Hudec, 2010). Consequently, the third essay analyzes the final research question:

(3) What are alternative terms for buyout funds that realign the interests of LPs and GPs?

Although various studies have already explored fund terms in the PE context (see Chung et al., 2012; Gompers and Lerner, 1999; Litvak, 2009; Metrick and Yasuda, 2010; Robinson and Sensoy, 2013b), no study so far has investigated the development of fund terms over time nor has any study differentiated between its drivers. Chapter 2.3 provides a profound understanding of the changes in fund terms over time making it possible to evaluate whether the alignment of interests between LPs and GPs has deteriorated in recent years. Furthermore, in contrast to existing publications, this analysis provides guidance towards new fund terms by employing a simulation of GP compensation depending on different fund term scenarios.

²⁰ The calculation assumes a management fee 2%, which is most likely too high for such a large fund. A more realistic fee of 1.5% would drive down yearly fees during the investment period of the fund to USD 264 million.

Overall, this dissertation analyzes the relation between LPs and GPs in the framework of the principal-agent theory and the related problems of adverse selection and moral hazard. It becomes clear that the evolution of the buyout segment leads to changes in this relation. The most obvious and relevant change for LPs is the reduced returns generated by GPs. However, there are also changes in the industry mechanics beyond lower average industry returns. These changes increase the challenges for LPs, which they have to address when considering the allocation of money to GPs: Do I invest in a fund of a GP? How much do I invest? On which terms do I invest?

This dissertation sheds light on these questions and provides useful implications for LPs to adapt their way of doing business to the new realities in the buyout industry. Ultimately, this dissertation facilitates the realignment of interests between LPs and GPs. In the long-run, low agency costs based on a sustainable relation between LPs and GPs will benefit the asset class as a whole.

The following Chapter 2 presents three self-contained essays in three different sections, which display individual academic papers. Each section addresses one of the three research questions. Section 2.1 investigates whether the past fund performance of a GP is still a good predictor for the performance of the successor fund. Section 2.2 analyzes the cost and drivers of dry powder in buyout funds, followed by an evaluation of alternative terms for buyout funds for realigning the interests of LPs and GPs in section 2.3. Chapter 3 of this dissertation summarizes the aggregated findings, develops practical implications as well as potential topics for future research, and concludes with an outlook in this field.

2 Essays

2.1 Essay 1 - How Persistent is Private Equity Performance? Evidence from Deal-Level Data

Abstract

The persistence of returns is a critical issue for investors in their choice of private equity managers. In this analysis we approach performance persistence in a new way, using a unique database containing cash-flow data on 10,637 portfolio company investments by 236 buyout funds. We confirm earlier findings of top-quartile persistence until the late-1990s, which we find was driven by experienced managers. However, as the market has grown and matured, fund managers no longer exhibit performance persistence. Private equity has conformed to the pattern found in most other asset classes where past performance is a poor predictor of the future.

Keywords: Private equity, returns, persistence, portfolio companies

JEL Classification Code: G11, G23

Authors: Braun, Reiner; Jenkinson, Tim; Stoff, Ingo

First Author: Braun, Reiner

Current Status: Working paper

Conferences: Global Private Investing, University of North Carolina, Dec. 6th 2012
Oxford Private Equity Forum, SAID Business School, Mar. 5th 2013
Private Equity: Science & Practice, Stanford University, Mar. 11th 2013
Internal Seminar, SAID Business School, Oct. 17th 2013
Finance Roundtable, University of North Carolina, Nov. 22th 2013
7th PE Findings Symposium, London Business School, Jun. 3rd 2014
Various practitioner conferences

2.1.1 Introduction

Persistent out-performance by financial intermediaries is rare. Although such claims are often heard, academic research across a range of assets has generally found past performance to be a poor predictor of future returns. An asset class where the claims of persistence have, to date, withstood academic scrutiny is private equity (PE) – including both venture capital (VC) and buyouts (Chung, 2012; Hochberg et al., 2012; Kaplan and Schoar, 2005). In this analysis we re-visit the issue of persistence of PE returns taking advantage of a new, extensive, deal-level data set of buyout fund investments. While confirming the finding of performance persistence from early papers, which focused on funds raised until the late-1990s, we find that persistence has disappeared as the market has expanded and matured. We also distinguish between fund manager experience and market maturity, and find that the drop in persistence is particularly pronounced among the top quartile more experienced fund managers.

Since PE investments are typically organized into limited partnership funds, most analyses of returns consider the persistence in performance across the successive funds of a particular PE manager. This creates challenges: funds have lives that are often in excess of ten years, reflecting the time taken to acquire, work with, and exit the portfolio companies in which they invest. Ultimate performance is only known accurately at the point when all investments have been exited and the cash has been returned to investors, but the speed with which funds enter and liquidate their positions varies considerably. Until exit, funds report their estimates of asset values, but these may be biased predictors of future cash returns. There are also difficulties in benchmarking performance, since funds are conventionally defined by the date they start investing, rather than when they actually invest or divest their funds. Furthermore, even though better fund performance databases have become available, they often do not have information on each fund of a particular manager. Consequently, there

are often gaps in the fund sequence which complicates temporal comparisons of performance. Finally, recent research has raised serious questions about the reliability of some of the fund data that had been made available to researchers. For all these reasons, analyses of performance across funds are far from straightforward for both academics and investors.

Funds, however, are simply a legal wrapper around a sequence of underlying portfolio company investments, and in this analysis we investigate the performance persistence of PE managers through a different lens: by stripping off the legal wrapper and considering how performance evolves across the deals undertaken by each PE firm. We shall refer to the latter as the general partner (GP), and distinguish between a GP – which manages a whole sequence of funds – and any particular fund. For example, Blackstone Group is the GP of multiple funds, including their latest PE fund Blackstone Capital Partners VI, which raised around USD 16 billion by its final close in January 2011. If one looks at the Blackstone website they provide considerable information on the companies in which they have invested over the years, but make little mention of which legal fund(s) made the investment. This is our approach. We consider the performance of the GP as a sequence of individual portfolio company investments, and characterize how this evolves over time.

This is the first study to analyze persistence in this way, since data is usually only available at a fund-level. We are able to take this approach as we have timed cash-flow data at the portfolio company level for 10,637 investments by 236 GPs. While a few other papers have analyzed portfolio company returns (such as Lopez-de-Silanes et al. (2013) whose data derives from the private placement memoranda of funds), this is the first to have timed cash-flow data at the deal-level. This allows us to move beyond the simple performance measures – such as internal rates of return (IRR) and money multiples (MM) – still favored by many in the PE sector. We also compute, and focus on, PE returns relative to public markets using the

public market equivalent (PME) approach introduced by Kaplan and Schoar (2005), which is arguably a much more relevant measure of performance.

Furthermore, for the vast majority of the GPs in our sample we have their complete investment history. We source the data from three fund-of-fund (FoF) managers who required all GPs who sought commitments to provide this detailed deal-level information. Importantly, the sample includes all the GPs upon which the FoF managers performed due diligence, whether or not they actually chose to invest. Also, the data are not derived from marketing materials or private placement memoranda – which can be selective in the deals they include – but are from formal requests for the full track-records of all current and prior investments in a standardized form.

On the other hand, obtaining data from investor due diligence – rather than the more conventional source via investors' fund holdings – has certain limitations. In particular, only funds that sought a follow-on fund will show up in our sample, and so fund sequences end when a PE manager quits the market – which will tend to follow poor performance. While this would be critical for an analysis of returns across the PE asset class, our focus is exclusively on return persistence within a GP. In the context of analyzing return persistence, missing the historical performance of GPs that died is less of an issue. Viewed from a practical perspective, our analysis sheds light on the performance persistence of GPs that have been seeking capital from investors (whether or not they were successful).

One implication of using deal-level data is that we observe gross returns, since management fees and profit shares (otherwise known as carried interest) are levied, respectively, on committed capital and profits at the fund-level. In order to understand the overall value creation of a GP the gross returns are a clean measure – since fund terms can

differ across investors, which will impact on the reported net returns at the fund-level. In practice, fund terms change rather slowly over time, and so the persistence of gross and net returns should be rather similar. We confirm this by comparing our results to earlier studies that used fund-level net returns. Clearly, investors are ultimately interested in the net returns they earn. Recent papers by Harris et al. (2013a), Higson and Stucke (2012), and Robinson and Sensoy (2013a) provide such evidence for the asset class as a whole, and Metrick and Yasuda (2010) provide detailed information about management fees and carried interest arrangements that apply at the fund-level, and therefore bridge the gap between gross and net returns.

The data we use in this analysis comprise over ten thousand buyout portfolio companies, for which we have the cash flows (investments and realizations), their industry, exact date, location of the investment and certain other characteristics. The investments are split roughly equally between the U.S. and Europe with a few investments in other regions, and span the period 1974-2010, although there are few observations in the sample before 1980. For those investments that have not been fully exited, we also have net asset values as of the latest date the manager reported their track record to the FoF manager. However, one major advantage of using deal-level data is that it allows us to focus on realized investments – and thereby avoid issues arising from potentially biased accounting values for remaining portfolio companies – without excluding entire funds that still retain a significant amount of unrealized investments. Recent papers by Brown et al. (2013) and Jenkinson et al. (2013) find that, across the lives of funds, remaining portfolio companies are valued conservatively. Since the extent of conservatism may well vary across GPs, this could impact on analyses of persistence at the fund-level. The main focus of our analysis is on the 6,195 fully-realized unique portfolio company investments in our sample, and so we avoid such valuation issues entirely.

There are some other significant advantages of using deal-level data. First, any analysis of persistence needs to define the frequency with which returns are observed. For instance, for mutual funds and hedge funds persistence is typically analyzed by observing monthly or annual returns, but clearly the analysis could be performed at higher or lower frequencies. For PE, researchers have been more constrained, as returns have only typically been available at the fund-level. The legal lives of funds are around 10 years, although (as we discuss below) successive funds overlap. In terms of the number of investments, funds typically comprise about 10-15 portfolio companies. So focusing on fund returns, while interesting in its own right, inevitably results in a low frequency analysis. In contrast, we are able to choose our frequency. We could, for example, analyze persistence from one deal to the next. However, there is a large amount of idiosyncratic volatility in returns, and it is unlikely that clear patterns would emerge. In this analysis we measure frequency at a medium frequency. We sequence each investment by the date of the first cash flow, and then construct “synthetic funds” comprising 8 deals. As we explain in more detail, we choose 8 as this is our best estimate for the number of investments a GP completes before starting to raise a successive fund, and so is the sort of track record investors would observe when considering the next pitch for money by a given GP. However, similar results are obtained for different sized synthetic funds. An advantage of this approach is that it generates many more (synthetic) funds than have been used in prior studies: our sample generates 755 synthetic funds of which 560 have a prior fund, and so can be included in the persistence analysis. By comparison, the sample in the seminal Kaplan and Schoar (2005) study comprised 76 U.S. buyout funds with a prior fund.

Second, there are potentially important issues regarding overlapping funds, as noted by Korteweg and Sorenson (2013). Typically, GPs start to raise a new fund around 3 years into

the life of the current fund. The overlapping nature of successive funds can create correlations in returns, due to exposure to common factors or even, in some cases, investments that are split across successive funds. This could produce spurious patterns of persistence between funds. By taking a deal-level approach, we avoid the problem of deals being included in adjacent funds. The synthetic portfolios are created by slicing up the temporal sequence of deals undertaken by each GP, irrespective of which fund actually made the investment. Furthermore, we are able to estimate the impact of this issue on measuring performance persistence by observing the impact of re-introducing deals that are split between funds as multiple observations of the same deal. We find relative performance persistence to be 2 percentage points lower if common deals are excluded from the analysis. The average probability of staying in the same performance quartile with two subsequent funds is 31% including, and 29% excluding, common deals.

A third important advantage of using deal-level data is that it enables far more accurate benchmarking of performance. It is clearly important to control for the date of investment when comparing performance, given the changing macro-economic and financial market conditions, and this is not straightforward in the case of PE funds. The complexity derives from the fact that investors commit capital to funds but the fund managers only call the capital when they find an investment. The distribution of investment over the life of the fund is not normally known to researchers, and so there is little alternative but to adopt the industry practice of using the reported “vintage year” of the fund. The vintage year itself can be defined in different ways – with the most common definitions being the date of the first investment and the date of the first closing of the fund. However, two funds with the same reported vintage year can invest their committed capital at a very different pace: one fund may not make its first investment for 2 years, whereas another fund might have invested most of

its capital within this period. Using reported vintage years therefore introduces noise into the benchmarking process. We are able to address this issue directly: we define a date for our synthetic portfolios which is the (investment-size) weighted average date of the investments. In this way, when measuring the relative performance of GPs, we are able to control for the timing of investments in a much more precise way.

Finally, the additional information we have on the industry, location, size etc. of each investment allows us to construct controls for such factors for each synthetic fund, which again allows more accurate benchmarking performance between GPs.

Our main findings are as follows. We first analyze performance persistence in absolute terms (rather than relative to other GPs) focusing on PME (which controls for movements in the local broad stock market index). To begin with we split our sample roughly in half, with the early sample – up to the end of 1998 – corresponding to the sample period analyzed in the Kaplan and Schoar (2005) paper. This allows a direct comparison with their results (although they studied both buyouts and VC, whereas we only study buyouts). For this early period we find very similar levels of absolute performance persistence. This is interesting given the fact that Kaplan and Schoar (2005) use fund-level net returns, whereas we use synthetic portfolios built up from unique and fully-realized deals, and gross data. Overall, our results of absolute performance persistence provide strong confirmation for the results found by Kaplan and Schoar (2005) up to the late 1990s.

Having replicated, as far as possible, the Kaplan and Schoar (2005) results, we then extend the analysis to include synthetic funds with vintage years from 1999 onwards. This nearly doubles the total number of synthetic funds (with prior funds) to 560. As we expand the sample to include the more recent deals we find a noticeable reduction in absolute return

persistence across the whole sample. Indeed, if we focus on just the post-1998 synthetic funds we find no statistically significant relation between prior and current fund PME's.

However, while looking at absolute return persistence for the whole sample is informative, it is the relative performance of GPs that is probably of more interest. We start by putting funds into performance quartiles according to their vintage years. Note, that at this point our ability to unambiguously assign vintage years – according to the capital-weighted date of the investments – is particularly useful. We then compute transition matrices to map how successive funds move between quartiles. For the whole sample we find marked top and bottom (“flop”) quartile relative persistence: using both PME and MM we find that around 34% of top and 36% of flop funds stay in the same quartile in their next fund. There is somewhat less top quartile persistence when using IRR as a performance measure.

To explore this further we switch to a logistical regression framework to analyze the probability of staying in the same quartile in successive funds. We find that relative performance persistence is only apparent for the pre-1999 synthetic funds, which echos the absolute performance results. To understand this finding, we also investigate the impact of GP experience.

At the same time as the PE industry has grown and become more competitive, some GPs have become a lot more experienced, and we explore whether such experience impacts independently on persistence. On the one hand it could be that the entry of inexperienced GPs lies behind the fall-off in persistence. On the other hand, it could be that experienced GPs lose their edge over time. While experience of completing a series of successful deals may have many benefits, there are also reasons to doubt whether success necessarily breeds success. The personal rewards that accrue to successful PE managers may reduce the drive to produce

exceptional performance, or reduce risk appetite, and there will be various organizational issues that larger GPs face – such as keeping talented individuals who lack influence or rewards in a large GP, and may prefer to raise their own fund.

To disentangle the impact of experience and market maturity, we estimate a multivariate model that allows us to compute the marginal effects on performance persistence of low/high experience and the early/late market conditions. Our overall finding is that in the early period, most of the persistence in returns was associated with experienced and top quartile GPs. For instance, the probability of repeating top quartile performance (as measured by PME) was around 50% in the pre-1999 period for those GPs with above median experience. However, the impact of experience on persistence has disappeared as the market has matured: in the more recent years this top quartile persistence probability has dropped to below 25% (i.e., what we would expect from a random draw). We show that this conclusion holds whether performance is measured by PME, MM or IRR. Even experienced GPs have, therefore, not been able to defy the general trend towards past performance becoming a poor indicator of future performance. PE appears to have conformed to the pattern of mean reversion found in most other areas of asset management.

Our analysis relates to various strands of the academic literature. Most directly, a contemporaneous analysis of U.S. buyout and VC funds by Harris et al. (2013b) finds a similar fall-off in performance persistence in recent years for buyout funds but continued and significant persistence for VC GPs. Our findings are complementary – we use deal-level data for both U.S. and European buyout fund managers and find similar results with our synthetic fund approach. Others who have studied persistence in PE returns include Chung (2012) who finds that persistence decreases as fund size increases. Using our deal-level data we find no significant impact of deal size on persistence.

These findings about GPs may help to explain similar results that have been found about ability of limited partners (LPs) to consistently produce superior relative returns. Lerner et al. (2007) found that endowments earned significantly higher returns from their PE investments than other classes of investors such as pension funds or insurances. Sensoy et al. (2013) update this work and find that LP performance persistence has disappeared in more recent years. The disappearing GP and LP performance persistence may plausibly be linked. Those LPs who were early investors in PE – such as endowments – established relationships which were valuable when the market was developing. We find the more experienced funds had significant top-quartile persistence in the early years. However, those relationships, and access to funds – at least on the buyout side – are now much less valuable and are no longer a source of LP out-performance. Although Harris et al. (2013b) find continued performance persistence for VCs, the flows of money into VC have fallen significantly in recent years, in direct contrast to the vast increase in money flowing into buyout funds. Consequently, the value of those VC relationships may be much less significant for overall portfolio returns than in the early years, when VC captured a much larger proportion of the overall PE investments.

The remainder of the essay is structured as follows. In the next section we explain carefully our deal-level data, sample selection issues and how we construct our synthetic portfolios. We present results on for the persistence of absolute performance in section 2.1.3. In section 2.1.4 we classify funds according to their performance relative to funds investing in the same period, and analyze – for the whole sample – whether there is relative performance persistence. Section 2.1.5 then focuses on relative persistence at the top (and bottom) quartiles, as it is the belief in top quartile persistence that drives many investors. Section 2.1.6 concludes.

2.1.2 Data

In this study, we use an extensive data set of PE buyout investments. The data set is derived from three large fund-of-fund (FoF) managers as an integral part of their due diligence process when selecting funds to invest in. One very important feature of the data set is that all GPs who sought capital from one or more of the FoF managers were required to provide information on all their prior fund investments. Therefore, if a GP sought capital for its fifth fund, detailed information would be required on funds 1 to 4. Of course, the FoF managers only chose to invest in a small fraction of the funds. Our data set comprises the full sample provided by all the GPs that manage buyout funds, and is in no way affected by the selections of the FoF managers.

2.1.2.1 Funds and Fund Families

The data from each of the three FoF managers is merged and duplicate GPs are removed. In total, there is information from 236 GPs and 726 buyout funds. Each time a GP approached the FoF managers, their track records were updated. The latest updates therefore vary by GP, with the most recent being at the end of 2010. For fund investments that had not been realized, the latest net asset valuations (NAVs) were provided. However, in this study we focus on realized investments thereby avoiding any issues regarding the accuracy of the NAV estimates, or timing issues regarding when the NAVs were reported.

A small number of GPs run more than one “family” of buyout funds – usually distinguished by geographic area. However, this is not a major issue in our data: for the 236 GPs, 258 fund families are identified. An attractive feature of the data is that it is balanced between funds with their origin in the U.S. (362 funds) and Europe (325 funds). It also

includes a limited number of funds from Asia (39 funds). Most previous academic studies have focused on U.S. funds only, whereas European PE transactions comprised 47% of global deals by value between 2000 and 2007 (Kaplan and Strömberg, 2009).

Details on the number of funds and families are provided in Table 2-1. Panel A shows that for 88% of the fund families we have complete information back to the first fund. In most of the remaining cases the sequence is complete but does not go back to the first fund (as such detailed records were not always kept for early investments). We have gaps in the sequence for only four fund families. When we are missing one or more funds in the sequence, we interpolate the deal sequence numbers using the median number of deals performed by the GP in the funds for which we have information. This completeness is a major strength of the data set, given our focus on performance persistence. No other data set that we are aware of combines a large number of funds with such a high level of completeness over the life of the GP/fund family. The completeness, and deal-level granularity, of our sample reflects the fact it was gathered from due diligence, rather than the actual holdings of investors. Balanced against these advantages, there are inevitable sample selection issues as GPs who disappear, often because of poor prior performance, will not be part of our data set. We address these selection issues first, before describing our sample.

Table 2-1: The Sample of Funds and Fund Families

This table illustrates the coverage of fund families with our sample of 10,637 realized and unrealized transactions and 726 funds. We assigned each fund to a fund family based on the GP managing the fund. In case a GP runs more than one fund sequence we formed more than one fund family for this GP, based on the countries in which the respective fund sequence operates. This procedure results in a total of 258 fund families. In Panel A we indicate the fund generation of the first fund in each fund family that we observe in our sample and the number of subsequent funds that we cover. For example, our sample includes 45 fund families for which we observe the complete sequence of the first, second, and third fund. At the right-hand side, instead of length of fund sequence, we show the number of incomplete sequences within a fund family. These numbers indicate that only four fund families in our sample are incomplete in the sense that we do not observe all funds in between the first and last fund of the respective fund family. In Panel B we provide the same information at the fund-level. For example, the 45 fund families for which we observe the complete sequence from the first to the third fund are made up by 135 funds.

Panel A: Number of Fund Families

	Length of Fund Sequence										Incomplete Sequences			Total	(%)
	1	2	3	4	5	6	7	8	9	10	4	5	6		
First Fund in Sequence	58	61	45	24	23	6	3	2	1	2		1	1	227	(88%)
2	7	1	4	1								1		14	(5%)
3	2	2	1	2	1									8	(3%)
4		1	2	1	2									6	(2%)
5											1			1	(0%)
6			1			1								2	(1%)
	67	65	53	28	26	7	3	2	1	2	1	2	1	258	(100%)

Panel B: Number of Funds

	Length of Fund Sequence										Incomplete Sequences			Total	(%)
	1	2	3	4	5	6	7	8	9	10	4	5	6		
First Fund in Sequence	58	122	135	96	115	36	21	16	9	20		5	6	639	(88%)
2	7	2	12	4								5		30	(4%)
3	2	4	3	8	5									22	(3%)
4		2	6	4	10									22	(3%)
5											4			4	(1%)
6			3			6								9	(1%)
	67	130	159	112	130	42	21	16	9	20	4	10	6	726	(100%)

2.1.2.2 Sample Selection Issues

There are broadly three ways to source data on PE – from investors’ actual holdings of funds, from the funds themselves, or from those funds that sought investors. Each involves potential biases. Most analyses of PE returns have relied on data sourced from some combination of the first two – that is from the holdings of LPs or the returns reported by GPs. Harris et al. (2012) discuss the main sources of PE data. A number of important previous papers that looked at PE returns and persistence, in particular Kaplan and Schoar (2005) and Phalippou and Gottschalg (2009), used Thomson Venture Economics (TVE) data, although this data has subsequently been shown by Stucke (2011) to have serious problems. For any

data source that relies on GPs to reveal their returns, biases are likely to arise from the lower probability of revealing the past performance of funds that performed poorly. Similarly, to the extent that data from GPs is back-filled, there will be survivorship biases.

Consequently, at the fund-level, LP-sourced data has many attractions. A number of papers, such as Ljungqvist and Richardson (2003) and Robinson and Sensoy (2013a) source return data from single large individual investors. However, such data inevitably raises questions about whether the particular investor was representative in their choices and/or success. Data sourced from multiple investors is likely to be the most representative, such as that used by Harris et al. (2013a) and Harris et al. (2013b), who use Burgiss data, sourced from over 200 institutional investors. Such data is unlikely to suffer any major biases, but will inevitably have gaps in the fund sequences, reflecting investors' choices about which funds to invest in. This is less important for analysis of PE returns, but clearly more constraining when analyzing persistence.

Regarding our data, the completeness of the fund sequences is exceptional, reflecting the requirement imposed by the FoF managers for GPs to reveal all their past deals. However, there are two potential biases in our sample. First, not all funds will have sought investments from the three FoF managers from whom we derived the data. This would be a problem if funds with particular performance persistence characteristics avoided our investors. This is a theoretical possibility, as some GPs may have a bias against having FoFs as investors (preferring direct relationships with the ultimate investors). However, it is worth remembering that this data was gathered right up until the end of 2010. The balance of power shifted significantly from GPs to LPs after the financial crisis. Many investors became over-allocated to PE and so ceased making any new commitments. Consequently, fund-raising became exceptionally challenging. Beggars can't be choosers, and so we doubt whether this potential

bias is significant. This belief is reinforced by the fact that we can replicate earlier performance persistence findings using our data set.

Second, given the nature of our data, there will inevitably be some survivorship bias in the sample. Only GPs who existed and sought funds from our FoF managers can be included in the sample. Consequently, our sample will exclude GPs who had raised funds in the past (from other investors) but subsequently quit the sector – although if they tried unsuccessfully to raise a fund, they could still appear in our sample. Little is known about the extent of such attrition among GPs focused on buyouts. It is certainly true that several GPs have closed their doors since the financial crisis; however, earlier cases of GPs disappearing from the market are likely to be dominated by unsuccessful first time funds, who were unable to raise a subsequent fund. Such GPs would clearly be relevant for an analysis of the returns earned by investors, but are not relevant to a study of performance persistence.

2.1.2.3 Deals

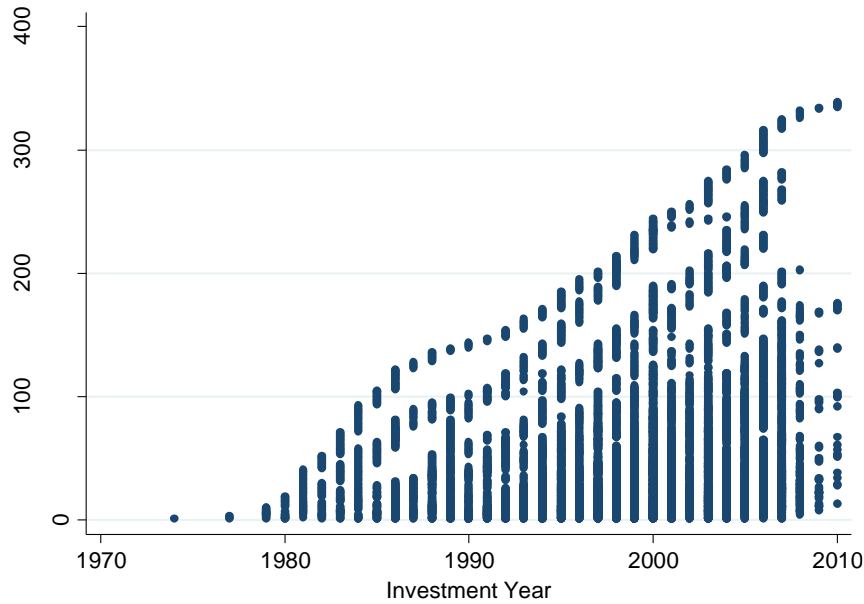
The 726 funds in our sample made 10,637 portfolio investments. At the time the data were provided, 6,600 of these investments had been fully exited, and 4,037 remained either partially or completely unrealized. The data set contains all monthly deal-level gross cash flows (before management and performance related fees) between the GP and the portfolio company as well as NAVs for all unrealized deals. This information enables us to calculate the exact gross PME, MM, and IRR for each deal. We also have full information of the industry, country, realization status, and investment date approximated by the first deal cash flow. For each fund, we know the fund type, origin, vintage year, sequence, size, and number of investments.

Figure 2-1 provides a visual overview of various characteristics of the deal-level data. Panel A displays the deal sequences per investment year for all 10,637 investments. We construct deal sequences by sorting the first to the last deal conducted within a fund family according to the deal investment date. As noted earlier, we ignore which legal fund(s) the deal was included within. The scatter in Panel A shows that our data set contains several fund families that started around 1980 and, over the subsequent three decades, have invested in more than 300 deals. New fund families are starting throughout the sample period, which is important for our research setting as we analyze separately how persistence has developed over time, as the PE sector has grown, and the effect of GP experience. Panel B plots equity investment (by the GP in our sample) per individual deal in millions of USD made by the funds over time, which shows the expected steady upward trend in deal size until the mid-2000s. This data set is the largest buyout deal-level sample yet constructed. Lopez-de-Silanes et al. (2013) construct a data set containing 7,453 deals (5,106 fully realized), although they do not have standardized cash-flow data for each portfolio investment.

Figure 2-1: Sample Distribution

This figure shows in Panel A the deal sequence number of individual deals plotted against the year in which each investment was made. Our sample covers 10,637 realized and unrealized individual transactions made between 1974 and 2010. We obtain the deal sequence variable by, first, sequencing all deals done within a certain fund family based on their investment date and, second, counting the number of deals starting with the first deal ever made within the fund family. In Panel B we show how the equity investment size (in millions of USD) of the buyout deals in our sample evolved over time (by investment years). We winsorize equity investment size at the 99th percentile.

Panel A: Deal Sequence by Investment Year



Panel B: Deal Size by Investment Year




Table 2-2 presents deal descriptives for all 10,637 deals in Panel A and for realized deals only – which are the focus of our attention – in Panels B to E. For each Panel, we provide median values of the three gross performance measures as well as deal characteristics, such as investment year, equity investment, and holding period. Panel A shows a median gross PME of 1.2, MM of 1.5, and IRR of 15.4% for the whole sample. Panel A also differentiates performance and deal characteristics in our sample by realization status. Deal returns for the 6,600 fully realized deals are substantially higher than the (interim) reported returns based on valuations. This is in line with recent findings on sticky accounting valuations of unrealized deals in the PE industry (Brown et al., 2013; Jenkinson et al., 2013).

To circumvent problems of biased accounting valuations and, thus biased returns, we exclude all (partially) unrealized deals throughout our analyses. We adjust all common investments of two funds from the same fund family. For example, if Fund III and IV of a given fund family financed the same buyout transaction, we remove the latter investment by Fund IV. By removing 405 such investments, we are sure that the remaining deals are unique transactions and investment decisions. Consequently, we present more detailed deal descriptives in Panel B to Panel E for the 6,195 realized unique deals that are the focus of our analysis.

Table 2-2: The Sample of Deals

This table reports deal-level characteristics of our sample. Generally, we provide median values for all deal characteristics. All performance measures show gross returns. First, we present deal characteristics for the 10,637 transactions differentiated by deal realization status (Panel A). A GP can either have sold all (realized) or parts (partially realized) of its shareholdings in a portfolio company. In case of some transactions, the GP still holds all of its shareholdings in the portfolio company (unrealized). Second, we display the same deal characteristics only for the 6,195 unique (so adjusted for common deals within a fund family) realized deals in our sample in four different dimensions. Panel B exhibits deal-level information within time categories (derived from Kaplan and Strömberg, 2009). Transactions are assigned to each of these time categories according to investment year. In Panel C we build groups of deals based on the region in which the deal took place. Panel D presents deal-level information by industry classification benchmarks (ICB). We only include four industries in which most of the deals have taken place. All remaining industries are summarized into “Other”. Finally, Panel E provides deal characteristics by deal sequence categories. Within each fund family we sorted investments from the first to the most recent transaction based on the dates of investment. Deal sequence counts these investments for each fund family starting with one for the first investment ever made within a fund family.

	Median						
	Obs.	PME	MM	IRR	Investment Year	Equity Investment	Holding Period
Panel A: Status Categories							
Realized	6,600	1.5	2.0	23.0%	1997	8.4	4.3
Partially	2,063	1.3	1.6	16.7%	2004	29.3	-
Unrealized	1,974	1.0	1.0	0.0%	2006	29.8	-
Total All Deals	10,637	1.2	1.5	15.4%	2000	13.9	4.3
Only Realized Deals (Adjusted for Common Deals within a Fund Family)							
Panel B: Time Categories							
1974-1989	807	0.9	1.7	13.2%	1986	1.7	5.6
1990-1994	1,177	1.4	2.3	25.3%	1992	4.8	4.8
1995-1999	2,116	1.5	1.9	20.1%	1997	9.8	4.8
2000-2004	1,796	1.6	2.0	28.1%	2001	15.1	3.7
2005-2010	299	1.7	2.0	49.1%	2005	18.7	2.2
Panel C: Region Categories							
Asia	132	1.6	1.7	19.9%	2001	12.6	3.8
Europe	2,988	1.2	2.0	25.7%	1998	7.2	4.0
North America	2,859	1.4	2.0	20.4%	1996	9.5	4.5
Other	216	1.3	1.6	19.4%	1998	16.7	4.0
Panel D: Industry Categories (Only Major Industries)							
Industrials	1,572	1.6	2.1	25.3%	1997	8.3	4.3
Consumer goods	920	1.4	2.0	24.4%	1997	9.1	4.3
Consumer services	855	1.6	2.1	25.3%	1998	13.6	4.0
Technology	891	1.1	1.5	10.6%	1998	4.9	4.3
Other	1,957	1.5	2.0	24.4%	1997	8.9	4.0
Panel E: Sequence Categories							
<= 20	2,803	1.6	2.1	26.7%	1997	7.4	4.2
21-40	1,328	1.4	1.8	19.8%	1997	8.5	4.3
41-60	773	1.4	2.0	21.3%	1998	9.6	4.3
61-80	465	1.4	2.0	20.2%	1998	7.6	4.0
81-100	288	1.3	1.7	14.8%	1998	7.6	4.1
101-150	378	1.7	2.0	22.7%	1999	13.6	3.8
>150	160	1.6	1.9	17.8%	2000	18.7	4.3
Total Only Realized Deals	6,195	1.5	2.0	23.2%	1997	8.4	4.2

Panel B provides the distribution of deals over time based on categories introduced by Kaplan and Strömberg (2009). Our total sample shows a similar distribution compared to their study; however, for the 6,195 realized deals we find a natural bias towards earlier years. The returns vary strongly for the different time categories, which is in line with the respective boom and bust periods. Yet, the last time category covering investments from 2005 until 2010, including 299 realized deals, indicates a potential bias. The time sensitive return measure IRR is very high with 49.1% and the median deal holding period amounts to only 2.2 years. Potentially, the last time category might mainly consist of the best deals of the most recent funds. GPs often quickly exit the most successful deals within their current portfolio in order to support their fundraising activities (Braun and Schmidt, 2014). Thus for each GP, the last realized deals we observe might be relatively good deals with high returns. In our econometric analysis we account for this potential bias by controlling for holding period and by conducting various robustness checks.

Panel C presents the regional distribution of deals and shows a nearly even split between deals conducted in Europe and North America. Panel D shows a selection of the major industries in which the targets are active based on ten industry classification benchmarks (ICB). All six remaining industries as well as some deals for which the industry is unknown are summarized into the category “Other”. We use the four region categories as well as eleven industry categories as control variables when conducting regression analysis in section 3. Panel E displays the 6,195 deals by deal sequence categories. The variation between sequence categories in median deal returns, investment years, equity investments, and holding period is quite low, which underlines that our data set consists of new GP fund families starting business nearly equally along different vintage years.

2.1.2.4 Synthetic Funds

In order to analyze performance persistence we build synthetic funds of 8 consecutive unique realized deals within a fund family. We do not analyze performance persistence at individual deal-level as the large amount of idiosyncratic volatility in deal returns would hamper the identification of clear patterns. To build the synthetic funds, we employ the deal sequence number of all unique realized investments within a fund family. The first synthetic fund (fund sequence number 1) consists of realized deals 1 to 8, fund 2 by deals 9 to 16, and so on. At the end of the deal sequence, we follow one of two procedures: First, if the number of remaining deals is 5 or higher, we build another synthetic portfolio. For example, if a fund family contains 45 realized investments, we build 5 synthetic funds from the first 40 investments and a 6th fund contains the 41th to 45th deals. Second, if the number of remaining deals is 4 or lower, they are excluded from the analysis.

The choice of how many deals to include in our synthetic portfolios is a matter of judgment. At a practical level, we want to measure performance at a medium frequency – so that we have enough observations to observe trends over time – while smoothing out some of the inevitable idiosyncratic volatility at the deal-level. We take 8 investments to build a synthetic fund as it is the average number of deals that a GP in our sample does through a fund before starting to invest through a next fund. GPs raise a new fund every 2 to 5 years (Gompers and Lerner, 2000). On average, GPs in our sample invest every 131 days (resulting in an average number of deals per fund of 14). Assuming a time of about 3 years in-between two funds, a GP does approximately 8 deals within this period of time.

We also compute the (equity investment size weighted) average investment date for each synthetic fund. For example, the unweighted average investment date of two investments undertaken at January 1st 1993 and January 1st 1994, respectively, is July 1st 1993. If, for

example, the first deal's equity investment size is double the one from the more recent deal, the size weighted investment date changes to May 1st 1993. This procedure enables us to assign a synthetic vintage year to each synthetic fund that indicates when most of the money was invested. This increases the accuracy of return benchmarking in contrast to simple fund vintage years.

Table 2-3 presents the distribution of all synthetic funds by vintage year and fund sequence. Overall, the 6,195 realized deals translate into 755 synthetic funds. These 755 synthetic funds exhibit synthetic vintage years ranging up to 2006. As expected, the number of funds per fund sequence number decreases with increasing sequence numbers. For example, our sample contains 195 fund families with at least one fund and 137 fund families with at least two funds and so on. Next, we focus on the 560 funds for which a previous synthetic fund exist, which we define as a fund sequence. The difference of 195 between 755 synthetic funds and 560 fund sequences originates from the nature of sequencing as, by definition, all 195 first time funds do not exhibit a previous fund. The 560 fund sequences, in which a previous fund is related to the current fund, are the basis for our analysis regarding performance persistence. Consequently, we only present median values of aggregated deal characteristics for these 560 fund sequences, which logically show similar patterns as the deal-level descriptives from Table 2-2.

Table 2-3: Distribution and Characteristics of Synthetic Funds

This table reports the distribution and characteristics of synthetic funds in two dimensions: synthetic vintage years (left side) and fund sequences (right side). For both dimensions, the number of funds is presented. Overall, the data set contains 755 synthetic funds consisting of five to eight fully realized deals. Out of these 755 synthetic funds, 560 funds have a previous fund; thus, the data set comprises 560 fund sequences. Furthermore, for both dimensions and their respective categories, median values of the two fund characteristics average deal equity investment (in millions of USD) and average deal holding period (in years) per fund are presented (for the 560 funds for which a previous fund can be observed).

Synthetic Vintage Year	Number of Funds		Median per Vintage Year (for 560 Fund Sequences)			Number of Funds		Median per Fund Sequence (for 560 Fund Sequences)	
	All	With previous Fund	Average Equity Investment by a Fund	Average Holding Period in a Fund	Fund Sequence Number	All	With previous Fund	Average Equity Investment by a Fund	Average Holding Period in a Fund
<1982	8	4	1.3	5.5	1	195	0		
1982	5	4	1.7	6.6	2	137	137	9.6	4.6
1983	8	6	1.4	5.5	3	92	92	10.2	4.6
1984	9	8	1.7	5.1	4	65	65	11.3	5.1
1985	7	4	2.2	5.8	5	50	50	12.5	5.1
1986	20	14	2.4	6.5	6	41	41	13.3	4.8
1987	9	7	1.7	6.3	7	33	33	13.4	4.6
1988	15	12	2.0	6.2	8	29	29	20.5	4.2
1989	16	13	3.7	5.4	9	18	18	6.3	4.1
1990	14	7	3.1	5.4	10	13	13	9.2	3.9
1991	32	15	3.7	5.7	11	9	9	6.1	4.9
1992	22	19	4.7	5.6	12	9	9	10.3	5.2
1993	31	23	5.2	5.1	13	9	9	8.9	4.8
1994	41	26	7.7	5.3	14	9	9	13.8	4.3
1995	42	33	8.2	4.5	15	8	8	15.3	5.4
1996	52	38	11.6	4.8	16	5	5	16.1	4.7
1997	53	39	13.2	5.1	17	5	5	22.8	4.2
1998	55	44	21.8	5.4	18	5	5	20.4	4.8
1999	64	50	20.7	4.8	19	5	5	36.1	4.1
2000	75	58	14.5	4.6	20	4	4	36.9	3.5
2001	36	22	14.9	4.2	>20	14	14	17.3	4.8
2002	46	37	22.0	3.8					
2003	45	34	25.3	3.3					
2004	29	26	24.7	3.2					
2005	13	10	33.7	2.9					
2006	8	7	7.8	2.5					
Total	755	560	11.8	4.7		755	560	11.8	4.7

In the next section we analyze whether the past predicts future performance in absolute terms, before turning to relative performance in section 2.1.4. As LPs are particularly interested in persistent superior performance, we then, in section 2.1.5, take a closer look at top quartile funds.

2.1.3 Absolute Performance Persistence

In this study, absolute fund returns are calculated from deal-level monthly cash flows between GPs and the portfolio companies, gross of fees and carried interest. As all deals in our sample are realized, returns stem from real cash flows only. We aggregate the 8 investment cash flow streams within each synthetic fund. This information is used to compute three absolute performance measures common in the PE asset class: (1) PME, which is the ratio between the present value of all cash flows to the GP from the fund portfolio divided by the present value of an alternative strategy to invest the same initial amount in a stock index. We use the three regional MSCI Performance Indices (Asia, Europe, North America) in local currency as benchmarks for our calculations and use the regional index where the GP is located; (2) MM, which is the ratio of the sum of all positive cash flows and the absolute value of the sum of all negative cash flows; and (3) IRR, which is the discount rate at which the net present value of all negative cash flows equals the net present value of all positive cash flows (from a GP's perspective). Our primary performance measure is the PME, as it indicates investment performance relative to public markets over the same period of time.

In column 1 of Table 2-4 we report results of regressing absolute PME performance of the current fund (t_0) on lagged absolute performance ($t-1$) for all 560 synthetic funds for which we observe a previous fund within the same fund family. Synthetic fund vintage years in this sample range from 1981 to 2006. We control for time (fund vintage years), industry (ten ICB industry categories), and region (share of capital invested in Asia, Europe, North America, and Other). Standard errors are adjusted for heteroskedasticity and clustered at the fund family level. The coefficient on lagged PME is positive and strongly significant. In economic terms, if the previous fund within a fund family has a 1% higher PME, the current fund returns a 17 basis points higher PME.

Table 2-4: Absolute Performance Persistence of PME

This table presents the results from ordinary least squares (OLS) regressions of gross Public Market Equivalents (PME) performance of the current fund (t_0) on lagged PME performance and various other controls. Independent variables are gross PME of the previous fund within the same fund family ($t-1$) and the fund before the previous one ($t-2$). The most comprehensive sample for this analysis contains 560 synthetic fund portfolios having a previous fund in the fund family. Columns 4 to 6 exhibit regressions on the subsample of 316 synthetic funds with an (equity investment size weighted) average investment date before December 31st 1998. Columns 7 to 9 show regressions based on the subsample of 244 synthetic funds invested after this date. In all regressions we control for synthetic fund vintage year, industry, and region. Synthetic fund vintage year fixed effects comprise the 26 years displayed in Table 2-3. Industry controls are based on ten basic ICB codes ranging from 1 (Oil & Gas) to 9000 (Technology). For deals with an unreported/unknown industry we form another category "unknown". Region controls consist of the four categories shown in Table 2-2 in Panel C. Industry and region controls are built for each synthetic fund by calculating the equity investment size weighted share for each category within a fund. For example, if 20% of the total equity of a synthetic fund are invested in Europe, the region control for Europe amounts to 0.2. In some regressions we also control for the log average equity investment size in millions of USD within the synthetic fund and the log average holding period in years. Standard errors are clustered at the fund family level and exhibited in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels respectively.

Variables	Gross PME								
	All Funds			Early PE Market (1981-1998)			Late PE Market (1999-2006)		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
PME _{t-1}	0.167*** (0.042)	0.107** (0.048)	0.110** (0.047)	0.222*** (0.056)	0.157*** (0.052)	0.155*** (0.053)	0.083 (0.058)	0.027 (0.065)	0.027 (0.066)
PME _{t-2}		0.047 (0.063)	0.041 (0.064)		0.029 (0.101)	0.031 (0.104)		0.040 (0.070)	0.032 (0.070)
LN Average Equity Investment			0.057 (0.067)			-0.004 (0.075)			0.119 (0.099)
LN Average Holding Period			0.208 (0.236)			-0.065 (0.369)			0.123 (0.348)
Synthetic Vintage Year Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes	yes
Industry Controls	yes	yes	yes	yes	yes	yes	yes	yes	yes
Region Controls	yes	yes	yes	yes	yes	yes	yes	yes	yes
Constant	1.562*** (0.250)	1.571*** (0.340)	1.473** (0.566)	2.396*** (0.363)	2.279*** (0.484)	2.376*** (0.729)	1.466*** (0.311)	1.535*** (0.406)	1.775** (0.723)
Observations	560	423	423	316	238	238	244	185	185
Adj. R-squared	0.132	0.172	0.176	0.192	0.259	0.259	0.157	0.266	0.276

In the second column we also include the PME of the fund before the previous one (t-2) within the same fund family. The coefficient on the previous fund (t-1) is still significant but the effect is reduced from 17 to 11 basis points. The coefficient on the fund before the previous fund (t-2) within the same fund family is weaker (0.047) and statistically insignificant. Together, the combined effect of a 1% higher past performance (in the two previous funds) is about 16 basis points. Hence, the performance of the fund before the previous fund does not add much information about current fund performance. In the third column, we test the robustness of this empirical pattern if controlling for the effects of average equity investment size and average holding period (both logarithmized) within the synthetic fund. The coefficients on lagged absolute performance remain basically unchanged in statistical and economic terms.

In order to replicate the analysis of Kaplan and Schoar (2005) and Sensoy et al. (2013), we split our sample of 560 synthetic funds for which we observe a previous fund at an investment weighted average investment date of December 31st 1998.²¹ In the fourth column of Table 2-4 we report regression results for this early PE market subsample of 316 fund sequences. The relation between previous and current fund PME is 0.22, that is 22 basis points for a 1% higher previous PME, and is statistically significant. This effect is slightly higher than the 0.17 reported by Kaplan and Schoar (2005). Column 5 shows the same regression for 238 funds for which we are also able to include the PME of the fund before the previous one. Here, our coefficient is around 0.03 and insignificant; Kaplan and Schoar (2005) also found a small and insignificant effect of the 2nd prior fund. Column 6 shows robustness of these results to the inclusion of average investment size and holding period. The

²¹ Kaplan and Schoar (2005) include mature funds up to a vintage year of 1996. Hence, the most recent funds in their sample started investing in 1996. As we use deal-level investment dates (equity investment weighted), it is reasonable to assume that these 1996 funds would have vintage years up to 1998 according to our (more accurate) procedure.

results from regressions 4 to 6 provide strong support for the findings and conclusions presented by Kaplan and Schoar (2005) for the buyout market until the late-1990s, despite our use of synthetic portfolios based on realized returns at the gross level as well, no doubt, as differences in the sample of underlying funds.

The results for the earlier period, compared with the full sample, suggest that persistence has decreased in more recent years. To test this, columns 7 to 9 of Table 2-4 exhibit the same specifications for our sample of 244 funds with a prior fund and a vintage year later than December 31st 1998. When regressing current PME on lagged PME (column 7), the effect is only 8 basis points for a 1% higher previous PME and statistically insignificant. The coefficient on the PME of the 2nd previous fund is 4 basis points and also insignificant (columns 8 and 9). These results show a clear pattern of decreased absolute performance persistence for GPs in more recent years. Based on synthetic funds containing investments later than 1998, when the Kaplan and Schoar (2005) sample ends, we no longer find any evidence of absolute buyout performance persistence.²²

2.1.4 Relative Performance Persistence

Investors in PE are not only interested in prior absolute performance, which is strongly influenced by factors like the macroeconomic environment, but also by the ability of a GP to persistently outperform their peers. A common way in which such relative performance is measured is by ranking funds according to their performance quartiles. Indeed, given the multiplicity of data providers, market definitions, vagaries over the definition of vintage years

²² Very similar results are found using MM instead of PME. The pattern of results using IRR is also similar, but statistically less significant in the earlier period.

and using a judicious choice of sample periods, it is often joked that all PE funds are (or at least claim to be) “top quartile”.²³

Using our detailed data, we are able to measure relative performance in a simple and precise way. We sort all 755 synthetic funds into vintage year performance quartiles, using our vintage years derived from the pattern of investment (as described earlier). Table 2-5 exhibits conditional probabilities that the current fund stays in the same performance quartile as the previous fund for our three measures of performance. For example, the probability that a GP will produce top quartile performance, according to PME, for two successive synthetic funds is 34.3% (Panel A). Irrespective of the performance measure, these transition matrices draw a coherent picture. In line with Kaplan and Schoar (2005) they show that persistence is particularly strong in the top and the flop quartile. The conditional probabilities at both ends of the performance dichotomy are in a range of 32.3% to 36.8% across all performance measures. Measured by PME, if the previous fund is top quartile the probability that the current fund’s return is above average is 56.7% (34.3% + 22.4%). At the negative end, if a fund was in the flop quartile the probability of below average subsequent performance is 63.1% (27.4% + 35.7%). Interestingly, the probability of moving from flop to top quartile is low (ranging from 12.7 to 16.5 percent) whereas the probability of moving from top to flop is noticeably higher (ranging from 20.8 to 23.9 percent).

²³ Harris et al. (2012) analyze some of the degrees of freedom GPs exploit when presenting their relative performance.

Table 2-5: Transition Matrices

This table presents the conditional probabilities that the current synthetic fund (t0) either stays in the same performance quartile as the previous one (t-1) within the fund family or moves into another quartile. We include only 560 funds for which we observe a previous fund in the fund family. Fund performance quartiles are built across all funds of a given fund vintage year, for example, we rank all funds of vintage year 1999 according to PME performance and build four quartiles, containing top funds, above average funds, below average funds and flop funds, respectively. Panel A, B, and C show the conditional Markov probabilities for the different quartile rankings by the gross return measures Public Market Equivalent (PME), Money Multiple (MM), and Internal Rate of Return (IRR).

Panel A: PME

		Quartile Fund _{t0} (%)			
		Top	Above Average	Below Average	Flop
Quartile Fund _{t-1}	Top	34.3	22.4	20.2	23.1
	Above Average	32.1	23.9	22.4	21.6
	Below Average	20.0	28.9	25.2	25.9
	Flop	12.7	24.2	27.4	35.7

Panel B: MM

		Quartile Fund _{t0} (%)			
		Top	Above Average	Below Average	Flop
Quartile Fund _{t-1}	Top	33.9	22.3	20.0	23.9
	Above Average	29.9	23.4	23.4	23.4
	Below Average	19.9	29.1	28.4	22.7
	Flop	16.5	18.4	29.0	36.2

Panel C: IRR

		Quartile Fund _{t0} (%)			
		Top	Above Average	Below Average	Flop
Quartile Fund _{t-1}	Top	32.3	25.4	21.5	20.8
	Above Average	28.2	21.1	24.7	26.1
	Below Average	21.3	25.7	28.7	24.3
	Flop	15.1	24.3	23.7	36.8

However, as with absolute performance persistence, we find lower performance persistence than Kaplan and Schoar (2005) in our entire sample of funds from 1981 to 2006. For example, for their sample of venture capital and buyout funds Kaplan and Schoar (2005) report a probability of top tercile PME performance persistence of 55%, which is 67% higher than under a random walk assumption $((55\% - 33\%) / 33\% = 67\%)$. In our sample, the top

quartile PME performance persistence probability is about 34%. This conditional probability is only 36% higher than under a random walk assumption $((34\% - 25\%) / 25\% = 36\%)$.

We further investigate relative performance persistence in a multivariate setting and estimate logit regressions of a binary relative persistence variable on several explanatory variables and controls. The results for the PME are presented in Table 2-6.²⁴ The relative performance persistence variable adopts a value of 1 if the previous and the current fund are in the same performance quartile and 0 otherwise. In the first column, we test whether the differences in persistence across quartiles (reported in Table 2-5) hold in a multivariate setting. Therefore, we include dummy variables for the previous fund's performance quartile. The reference group is the above average quartile. The coefficients confirm that persistence in the top and flop quartiles is statistically stronger than in the above average quartile.

In the second column of Table 2-6 we test, using the whole sample of 560 funds with a prior fund again, whether the decrease in absolute performance persistence in more recent years also holds for relative performance. To this end, we include a binary time variable that adopts a value of 0 for early synthetic funds with a vintage year before December 31st 1998 and a value of 1 for later funds. The coefficient on the time dummy is negative but statistically insignificant.

²⁴ All controls and standard errors are identical to those in the absolute performance persistence regressions reported in Table 2-4.

Table 2-6: Fund Performance Quartiles Persistence

This table presents the results from logit regressions of Public Market Equivalents (PME) performance persistence on a binary time, a binary fund sequence variable, and various other controls. The dependent variable of performance persistence adopts a value of one if the current fund (t0) is in the same performance quartile as the previous one (t-1) within the fund family, and zero otherwise. We include 560 synthetic funds that have a previous fund in the fund family. The time dummy splits this sample of funds of the (equity investment size weighted) average investment dates into early time, indicating funds invested before December 31st 1998, and late time all funds invested thereafter. The fund sequence dummy splits this sample of relevant funds at the median fund sequence. Accordingly, low fund sequence indicates fund sequences below 5, while high represents fund sequences of 5 and higher. We run separate regressions on the subsample of 314 early vintage year funds (column 5) and 244 late vintage year funds (column 6). In all regressions we control for the performance quartile in which the previous fund has been (reference category is the above average quartile), the log of the average equity investment size in millions of USD within the synthetic fund, the log of the average holding period in years, and industry as well as region controls. Standard errors are clustered at the fund family level and exhibited in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels respectively.

Variables	Dummy: Change in (PME) Quartiles from Fund _{t-1} to Fund _{t0} (no change=1)					
	All				Time	
	(1)	(2)	(3)	(4)	Early 1981-1998 (5)	Late 1999-2006 (6)
Time Dummy (early=0; late=1)		-0.377 (0.233)		-0.419* (0.233)		
Fund Sequence Dummy (low=0; high=1)			0.539** (0.210)	0.561*** (0.209)	0.603** (0.293)	0.386 (0.323)
Fund _{t-1} Quartile (ref: above average)						
Top Quartile	0.575* (0.329)	0.601* (0.333)	0.578* (0.326)	0.606* (0.328)	1.045*** (0.380)	-0.111 (0.437)
Below Average Quartile	0.130 (0.267)	0.144 (0.271)	0.136 (0.262)	0.150 (0.265)	0.254 (0.351)	-0.025 (0.421)
Flop Quartile	0.629* (0.324)	0.643** (0.327)	0.618** (0.314)	0.634** (0.317)	0.982*** (0.360)	0.145 (0.479)
LN Average Equity Investment	-0.083 (0.080)	-0.040 (0.087)	-0.123 (0.081)	-0.076 (0.088)	-0.095 (0.104)	-0.059 (0.157)
LN Average Holding Period	-0.125 (0.305)	-0.341 (0.358)	-0.143 (0.303)	-0.377 (0.352)	-0.653 (0.528)	-0.120 (0.538)
Industry Controls	yes	yes	yes	yes	yes	yes
Region Controls	yes	yes	yes	yes	yes	yes
Constant	-1.320* (0.726)	-0.679 (0.868)	-1.637** (0.744)	-0.951 (0.864)	-0.555 (1.255)	-1.362 (1.109)
Observations	560	560	560	560	314	244
McFadden's Pseudo-R-squared	0.022	0.026	0.033	0.037	0.080	0.028

Next, we test whether the probability of relative performance persistence is different between inexperienced and experienced GPs.²⁵ We construct a fund sequence dummy that splits the sample at the median synthetic fund sequence of 4. A value of 0 (low fund sequence) indicates that the current fund has a fund sequence below 5. Accordingly, an inexperienced GP has done up to 32 investments. The fund sequence dummy adopts a value of 1 (high fund sequence) if the current synthetic fund has a sequence within the fund family of 5 or higher. Column 3 in Table 2-6 shows that the coefficient on the fund sequence dummy is positive and significant.

The relation remains statistically significant and at the same level of economic relevance when adding the time dummy to the regression (column 4). The marginal effect of the fund sequence dummy on relative performance persistence across all performance quartiles is 11.4%-points.²⁶ In other words, based on our entire time period (1981-2006) the probability of staying in the same quartile as the previous fund is 11.4%-points higher for synthetic fund sequences of 5 or higher (36.1%) than for those below 5 (24.7%). When controlling for GP experience, the time dummy becomes slightly stronger and statistically significant as well. In economic terms, the marginal effect of this time dummy on the probability of relative performance persistence in the fourth column is -8.3%-points (estimated conditional probabilities of being persistent are 33.7% prior to 1999 and 25.4% thereafter). This means that, *ceteris paribus*, relative performance persistence across all performance quartiles is 8.3%-points lower after 1998 and has fallen to a random walk level (of about 25%).

²⁵ We use fund generation as a proxy for a GP's experience as in the study of Kaplan and Schoar (2005).

²⁶ Throughout this analysis, we report average effects for estimated conditional probabilities and marginal effects obtained by using Stata's margins command.

Since the results for the later years appear very different to the early period we also estimate separate models for early funds (column 5) and late funds (column 6). This allows us to test whether the relation between GP experience and relative performance persistence has changed over time. The coefficient on the fund sequence dummy is only significant and much stronger in the early phase before 1999. The estimations yield a 12.0%-points higher probability of relative persistence for high fund sequences in the early period (26.9% vs. 38.9%) and 7.4%-points in the later period after 1998 (23.6% vs. 31.0%). This suggests that the positive effect of GP experience on relative performance persistence across all quartiles was much stronger before 1999. GP experience appears to have mattered much more in the days before the PE industry matured.

Another empirical pattern that these regressions in Table 2-6 reveal is that the heterogeneity of relative performance persistence among quartiles in our overall sample is driven by the early vintage years. The coefficients on the top and flop quartile dummies are (in line with the pattern shown in Table 2-5) strong and significant before 1999 (column 5). Both relations become insignificant and decrease to lower levels in the later time period (column 6). This shows that relative performance persistence, in particular in the top quartile, has substantially decreased in more recent years and is no longer statistically different among performance quartiles.

These results directly raise further questions related to persistent outperformance in the PE asset class and LPs' fund investment decisions: If persistence in the top quartile is not statistically stronger than in the other quartiles anymore, has it totally disappeared? If GP experience is not much associated with relative performance persistence per se anymore, does this also hold for outperformance, that is the top quartile? Or are established GPs able to defy the general trend of weakened persistence?

2.1.5 Top Quartile Performance Persistence

To answer these questions we run several logit regressions to explore the determinants of top quartile performance. In Table 2-7 the dependent variable is a top quartile performance dummy that adopts a value of 1 if the current fund (t_0) returns are top quartile among all synthetic funds of the same vintage year and 0 otherwise.²⁷ All controls and standard errors are identical to those in Table 2-4 and Table 2-6.

In the first column we report results of regressing the top quartile dummy (t_0) on an equivalent top quartile dummy for the previous fund ($t-1$) within the same fund family. The coefficient on lagged relative top quartile performance is strongly positive and statistically significant. The estimated marginal effect is 10.7%-points and indicates that the probability of being top quartile with the current fund in the entire sample period is 10.7%-points higher if the previous fund was also top quartile. The relation between current and lagged relative outperformance remains basically unchanged if controlling for time and fund sequence as introduced in the previous section (column 2). The probability of staying in the top quartile – as estimated within this multivariate model – in the period between 1981 and 2006 is 32.4%, which is close to the value of 34.3% from the simple transition matrix in Panel A of Table 2-5.

²⁷ In this set-up the dependent variable is conceptually slightly different from the one that we used to examine relative performance persistence per se in the previous section. Now, the binary dependent variable just indicates top quartile persistence of the current (t_0) fund, while in the previous section the relative performance variable was made up by performance of the prior ($t-1$) and the current (t_0) fund. The set-up used for top quartile persistence allows us to interact prior relative performance with time and GP experience.

Table 2-7: Fund Top Performance Quartile Persistence

This table presents the results from logit regressions of Public Market Equivalents (PME) top quartile performance with the current fund (t0) on top quartile performance with the previous fund (t-1), a binary time variable, a binary fund sequence variable, and various other controls. The dependent variable of top quartile performance with the current fund adopts a value of one if the current fund is top quartile and zero otherwise. We include 560 synthetic funds that have a previous fund in the fund family. The time dummy splits this sample of funds into early time, indicating funds invested before December 31st 1998, and late time all funds invested thereafter. The fund sequence dummy splits this sample of relevant funds at the median fund sequence. Accordingly, low fund sequence indicates fund sequences below 5, while high represents fund sequences of 5 and higher. We run separate regressions on the subsample of 314 early vintage year funds (column 3) and 244 late vintage year funds (column 4) as well as on the subsample of low and high fund sequence containing 294, respectively, 266 funds each (columns 5 and 6). In column 7 we also include a three-way interaction term of these variables. In all regressions we control for the log of the average equity investment size in millions of USD within the synthetic fund, the log of the average holding period in years, and industry as well as region controls. Standard errors are clustered at the fund family level and exhibited in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels respectively.

Variables	Dummy: Top Quartile (PME) in Fund _{t0} (top=1)						
	All		Time		Fund Sequence		All
	(1)	(2)	Early 1981-1998 (3)	Late 1999-2006 (4)	Low 1-4 (5)	High >4 (6)	(7)
Fund _{t-1} Top Quartile Dummy	0.562** (0.252)	0.568** (0.250)	0.824*** (0.277)	0.085 (0.376)	0.342 (0.321)	0.667* (0.391)	0.351 (0.451)
Time Dummy (early=0; late=1)		-0.147 (0.254)			0.070 (0.311)	-0.281 (0.397)	0.017 (0.360)
Fund Sequence Dummy (low=0; high=1)		0.100 (0.220)	0.115 (0.324)	0.003 (0.286)			-0.123 (0.317)
Fund _{t-1} Top Quartile Dummy*Time Dummy							-0.055 (0.647)
Fund _{t-1} Top Quartile Dummy*Fund Sequence Dummy							1.198* (0.646)
Time Dummy*Fund Sequence Dummy							0.210 (0.476)
Fund _{t-1} Top Quartile Dummy*Time Dummy*Fund Sequence Dummy							-1.652* (0.852)
LN Average Equity Investment	-0.007 (0.091)	0.003 (0.096)	-0.003 (0.111)	0.033 (0.166)	-0.115 (0.115)	0.176 (0.178)	0.002 (0.095)
LN Average Holding Period	-0.111 (0.319)	-0.192 (0.353)	-0.255 (0.544)	0.039 (0.573)	-0.107 (0.471)	-0.366 (0.494)	-0.125 (0.355)
Industry Controls	yes	yes	yes	yes	yes	yes	yes
Region Controls	yes	yes	yes	yes	yes	yes	yes
Constant	-0.657 (0.726)	-0.467 (0.857)	-0.197 (1.155)	-0.988 (1.241)	-0.934 (1.030)	-0.026 (1.419)	-0.567 (0.858)
Observations	560	560	314	244	294	266	560
McFadden's Pseudo-R-squared	0.043	0.038	0.102	0.060	0.071	0.071	0.050

To test the pattern that relative outperformance persistence changes over time, in columns 3 and 4 we run separate regressions for the early and the late periods, respectively. In column 3, for the early vintage year funds, the coefficient on the top quartile dummy for the previous fund is 0.824 and statistically significant. Therefore, for funds before 1999 the estimated conditional probability that the current fund is top quartile when the previous fund

was top quartile is 36.0%. For the later vintage year synthetic funds (column 4), we obtain an insignificant coefficient of only 0.085 on lagged relative top performance. The estimated probability of staying top quartile in the later period has decreased to 25.2%. These results confirm that top quartile relative performance persistence substantially decreased as the PE industry matured.

In columns 5 and 6 of Table 2-7 we assess whether top quartile persistence is contingent on GP experience. These regressions on the low and high fund sequence subsamples both yield positive coefficients on lagged top quartile performance. However, only in the high experience sample is the coefficient statistically significant (and then only at the 10% level). The estimated conditional probability of persistent relative outperformance is 28.0% in the low and 34.3% in the high fund sequence subsample.

Altogether, the results reported in columns 1 to 6 in Table 2-7 reveal that persistent top-quartile performance among buyout funds existed in the early years, but has disappeared as the market has matured. At the same time top quartile persistence in the entire time period is somewhat stronger for experienced GPs. These findings raise the question whether GP experience can defy the general market trends. In other words: Do experienced GPs remain persistent outperformers?

We test this aspect by building a three-way interaction model reported in column 7 of Table 2-7. Besides the two interaction terms of the top quartile dummy for the previous fund and the time and fund sequence dummy, respectively, we include a three-way interaction term of all three variables. As (three-way) interaction terms in non-linear models are very difficult to interpret, we compute the probabilities of current relative outperformance contingent on past relative outperformance over time and fund sequence. Table 2-8 displays these

probabilities. While the interpretation as conditional probabilities is identical to the numbers presented in Table 2-5, the values in Table 2-8 are estimated and result from four separate estimations of model 7 in Table 2-7 for all four performance quartiles (columns 1 to 4 in Table 2-8).²⁸ We also repeat this procedure for PME (Panel A), MM (Panel B), and IRR (Panel C).

The interpretation of the black numbers in Table 2-8 can be illustrated by considering the top left-hand block. The probability of staying top quartile measured by PME for experienced GPs in the early PE market is 52% (Panel A, column 1). Furthermore, for each model we estimate the linear combinations of the GP experience, the time, and the previous fund top quartile performance coefficients. Thereby, we are able to compute the differences (grey numbers in Table 2-8) and obtain standard errors as well as significance levels. For example, the fall in the estimated probability of staying top quartile for experienced GPs in the later period (compared with the early period) is 32%-points and highly statistically significant.

²⁸ Please note that Table 2-8 only displays the estimated conditional probabilities of staying in the same performance quartile. This corresponds to the values in the diagonals of the matrices shown in Table 2-5. As the probabilities of migrating to another performance quartile are omitted in Table 2-8, the probabilities of each time and experience pair (e.g., early & low funds) across all quartiles do not, and do not need to, sum up to one.

Table 2-8: Probability of Quartile Persistence (Current and Next Fund)

This table presents the predicted probabilities of being in the same quartile with the current (t0) and the previous fund (t-1) within a fund family based on the regression set-up shown in model 7 of Table 2-7 for top quartile fund performance persistence (column 1) as well as unreported, identical regressions for the quartiles above average (column 2), below average (column 3), and flop (column 4). Column 5 displays average probabilities across all quartiles, column 6 compares average values of all quartiles over time and funds sequence. Panel A, B, and C report performance quartile persistence based on Public Market Equivalent (PME), Money Multiple (MM), and Internal Rate of Return (IRR), respectively. For each quartile and performance measure we display a two by two matrix containing the average predicted probabilities of being in the same quartile with the current and previous fund (black numbers) over time (early vs. late, i.e., before December 31th 1998 vs. after) and fund sequence (low vs. high, i.e., fund sequence below 5 vs. 5 & beyond). All predictions of the point estimations are highly significant. Furthermore, we show the differences in the point estimates for these four scenarios (grey numbers). *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels respectively.

QUARTILES in t-1 & t0:		1. Top			2. Above Average			3. Below Average			4. Flop			5. Averages Across All Quartiles			6. Comparison Along Time and Fund Sequence	
		Time												Time			Time	
		Early	Δ	Late	Early	Δ	Late	Early	Δ	Late	Early	Δ	Late	Early	Δ	Late	Early	Late
Panel A: PME																		
Fund Sequence	Low	0.28	-0.01	0.28	0.25	-0.05	0.20	0.25	-0.04	0.21	0.30	-0.09	0.21	0.27	-0.05	0.22	0.25	
	Δ	0.24 *		-0.07	-0.04		0.10	-0.03		0.08	0.18 **		0.17	0.09		0.07	0.32	0.29
	High	0.52	-0.32	0.21	0.21	0.09	0.30	0.22	0.06	0.28	0.48	-0.11	0.38	0.36	-0.07	0.29	0.33	
Panel B: MM																		
Fund Sequence	Low	0.35	-0.04	0.31	0.32	-0.20 **	0.11	0.28	-0.04	0.24	0.33	-0.16 *	0.17	0.32	-0.11	0.21	0.26	
	Δ	0.05		-0.15	-0.18 **		0.11	0.02		0.04	0.21 **		0.17	0.02		0.04	0.33	0.28
	High	0.39	**	0.16	0.14	0.09	0.22	0.29	-0.02	0.27	0.53	-0.20 *	0.34	0.34	-0.09	0.25	0.29	
Panel C: IRR																		
Fund Sequence	Low	0.29	-0.11 *	0.18	0.25	-0.11	0.13	0.25	0.08	0.33	0.33	-0.09	0.24	0.28	-0.06	0.22	0.25	
	Δ	0.16 **		0.02	0.04		0.04	0.01		-0.10	0.20 **		0.09	0.10		0.01	0.33	0.28
	High	0.45	**	0.20	0.29	-0.12 *	0.17	0.26	-0.03	0.23	0.53	-0.20 *	0.32	0.38	-0.15	0.23	0.31	

These estimated probabilities of PME top quartile persistence summarize the patterns shown in the regressions of Table 2-7. The experience-related difference in top-quartile performance persistence in the overall sample is driven by the extraordinary top quartile persistence of experienced GPs before 1999. However, the performance persistence of experienced GPs has disappeared in recent years. Performance persistence in the top quartile for less experienced GPs has been limited and stable over time. Overall, we find no evidence of top-quartile performance persistence from 1999 onwards irrespective of GP experience, as the average probability of about 24% (average of 28% for inexperienced and 21% for experienced GPs) is almost exactly at the expected value of 25% under a random walk assumption. This is in sharp contrast to the relative top quartile persistence level prior to 1999 of 40% (28% for inexperienced and 52% for experienced GPs).

This pattern of relatively stable probabilities of staying top quartile for funds managed by inexperienced GPs over time (at levels only slightly above or actually at random walk level) and a sharp drop for experienced GPs also holds for measuring performance by MM (Panel B) and IRR (Panel C).

The other results of note in Table 2-8 relate to the poorly performing GPs. Interestingly, there is a significantly higher probability of repeating bottom-quartile performance for experienced GPs. This was particularly true in the earlier period, where the difference in probability for low and high sequence funds is statistically significant using all three performance measures. The difference in bottom-quartile persistence has narrowed in recent years, but it still remains a mystery as to why such GPs survive. It may be that investors are too loyal to the GPs they have previously invested with, or do not probe their relative performance sufficiently.

Finally, in column 5 the probabilities are aggregated over quartiles, and, in column 6, by fund sequence and time. These confirm, within this multivariate setting, the main results obtained across the whole sample earlier using simpler models. We also investigate the impact of re-introducing deals that are split between successive funds, to quantify the impact on the estimated persistence. Without such deals the results in the final column of Table 2-8 indicate that the probability of staying in the same quartile, as measured by PME performance, across the whole sample is 29%. If we re-introduce deals that were split between funds – so that they appear more than once in the synthetic funds – the estimated persistence rises to 31%. This suggests that not controlling for common deals increases, as expected, the estimated persistence somewhat.

The results in this section demonstrate that while GPs who were early entrants and performed well managed to replicate their top-quartile performance for some time, whatever advantages they had evaporated as the market matured. In recent years, the probability of repeating top-quartile performance is no higher for experienced GPs, and close to random overall. Experience appears to have been valuable in the early development of the PE market, but, in common with most other asset classes, private equity has succumbed to mean-reverting tendencies. Understanding the organizational, competitive, incentive or succession challenges faced as GPs mature is beyond the scope of this study but represents a fertile area for future research.

2.1.6 Conclusion

In this study we have re-visited the performance persistence of PE fund managers. Whereas previous studies have considered performance across funds, our unique sample of

detailed cash flow data on over ten thousand buyout transactions allows us to strip away the legal wrapper of the fund and analyze performance persistence of GPs in new ways. In particular, we avoid all issues regarding unrealized investments by focusing on fully exited transactions. We also avoid problems associated with overlapping funds by constructing portfolios that are strictly non-overlapping temporally, and contain no common deals. For these synthetic portfolios we are able to construct precise vintage years based on the allocation of capital to the individual portfolio companies, which allows for more accurate benchmarking of relative performance. Using synthetic portfolios also allows us to choose the frequency at which to measure performance persistence, rather than to accept the 10+ year life of the typical fund. Our choice of a medium frequency (8 exited deals) for the portfolios generates a large sample of 755 synthetic funds in our sample; of these 560 have a prior fund.

Our main findings are as follows. First, our analysis confirms the findings of Kaplan and Schoar (2005), who found performance persistence across funds raised until the late-1990s. The similarity in our results is remarkable given that we have gross returns at the deal-level that is derived from investor due diligence, whereas Kaplan and Schoar (2005) used net returns at the fund-level derived from Thomson Venture Economics. When we extend the sample with deals that exited until 2010 we find that absolute performance persistence has disappeared in recent years. We focus in the main on performance relative to public markets (PMEs) but similar results are found for money multiples and IRRs.

Second, we consider the performance of GPs relative to other GPs who were investing at the same time, by constructing performance quartiles by synthetic fund vintage year. Using simple transition matrices, across the whole sample we find considerable top and bottom relative persistence: around 34% of top and bottom quartile funds stay in the same quartile in the next fund. We explore this further in a logistical regression framework, which allows us to

control for various deal, industry, and GP characteristics. In line with the results obtained for absolute performance persistence, we find that relative performance persistence is only apparent until the late-1990s, and has essentially disappeared for funds investing in the second half of the sample.

Third, we explore whether the disappearing performance persistence is associated simply with the industry becoming mature with the passage of time, as more managers have entered the industry, or whether the GP experience has a separate role to play in the explanation. We find strong top quartile performance persistence until the late-1990s that was driven by experienced GPs. This experience advantage has disappeared in recent years. Across GPs, the probability that a GP will repeat top quartile performance is now no better than random.

For investors this research has clear, but uncomfortable, implications. The often-heard mantra of sticking with top-performing managers is no longer a recipe for success. If experience, or past performance, provides little guidance on the choice of GP, it remains to be seen whether alternative reliable predictors can be found.

2.2 Essay 2 - The Cost of Private Equity

Abstract

Private equity (PE) has developed into a well-established asset class with strong growth in capital commitments over the last decades. Consequently, fund returns have decreased over time and investors have become more cost conscious. Based on a unique data set of 358 PE buyout funds with vintage years between 1983 and 2007, we analyze whether the maturing PE asset class has become less costly over time. We define costs as the difference between gross and net returns (return spread) and provide a spread benchmark useful for investors to evaluate a fund's costliness. Next, we show that, in line with our expectations, return spreads have decreased over time. However, when we control for falling gross returns causing lower performance-based fees, surprisingly, the cost of PE investing has increased. We relate the higher costs to increased levels of dry powder due to swelling capital flows into the industry. We conclude that the PE industry is a victim of its own success, suggesting that investors in the asset class should consider a more anti-cyclical investment approach.

Keywords: Private equity, fund return, capital deployment, investment behavior

JEL Classification Code: G11, G23, G24

Authors: Stoff, Ingo; Braun, Reiner

First Author: Stoff, Ingo

Current Status: Submitted to the Review of Financial Economics

Conferences: Accepted at Midwest Finance Association Conference 2014

2.2.1 Introduction

Private equity (PE) has developed from a niche to a well-established asset class with a record fundraising volume of USD 550 billion in 2007 (Thomson ONE). However, with increasing fund sizes, PE returns have dropped from their staggering levels of the past (Harris et al., 2013a). Obviously, PE has gone through a similar development as the related hedge fund industry, in which increasing asset allocations to the industry have driven down existing alpha returns (Fung et al., 2008). Finding attractive deals for general partners (GPs), who invest money committed by limited partner (LP) investors, and generating outstanding fund returns have become increasingly difficult. With lower returns, LPs such as pension funds and insurance companies have taken a more critical view of GP compensation.²⁹ In turn, the academic literature has started to critically examine the complex compensation structure between GPs and LPs (Gompers and Lerner, 1999; Litvak, 2009), the split between fixed and variable compensation (Chung et al., 2012; Metrick and Yasuda, 2010), and the relation between the costliness of fund terms and net performance (Robinson and Sensoy, 2013b).

Our research is the first to provide empirical evidence of the cost of the PE asset class, more specifically the buyout segment, over time and based on actual data.³⁰ We can therefore verify whether buyout fund investing has become less costly for LPs, as expected with increasing market maturity. At first glance, we indeed find that the cost of PE has dropped for fund vintages between 1983 and 2007. However, when lower gross fund returns (i.e., lower performance-based fees) are controlled for in the maturing industry, the cost of the buyout asset class has actually increased. We relate this finding to a vastly overlooked determinant of

²⁹ See Financial Times, Feb. 12th, 2012, “Investors push for private equity fee cuts.”; See The Economist, May 6th, 2010, “Private inequity.”

³⁰ This analysis focusses explicitly on later-stage buyout funds investments and neglects VC fund investments.

costs: a fund's unused callable capital ("dry powder"), driven by swelling capital flows into the asset class. Obviously, the buyout industry is a victim of its own success.

Throughout our research we define the cost of buyout fund investing as the difference between actual gross and net fund returns, the so-called return spread. We focus on the return spread because it expresses the return extracted by GPs for managing a specific fund. The spread thus displays the cost of buyout fund investing from an LP perspective. For example, the return spread of a fund generating a gross money multiple (MM) of 2.0 and a net MM of 1.7 amounts to an MM of 0.3, which translates to costs of USD 3 billion for a mega fund consisting of USD 10 billion committed capital.³¹

We first provide an overview of the development of return spreads over time. In a second step, we then analyze the impact of a fund's dry powder on its return spread. So far, the academic literature has predominantly focused on two fund-level drivers of the cost of buyout investing: fund gross returns and fund terms. In contrast, we also consider the deal-level perspective by analyzing how the investment behavior of GPs influences the return spread. We measure GP investment behavior by a fund's deployment rate, which is the ratio of all equity investments the fund makes to its size (committed capital). Thus, the deployment rate expresses the dry powder of a buyout fund, which is quite costly. When charging a 2% management fee for a USD 10 billion fund, respectively, USD 200 million in fees per year, the effective management fee increases to 4% if the GP is only able to invest 50% of the total committed capital of the fund.

³¹ We define the money multiple (MM) as the ratio of all positive to all negative cash flows. Thereby, the gross MM from a GP perspective includes all cash flows between a GP and its portfolio companies, whereas the net MM from an LP perspective uses all cash flows between LPs and GPs.

Dry powder has become an important issue, especially since 2006, because excessively large funds were raised in 2006 and 2007 that faced low Merger and Acquisition (M&A) activity after the financial crisis.³² Consequently, the global volume of callable capital reserves of buyout funds more than doubled from 2004 to 2008, reaching a record volume of USD 485 billion, remaining high in 2012 with about USD 360 billion.³³

The resulting lower fund deployment rates should increasingly drive up return spreads and, thus, lower investor net returns. Extreme cases of this issue are so-called zombie funds, which virtually stop investing and simply extract annual management fees out of their funds.³⁴ Overall, besides fund gross returns and fund terms, a fund's dry powder, now more than ever, displays a third and mainly overlooked driver of the return spread.

Lastly, this study analyzes the underlying determinants of the deployment rate to extract important implications for LPs by asking the following questions: Is the GPs' investment behavior driven by factors related to the GP or by market factors? If the former, such as GP experience or gaming for fees, drive the deployment rate, LPs should increasingly consider a GP's ability to deploy capital or the structure of compensation agreements in their GP selection. In contrast, if the underlying driver of the deployment rate is mainly rooted in external market parameters, LPs should account for these parameters to correctly adjust their return expectations and possibly reconsider their investment decisions.

³² See The Economist, Mar. 2nd, 2013, "When the music gets you."

³³ See Private Equity Trends, 2013 Q2 newsletter, Private Equity Growth Capital Council.

³⁴ See The Economist, Mar. 23rd, 2013, "Zombies at the gates."; See Financial Times, Dec. 11th, 2011, "Private equity trapped in zombie funds."

This study employs a unique and proprietary sample of 358 mainly realized buyout funds headquartered in North America, Europe, and Asia with vintage years between 1983 and 2007. Since the return spread is meaningful only for nearly realized funds, we exclude all funds younger than five years or in which less than 75% of all deals are at least partially realized. For each fund we have full knowledge of the fund gross and net returns, as well as gross cash flows between the GPs and their portfolio companies, which enables us to evaluate the GPs' investment behavior.

Our first univariate results show a mean (median) MM return spread of 0.57 (0.47) and an internal rate of return (IRR) spread of 10.3 (9.5) percentage points. To provide universal benchmarks and circumvent a potential positive selection bias, we create a simple return spread benchmark by grouping funds into gross performance groups and calculating the return spread for each group. For example, our benchmark provides an MM spread of around 0.42 for funds with a gross MM between 2.2 and 2.4 and a lower spread of 0.25 for funds with a gross MM between 1.6 and 1.8, due to the lower variable compensation of GPs.

The multivariate results show that the return spreads have decreased over time, which at first glance implies a reduction in the cost of buyout fund investing. However, this reduction is solely caused by decreasing performance-based fees due to the falling gross returns over time. When controlling for reduced gross returns, we find that the return spreads have actually increased between 1983 and 2007, indicating that it has become more expensive for LPs to invest in the buyout segment.

Next, we analyze the impact of GP investment behavior on return spreads. We find that a fund's deployment rate shows a negative relation to return spreads, which is both statistically and economically highly significant. Furthermore, our results show that falling deployment

rates over time are partly responsible for the increasing cost of buyout fund investing. Thus, LPs should consider a GP's investment behavior as another and important cost driver of PE investing, which has so far been widely overlooked.

Lastly, we find that the deployment rate is mainly unaffected by internal GP-related matters, so that we can negate potential agency conflicts between LPs and GPs in our research context. Instead, deployment rates are predominantly driven by external market factors such as overall PE fundraising and M&A activity. Increased fundraising volumes are related to lower deployment rates, whereas higher M&A activity relates positively to fund deployment rates.

This analyzes contributes to the literature in several dimensions. First, it describes return spreads based on actual data useful to benchmark a fund's costliness. Second, the study analyzes the cost of buyout fund investing over time. Third, it introduces GP investment behavior as another driver of the cost of this asset class, which has so far been widely overlooked. Fourth, the paper enriches the discussion on the overfunding of the maturing buyout industry with a new aspect. Increased capital commitments not only drive down returns, but also increase the problem of dry powder, making the buyout industry more expensive.

The results also provide important implications for investors in the buyout asset class. First, buyout overfunding and the negative effect on returns and costs call into question the overall amount of capital allocated to the asset class, since funds' net returns may decrease below a level that does not reflect their risk and illiquidity. Thus, LPs should be aware of the consequences of overfunding and should verify and potentially adjust their return expectations.

Second, each LP should also reconsider the timing of allocating money to the buyout industry, since fund deployment rates and their impact on return spreads reinforce the relation between buyout fund net returns and general market cycles. For example, LPs investing in boom phases such as 1999 and 2000 or 2006 and 2007 suffer twice, namely, from lower gross returns and from higher return spreads due to low deployment rates, potentially leading to disappointing fund net returns. Of course, LPs also benefit twice when investing in rough fundraising environments in which better gross returns and higher deployment rates can be expected. Overall, the acknowledgment of GP investment behavior favors a more anti-cyclical investment approach.

The study is organized as follows. Section 2.2.2 clarifies the theoretical background and section 2.2.3 introduces the data set and describes the sample. Section 2.2.4 presents the multivariate analyses and section 2.2.5 concludes.

2.2.2 Theoretical Background

2.2.2.1 Literature Review

As buyout fund returns have fallen over time and fundraising has become more difficult since the financial crises starting in 2007, LPs and academics have increasingly shifted their attention toward complex compensation schemes in the PE industry (Gompers and Lerner, 1999; Litvak, 2009). Several studies have taken a rather critical view of the GP-LP relation and related compensation terms, which are ultimately intended to align the interests of both parties. Phalippou (2009) argues that GPs charge excessive fees and asks how investors can accept such contracts? Metrick and Yasuda (2010) question whether the interest of GPs and LPs are actually aligned, since only one-third of all fee payments are performance-based,

whereas two-thirds are mainly due to fixed management fees. The study of Chung et al. (2012), however, relaxes this concern, since the performance of the current fund affects the fund size and management fees of the next fund, so that the present value of performance-based fees is much higher than one-third. Robinson and Sensoy (2013b) verify that a GP's fund terms are largely unrelated to its net performance.

One reason for the ambiguous findings in the previous literature is the lack of comprehensive and complete data. Ideally, in our research setting we would like to observe all gross cash flows between GPs and their portfolio companies, all net cash flows between LPs and GPs, and all fund terms. All previous studies have omitted at least one bit of information and take different approaches to circumvent the missing data.

Gompers and Lerner (1999) and Litvak (2009) focus on analyzing fund terms. Metrick and Yasuda (2010) employ an option framework and various assumptions regarding fee terms, gross returns, and GP investment behavior to simulate fixed and variable GP compensation. Chung et al. (2012) establish a theoretical model and estimate the variable compensation arising from a GP's future fundraising, also based on the option framework of Metrick and Yasuda (2010). Robinson and Sensoy (2013b) work with a rich data set that includes quarterly net cash flows between GPs and LPs, as well as related fund terms, yet they lack the related gross cash flows.

In contrast, this analyzes benefits from the availability of actual net and gross returns and gross cash flows. Based on this sample, this study is the first to provide large-sample empirical evidence on MM and IRR spreads over time. Second, it explicitly analyzes the influence of GP deal-level investment behavior on return spreads. Thus, this study adds to the existing literature, which considers only fund terms and gross performance, by introducing

GP investment behavior as a third driver of the return spread and a third driver of the buyout asset class.

2.2.2.2 Fund Terms and Drivers of the Return Spread

Generally, there are three (partly interacting) drivers of the cost of buyout investing: fund terms, fund gross performance, and GP investment behavior.³⁵ Fund terms are the most obvious driver of final GP compensation and, thus, the return spread. Fund terms are negotiated between a GP and an LP and documented in the limited partnership agreement (LPA).

Standard – at least historically – fund terms are described as 2/20 (Robinson and Sensoy, 2013b), which means that a GP receives 2% management fees per year and a performance-related 20% carried interest (carry) of excessive fund profits beyond a preferred return to LPs. While these fund terms may appear simple at first glance, Metrick and Yasuda (2010) describe their high complexity. They broadly differentiate fund terms into three categories: fee terms, carry terms, and other transaction fee revenues.

In terms of management fees, the two decisive parameters are the size of the yearly fee and the basis on which the fee is calculated. Predominantly for buyout funds, the percentage management fee is paid on the committed capital in the investment phase and, afterward, on the fund's net invested capital. Usually the investment period ends when either five years have passed, commitments have been drawn to a predetermined extent, or a GP starts investing in a new successor fund.

³⁵ We neglect a fourth cost driver applicable only to the IRR measure, which is lagged cash flow timing between LPs and GPs, in comparison to GPs and their portfolio companies. We exclude this fourth driver because we focus our multivariate analyses solely on the more intuitive MM return measure.

Besides the possibility of switching the basis of the calculation from committed capital to net invested capital, a common procedure is to reduce the percentage management fee after the investment period. Second, the two most important carry terms comprise the carried interest level and preferred return, which usually implies that the GP receives 20% of all fund profits after meeting the preferred return (called hurdle rate often being 8%) above the carried interest is actually paid to the GP. However, various parameters affect the overall amount and timing of carry payments, such as waterfall distribution terms, catch-up clauses, and clawback provisions, well described by Metrick and Yasuda (2010).

Third, transaction, monitoring, and other deal-related fees can be significant fund term components and are usually stipulated as a percentage of the enterprise value at deal entry or a profitability measure.

Throughout this study we define gross returns as returns before management fees and carry payments, whereas transaction and monitoring fees have reduced gross returns, since they are already deducted at deal-level. Net returns are defined as returns to LPs, from which management fees and carried interest have been deducted, as well as the LPs' share of deal-related fees have been included.

After describing the standard fund terms in the buyout industry, we now provide an example of how fund terms affect the cost of fund investing and, thus, the return spread. Imagine a USD 100 million fund with a yearly management fee of 2%, which would result in yearly fees of USD 2 million translating into an MM return spread of 0.02 per year (USD 2 million/USD 100 million). Assuming an investment period of five years, the return spread caused by the fixed management fee would add up to 0.10 (5 years*0.02 return spread per year) plus additional management fees after the investment period, depending on the (net)

invested capital and management fee applied. Over a complete fund life, yearly management fees can amount to around USD 15 million resulting in an MM return spread of 0.15.

Next, fund terms also influence the return spread through the interaction with gross returns, which is related to the carried interest. Assuming a standard carry agreement, the GP receives 20% of all fund profits, as long as the hurdle rate is reached (usually after costs). Thus, higher gross returns usually translate into higher performance-based compensation for the GP and consequently also into a higher return spread. For example, a GP would receive USD 40 million in carried interest (with a 20% rate) while using USD 100 million of committed capital and receiving USD 300 million from divestments over a fund's lifetime ($(\text{USD } 300 \text{ million} - \text{USD } 100 \text{ million}) * 20\% = \text{USD } 40 \text{ million}$). Each additional USD 100 million in fund profits increases the variable GP compensation by another USD 20 million. Overall, due to the carried interest, increasing fund profits also increases the return spread.

The third and widely overlooked PE cost driver is rooted in the interaction of fund terms and GP investment behavior. This mechanism is far less obvious than in the cases of fund terms and gross returns. Throughout our research, we focus on one dominant proxy of GP investment behavior that is commonly used in the overall PE industry: a fund's deployment rate. We define a fund's deployment rate by the maximum ratio of cash investments into portfolio companies financed by the fund divided by the fund's committed capital. Thus, if a fund with committed capital of USD 100 million invests USD 90 million over its lifetime, the deployment rate amounts to 90%. This definition does not include capital calls due to, for example, management fee payments and so displays the gross deployment rate, which is therefore lower than the net deployment rate also used in practice.

A simple example shows the economic relevance of the deployment rate. Again, assume a USD 100 million fund with a 2% management fee for a five-year investment period, which consequently receives USD 10 million in fees for this period. If the GP is able to also invest USD 100 million (equivalent to a gross deployment rate of 100%), the return spread of USD 10 million fees amounts to 0.1 MM (USD 10 million/USD 100 million). However, if the deployment rate is only 50%, the spread doubles to 0.2 MM (USD 10 million/USD 50 million). Effectively, a deployment rate of only half the committed capital implies a doubling of the effective management fee percentage to 4% for the capital employed (USD 2 million fees per year/USD 50 million deployed capital). This example is far from being hypothetical, especially for the fund vintage years 2006 and 2007. During these boom years, GPs were able to raise excessive amounts of committed capital; however, they later faced recessionary markets. Thus, the buyout industry had to cope with record levels of dry powder after 2007, which is expressed in low deployment rates, especially for funds starting in mid-2000 with expected fund maturity in 2015 and later.

The economic relevance of the deployment rate points out the importance of GP investment behavior as a third driver of the cost of buyout fund investing. It has been overlooked in the economic literature so far, which mainly focuses on fund terms and net returns due to the unavailability of data. Now, the relevant question to LPs is whether GPs are to blame for these costs. To answer this question, we must first discuss the potential underlying determinants of GP investment behavior. More precisely, is a fund's deployment rate determined by GP characteristics and choices or mainly driven by external market conditions?

2.2.2.3 GP-LP Relation and a Fund's Deployment Rate

The PE industry's legal setup provides grounds for potential agency conflicts: LPs commit capital to GPs with little participation rights, whereas GPs have more or less complete discretion over the final investment decisions. In this respect, the academic literature provides various examples of agency conflicts between GPs and LPs (e.g., Braun and Schmidt, 2014; Metrick and Yasuda, 2010), studies on GP investment behavior (e.g., Ljungqvist et al., 2008), and ways to circumvent potential agency problems (Fang et al., 2013). The deployment rate also presents a suitable example for the different perspectives and interests of LPs and GPs and GP investment behavior. On the one hand, LPs would like to see their committed capital fully employed, since dry powder increases LPs' opportunity costs and effectively increases the cost of buyout fund investing, namely, the return spread. On the other hand, LPs should be hesitant to pressure GPs to fully invest their money, since GPs might, in turn, make unprofitable investments they might not make otherwise. In this case, the return spread might be minimized, but at the cost of much lower gross returns and consequently lower net returns.

Altogether, the question is whether a fund's deployment rate is predominantly determined by GP choices or external market conditions. If the deployment rate is driven internally or even intentionally by GPs, LPs should be careful in their GP selection or even enlarge LPAs to mitigate potential agency conflicts. In contrast, if GP investment behavior regarding the deployment rate is mainly driven by external market conditions, LPs should simply account for these costs in their investment decisions.

2.2.3 Data Set and Descriptives

Data availability is a key determinant of PE research. In this respect, our study benefits from a unique and proprietary buyout sample, which is – to the best of our knowledge – the first data set that includes actual gross and net returns and, thus, also the exact return spread of each fund. Furthermore, full knowledge of monthly gross cash flows provides a clear picture of GP investment behavior.

2.2.3.1 Data Sampling and Selection Bias

The anonymized data set was provided by two LPs investing in hundreds of buyout funds. The LPs extensively collect data from GPs who raise new funds and search for committed capital. First, during fundraising, a GP is usually asked to provide data on all past investments including gross cash flows, so that the LPs have full knowledge of the historic investment activity. Based on this information, the LPs make the final decision of whether to invest in the GP's new fund. Second, after committing capital to a GP, LPs thoroughly monitor and document the ongoing net cash flows, however, by definition, only for the funds in which they are invested.

Thus, our sample faces several selection biases. An obvious instance is that we only observe gross and net return pairs for funds the LPs are actually invested. Assuming a superior fund selection capability of these two professional LPs, the sample predominantly consists of well-performing funds. Another selection bias might be that especially poorly performing GPs simply do not search for committed capital from professional LPs, again restricting our sample to well-performing funds. However, there are two reasons why such biases should not be relevant. First, we are not interested in the mean or median return spread levels in the buyout industry per se. Instead, we assess return spreads along different levels of

gross returns, namely, a spread benchmark. Second, we analyze return spreads over time and therefore evaluate general time trends, independent of the absolute level of return spreads.

Our initial data set includes 586 buyout funds with gross and net return pairs. We exclude 13 funds with missing fund vintage year and size information. Next, we exclude all mainly unrealized funds, since the return spread is only meaningful for mostly realized funds. The parameters for defining funds as mainly unrealized are arguably arbitrary. We exclude all 91 funds younger than 60 months and with vintage years later than 2007 to ensure that the investment period of five years is expired and fund investments are carried out. Furthermore, we exclude all funds in which less than 75% of all initial investments are partially or fully realized. This procedure leaves us with 358 buyout funds (184 from North America, or 51.4%; 151 from Europe, or 42.2%; and 23 from Asia, or 6.4%) for which we have reliable return spreads, which we also verify in various robustness checks.

2.2.3.2 Summary Statistics

The data sampling process of excluding mainly unrealized funds is also displayed in the fund descriptives in Table 2-9, since the mean (median) vintage year of 1997 (1998) for the sample is rather early.

Table 2-9: Summary Statistics

This table displays the fund-level statistics of our sample for 358 buyout funds between 1983 and 2007. Generally, we provide for each variable the number of funds, as well as values for the mean, standard deviation, and five percentiles. The variables cover fund characteristics such as vintage year, generation (indicating the sequence number of a fund per GP), and size (in USD millions). Furthermore, the table presents gross and net money multiple (MM) and internal rate of return (IRR) returns as well as the related return spreads. Please note that the difference between gross and net returns gives the mean return spread, but rightly exhibits small deviations for the different percentiles. The table also displays the gross deployment rate, which is the ratio of all equity investments from a fund in portfolio companies, divided by fund size. The last five variables are potential determinants of the deployment rate. The intermediate MM return is defined as the gross MM of all realized deals within a fund at the end of the investment period. The absolute fund growth is the difference in fund size (in USD millions) between the current and previous funds. The relative fund growth displays the percentage growth between a GP's two successive funds. The fund growth variables only show 222 observations, since only funds with a previous fund can be included. Fundraising and M&A activity are defined as the three-year average change in yearly fundraising volumes and M&A activity (as the number of buyout deals) per GP origin, respectively (from Thomson ONE). We winsorize all variables at the 99th percentile.

Variables	# of funds	Mean	SD	10th Percentile	25th Percentile	Median	75th Percentile	90th Percentile
Fund Vintage	358	1997	5	1990	1994	1998	2000	2003
Fund Generation	358	2.8	1.9	1.0	1.0	2.0	4.0	6.0
Fund Size [USD million]	358	752.6	1041.0	58.0	129.0	329.0	870.0	2061.8
Gross MM	358	2.79	1.22	1.56	1.95	2.56	3.22	4.15
Gross IRR	358	35.1%	22.9%	13.1%	19.1%	30.3%	46.1%	65.6%
Net MM	358	2.22	0.90	1.33	1.62	2.03	2.49	3.42
Net IRR	358	24.8%	17.8%	8.3%	12.8%	21.5%	32.8%	48.1%
MM Return Spread	358	0.57	0.40	0.19	0.28	0.47	0.74	1.05
IRR Return Spread	358	10.3%	7.2%	3.7%	5.3%	8.4%	13.0%	19.1%
Deployment Rate	358	91.3%	18.4%	70.9%	83.8%	93.3%	100.3%	109.2%
Intermediate MM Return	358	2.82	1.25	1.55	1.99	2.59	3.28	4.49
Absolute Fund Growth [USD million]	222	466.8	807.8	-14.5	58.4	211.9	550.0	1265.0
Relative Fund Growth	222	218.6%	648.0%	-9.5%	35.2%	94.9%	191.8%	330.6%
Fundraising Activity	358	43.1%	30.1%	4.0%	32.2%	44.5%	54.6%	69.6%
M&A Activity	358	18.5%	16.8%	-5.0%	5.1%	18.1%	28.2%	39.8%

The mean (median) fund generation number amounts to 2.8 (2.0). The mean (median) fund size of USD 753 million (USD 329 million) is quite comparable to the buyout sample of Robinson and Sensoy (2013b), USD 988 million (USD 313 million).

The high gross and net returns also confirm a positive selection bias. The mean (median) net returns of 2.22 MM (2.03 MM) and 24.8% IRR (21.5% IRR) are considerably higher than the benchmark values provided by Harris et al. (2013a), with mean net MMs between 1.68 and 2.13 and mean net IRRs ranging from 10.9% to 17.1% for different data

sets. Consequently, the return spreads reach tremendous levels, with a mean (median) MM spread of 0.57 (0.47) and a mean (median) IRR spread of 10.3 percentage points (8.4 percentage points). Importantly, these return spreads are not representative of the buyout industry as a whole and must be interpreted carefully, since they are driven by the high gross performance within our sample. To allow for a more objective interpretation, we present spread benchmarks along different levels of gross returns in section 2.2.3.3.

The key variable throughout this analyzes, a fund's gross deployment rate, exhibits a mean (median) of 91.3% (93.3%) and wide distribution, with a standard deviation of 18.4%. Interestingly, some funds exhibit deployment rates above 100%, which occurs when GPs directly reinvest proceeds from selling a portfolio company. These descriptives are in line with the mean net deployment rate of 90.0% after five years and 93.6% in year ten over a fund's lifetime (calculated from cumulative drawdowns from LPs) from Ljungqvist and Richardson (2003).

The last five variables of Table 2-9 are used later to verify whether the deployment rate is predominantly driven by GP-related internal or external market factors. The intermediate MM is defined as the gross MM of all realized deals within a fund at the end of the investment period. Surprisingly, the intermediate gross MM is even higher than the final gross MM, supporting the argument that GPs tend to keep so-called "living dead" investments in their fund (Phalippou and Gottschalg, 2009). Next, the variable absolute fund growth is simply calculated by subtracting the size of the previous fund from that of the current fund. Similarly, relative fund growth displays the percentage growth between a single GP's two successive funds. Lastly, Table 2-9 presents two external potential market drivers of the deployment rate, fundraising and M&A activity, in line with Ljungqvist and Richardson (2003). Fundraising activity is defined as the three-year average change in yearly fundraising

volumes per GP origin, whereas M&A activity is calculated as the three-year average change of yearly M&A activity per GP origin. To mitigate the problem of multicollinearity between fundraising and M&A activity, we base the fundraising proxy on the three previous years of a fund's vintage year, whereas we take the vintage year and the two following years for the calculation of M&A activity.


2.2.3.3 Return Spread Benchmark

Panels A and B of Figure 2-2 illustrate the close relation between gross fund MMs and IRRs to their return spreads, especially at lower gross return levels. The observations vastly cluster around the best-fitting line, whose ascending slope lies slightly above 0.25, for example, for MMs. This slope basically corresponds to the carried interest of 20%.

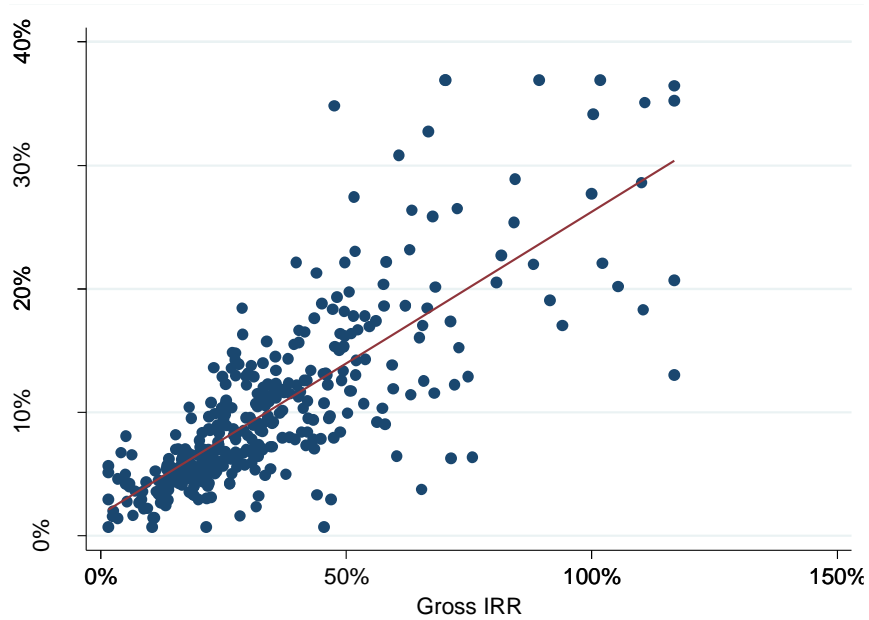
Figure 2-2: Relation between Fund Gross Returns and Return Spreads

Panel A shows the MM return spreads plotted against the gross MM returns for the 358 funds in our sample. Panel B displays the IRR return spreads plotted against the gross IRR returns. The straight line describes the best-fitted relation between both variables.

Panel A: MM Return Spread, by Gross MM



Panel B: IRR Return Spread, by Gross IRR



In addition to Figure 2-2, Table 2-10 presents the related descriptive statistics by differentiating the 358 funds into performance groups. For example, in Panel A, all funds with a gross MM below 1.4 are assigned to the first group; thereafter, the performance groups categorize funds in steps of 0.2 MM until a gross MM of 3.0 is reached; it then categorizes funds in steps of 0.5 MM until a gross MM of 4.0 is attained, and so on and for all the remaining funds.

Table 2-10: Return Spread Benchmarks

This table presents the MM (Panel A) and IRR (Panel B) return spread benchmarks for the 358 buyout funds. The observations are clustered in stepwise gross return categories starting, for example, with 20 funds with gross MM fund returns below 1.4, followed by 21 funds with performance between 1.4 and 1.6. For each gross return cluster, Panels A and B present the number of funds and their mean and median return spreads, as well as the standard deviation.

Panel A: MM Return Spread					Panel B: IRR Return Spread				
Gross MM	# of Funds	Mean	Median	SD	Gross IRR	# of Funds	Mean	Median	SD
< 1.4	20	0.16	0.16	0.06	< 8%	20	3.9%	4.1%	2.0%
1.4-1.6	21	0.21	0.20	0.05	8% -12%	12	2.9%	3.1%	1.4%
1.6-1.8	26	0.25	0.25	0.09	12%-16%	27	4.9%	4.9%	1.3%
1.8-2.0	32	0.32	0.32	0.09	16%-20%	38	5.4%	5.5%	1.5%
2.0-2.2	31	0.36	0.33	0.12	20%-24%	36	6.5%	6.4%	2.4%
2.2-2.4	26	0.42	0.42	0.15	24%-28%	33	8.8%	8.3%	3.1%
2.4-2.6	30	0.49	0.47	0.21	28%-32%	30	9.3%	9.0%	3.8%
2.6-2.8	24	0.53	0.52	0.19	32%-36%	32	9.9%	10.1%	2.8%
2.8-3.0	30	0.64	0.62	0.15	36%-40%	13	11.4%	11.5%	4.3%
3.0-3.5	41	0.74	0.76	0.29	40%-50%	48	12.9%	12.6%	5.6%
3.5-4.0	30	0.81	0.79	0.32	50%-60%	24	15.6%	15.3%	4.8%
>=4.0	47	1.17	1.05	0.56	>=60%	45	21.9%	20.5%	9.3%
Total	358	0.57	0.47	0.40	Total	358	10.3%	8.4%	7.2%

Furthermore, Table 2-10 shows the return spread descriptives by performance group, which can be used to benchmark the cost of individual funds in the buyout industry. For example, a fund with a gross MM performance between 1.8 and 2.0 has mean and median MM return spreads of 0.32, whereas a fund with a gross return between 2.2 and 2.4 exhibits MM return spreads of 0.42. Although these benchmarks do not differentiate between different GP or fund characteristics, they provide novel insight into the costs of buyout fund investing.

This analysis can be used by LPs to conduct an indicative benchmarking of a fund’s costliness.

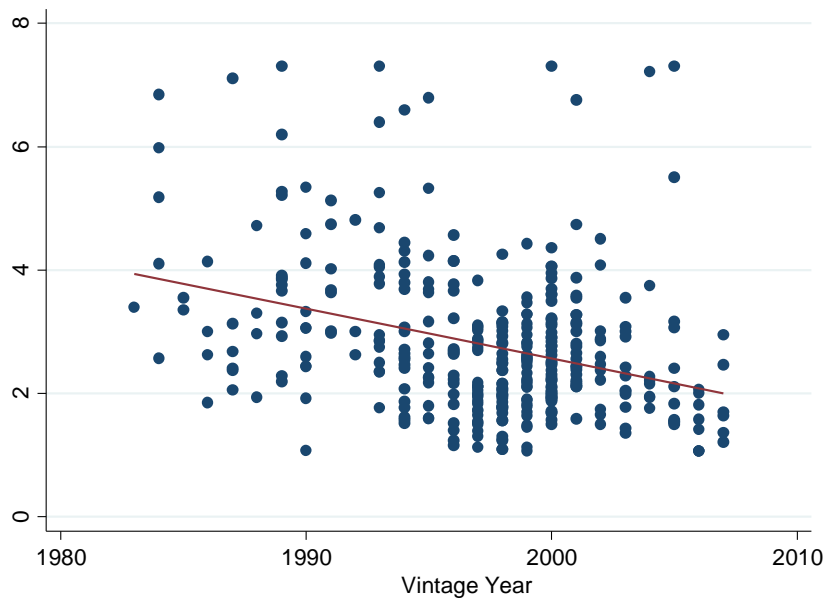
2.2.3.4 Return Spreads over Time and Compared to Fund Characteristics

Figure 2-3 illustrates the development of gross returns and return spreads over time. Panel A presents the scatter plot of 358 fund gross MM returns per vintage year. In line with Harris et al. (2013a), Kaplan and Schoar (2005), and Stucke (2011), we observe strong cyclical return patterns. Fund vintage years during the economic boom periods between 1996 and 1999 or 2006 and 2007 show low returns, whereas funds raised in the bust periods of 2001 or 2002 show stronger returns.


Figure 2-3: Gross Returns and Return Spreads over Time

Panel A shows the gross fund MM plotted against the fund vintage years for the 358 funds in our sample. Panel B displays the MM return spreads plotted against the fund vintage years. The straight line describes the best-fitted relation between both variables.

Panel A: Gross MM, by Fund Vintage Year



Panel B: MM Return Spread, by Fund Vintage Year



However, imposing a linear fitted line, Panel A displays a general downward trend in fund returns. Panel B presents the MM return spread per fund vintage year and exhibits similar cyclical patterns and also a general downward trend. Thus, a preliminary conclusion is that the return spread and thus the cost of investing in buyout funds has decreased over time.

The large influence of gross returns on the return spread is also displayed in Table 2-11. Panel A shows all 358 funds separated into five boom and bust vintage year groups according to Kaplan and Strömberg (2009). Again, the return spreads display cyclical movements with low spreads in the boom periods 1997 to 2000 and from 2004 on. Panel B of Table 2-11 displays the relation between fund size and return spreads when sorting all funds into fund size quartiles. At first glance, increasing fund size seems to be related to decreasing mean return spreads, since the MM spreads decrease from 0.62 to 0.47 and the IRR spreads decrease from 11.2 percentage points to 9.2 percentage points. However, since fund gross performance also decreases with increasing fund size quartiles – for example, from an MM of 3.09 to an MM of 2.34 – the relation might simply be driven by the variation in gross

performance. Finally, Panel C shows that return spreads also decrease with increasing deployment rates, with MM spreads decreasing from 0.65 to 0.45 and IRR spreads decreasing from 11.8 percentage points to 9.6 percentage points. In this case, the relation is hardly driven by gross performance, which actually varies between quartiles without a clear pattern.

Table 2-11: Return Spread Descriptives

This table presents different measures of MM and IRR return spreads and mean MM and IRR gross returns for vintage year groups (Panel A), fund size quartiles (Panel B), and deployment rate quartiles (Panel C). The vintage year groups in Panel A are based on those of Kaplan and Strömberg (2009) and divide the 358 funds in our sample into boom and bust time periods. Panel B ranks all 358 funds by fund size (in USD millions) and separates them into quartiles, where, for example, the largest 25% of all funds is grouped into the high category, the following 25% into the above average category, and so on. Accordingly, Panel C groups all funds by their deployment rates into quartiles ranging from low (the lowest 25%) to high (the highest 25%).

Variable Groups	# of Funds	MM Return Spread			IRR Return Spread			Mean Gross Returns	
		Mean	Median	SD	Mean	Median	SD	MM	IRR
Panel A: Vintage Year Groups									
<=1990	46	0.74	0.58	0.51	10.7%	9.2%	7.0%	3.62	42.2%
1991-1996	83	0.64	0.54	0.44	12.1%	10.7%	7.7%	3.17	42.6%
1997-2000	148	0.49	0.45	0.32	8.3%	6.7%	6.1%	2.43	26.9%
2001-2003	48	0.56	0.50	0.33	13.1%	12.1%	7.5%	2.72	43.4%
>=2004	33	0.48	0.32	0.47	10.5%	7.7%	7.7%	2.37	31.3%
Panel B: Fund Size Quartiles									
Low	90	0.62	0.50	0.43	11.2%	10.0%	7.7%	3.09	40.8%
Below Average	89	0.64	0.52	0.45	10.9%	9.7%	6.8%	3.06	37.1%
Above Average	90	0.53	0.46	0.37	10.0%	7.4%	7.8%	2.65	33.5%
High	89	0.47	0.40	0.32	9.2%	7.0%	6.1%	2.34	29.1%
Panel C: Deployment Rate Quartiles									
Low	90	0.65	0.54	0.43	11.8%	9.4%	8.0%	2.66	34.1%
Below Average	89	0.61	0.50	0.46	10.1%	7.8%	8.0%	2.89	36.5%
Above Average	90	0.55	0.48	0.35	9.8%	8.7%	5.9%	2.96	36.0%
High	89	0.45	0.36	0.34	9.6%	7.8%	6.4%	2.63	33.8%

To clarify the relations between gross returns, return spreads, fund size, and deployment rate, we present a correlation matrix in Table 2-12. As expected, we find strong correlations between gross MM and gross IRR returns as well as between gross returns and return spreads. Furthermore, fund size is negatively related to gross performance and therefore also to return spreads. In contrast, the deployment rate is generally not significantly related to fund gross performance; however, it exhibits a significant correlation to return spreads. These results already indicate that the relation between the deployment rate and

return spreads might not be driven by the variation in gross returns. However, since gross returns have a dominant influence on the return spread, we now turn to a multivariate setting. Lastly, the fund deployment rate is also positively related to a fund's size, which is discussed in section 2.2.4.3

Table 2-12: Correlation Matrix

This table presents the correlation matrix for the variables used in the regressions in Table 2-13 and Table 2-14. The superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

	MM Return Spread	IRR Return Spread	Gross MM	Gross IRR	Fund Size	Deployment Rate
MM Return Spread	1					
IRR Return Spread	0.4651***	1				
Gross MM	0.8309***	0.3954***	1			
Gross IRR	0.5825***	0.7851***	0.7047***	1		
Fund Size	-0.1447***	-0.1346**	-0.2120***	-0.1780***	1	
Deployment Rate	-0.2172***	-0.1558***	-0.031	-0.0117	0.1195**	1

2.2.4 Analysis

In this section, we will answer three questions. First, how have return spreads of our buyout fund sample developed over time? Second, how does GP investment behavior determine the return spread? Third, is GP investment behavior internally or externally driven?

Our regression setup throughout section 2.2.4 employs OLS regressions with heteroskedasticity-robust standard errors, which are clustered at the GP level. Furthermore, we include fixed effects for the GP origin being North America, Europe, and Asia as well as time fixed effects (vintage year groups). We winsorize all continuous variables at 1% to minimize the distorting effect of outliers. We also focus our multivariate analyses on MM return spreads only, for three reasons. First, the IRR measure, although widely used in the PE industry, faces increasing criticism in the academic literature (Braun et al., 2013; Harris et al., 2013a). Second, the IRR spread is not only driven by gross returns, fund terms, and GP

investment behavior, but also by the time lag of cash flow streams between LPs and GPs (net cash flows) and GPs and their portfolio companies (gross cash flows). Although this time lag is usually kept at a minimum in practice, we prefer to employ the MM measure to exclude the noise from cash flow timing. Third, the results are robust to using IRR returns. Fourth and most importantly, the return measure MM allows simple and intuitive economic interpretations of our analyses. A change of 1 percentage point in IRR return spread is difficult to interpret in economic terms, whereas an MM return spread change of, for example, 0.1 can be easily translated into an effect of USD 10 million for a fund size of USD 100 million.

2.2.4.1 How Have Return Spreads Developed over Time?

Table 2-13 displays five regressions on the MM return spread, analyzing its historic development. Regression 1 includes only GP origin fixed effects and a time variable that is the difference between a fund's vintage year and 1982 (the vintage year of our earliest fund minus one) so that the variable ranges between one and 25. The highly significant coefficient indicates that the MM return spread decreases by 0.017 per vintage year, which indicates at first glance that the cost of investing into buyout funds has decreased over time. The results are robust to unreported regressions with variations in the fund vintage years when we, for example, exclude all fund vintage years before 1988 or after 2002 (the five earliest and latest vintage years in our sample).

Table 2-13: Return Spreads over Time

This table presents the results from OLS regressions on the MM return spread. The regressions include two time-related variables. First, regressions 1 and 2 contain a time variable defined as the fund vintage year minus 1982. Since the vintage years in our sample range from 1983 to 2007, the time variable exhibits values from one to 25 and can therefore cover linear trends over fund vintage years. In contrast, regressions 3 to 5 use vintage year groups based on those of Kaplan and Strömberg (2009), presented in Panel A of Table 2-11, to allow for non-linear relations between time periods and MM returns spreads. We chose the vintage year group 3, which ranges from 1997 to 2000, as the base group, because it contains the highest number of funds. All regressions include GP origin fixed effects. Heteroskedasticity-robust standard errors are clustered at the GP level and shown in parentheses. The superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	MM Return Spread				
	(1)	(2)	(3)	(4)	(5)
Gross MM		0.284*** (0.014)		0.286*** (0.014)	0.289*** (0.015)
Fund Size					0.008 (0.020)
Hurdle Dummy					0.291*** (0.080)
Gross MM*Hurdle Dummy					-0.202*** (0.059)
Time [Vintage Year - 1982]	-0.017*** (0.005)	0.006** (0.003)			
Vintage Year Group 1 [Vintage < 1991]			0.246*** (0.086)	-0.097** (0.049)	-0.097* (0.052)
Vintage Year Group 2 [1991 <= Vintage < 1997]			0.149** (0.058)	-0.067** (0.034)	-0.064* (0.036)
Vintage Year Group 4 [2001 <= Vintage < 2004]			0.067 (0.052)	-0.024 (0.028)	-0.026 (0.029)
Vintage Year Group 5 [Vintage >= 2004]			-0.027 (0.080)	0.001 (0.025)	-0.009 (0.026)
GP Origin Fixed Effects	yes	yes	yes	yes	yes
Constant	0.836*** (0.092)	-0.326*** (0.062)	0.511*** (0.039)	-0.208*** (0.036)	-0.223*** (0.046)
Observations	358	358	358	358	358
Adjusted R-squared	0.053	0.697	0.061	0.699	0.700

Next, we include the funds' gross MM returns in regression 2. As expected, fund gross performance shows strong explanatory power for the return spread, with a highly significant coefficient of 0.284 and a tremendous increase in adjusted R-squared, from 5.3% to 69.7%. Thus, a gross return MM increase of 1.0 leads to an increased spread of 0.284, mainly related to the carried interest of at least 20% and other deal-related fees. Surprisingly, the time variable no longer indicates decreasing spreads over time in regression 2. Instead, the coefficient even reverses to a significant positive value of 0.006. So, controlling for falling gross MM returns, buyout fund investing has actually become more expensive over time. The return spread has increased by 0.006 for each year since 1983, independent of fund gross performance. For a period of 10 years, the coefficient of 0.006 results in economically relevant costs of buyout fund investing of a 0.060 MM spread, compared to the mean MM return spread of 0.57. Again, these results remain robust when the five earliest and latest vintage years are omitted.

In regressions 3 and 4, we verify the time trend by replacing the time variable by the five boom and bust categories from Table 2-11. We choose the time period between 1997 and 2000 as a reference category because this time category contains the most funds. The remaining time category dummies support the conclusion of the increasing cost of buyout fund investing, since the first two time categories show significant negative coefficients. Thus, the spread before 1991 was 0.097 MM lower and the spread from 1992 to the end of 1996 was 0.067 MM lower than in the base time category, controlling for gross performance. Thereafter, the two time categories after 2000 show a similar relation to the return spread. We can therefore conclude that the major break in the return spread took place after 1996, since the return spreads were higher than before and remained at that level since then. Regression 5 additionally includes fund size, a dummy variable that equals one if the fund's gross

performance lies below the standard hurdle rate of 8% and zero otherwise, and an interaction of this dummy and gross returns. By including the dummy and the interaction term, we want to control for the non-linear relation between gross returns and the return spread, since the performance-based carry is usually only paid if the gross return after costs lies above an IRR of 8%. Overall, including the three variables adds hardly any explanatory power to the regression model in terms of adjusted R-squared. Furthermore, when controlling for decreasing fund returns, the conclusion that investing in buyout funds has become more expensive, especially since 1996, remains robust.

What might have caused this development? As we already control for gross fund performance, there remain only two potential drivers of PE investing costs: changes in fund terms and/or changes in GP investment behavior. Unfortunately, our data lacks fund terms. However, Stoff and Braun (2014) find that fund terms have been quite stable over time. Therefore, fund terms can be ruled out as an explanation for the increase in buyout investing costs. Consequently, only GP investment behavior remains to explain the increased costs, which we analyze in the next section.

2.2.4.2 How does GP Investment Behavior Determine the Return Spread?

In Table 2-14 we analyze the relation between GP investment behavior and return spread by using the regression setup from Table 2-13 and including a fund's gross deployment rate. Again, we start with a simple regression including GP origin, as well as time fixed effects, and the deployment rate. The coefficient of the deployment rate is highly significant and displays a negative relation with the return spread. This relation holds even when controlling for a fund's gross performance in regression 2 with a highly significant coefficient of -0.415. More precisely, for example, a deployment rate 10 percentage points higher, such as 100%

instead of 90%, decreases the MM return spread by 0.042 ($10\% \cdot 0.415$). This result has remarkable economic relevance. First, variations of 10 percentage points seem quite common in the PE industry, since the deployment rate shows considerable variation, with a mean (median) of 91.3% (93.3%) and a standard deviation of 18.4% (Table 2-9). Second, the effect of a 0.042 change in the return spread amounts to a share of 7% of the total mean MM return spread of 0.57 in our sample ($0.042/0.566 = 7\%$).

Table 2-14: Influence of the Deployment Rate on Return Spreads

This table presents the results from OLS regressions on the MM return spread. Generally, the regressions contain time and GP origin fixed effects, where we use the vintage year groups (from Panel A in Table 2-11) as the time fixed effects in regressions 1 to 4. In regressions 3 and 4 we control the robustness of the results by excluding all funds from the five earliest and latest vintage years. Therefore, regression 3 includes only 340 funds raised after 1987 and regression 4 contains only 313 fund vintages raised before 2003. In regressions 5 and 6 we replace the vintage year groups by the time variable as a robustness check and to verify the relation between the change in deployment rates and the change in MM return spreads over time. Heteroskedasticity-robust standard errors are clustered at the GP level and shown in parentheses. The superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	MM Return Spread					
	All Funds		Vintages	Vintages	All Funds	
	(1)	(2)	> 1987 (3)	< 2003 (4)	(5)	(6)
Deployment Rate [in %]	-0.510*** (0.144)	-0.415*** (0.096)	-0.434*** (0.101)	-0.428*** (0.108)		-0.407*** (0.096)
Gross MM		0.283*** (0.014)	0.281*** (0.015)	0.276*** (0.017)	0.284*** (0.014)	0.280*** (0.014)
Time [Vintage Year - 1982]					0.006** (0.003)	0.005 (0.003)
Time Fixed Effects	yes	yes	yes	yes	no	no
GP Origin Fixed Effects	yes	yes	yes	yes	yes	yes
Constant	0.984*** (0.137)	0.186* (0.096)	0.207** (0.103)	0.217* (0.110)	-0.326*** (0.062)	0.085 (0.114)
Observations	358	358	340	313	358	358
Adjusted R-Squared	0.114	0.733	0.736	0.698	0.697	0.730

Notably, since this research concentrates on the major parameter of GP investment behavior, namely, a fund's deployment rate, we neglect less important factors such as deal timing or holding periods over a fund's life. Based on unreported regressions, we find that the results are also robust to the inclusion of these GP investment behavior variables. A further concern about the validity of the results might be the lack of fund terms in our regressions and, thus, potential unobserved variable bias. For example, funds with higher deployment rates might simply have more LP-friendly fund terms, leading to lower return spreads. This concern might be especially relevant for two reasons. First, the correlation matrix (Table 2-12) indeed shows that deployment rate and fund size are positively correlated. Second, Stoff and Braun (2014) show that a fund's management fee is very strongly and negatively correlated to its fund size. Thus, part of the results may well be driven by larger funds with high deployment rates and lower fees. To verify this potential unobserved variable bias, we include fund size as a proxy for fund terms in unreported robustness checks and find our results nearly unchanged. As in regression 5 of Table 2-13, fund size is not significant and the coefficient of the deployment rates remains strongly significant.

The economic interpretations are especially relevant for LPs invested in funds with vintage years 2006 and 2007 with low expected deployment rates or, in extreme cases, for investors invested in zombie funds. These LPs suffer for two main reasons: First, on average, they are adversely affected by lower gross returns, since funds started in 2006 or 2007 might have conducted overly expensive deals in those years and have faced the economic recession thereafter (for the general relation between fund in-flows and deal valuations, see Gompers and Lerner, 2000). Second, especially for fund vintage years 2006 and 2007, lower gross returns are also accompanied by low deployment rates, which further reduce net returns. Thus, buyout fund investors might benefit from an anti-cyclical investment approach in line

with the related venture capital research of Achleitner et al. (2013). In summary, the negative relation between deployment rates and return spreads urges investors to consider GP investment behavior as a third parameter driving the cost of the buyout asset class.

So far, we have only analyzed the cross-sectional relation between GP investment behavior and the return spread; however, we have not discussed the potential reasons behind the increasing cost of buyout fund investing over time reported in Table 2-13. Are changing deployment rates over time responsible for the change in buyout investing costs? To test this, we again replace the time category dummies by the time variable ranging from one to 25. Regression 5 in Table 2-14 simply replicates regression 2 from Table 2-13 for ease of comparison and provides a significant coefficient of 0.006 for the time variable. In contrast, when the deployment rate is included in regression 6 of Table 2-14, the time variable coefficient loses its significance and is reduced to 0.005. These results are again robust to excluding the five first or last vintage years from the sample (unreported tests). Thus, we can conclude that the reduced deployment rates – or, rather, the increasing costs of dry powder – are at least partly responsible for the increasing cost of investing in buyout funds. However, a clear limitation of this analysis is that the deployment rate might not solely explain the higher costs over time. Note that the time variable coefficient in regression 6 is close to being significant at common levels, with a p-value of around 0.12. Furthermore, this analysis cannot verify the effect of changes in fund terms on the spread over time due to data constraints, although it is reasonable to assume that a decrease in fund terms would result in lower, not higher, spreads.

Now, how should LPs cope with the highly relevant third cost driver of buyout fund investing? The answer depends on the underlying reasons for the variation in deployment rates. If the rate is purely affected by market-driven factors, LPs cannot blame GPs and can

only account for these market parameters and incorporate them into their investment decision. However, if the variation in deployment rates is also related to internal GP-related matters, this finding could imply, for example, some kind of GP gaming behavior. In this case, LPs should account for these factors in their GP selection and when setting up LPAs with specific GPs. In this context, in the next section we analyze potential market- and GP-related reasons for the variation of the deployment rate.

2.2.4.3 Is GP Investment Behavior Internally or Externally Driven?

Table 2-15 verifies the five potential reasons for the variation in the deployment rate we consider the most relevant, differentiated by three GP- and two market-related reasons. The three GP-related reasons are GP experience, intermediate fund performance, and fund growth. First, GP experience might be positively related to the deployment rate, since high GP experience usually includes longstanding and fruitful deal sourcing networks (Kaplan and Schoar, 2005). Thus, high experience might ensure a steady deal flow and consequently the ability to allocate all of the committed capital. We proxy for GP experience and deal sourcing ability with the variables fund size and fund generation, respectively, in line with Demiroglu and James (2010) and Kaplan and Schoar (2005). Second, intermediate fund performance might be negatively related to the deployment rate. We argue that high and thus satisfying performance for LPs at the time a GP is allowed to raise a new fund might induce the GP to concentrate on raising this new fund while neglecting to fully invest the current fund. In case of the old fund's outstanding performance, LPs tend to overlook this GP behavior and might focus on committing money to the successful GP's new fund. Third, fund growth, meaning the absolute or relative change in fund size from the previous to the current fund, might be negatively related to the deployment rate. GPs that have been able to receive much larger

funds might not be able to follow their past, potentially more focused investment strategy and might lack the necessary market knowledge and deal sourcing networks to fully allocate their funds.

The two external potential drivers of the deployment rate are the common market parameters of demand and supply. On the one hand, fundraising activity is a proxy for the demand for buyout transactions; on the other hand, we use M&A activity as a measure of the supply of potential buyout transactions. We expect increasing fundraising activity to decrease the deployment rate, whereas increasing M&A activity will potentially increase it.

The regression analyses in Table 2-15 employ the same setup as in the previous regressions that include GP origin and time fixed effects; however, the dependent variable is now the deployment rate. Regressions 1 and 2 indicate with highly significant coefficients of 0.022 and 0.015 that GP experience, approximated by fund size and fund generation number, is positively related to the deployment rate. This finding is in line with the expectation that high-experienced GPs have better access to more deals in the market. Consequently, this result provides another rationale for why high-experienced GPs have been able to charge slightly more attractive fund terms (Robinson and Sensoy, 2013b), as the demand for such GPs is higher than for low-experienced GPs. Regression 3 shows that the intermediate performance of a fund at the end of the investment phase is not significantly related to the deployment rate. Thus, GPs do not exploit their position of impressing investors with attractive intermediate returns while being actually focusing on raising the next fund. In this regard, we do not find evidence of GPs gaming LPs. Regressions 4 and 5 also do not reveal any relation between the relative or absolute growth in fund size and the deployment rate. So, GPs are actually able to adapt to the new market segment they grow in, which underlines the scalability of the buyout business model (Chung et al., 2012).

Table 2-15: Determinants of the Deployment Rate

This table presents the results from OLS regressions on a fund's gross deployment rates. The regressions present the relation between the deployment rate's internal determinants (regressions 1 to 5) and external determinants (regressions 6 to 7). See Table 2-9 for a more detailed definition of the variables. Specification 8 displays a more comprehensive regression by including all relevant variables. We leave out fund generation in favor of fund size and relative fund growth in favor of absolute fund growth to avoid the issue of multicollinearity. All regressions include time and GP origin fixed effects. Heteroskedasticity-robust standard errors are clustered at the GP level and shown in parentheses. The superscripts *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively.

Variables	Deployment Rate in %							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Fund Size [USD million]	0.022*** (0.008)							0.017** (0.008)
Fund Generation		0.015** (0.006)						
Intermediate MM Return			-0.003 (0.010)					0.002 (0.009)
Relative Fund Growth [%]				0.000 (0.001)				-0.001 (0.001)
Absolute Fund Growth [USD million]					0.000 (0.000)			
Fundraising Activity						-0.101** (0.040)		-0.145*** (0.048)
M&A Activity							0.127** (0.058)	0.125* (0.071)
Time Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes
GP Origin Fixed Effects	yes	yes	yes	yes	yes	yes	yes	yes
Constant	0.904*** (0.019)	0.876*** (0.029)	0.936*** (0.031)	0.937*** (0.019)	0.932*** (0.020)	0.969*** (0.024)	0.914*** (0.019)	0.959*** (0.042)
Observations	358	358	358	222	222	358	358	222
Adjusted R-squared	0.055	0.061	0.041	0.063	0.065	0.062	0.049	0.141

Regressions 6 and 7 analyze the potential external market drivers of the deployment rate. Regression 6 shows a significant negative relation between the average change in fundraising in the three years before the fund starts investing and, thus, confirms our expectations. Increasing fundraising activity increases the demand and competition for PE deals and consequently lowers deployment rates. In economic terms, we find that a 10 percentage point change in average fundraising growth results in a deployment rate around 1.0 percentage point lower, whereas commonly observed variations of this variable are expected to be higher, with a standard deviation of 30.1% (Table 2-9). Furthermore, regression 7 also confirms our expectations, since increased M&A activity is positively related to the deployment rate. A 10 percentage point change in average M&A activity (with a standard deviation of 16.8%) results in a 1.3 percentage point higher deployment rate. Regression 8 summarizes all five internal and external variables while only excluding the redundant variables fund generation and absolute fund size growth (to avoid multicollinearity) and presents similar results as in the individual regressions. Overall, we can conclude that, besides GP experience, the deployment rate predominantly depends on external market-driven variables. This conclusion underlines the need of LPs to account for the deployment rate in their investment decision as a third exogenous driver of the cost of buyout fund investing.

2.2.5 Conclusion

In the last decades, buyout investing has emerged from a small niche to a well-established asset class. Both GPs and LPs have professionalized so that the PE market has become more mature. However, the lack of data has prevented a comprehensive view of the cost of the PE asset class, as well as their drivers. This analyzes is the first to provide large-

sample empirical evidence on the cost of buyout fund investing, namely, the spread between gross and net performance, based on actual data. We provide a spread benchmark that indicates, for example, an MM spread of around 0.3 for a fund with a gross MM performance between 1.8 and 2.0. LPs can use this benchmark for a rough evaluation of fund costliness, although it neglects all other fund-specific characteristics, as well as changes over time.

In our multivariate analyses, we find that return spreads have decreased between the vintage years 1983 and 2007, indicating lower costs of buyout fund investing, at first glance. However, when the dominant driver of the return spread, which is a fund's gross performance, is controlled for, return spreads have actually increased over time. Since fund terms, on average, have been quite stable over time, we identify GP investment behavior as a third, vastly overlooked driver of buyout investing cost. We proxy for investment behavior with the deployment rate, which is especially low for fund vintage years 2006 and 2007 and is of particular relevance in so-called zombie funds. The deployment rate, which reflects the costs of dry powder, exhibits a strong correlation to return spreads in our cross-sectional analyses and also partly explains the increasing cost of buyout fund investing over time. Importantly, we verify that the deployment rate is not driven by GPs gaming LPs to maximize their fees but, rather, mainly by external market drivers such as fundraising and M&A activity.

Overall, we interpret the increased cost of fund investing after controlling for fund gross performance as a consequence of the evolution of the buyout industry from a niche to a nearly commoditized asset class. The success of the buyout segment has led to increasing, although very cyclical, fund flows into the industry, with two broad return-related effects. First, gross fund returns have fallen. Second, it has become more and more difficult for GPs to allocate all of the committed capital, leading to lower deployment rates and, thus, to higher return

spreads. Our analysis shows that the cost of buyout investing, when falling gross returns are controlled for, has especially increased since 1996. This finding is well in line with Braun et al. (2013) and Sensoy et al. (2013), who find that GP and LP performance persistence disappeared around 1996. Obviously, the buyout industry seems to have undergone a fundamental change in this period. Since then, LPs not only exhibit lower performance persistence, but also must consider GP investment behavior as an increasingly important cost driver of buyout investing, also explaining lower net returns.

In this respect, this study offers important implications for investors allocating funds to the buyout industry. First, LPs could intensify their efforts to change fund terms in their favor by introducing a mechanism in LPAs that compensates LPs for low deployment rates. Thereby, LPs could effectively transfer a share of the downside of low deployment rates to GPs and would incentivize GPs to reach a deployment rate of 100%.³⁶ However, such a mechanism is controversial, since it might have ambiguous effects. On the one hand, such a legal agreement potentially creates a new agency conflict between GPs and LPs, since GPs might be induced to sacrifice rigorous deal selection for a high deployment rate, leading to overall lower net returns. On the other hand, the mechanism might also educate GPs to set realistic target fund sizes for their next fund, which can potentially mitigate the fundamental issue of overfunding in the buyout industry.

Second, it is not only relevant to reduce GPs' incentives to raise ever-increasing (although with strong cyclicity) fund sizes, but also fundamental for LPs to reconsider their overall allocation of funds to the buyout industry. From a long-term perspective, overfunding causes decreasing gross returns and deployment rates. Thus, LPs have to evaluate how much

³⁶ LPAs can, for example, include an agreement on management fee cuts after the investment period, if the deployment rate does not reach a minimum level.

committed capital the buyout industry can cope with before net returns are overly depressed below a level that does not reflect the risk and illiquidity of the asset class. A similar development has already taken place in the smaller (in terms of committed capital) venture capital industry, from which some LPs have turned away completely (Lerner, 2011). Truly, a sustainable allocation of funds to PE buyout funds and the availability of buyout deals can hardly be influenced by single LPs but lies in the interest of all market participants.

Therefore, third, individual LPs should more rigorously evaluate the particular cyclicity of the buyout industry. Investments in boom periods such as 2006 and 2007 most likely generate disappointing gross returns and high costs due to low deployment rates so that LPs suffer twice. In contrast, committing capital in bust periods in a tough fundraising environment can potentially also benefit LPs twice: they can expect higher gross returns and lower costs of buyout fund investing. Thus, LPs should consider taking a more anti-cyclical investment approach.

2.3 Essay 3 - The Evolution of Private Equity Fund Terms Beyond 2 and 20

Abstract

Fund terms in the private equity (PE) industry, historically described as 2% fixed management fee and 20% success-based carried interest (2/20), have raised concerns of misaligned interests between limited partners (LPs) and general partners (GPs). The traditional 2/20 fund terms do not seem to reflect the new realities of the PE industry anymore with larger fund sizes and lower gross returns, so that the major share of GP income increasingly stems from performance-unrelated management fees. Based on a proprietary data set covering the detailed fund terms of 210 PE buyout funds with vintage years between 1989 and 2012, this study is the first to empirically analyze the evolution of fund terms and underlying drivers. We find that fund terms have been surprisingly stable over time despite concerns and pressure from LPs. In recent years LPs have only been able to reduce the heavily criticized uses of deal-by-deal carry and deal-related fees benefiting GPs. Instead, the major driver of more investor-friendly fund terms offered by GPs has been growing fund sizes, which has however not been sufficient to keep interests between LPs and GPs aligned. Obviously, GPs have accepted moderately lower management fee percentages with the aim of maximizing fund sizes and thus management fees in absolute (dollar) amounts. Our research provides guidance towards realigning the interests between LPs and GPs with a new – more success-based – fund term structure in the PE industry in which the share of fixed to total GP compensation decreases from around 60% to 20%.

Keywords: Private equity, fund terms, misalignment of interests

JEL Classification Code: G11, G23, G24

Authors: Stoff, Ingo; Braun, Reiner

First Author: Stoff, Ingo

Current Status: Published in *Journal of Applied Corporate Finance*
(Vol. 26, No. 1, pp. 65-75)

2.3.1 Introduction

Most private equity (PE) funds are organized as limited partnerships, and there is a good deal of controversy about the terms that such PE funds have struck with their limited partners (LPs). In the typical arrangement, the LPs are pension funds, insurance companies, or other institutional investors that commit capital to specific funds that are managed by general partners (GPs) like KKR or Bain Capital. The GPs invest this capital at their own discretion in individual target companies. For their efforts, most GPs have received compensation that has historically consisted of two main components: (1) a yearly fixed management fee of about 2% of the committed capital and (2) a success-based fee equal to 20% of the fund's profits above those needed to provide LPs with a minimum threshold rate of return (often referred to as "the hurdle rate").

In the early days of the PE industry, the mixture of fixed and variable compensation provided by this 2 and 20 ("2/20") arrangement was widely believed to ensure a congruence of interests between GPs and LPs. But two fairly recent developments have raised questions about the effectiveness of the compensation scheme in the buyout segment.³⁷ First, the amount of capital committed to buyout funds in aggregate has risen sharply since the 1980s, causing performance-unrelated management fees for large funds to reach staggering levels of hundreds of USD millions per year. Second, and perhaps partly as a consequence of the first, the average returns of buyout funds have been falling over time (Harris et al., 2013a; Higson and Stucke, 2012), reducing the relative importance of variable compensation for GPs. The resulting increase in the fixed relative to the variable compensation of GPs has fueled concerns of industry participants that GPs have become increasingly focused on maximizing

³⁷ Although fund terms are relatively similar between early-stage venture capital and later-stage buyout funds, this analysis explicitly focusses on buyout funds.

fund sizes and the associated management fees, while neglecting their original mandate to create above-market returns for their investors (Hudec, 2010).

The growing awareness of this misalignment of interests between LPs and GPs, together with the increasing “professionalization” of LPs, has led to a joint initiative of several large LPs – known as “ILPA Principles”³⁸ – whose aim is to adjust fund terms to reflect the new realities within the buyout industry. Furthermore, the difficult fund-raising environment since 2007 has shifted bargaining power to LPs from GPs seeking capital commitments for new funds.³⁹ Both of these factors have pressured GPs into making concessions to LPs. The goal of this study is to analyze whether, and the extent to which, buyout fund terms have evolved beyond the classical 2/20 structure in recent years to reflect the new reality of lower returns and the increasing number of “mega” funds with a USD 1 billion or more of capital to invest.

This study also attempts to determine whether more “LP-friendly” fund terms are related to increases in fund size (which we later refer to as the “fund size effect”) or, more generally, to all funds of recent vintage years (“market effect”) in a way that reflects the increased bargaining power of LPs. The study concludes with a closer look at the new fund terms and their shift away from fixed management fees. Using our findings as a basis, we also offer some suggestions for the buyout industry in its ongoing effort to design more investor-friendly terms with the promise of restoring the unity of goals and interests between LPs and GPs.

Our findings are based on an analysis of a proprietary data set of detailed fund terms from 210 PE buyout limited partnerships established between 1989 and 2012. In what came

³⁸ ILPA is the Institutional Limited Partners Association, which recommends best practices for PE fund terms.

³⁹ See *The Economist*, May 6th, 2010, “Private inequity.”; See *Financial Times*, Feb. 28th, 2012, “Investors push for private equity fee cuts.”

as a bit of a surprise to us, we find that fund terms in the buyout industry have been remarkably stable over time. We identify a clear industry standard for most fund terms, with a modest amount of variation that appears to reflect mainly just differences in fund size. The standard management fee charged on the committed capital during the investment period is 2.00% for funds smaller than USD 1 billion, dropping to 1.75% for funds between USD 1 billion and USD 5 billion, and to 1.50% for funds larger than USD 5 billion.

For the majority of more recent funds, we find that LPs and GPs agreed to a change in the *basis* for calculating the management fee from committed capital to (net) invested capital that take effect *after the end of the investment period*. The most common arrangement also includes a carried interest (carry) level of 20% for profits above an 8% hurdle rate, with the carry distribution varying with the origin of the GP (predominantly deal-by-deal carry for American based funds and whole-fund carry for European funds). Finally, it includes the stipulation that all deal-related fees accrue directly to the funds and their LPs.

Overall, we conclude that the fund size effect – the tendency for changes in fund terms to reflect the increasing sizes of the funds – dominates the market effect. During the period we study, only the most heavily criticized fund terms, such as the deal-by-deal carry distribution and deal-related fees, have become more investor-friendly. Yet, this success of LPs seems to only amend the worst liabilities from the beginning of the buyout industry. At the same time, the main driver of reduced fund terms appears to be the growth in fund size, and its effects on management fees. Our results show that a 100% increase in fund size is associated with a reduction of the management fee by 0.16 percentage points (e.g., from 2.00% to 1.84%). The clear message here is that, as fund sizes have soared, GPs have accepted moderately lower management fee *percentages* with the aim of generating higher management fees in *absolute* dollars.

The increasing importance of fixed management fees in the buyout industry supports the concern about misalignment of interests between LPs and GPs. With this concern in mind, we evaluate the expected effects on goal congruence and GP incentives of two new fund term structures that contain lower management fees (0.50% and 1.00%) and a higher carried interest (30%). We conclude that these new arrangements, which have been introduced by Bain Capital in their latest fund (Bain Capital XI), have little effect on the total fee compensation received by GPs. But more important, these new terms would have the beneficial effect of reducing the share of GP income from fixed fees below 20% of their total compensation. For this reason, the roll-out of these new terms could well be a big step toward a buyout industry with compensation terms that realign the interests of LPs and GPs.

2.3.2 Theoretical Background

The academic literature provides a comprehensive view of the complexity of fund terms and categorizes them into three classes: management fees, carried interest, and other deal-related fees (see for example Chertok and Braendel, 2010a/2010b; Gompers and Lerner, 1999; Litvak, 2009; Metrick and Yasuda, 2010; Robinson and Sensoy, 2013b). In the next few paragraphs, which are mainly based on Metrick and Yasuda (2010), we will briefly explain the most important fund terms.

First, since the initial stated purpose of the fixed management fee has been to cover the operating expenses of a fund, the calculation of these fees varies during the lives of the funds to reflect the typical pattern of intensive GP involvement at the start followed by diminished activity as the GP shifts from deal-making to a mainly monitoring and selling function. In recognition of this pattern, the fee structure of most funds is divided into an investment period

of from three to five years, during which the committed capital is expected to be put to work in individual investments, and a divestment period for the remainder, generally five to seven years, given the typical fund life of ten years. Because GPs do most of the work of sourcing and executing deals during the investment period, they typically receive the management fee on the full amount of committed capital during that time, regardless of whether the capital has been called from the LPs to finance a transaction. During the divestment period that follows, however, the GPs earn lower management fees that reflect the generally less time-consuming activities of monitoring and exiting their portfolio companies.

In recent years, however, there have been two widely adopted changes in assessing management fees. At the end of the investment period, either the *basis* on which the management fee is charged switches from the committed capital to a lower measure often expressing the (net) invested capital or the fee *level* itself is reduced (or both).

Next, we explain the terms of the carried interest, which are intended to align interests between LPs and GPs. The standard arrangement is for GPs to receive 20% of all fund profits (mainly after deducting all costs, e.g., management fees) after the fund returns satisfy the specified hurdle rate, which for most funds is about 8%. Other more detailed terms of the carry include the waterfall distribution and catch-up provisions. The waterfall distribution describes the timing of carry payments to GPs, which can be loosely described as falling into two categories: deal-by-deal carry (also known as the “American waterfall”) and whole-fund carry (“European waterfall”). The American waterfall, which has received considerable criticism from LPs, allows GPs to extract carried interest for successful deals without regard for the performance of the rest of the portfolio, thus creating the potential for larger carried interest when exits from high-return deals are followed by the resolution of deals with low returns (Hudec, 2010). The European waterfall avoids this undesirable outcome – and the

incentives to do lower quality deals whose exits can be deferred – by making carry payments a function of the return on the entire portfolio of deals, and by delaying the carry payments to GPs until LPs get back their initial investment in the fund plus the hurdle rate return.

Finally, the catch-up provision specifies the speed with which GPs receive their carry after meeting the hurdle requirement. For example, a catch-up of 100% implies that the GP will receive all profits above the hurdle until its share of profits reaches the specified carry level – again, typically 20%.

Third is the category of deal-related fees, which have also been strongly criticized by industry practitioners.⁴⁰ We distinguish between transaction fees, monitoring fees, and abort fees. Transaction fees are in fact paid by the companies acquired by the GPs, and typically amount to 1% to 2% of a deal's enterprise value. Monitoring fees are also paid by the portfolio companies, and generally range between 1% and 5% of Earnings before interest, taxes, depreciation, and amortization (EBITDA) and display an additional compensation for GPs for performing what is part of their normal activity, namely the ongoing monitoring of their portfolio companies (Metrick and Yasuda, 2010). Abort fees result from cases in which GPs do not win a bidding contest; such fees are paid by the target companies, and are intended to provide compensation for the GPs' due diligence costs. In most cases, these additional deal-related fees flow back to the fund, which has initially also paid the due diligence costs, and so benefit mainly the LPs. However, part of such fees depending on the agreement between GPs and LPs – and this is the source of the controversy – also flow directly to GPs as a third source of income.

⁴⁰ See *The Economist*, Nov. 12th, 2011, "Fee high so dumb."

2.3.3 Data Set and Descriptive Statistics

2.3.3.1 Data Set

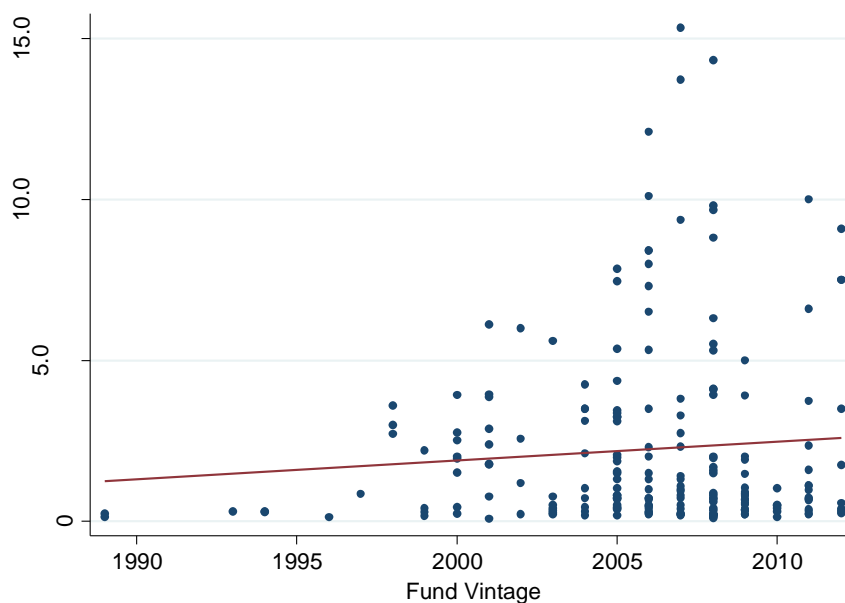
Our sample consists of 210 buyout funds with vintage years between 1989 and 2012. For each of these 210 funds, we collected information on management fees and carried interest. But more detailed terms, such as deal-related fees, were accessible for only about 170 funds. The sources of our information on fund terms were private placement memoranda (PPM) and the limited partnership agreements (LPAs) collected by one large international LP. Of the 210 funds, 116 (or 55%) are based in Europe, 86 (41%) in North America (41%), and 8 (4%) in Asia.

The sample is particularly useful for assessing recent developments since it consists mainly of funds launched after the year 2000. Furthermore, the sample includes a broad variety of fund sizes, ranging from mid-cap funds to funds larger than USD 10 billion, and includes funds representing USD 468 billion of total committed capital.⁴¹ Furthermore, as can be seen in Figure 2-4, which displays fund sizes over fund vintage years, the private equity industry has experienced pronounced boom and bust market cycles with, for example, low fund volumes around 2003 and record fund sizes in 2007.

⁴¹ In respect to fund terms, our sample displays similar descriptives in comparison to existing academic studies dealing with PE buyout funds. See Metrick and Yasuda (2010) as well as Robinson and Sensoy (2013b).

Figure 2-4: Fund Sizes over Fund Vintage Years

This scatter plot displays the relation between fund sizes and fund vintage years for all 210 buyout funds in our sample between 1989 and 2012. Fund size is expressed in USD billions. The line indicates the linear best fitted values to explain the relation between both variables. The straight line describes the best-fitted relation between both variables.



2.3.3.2 Descriptive Statistics

Table 2-16 presents fund term descriptives for the entire sample as well as categorized by time and fund size. Following existing research that differentiates fund vintage years into boom and bust periods (Kaplan and Strömberg, 2009), we identify four distinct time periods: (1) all vintages before 2001, (2) the “bust” vintage years of 2001 to 2003, (3) “boom” vintages of 2004 to 2007, and (4) “bust” vintages from 2008 to 2012. By clustering funds into these four groups, we are able to detect the relation between the fund-raising environment of a certain period and fund terms as well as any changes in fund terms over time. As can be seen in Panel A of Table 2-16, the lion’s share of our 210 buyout funds are drawn from the two more recent time categories, with 86 funds launched during the boom years from 2004 through 2007, and 82 funds raised in the period after 2007.

Table 2-16: Fund Term Descriptives over Time and Fund Size Categories

This table shows the sample distribution (Panel A), fund characteristics (Panel B), and fund term descriptives (Panel C) for vintage year and fund size categories. Panel A reports the sample distribution in the four time and three fund size categories for the complete sample of 210 funds. Panel B displays the mean and median fund sizes in USD million as well as the share of funds with headquarter in Europe along the categories. Panel C presents fund term descriptives related to management fees, carried interest, and other deal fees.

	# of Funds	All Funds	Vintage Year Categories				Fund Size Categories [USD million]		
			<2001	2001-2003	2004-2007	>2007	<1,000	1,000-5,000	>5,000
Panel A: Sample Distribution									
Number of Funds	210		22	20	86	82	109	71	30
Panel B: Sample Fund Characteristics									
Mean Fund Size [USD million]	210	2,231	1,358	2,078	2,580	2,135	452	2,406	8,279
Median Fund Size [USD million]	210	894	642	1,473	1,063	757	397	2,100	7,675
Share of Funds with Headquarter in Europe	210	55%	55%	75%	45%	61%	69%	44%	33%
Panel C: Fund Term Descriptives									
Mean Management Fee	210	1.85%	1.97%	1.74%	1.83%	1.86%	1.99%	1.75%	1.54%
Median Management Fee	210	2.00%	2.00%	1.57%	1.75%	2.00%	2.00%	1.75%	1.50%
Share of Funds for which Fee Basis Changes after Investment Period	210	90%	82%	90%	91%	93%	87%	93%	97%
Share of Funds for which Fee Level Changes after Investment Period	210	37%	27%	35%	42%	35%	20%	49%	70%
Mean New Fee Level after Investment Period	78	1.17%	1.38%	1.00%	1.12%	1.24%	1.50%	1.18%	0.82%
Median Carry Level	210	20%	20%	20%	20%	20%	20%	20%	20%
Median Hurdle Rate	207	8%	8%	8%	8%	8%	8%	8%	8%
Share of Funds with European Waterfall Distribution Rule	197	53%	26%	65%	45%	65%	70%	43%	20%
Mean Catch-up	157	93%	89%	92%	94%	94%	94%	92%	93%
Mean Transaction Fees Benefiting the Fund	168	75%	56%	67%	74%	80%	72%	75%	84%
Mean Monitoring Fees Benefiting the Fund	167	77%	56%	71%	75%	84%	74%	78%	86%
Mean Abort Fees Benefiting the Fund	170	86%	65%	79%	85%	91%	82%	87%	96%

As reported in Panel B of Table 2-16, the average fund size for the entire sample (and thus spanning all four time periods) is USD 2.231 million (while the median size is USD 894 million). Among the four different periods, the average fund size is largest (USD 2.580 million) during the boom years between 2004 and 2007 (when the median fund is just over USD 1 billion). The relatively small fund sizes in the period before 2001 can be explained by some older funds included in this group with volumes below USD 500 million. The proportion of funds located in Europe varies considerably over time, and decreases with fund size as expected, since the well-known mega buyout funds are still located mainly in North America.

Panel C summarizes our findings about management fees, carried interest, and deal fees. First, the average management fee for the 210 buyout funds turns out to be 1.85% – while the median was 2.00% – which is mainly consistent with the standard 2/20 terms. As expected, we find that management fees followed patterns of economic cycles for the first three time categories. The average (median) fees were 1.97% (2.00%) before 2001 and then fell to 1.74% (1.57%) between 2001 and 2003, when fund-raising had become very difficult for GPs. When the economic environment improved between 2004 and 2007 and LPs increasingly allocated funds to the buyout industry, GPs have been able to raise fee levels again to 1.83% (1.75%). Surprisingly, after 2007 the average (median) management fee increase further to 1.86% (2.00%), despite the drying up of capital flows to the industry, and growing criticism of excessive fund terms.

Thus, we find that fixed management fees, as a key driver of GP income, have been surprisingly stable over time. In contrast, a comparison of management fees among the three fund size categories shows a clear negative relation. The standard fee for funds below USD 1 billion in committed capital is about 2.00%, but for funds with capital of between USD 1 and

USD 5 billion, the fee is about 1.75%, and funds larger than USD 5 billion earn a management fee of about 1.50%. Table 2-17 provides more information about the relation between time and management fees. Panel A reports the number of funds differentiated by the vintage year and fund size categories. Panel B shows for each cell the median management fee during the investment period of a fund (means shown in Panel C). Strikingly, the median management fee is constant over time at 2.00% for funds below USD 1 billion committed capital and for funds above USD 5 billion it remains essentially unchanged at 1.50%. Funds with volumes between USD 1 billion and USD 5 billion also exhibit constant fees of 1.75% after 2003. Even the somewhat more variable average management fee reported in Panel C confirms the stability of the pattern over time.

Table 2-17: Management Fee by Time and Fund Size Categories

This table shows the sample distribution (Panel A), median management fee (Panel B), and mean management fee (Panel C) in the investment period for all 210 funds in a matrix structure by vintage year and fund size categories.

Vintage Year Categories	Fund Size Categories [USD million]		
	<1,000	1,000-5,000	>5,000
Panel A: Number of Funds			
<2001	12	10	
2001-2003	9	8	3
2004-2007	40	32	14
>2007	48	21	13
Total	109	71	30
Panel B: Median Management Fee			
<2001	2.00%	1.88%	
2001-2003	2.00%	1.50%	1.50%
2004-2007	2.00%	1.75%	1.50%
>2007	2.00%	1.75%	1.50%
Total	2.00%	1.75%	1.50%
Panel C: Mean Management Fee			
<2001	2.04%	1.89%	
2001-2003	2.00%	1.55%	1.50%
2004-2007	1.97%	1.76%	1.59%
>2007	2.00%	1.76%	1.50%
Total	1.99%	1.75%	1.54%

In sum, for funds of similar size, management fee percentages are remarkably stable. And the same is true of the other three management fee terms. In Table 2-16, we report evidence of a clear positive relation between fund size and the percentage of funds that experienced changes in both their fee basis as well as their level of fees. For the 78 funds that exhibited a change in the level of their fees, we observe lower fees associated with higher fund sizes. In contrast, there is no clear pattern of changes over time for these three fee terms with a single exception: after 2007 93% of the funds in our sample experience a change in the basis on which fees are calculated. Since that time the industry standard is to charge management fees on (net) invested capital after the end of the investment period.

When we analyze the development of carried interest terms, we also find a clear industry standard that consists of a carry level of 20%, a hurdle rate of 8%, and an average catch-up rate of slightly above 90% (but a median rate of 100%). The use of the European carry distribution might have also increased over time to 65% after 2007. This development might, however, be simply explained by the variation in fund location for the different time categories in our sample (e.g., whether a fund was raised in Europe or America) since European funds mostly incorporate a European carry distribution.

Finally, we find clear evidence of changes in the three kinds of deal-related fees for certain categories of fund size. Before 2001 about 60% of these fees benefited the fund (and thus LPs) while the other 40% provided additional income to GPs. After 2007 an average of from 80% to 91% of these fees flow back to the funds – though the fact that the median value was 100% indicates that 100% has become the new industry standard.

2.3.4 Analysis

When analyzing changes in the fund term descriptives, it is hard to differentiate whether these changes are driven by different market periods, differences in fund sizes, and variations in sample characteristics such as fund location. To explore the significance of market and fund size effects, we design a series of multivariate regressions in which the dependent variables measure the fund terms discussed above. More specifically, we use standard OLS regressions for continuous dependent variables such as management fee, carry level, hurdle rate, catch-up, and deal-related fees (all expressed as percentages). In contrast, we use logit regressions for the binary dependent variables fee basis and change in fee level (with 0 for no change and 1 for a change) as well as the kind of carry waterfall (0 for American and 1 for European). In each regression we control for fund location by including dummies for Europe and Asia, while treating North America as the base case.

Along with the regional controls, we use two variables to evaluate the fund size and market effects. First, we include fund size as the natural logarithm, because we expect the percentage change, in comparison to the absolute change, to be stronger related to fund terms. For example, we expect that doubling the fund size from USD 1 billion to USD 2 billion is associated with a reduction in the percentage fee, while an equal absolute increase from USD 10 billion to USD 11 billion has a negligible effect on fee percentages.

Second, to capture the general development of fund terms over time, independent of market cycles, we include a time trend variable for the fund vintage year. Specifically, it measures the vintage year minus 1988, which is the first vintage year in our sample minus one. Therefore, this time variable ranges from 1 to 24 as the vintages in our sample lie between 1989 and 2012 and it indicates the change in fund terms for one additional year.

2.3.4.1 The Drivers of Fund Terms

As summarized in Table 2-18, the findings of our multivariate regressions confirm that the effects of fund size on management fees mostly dominate the market effects related to increased bargaining power of LPs. Coefficients of the natural logarithm of fund size are strongly significant for seven fund terms, whereas the time variables are significant for only the carry waterfall distribution and deal-related fees.

Notably, our regression findings with respect to changes in fee basis, carry levels, hurdle rates, and catch-up percentages show no relation to fund size or time variables, which is consistent with the high values and relatively small variation of these fund terms reported earlier in Table 2-16. In the following paragraphs, we analyze our results for these three fund term categories in more detail.

The results of regression 1 confirm that management fees are nearly unrelated to the passage of time, whereas we find a strong statistical relation to fund size. The coefficient of -0.157 implies that a 100% increase in fund size is associated with a decrease of the management fee by 0.16 percentage points, for example from 2.00% to 1.84%. Regression 2 shows no significant relation between fund size or time on the possibility of change in the fee basis after the investment period. As already noted, a change from committed capital to (net) invested capital had already been become the industry standard by 1989, the time period with which our sample begins.

Table 2-18: Regressions of Fund Terms

This table presents the results of OLS and logit regressions on management fee related terms, carried interest terms, and deal fees. We conduct standard OLS regressions on the continuous dependent variables. We use logit regressions in which the dependent fund term variables adopt a value of zero or one. In regression 2 we have to exclude the Asia dummy as this dummy perfectly predicts the fee basis change (all 8 observations feature a change). Numbers in the upper rows represent the regression coefficients. The superscripts *, **, *** indicate the p-values of the 10%, 5%, and 1% significance levels, respectively. In the lower rows the standard errors are reported in parentheses. We cluster robust standard errors by GP. The last row indicates the adjusted R-squared for OLS regressions and McFadden's pseudo-R-squared for logit regressions.

Fund Term Type Regression Type	Management Fee				Carried Interest				Deal Fees		
	OLS	Logit	Logit	OLS	OLS	OLS	OLS	Logit	OLS	OLS	OLS
Variables	Management Fee [%]	Fee Basis Change [0 no, 1 yes]	Fee Level Change [0 no, 1 yes]	Management Fee after Change [%]	Carry Level [%]	Hurdle Rate [%]	Catch-up [%]	Waterfall [0 US, 1 EU]	Share of Fees Flowing Back to Fund [%]		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
LN(Fund Size)	-0.157*** (0.014)	0.246 (0.232)	0.698*** (0.183)	-0.251*** (0.030)	0.134 (0.108)	-0.153 (0.164)	-0.729 (0.822)	-0.645*** (0.235)	5.999*** (1.542)	5.252*** (1.613)	6.551*** (1.553)
Time [Fund Vintage Year - 1988]	-0.001 (0.006)	0.074 (0.062)	0.018 (0.043)	0.016 (0.013)	-0.029 (0.030)	-0.040 (0.026)	0.186 (0.341)	0.237*** (0.056)	1.799** (0.743)	2.107*** (0.708)	1.882* (0.997)
Headquarter Dummy: Europe	-0.083** (0.037)	0.004 (0.503)	-0.660 (0.414)	-0.015 (0.086)	-0.690** (0.320)	-0.028 (0.304)	-2.411 (2.733)	3.370*** (0.524)	9.007** (3.612)	5.219 (3.859)	12.605*** (3.886)
Headquarter Dummy: Asia	-0.084 (0.062)		-1.713** (0.793)	0.092* (0.049)	-0.714** (0.337)	0.291 (0.302)	1.281 (4.150)	2.110* (1.117)	12.161 (11.781)	7.191 (12.078)	11.705 (11.850)
Constant	3.018*** (0.139)	-0.713 (1.599)	-5.416*** (1.514)	2.798*** (0.271)	20.325*** (0.373)	9.433*** (1.053)	96.289*** (7.008)	-1.551 (1.712)	-6.289 (18.633)	-1.813 (18.549)	-2.997 (22.915)
Observations	210	210	210	78	210	207	157	197	168	167	170
Adj./McFadden R-squared	0.420	0.032	0.159	0.491	0.078	0.025	0.012	0.446	0.178	0.168	0.189

In contrast, we find that a change in fee levels was more likely to be associated with increases in the size of funds (as indicated by the strongly significant coefficient of 0.698 in regression 3). Consistent with this result, we also find that fund size is significantly and negatively related to the changes in management fees that become effective after the investment period (with a coefficient of -0.251 in regression 4). For example, doubling the fund size is associated with a reduction in the management fee of 0.25 percentage points – for example from 1.50% to 1.25% – after the investment period.⁴²

Overall, the findings show that growth in fund size is associated with lower management fees, a higher probability of changing the level of fees, and lower management fee after the end of the investment period. But none of these relations appear to have changed during the four different time periods in our sample.

Regarding the four major carried interest fund terms, we find that carry percentages, hurdle rates, and catch-up provisions in the buyout industry are largely, if not completely, standardized. They do not change with either fund size or over time. As a consequence, we focused entirely on one feature where there is some variation, namely the distribution of the carry (Table 2-18, regression 8).

As already noted, the European waterfall is more often used in Asia as well as Europe than in North America (as indicated by the positive regional dummies). Interestingly, the negative coefficient of the natural logarithm of fund size suggests that the probability of a European waterfall decreases for larger funds. But this result may reflect mainly the reality that larger funds are usually managed by older GPs that are likely to have started with an American waterfall. In contrast, the positive coefficient of 0.237 for the time variable in

⁴² This relation even remains significant in unreported regressions when we control for the initial management fee in the investment period. This result reflects the fact that the initial management fee also varies with fund size.

regression 8 shows that more recent funds are more likely to include the European waterfall in their LPAs.

The three deal-related fees also have become more investor-friendly, reflecting both a fund size as well as a general market effect (regressions 9 to 11). For example, a doubling of fund size increases the share of transaction fees flowing back to the fund in favor of LPs by 6.0 percentage points. In addition, assuming that fund sizes have remained constant, the share of these fees benefiting LPs has increased by 1.8 percentage points per year over time. Changes in monitoring and abort fees over time and fund size are remarkably similar to the changes in transaction fees, indicating that these terms are usually kept at the same level per fund.

2.3.4.2 Is There a New Normal of Fund Terms?

Overall, today's fund terms appear to have changed remarkably little over the life of the PE industry, deviating from the classical 2/20 compensation structure only in the case of funds with larger amounts of committed capital. The standard management fee is 2.00% for funds of up to USD 1 billion, 1.75% for funds between USD 1 billion and USD 5 billion, and 1.50% for funds with more than USD 5 billion. Furthermore, the buyout industry standard now effectively includes a provision that, at the end of the investment period, the basis for calculating management fees switches from total capital committed to a measure of (net) invested capital. And for most larger funds, the percentage management fees also go down at the end of the investment period. Carried interest fund terms are highly standardized, with 20% carry, hurdle rates of 8%, and catch-up rates of 100% almost universally employed by European as well as North-American-based funds. At the same time, however, whereas European funds follow the more LP-friendly practice of whole-fund distribution of carried

interest, the majority of North American funds continue to distribute carry on a deal-by-deal basis. But even among such funds, there appears to be a clear trend toward the European waterfall. And along with this pro-investor change, the new normal now is 100% of deal-related fees accruing to LPs.

Apart from these structural changes in general fund terms, some LPs have also succeeded in negotiating more favorable terms in their dealings with certain GPs. For example, in the case of about ten funds raised, we find provisions that specify a reduction in the management fee percentage during the investment period for “early closers” – that is, LPs that commit capital early in a fundraising cycle – or in the case of fund sizes above a certain level (for example, fees of only 1% for all funds above USD 8 billion) or for large anchor investors (LPs committing more than USD 300 million pay a 0.25 percentage points lower fee). Second, GPs are increasingly allowing an earlier reduction of the management fee basis or level in cases where successor funds are either set up or achieve a certain fund size. This arrangement is intended to prevent GPs from earning management fees during the investment period simultaneously from two funds. Third, although more often used in venture capital LPAs, we also observe staggered carry interest levels above 20% when GPs are able to generate returns above a certain money multiple (MM) threshold. In one arrangement, for example, carry increases to 25% for all profits above a MM of 3.

Still, the results from the regressions show that fund terms have been remarkably stable over time. LPs have only been able to renegotiate heavily criticized deal-related terms as well as American carry distributions to their advantage and appear to have been successful in persuading some GPs to adjust specific fund terms. Yet, when one looks at the evolution of all fund terms, these achievements seem fairly modest, more like cosmetic changes than structural improvements of fund terms in LPAs.

2.3.4.3 Alignment of Interests between LPs and GPs

Our results show that GPs have been willing (or forced) to reduce their management fees with increasing fund sizes. Yet the essential question remains: Has the sensitivity between changes in fund size and changes in management fees been sufficient to prevent a misalignment of interests between LPs and GPs?

Imagine a simple example of two successive funds from the same GP. Fund 1 has raised a USD 1 billion in capital and a fixed management fee of 2.00%. Assuming an investment phase of five years, this GP would earn USD 100 million in fees during this period. But now we assume that the GP is able to raise a successor fund with USD 2 billion in committed capital, but a lower management fee of 1.84% per year (in line with the sensitivity of Table 2-18, regression 1). In that case, the GP would receive USD 184 million in fixed fees, even though its fixed costs for operating the successor fund would most likely not have increased by 84%.

Thus, to ensure an effective alignment of interests between LPs and GPs, the industry would need a more negative relation between fund size growth and fixed management fee percentages. But, as this example suggests, fixed management fees that were originally justified as covering the ongoing operating costs of GPs appear to have evolved into a main source of GP income. This development has fuelled concerns of industry participants that GPs might be encouraged to accept lower-return deals that enable them to put more LP capital to work.

2.3.4.4 How to Realign the Interests of LPs and GPs?

The buyout industry has matured over time, leading to a proliferation of mega funds with large capital commitments and reductions in gross fund returns. In response to this

development, we propose that an effective realignment of interests can be achieved only by a significant rethinking of fund terms that reflects and responds to the new realities within the buyout industry. In the next section, we explore the possibility of experimenting with a different set of fund terms that go beyond 2/20. In an effort to evaluate different fund terms and their power to limit the conflict of interests between LPs and GPs, we show scenario calculations for various new terms in Table 2-19.

We start by reproducing a number of different cases from Bain Capital's latest fund XI. Bain Capital fund managers have offered their LP clients three different combinations of fund terms. For its mega funds, Bain Capital still offers terms – 1.50/20 and a hurdle rate of 7% – that are consistent with the industry standard. But in addition to the standard option, LPs can choose between a 1.00/30 with a 7% hurdle rate or a 0.50/30 term without any hurdle. Bain Capital advertised the latter by calling it “the ultimate” realignment of interests between LPs and GPs.⁴³

We follow the new fund terms by Bain Capital for two reasons. First, the lower fee percentages better reflect the initial idea of the role of management fees to cover a GP's ongoing operating expenses. Second, a higher carried interest increases the importance of variable compensation for GPs and thus counteracts the overall trend in which lower gross returns are reducing the carry payments to GPs. Furthermore, we suggest that eliminating the hurdle rate in the buyout industry, as offered in Bain Capital's third combination, might be worth considering (even though it appears to favor mainly GPs' interests). Like the 2/20 structure, the 8% hurdle rate may well reflect its origins in a decade with substantially higher interest rates and buyout industry returns. To the extent this is so, LPs should consider exchanging hurdle rates for more investor-friendly management fee terms.

⁴³ See Chicago Tribune, Apr. 15th, 2013, “Bain nears USD 3 billion fundraising close.”

2.3.4.5 Simulation of Three Different Fund Terms

To reduce the complexity of our analysis, we start by using the standardized parameters for fund terms, fund investment behavior, and fund gross returns described below. We use the most common fund term structure in our sample for fund sizes between USD 1 billion and USD 5 billion as our base case being a 1.75% management fee during the investment period on the committed capital and on net invested capital thereafter. Furthermore, the analysis employs a carried interest of 20%, an 8% hurdle rate, a 100% catch-up, and a whole-fund European carried interest distribution, whereas deal-related fees are neglected so that they do not benefit GPs or the fund and thus LPs. Besides these fund terms, the analysis assumes a fund size of USD 2.5 billion (close to our sample mean), an investment period of 5 years with an investment speed of 20% of the committed capital per year, and an average deal holding period of 4.5 years.

To analyze the effects of different fund gross returns on fund terms, we also introduce three fund return scenarios while holding constant the MM returns, whereas the internal rate of return (IRR) resulting from assumptions, such as the deal holding period. Forecasting future gross returns of buyout funds – or even just estimating their final gross returns – on the basis of recent vintage years is by no means a simple or straightforward task. We have our three scenarios below vetted by interviews with industry experts and believe them to be reasonable. At the same time, we also run robustness checks with alternative fund gross returns and reached similar conclusions. The three scenarios include gross returns of 2.25 MM and 20.7% IRR (high returns), 1.75 MM and 13.5% IRR (medium returns), and 1.25 MM and 5.2% IRR (low returns).

Our results, which are summarized in Table 2-19, amount to a simulation of the impact of the new fund terms on the sharing of gains between GPs and their LPs; and, as such, they

can be seen as useful – and indeed critical – for developing a guideline to unify the interests between LPs and GPs.

Table 2-19: Simulation Results of Three Fund Term Cases

This table presents the results of simulations of net money multiple (MM) fund returns and GP compensation for three gross return scenarios and three fund term cases. The three fund terms cases build on similar terms offered by Bain Capital in their fund Bain Capital XI. In slight contrast to their terms, we base our default case 1 on a fund with USD 2.5 billion in committed capital (nearly the mean value in our total sample), which has typically a management fee of 1.75%. Furthermore, we stick to the industry standard of a hurdle rate being 8%, although results will be the same to a hurdle of 7% by Bain Capital in our simplified return scenarios. We choose the three return scenarios so that the low return case lies below the hurdle rate, which is important to evaluate the differences between cases 2 and 3. Panel A presents the net MM fund returns per gross return scenario and fund term case. Panel B shows the total GP compensation consisting of management fee and carried interest payments in USD millions. Panel C displays the share of management fees in comparison to the total GP income and two related averages. The unweighted average is simply the mean of the three return scenarios as given in Panel C. The weighted average is calculated by multiplying the absolute GP income from Panel B with the percentage from Panel C (e.g.: USD 856 million * 33.7%) for each of the three return scenarios, adding them up, and dividing the sum by the sum of the three total GP incomes per case.

	Case 1	Case 2	Case 3
Management Fee	1.75%	1.00%	0.50%
Carried Interest	20%	30%	30%
Hurdle Rate	8%	8%	0%
Panel A: Net MM Return			
High (2.25 Gross MM)	1.91	1.83	1.85
Medium (1.75 Gross MM)	1.51	1.48	1.50
Low (1.25 Gross MM)	1.13	1.18	1.15
Average	1.52	1.50	1.50
Panel B: Total GP Income [USD million]			
High (2.25 Gross MM)	856	1,053	995
Medium (1.75 Gross MM)	606	678	620
Low (1.25 Gross MM)	289	165	245
Average	584	632	620
Panel C: Share of GP Income from Mgmt. Fee			
High (2.25 Gross MM)	33.7%	15.7%	8.3%
Medium (1.75 Gross MM)	47.6%	24.3%	13.3%
Low (1.25 Gross MM)	100.0%	100.0%	33.6%
Average (Unweighted)	60.5%	46.7%	18.4%
Average (Weighted)	49.5%	26.1%	13.3%

We start with Panel A in Table 2-19, which reports the net MM fund returns for the three return scenarios and three different fund terms (Cases 1, 2, and 3). To illustrate our findings, for example, a fund with a gross MM return of 2.25 (and in our case a 20.7% IRR) generates a net MM of 1.91 for LPs when applying the standard structure 1.75/20, and after deducting management fees and carried interest. As expected, in the high gross return scenario, net returns for LPs are lowest with the new fund terms that include a 30% carry. In contrast, the new terms benefit LPs in the low gross return scenario, with slightly higher net return MMs of 1.18 and 1.15 in comparison to 1.13 for the classical terms. The simple averages for the three sets of fund terms are very similar, with net returns about 1.50 MM, although the new terms cause LPs to forgo a slightly larger part of the upside.

Panel B displays the total GP income, including fixed and variable compensation in total dollars. By definition, the total GP income shows the exact opposite pattern as the net returns in Panel A, since net MM returns to LPs are calculated by deducting total GP income from gross MM returns.

More interestingly, Panel C presents the share of GP income from fixed management fees. As expected, the share decreases under the new fund terms since management fees are reduced but carry interest gains more importance. When taking the simple average of the results from the three return scenarios, GP income from management fees amounts to just over 60% when using the classical fund terms and it remains as high as 49.5% when weighting the share by the different levels of GP income per scenario.⁴⁴ By far the lowest

⁴⁴ Metrick and Yasuda (2010) find that two-thirds of all fees are related to fixed fees, which is slightly higher than our results of 50% to 60%. The gap can be explained by different assumptions: They assume a higher management fee of 2% in the investment period and also include deal-related fees in their calculations. For example, they classify the usually more important transactions fees as fixed but monitoring fees as variable compensation. Furthermore, they calculate the present value of all future income streams, which increases the importance of management fees as they are mainly paid early in a fund's life, whereas carry interest is usually paid later.

shares of fixed to total GP compensation – of 18.4% unweighted average, respectively, 13.3% weighted average – are found in case 3 of Table 2-19, Panel C, with a 0.50/30 term structure with no hurdle rate. These results clearly show the important effect of the new fund terms on the mix of GP compensation and their potential role in realigning the interests in the buyout industry. In brief, such terms could go a long way in ensuring that GPs' main source of income is fund returns and not fund sizes anymore.

2.3.5 Conclusion

Buyout fund terms have faced increased criticism from LPs, especially since 2007, and we accordingly have expected such terms to become more investor-friendly. Instead, our analysis contributes to the existing literature by providing evidence that fund terms have been surprisingly stable over time and by describing the standard fund terms in the buyout industry.

Nevertheless, during the same period, we have also observed a renegotiation of certain heavily criticized fund terms, as reflected in the growing use (by American GPs) of whole-fund carry distribution and the flow-through of 100% of deal-related fees to the fund and hence the LPs. These changes, while representing an end to some of the buyout industry's least defensible practices, fall well short of a significant reform of fund terms.

As evidence of this, we find that the recent increase in fund sizes are the only important force for bringing about reductions in percentage management fees. More specifically, in the past GPs have accepted lower fund terms, especially management fee terms, only when associated with increasing fund sizes. In such cases, GPs appear to have volunteered to sacrifice moderately lower percentage management fees for higher absolute (dollar) fees. We provide a simple example to illustrate this tradeoff, which shows that the

management fee reductions that have accompanied the increase in fund sizes have done little to address what appears to be a conflict of interest between LPs and GPs over the optimal buyout fund size. As evidence of this conflict, industry participants have criticized GPs for the excessive focus on increasing fund sizes, at the expense of net returns to LPs.

We shed light on this conflict of interest by evaluating innovative fund terms now being offered by one reputable GP Bain Capital. In an attempt to simulate the effects of the new fund terms on the division of gains between GPs and LPs – and hence on GP incentives to invest in only high-return portfolio companies – we apply three gross fund return scenarios (high, medium, low) to three different sets of fund terms, including variations of Bain Capital’s 1/30 structure with an 8% hurdle rate and a 0.50/30 structure with no hurdle.⁴⁵ Our results, while showing nearly similar net fund returns for LPs in all three term cases, show GPs earning a slightly higher total compensation under the new fund terms.⁴⁶ And most important, we find that the 0.50/30 structure reduces the share of management fees to total GP compensation from around 55% to below 20% on average. Our view is that by ensuring that as much as 80% of the total GP income depends on fund returns, the buyout industry will have gone a long way toward creating fund terms that reflect the new realities in the buyout industry while unifying the interests of both parties involved.⁴⁷

⁴⁵ Bain Capital XI, with a target fund size of USD 6 billion offered three different sets of fund terms: 1.50%/20 with a 7% hurdle, 1.00%/30 with a 7% hurdle, 0.50%/30 without a hurdle. We slightly deviate from these cases for two reasons. First, we choose a hurdle rate of 8% as this rate reflects the industry standard, although the results are unaffected by the choice between 7% or 8%. Second, we base the calculation on a hypothetical fund with USD 2.5 billion in size because this amount lies near our sample mean.

⁴⁶ Yet, this difference might be negligible as we do not incorporate the timing of cash flows to GPs (which is usually earlier in a fund life for management fees) and especially as the mix between management fee and carried interest income of GPs differs strongly.

⁴⁷ On the other hand, because the management fee should allow GPs to cover their operating costs, a 0.50% fee might only be feasible for mega funds, whereas smaller funds most likely need higher fees at about 1.00% or more for this purpose.

Whether GPs or LPs will benefit from these innovative terms will depend mainly, of course, on the funds' gross returns. By definition, lower fixed fees and a larger carried interest of 30% provide greater benefits for GPs when fund returns are higher, whereas LPs are effectively given more protection when fund returns are low. Yet, these optimization efforts overlook the broad purpose of the new fund terms: to provide GPs with stronger incentives to generate higher returns. To the extent the new terms succeed in bringing about this effect, both GPs and LPs will benefit in the end.

Thus, although fund terms have been quite stable in the past decades, we believe that terms like those offered by Bain Capital have a good chance to become the new industry standard, and for two reasons: First, LPs have become increasingly "professionalized," which has led to greater focus on GP compensation and ways of realigning their interests with LPs'. A second important consideration is that the "signaling" benefits for – in the form of the message of confidence sent by – those GPs willing to distinguish themselves by offering terms like 1.00/30 or 0.50/30 could encourage more GPs to move in this direction. Presumably, only those GPs most confident about their ability to deliver higher returns will offer these new terms. Conversely, for all but the most reputable and established PE buyout firms, those GPs that do not offer the new terms may well be seen as signaling little confidence in their abilities to do what they are being paid to do: namely, producing above-market returns.

One last point in closing: Our main target in this analyzes is not to challenge the total compensation of GPs, but rather the proportion of it constituted by fixed management fees. We remain agnostic about the total level of compensation, recognizing the possibility that it could well reflect the distinctive capabilities of GPs as well as the supply of and demand for committed capital. Our focus is instead what we view as a clear misalignment of interests

between LPs and GPs, and hence our prescription suggestions are meant to bear on only those fund terms that affect the mix between fix and variable GP compensation, not the total amount. Our working premise is the importance of rediscovering the original purpose of management fees: enabling GPs to cover the costs of running their daily operations. If one begins with and continues to uphold this premise, carried interest should be the only major source of GP income.

3 Conclusion

The buyout segment has developed from a niche to a well-established asset class. Historic outperformance has attracted very large amounts of institutional money that on the one hand has driven down fund returns and on the other hand has affected the fundamental industry mechanics. This dissertation analyzes whether institutional investors of the buyout asset class have to adjust their behavior to the new realities, and if so, in which way?

In the context of this overarching research topic, this dissertation develops three specific research questions, which are based on the most important questions LPs have to answer in their daily business: Should I invest in the new fund of a GP? If yes, how much money should I invest? To which fund terms should I invest?

This dissertation provides practical implications for LPs to help answer these questions in a buyout industry that has grown dramatically over the last two decades. Ultimately, the aim of this dissertation is to facilitate a sustainable relation between LPs and GPs that will benefit all market participants in the long-run.

Related to the question whether a LP should invest in the new fund of a GP, the first essay of this dissertation asks specifically whether the past fund performance of a GP is still a good predictor for the performance of the successor fund. In this respect, it deals with a distinct feature of the PE asset class, which was first documented by Kaplan and Schoar (2005). They reported performance persistence between two consecutive funds of one GP, in their sample of realized funds with vintage years until about 1996. Clearly, this performance persistence confirms the strategy of LPs who mainly try to allocate their capital to the follow-on funds of top performing GPs.

The results of the first essay confirm the early findings of Kaplan and Schoar (2005) and explain further that the strong top quartile performance persistence in the early days of the industry was mainly driven by experienced GPs with dozens of executed deals. In more recent years, the top quartile persistence of these experienced GPs decreased so that in the later period the probability of staying in the top quartile with a new fund is exactly 25%, the random walk probability. Thus, past performance no longer indicates future performance.

The essay adds to the emerging literature on performance persistence by using a new deal-level data set and a new deal-level approach to circumvent various methodological challenges. Furthermore, the analysis identifies the source of the past performance persistence, which is the GPs' experience.

The implications for LPs are expansive as their major fund selection criteria has been the past performance of GPs. This recipe has lost its power. In turn, LPs have to increasingly go beyond past performance of GPs to find the outperforming funds of the future. For example, LPs should better understand the elements of past performance. What is the distribution of deal returns, i.e., are deal returns homogeneously distributed or are fund returns driven by a few outliers that might reflect luck instead of skill? What was the value-creation strategy of the past deals? Is it sustainable? Who was the responsible partner or investment manager for particular successful deals? Is there a specific "employee-alpha"? Besides understanding the ingredients of past performance, LPs should also look more closely at the current setting for a new fund. For example, are the fund managers still "hungry"? Does their investment strategy still fit the current market situation and fund size?

All these questions are highly relevant for practitioners and also offer fruitful topics for future research. For example, researchers can empirically analyze these questions to improve our understanding of what drives performance in the PE industry. From a principal-agency

theory perspective, the better understanding of performance drivers will decrease information asymmetries and consequently agency costs.

Related to the question of how much money a LP should invest in a new fund, the second essay investigates the controversial topic of the costs of buyout fund investing by analyzing the costs and drivers of dry powder in buyout funds. Historically, LPs have been mainly interested in their final net returns, which until recently have usually been quite satisfying. However, as fund net returns have decreased and LPs have professionalized, the cost of investing into the buyout asset class has been increasingly questioned by LPs.

The results expand the existing literature in several ways. First, the essay provides a spread benchmark for different levels of gross returns. Second, the essay presents the return spread over time. At first glance, the return spread has decreased over time, suggesting that investing into buyout funds has become less costly. Yet, this development is mainly driven by lower gross returns and, consequently, lower performance-based GP compensation. When controlling for the lower gross returns, the cost of buyout investing has actually increased. Third, the essay shows that the increasing amounts of unused capital are an important driver of the increased costs over time, which has so far been mainly disregarded in academic research. Fourth, the analysis investigates the drivers of dry powder and differentiates between internal GP-related and external market-related drivers. The results show that dry powder is primarily influenced by market-related factors, as opposed to GP-related factors, such as the capital flows into the buyout industry and the availability of deals.

The results also provide important implications for the investors of buyout funds. On the one hand, the essay addresses the recent discussion on capital over-allocation to the buyout industry. Increasing capital flows do not only drive down gross (and net) returns, but also increase the issue of dry powder (depending on the market conditions) making the buyout

asset class more expensive, which further drives down its net returns. On the other hand, the essay provides implications for LPs not only regarding the *amount* of capital allocated to buyout funds, but also regarding the *timing*. For example, in the last decades allocating money to buyout funds in boom periods was accompanied by large amounts of costly dry powder, whereas dry powder has rarely been a problem when committing capital in bust periods and, consequently, usually in smaller funds. Overall, the results imply that LPs should – if possible with their overall asset allocation – follow a rather anti-cyclical investment strategy.

As this dissertation quantifies the substantial costs of dry powder, a consecutive step for future research might be the contractual design of LPAs with respect to unused capital. This topic exhibits a high level of complexity due to the problem of moral hazard. If the LPAs between LPs and GPs simply include a penalty for a low deployment rate, GPs might forego a rigorous deal selection, which would lower fund returns for LPs, in order to maximize a fund's deployment rate.

Related to the question to which fund terms LPs should invest in a new fund of a GP, the third essay analyzes alternative terms for buyout funds to realign the interests of LPs and GPs. Historically, fund terms mainly include a fixed yearly management fee of 2% on the committed capital and a success-based compensation of 20% of all fund profits. The initial idea of the 2% management fee has been to allow GPs to cover their ongoing operating expenses, while their major source of income should rest upon the success-based compensation. Yet, as fund sizes have increased over the decades while returns, and therefore success-based fees, have decreased, the major driver of GP income switched to the fixed management fee.

In this respect, the third essay shows that the fund terms have been surprisingly stable over time, the only exceptions being success-based fee distribution rules and deal-related fees

that have become more LP-friendly. In contrast, the major driver of changes in fund terms is changes in fund sizes. However, the sensitivity between an increase in fund sizes and a decrease in the fixed management fee percentages, as the most important fixed fund term component, has been far too low, resulting in increasing fixed management fees in absolute terms for GPs. This low sensitivity has caused a misalignment of interests between LPs and GPs over time.

To realign the interests, the third essay simulates the share of GP income from fixed and variable components for standard fund terms and also for two alternative terms with lower management fee percentages and higher success-based compensation. The results show that fixed fees with the standard terms, in line with Metrick and Yasuda (2010), make up around 60% of total GP compensation. In contrast, the alternative terms with a carried interest of 30% and a management fee up to 1% reduce the share of fixed to total GP compensation to around 20%.

The results advance the academic literature by providing an analysis of fund terms over time and of their respective drivers. This analysis makes it possible to exactly evaluate whether the conflict of interests between LPs and GPs has worsened over time.

Essay 3 has important implications for the LP-GP relation as it provides guidance towards a more balanced incentive structure that helps realigning the interests of LPs and GPs. The alternative fund terms should push GPs to do what they are paid for: generate decent returns for their investors opposed to maximize their assets under management and thus their fixed management fees.

As various academic studies have already analyzed controversial topics concerning PE fund terms, future research should focus on new developments in the industry that aim at reducing the high costs of investing into buyout funds. For example, McCahery and

Vermeulen (2013) show that LPs increasingly co-invest in LBO transactions and in this way circumvent GPs (and avoid paying their compensation). Furthermore, GPs increasingly offer separate accounts for their LPs with lower fees, depending on the level of the committed capital. Although these developments at first glance make buyout fund investing less costly, researchers can investigate whether new agency conflicts evolve.⁴⁸

Overall, the progress of market participants' behavior has dropped behind the development of the buyout industry in terms of capital commitments. To put it more clearly, today the global buyout industry manages more than one trillion USD⁴⁹, whereas many market participants still follow old behavioral patterns. LPs still aggressively commit capital to the top performing funds of the past, although this strategy displays an artefact of the early days. Furthermore, LPs follow the distinct boom and bust periods when committing capital to GPs, although a more anti-cyclical strategy promises better net returns. On the other hand, GPs still tend to maximize their assets under management in order to reap profits from fixed management fees. These points show that the buyout segment and its participants still have to advance before it can be called a "mature asset class". This dissertation sheds light on the necessary steps towards this goal.

First, LPs have to advance their fund selection. Interestingly, in contrast to top quartile performance persistence, so-called flop quartile persistence remains high, although these funds should actually drop out of the market. The road to success to an enhanced fund

⁴⁸ A study from Altius, a PE advisory and fund-of-funds investor, points out the problem of adverse selection for co-investments. LPs might be offered simply the least attractive deals as co-investments. In that sense, LPs would trade lower fees for lower returns (see PE Manager, Apr. 4th, 2014, "Co-invests likely to underperform.").

⁴⁹ The Global Private Equity Report 2014 of Bain & Company shows unrealized values of portfolio companies of USD 930 billion and USD 399 billion of liquid dry powder (p. 6 and p. 10), for buyout funds only.

selection lies in increased transparency by means of improved fund reporting and industry-wide databases, such as the Burgiss research consortium.⁵⁰

Second, LPs have to verify whether their return expectations for buyout investments are realistic. The historical attractive returns or the current positive market sentiment are most likely a misleading basis for forecasting buyout fund returns. For example, the currently increasing pricing multiples for LBOs benefit recent funds' exits; however, transaction pricing multiples will not increase forever.⁵¹ Thus, LPs have to carefully reconsider the cyclicity, risk, and costs of the asset class when forming their return expectations, which will ultimately also determine the amount and timing of their commitments.

Third and most importantly, the buyout industry should move towards fund terms that realign the interests of LPs and GPs. In this sense, performance-based fees should be the major source of GP income. As fund returns usually deteriorate with increasing capital commitments (Harris et al., 2013a), more performance-based terms might reduce the incentive for GPs to raise ever increasing funds. This would mitigate the key problem of capital over-allocation in the buyout industry.

To conclude, the evolution of the buyout industry in the last decades is impressive. Still, the buyout segment is a maturing, but not yet a mature industry.

⁵⁰ Burgiss, a provider of portfolio management software, data and analytics to PE investors, and academic researchers formed the Private Equity Research Consortium (PERC) to increase the data quality for PE research and industry professionals.

⁵¹ See Global Private Equity Report 2014 of Bain & Company (p. 35).

References

ACHARYA, V. V., GOTTSCHALG, O. F., HAHN, M. & KEHOE, C. 2013. Corporate governance and value creation: Evidence from private equity. *The Review of Financial Studies*, 26, 368-402.

ACHLEITNER, A.-K., BRAUN, R. & ENGEL, N. 2011. Value creation and pricing in buyouts: Empirical evidence from Europe and North America. *Review of Financial Economics*, 20, 146-161.

ACHLEITNER, A.-K., BRAUN, R., ENGEL, N., FIGGE, C. & TAPPEINER, F. 2010. Value creation drivers in private equity buyouts: Empirical evidence from Europe. *Journal of Private Equity*, 13, 17-27.

ACHLEITNER, A.-K., ENGEL, N. & REINER, U. 2013. The performance of venture capital investments: Do investors overreact? *Review of Financial Economics*, 22, 20–35.

AKERLOF, G. A. 1970. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84, 488-500.

ANG, A., AYALA, A. & GOETZMANN, W. N. 2013. Investment beliefs of endowments. Working paper, NBER.

ANG, A., PAPANIKOLAOU, D. & WESTERFIELD, M. M. 2012. Portfolio choice with illiquid assets. Working paper, NBER.

AUSTIN, S. 2009. Majority of VCs in Survey Call Industry 'Broken'. *The Wall Street Journal*.

AXELSON, U., JENKINSON, T., STRÖMBERG, P. & WEISBACH, M. S. 2013. Borrow cheap, buy high? The determinants of leverage and pricing in buyouts. *Journal of Finance*, 68, 2223-2267.

BAIN & COMPANY. 2014. *Global Private Equity Report*. Boston.

BERHOLD, M. 1971. A theory of linear profit-sharing incentives. *The Quarterly Journal of Economics*, 85, 460-482.

BOBELDIJK, Y. 2014. Co-invests likely to underperform. PE Manager.

BRAUN, R., JENKINSON, T. & STOFF, I. 2013. How persistent is private equity performance? Evidence from deal-level data. Working paper, Friedrich-Alexander University, Erlangen-Nuremberg.

BRAUN, R. & SCHMIDT, M. 2014. The limited partnership model in private equity: Deal returns over fund's life. Working paper, Friedrich-Alexander University, Erlangen-Nuremberg.

BROWN, G. W., GREDIL, O. R. & KAPLAN, S. N. 2013. Do private equity funds game returns? Working paper, University of North Carolina (UNC).

CHERTOK, S. & BRAENDEL, A. 2010a. Closed-end private equity funds: A detailed overview of fund business terms, part I. *Journal of Private Equity*, 13, Spring, 33-54.

CHERTOK, S. & BRAENDEL, A. 2010b. Closed-end private equity funds: A detailed overview of fund business terms, part II. *Journal of Private Equity*, 13, Summer, 1-27.

CHICAGO, T. 2013. Bain nears USD 3 billion fundraising close. *Chicago Tribune*.

CHUNG, J.-W. 2012. Performance persistence in private equity funds. Working paper, Korea University.

CHUNG, J.-W., SENSOY, B. A., STERN, L. H. & WEISBACH, M. S. 2012. Pay for performance from future fund flows: The case of private equity. *Review of Financial Studies*, 25, 3259–3304.

CUMMING, D., SIEGEL, D. S. & WRIGHT, M. 2007. Private equity, leveraged buyouts and governance. *Journal of Corporate Finance*, 13, 439-460.

DEMIROGLU, C. & JAMES, C. M. 2010. The role of private equity group reputation in LBO financing. *Journal of Financial Economics*, 96, 306-330.

DRIESSEN, J., LIN, T.-C. & PHALIPPOU, L. 2011. A new method to estimate risk and return of non-traded assets from cash flows: The case of private equity funds. Working paper, Tilburg University.

- ECONOMIST, T. 2010. Private inequity. *The Economist*.
- ECONOMIST, T. 2011. Fee high so dumb. *The Economist*.
- ECONOMIST, T. 2013. When the music gets you. *The Economist*.
- ECONOMIST, T. 2013. Zombies at the gates. *The Economist*.
- EISENHARDT, K. M. 1989. Agency theory: An assessment and review. *Academy of Management Review*, 14, 57-74.
- FAMA, E. F. & JENSEN, M. C. 1983. Separation of ownership and control. *Journal of Law and Economics*, 26, 301-325.
- FANG, L. H., IVASHINA, V. & LERNER, J. 2013. The disintermediation of financial markets: Direct investing in private equity. Working paper, INSEAD.
- FRANZONI, F., ERIC N. & PHALIPPOU, L. 2012. Private equity performance and liquidity risk. *Journal of Finance*, 67, 2341–2373.
- FUNG, W., HSIEH, D. A., NAIK, N. Y. & RAMADORAI, T. 2008. Hedge funds: Performance, risk, and capital formation. *Journal of Finance*, 63, 1777–1803.
- GOMPERS, P. A. & LERNER, J. 1999. An analysis of compensation in the U.S. venture capital partnership. *Journal of Financial Economics*, 51, 3–44.
- GOMPERS, P. A. & LERNER, J. 2000. Money chasing deals? The impact of fund inflows on private equity valuations. *Journal of Financial Economics*, 55, 281–325.
- GREDIL, O., GRIFFITHS, B. E. & STUCKE, R. 2014. Benchmarking private equity: The direct alpha method. Working paper, University of North Carolina Kenan-Flagler Business School.
- GUO, S., HOTCHKISS, E. S. & SONG, W. 2011. Do buyouts (still) create value? *Journal of Finance*, 66, 479-517.
- HARRIS, R. S., JENKINSON, T. & KAPLAN, S. N. 2013a. Private equity performance: What do we know? *Journal of Finance*, forthcoming.

- HARRIS, R. S., JENKINSON, T. & STUCKE, R. 2012. Are too many private equity funds top quartile? *Journal of Applied Corporate Finance*, 24, 77-89.
- HARRIS, R. S., JENKINSON, T., KAPLAN, S. N. & STUCKE, R. 2013b. Has persistence persisted in private equity? Evidence from buyout and venture capital funds. Working paper, UAI Foundation, University of Virginia.
- HIGSON, C. & STUCKE, R. 2012. The performance of private equity. Working paper, London Business School.
- HOCHBERG, Y. V., LJUNGQVIST, A. & VISSING-JORGENSEN, A. 2012, Informational hold-up and performance persistence in venture capital. Working paper, NBER.
- HUDEC, A. 2010. Negotiating private equity fund terms. *Business Law Today*, 19, 45-49.
- JENKINSON, T., SOUSA, M. & STUCKE, R. 2013. How fair are the valuations of private equity funds? Working paper, University of Oxford.
- JENSEN, M. C. & MECKLING, W. H. 1976. Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics*, 3, 305-360.
- KAPLAN, S. N. 1989. The effects of management buyouts on operating performance and value. *Journal of Financial Economics*, 24, 217-254.
- KAPLAN, S. N. & SCHOAR, A. 2005. Private equity performance: Returns, persistence, and capital flows. *Journal of Finance*, 60, 1791-1823.
- KAPLAN, S. N. & STRÖMBERG, P. 2009. Leveraged buyouts and private equity. *Journal of Economic Perspectives*, 23, 121-146.
- KASERER, C. 2011. Return attribution in mid-market buy-out transactions – New evidence from Europe. Working paper, Technical University Munich.
- KASERER, C. & DILLER, C. 2005. Private equity funds: Return characteristics, return drivers and consequences for investors. *Capco Institute Journal of Financial Transformation*, 2005, 107-119.

- KORTEWEG, A. & SORENSEN, M. 2013. Decomposing private equity performance persistence, Working paper, Stanford University.
- LELAND, H. E. & PYLE, D. H. 1977. Informational asymmetries, financial structure, and financial. *Journal of Finance*, 32, 371-387.
- LERNER, J. 2011. The future of private equity. *European Financial Management*, 17, 423-435.
- LERNER, J., SCHOAR, A. & WONGSUNWAI, W. 2007. Smart institutions, foolish choices: The limited partner performance puzzle. *Journal of Finance*, 62, 731-764.
- LITVAK, K. 2009. Venture capital limited partnership agreements: Understanding compensation arrangements. *University of Chicago Law Review*, 76, 161–218.
- LOPEZ-DE-SILANES, F., PHALIPPOU, L. & GOTTSCHALG, O. F. 2013. Giants at the gate: On the cross-section of private equity returns. *Journal of Financial and Quantitative Analysis*, forthcoming.
- LJUNGQVIST, A. & RICHARDSON, M. 2003. The cash flow, return and risk characteristics of private equity. Working paper, NBER.
- LJUNGQVIST, A., RICHARDSON, M. & WOLFENZON, D. 2008. The investment behavior of buyout funds: Theory and evidence. Working paper, NBER.
- MCCAHERY, J. A. & VERMEULEN, E. P. M. 2013. Recasting private equity funds after the financial crisis: The end of ‘two and twenty’ and the emergence of co-investment and separate account arrangements. Working paper, Tilburg University.
- METRICK, A. 2007. *Venture capital and the finance of innovation*. John Wiley & Sons.
- METRICK, A. & YASUDA, A. 2010. The economics of private equity funds. *Review of Financial Studies*, 23, 2303-2341.
- METRICK, A. & YASUDA, A. 2011. Venture capital and other private equity: A survey. *European Financial Management*, 17, 619–654.

MULCAHY, D., WEEKS, B. & BRADLEY, S. B. 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. Ewing Marion Kauffman Foundation.

PHALIPPOU, L. 2008. The hazards of using IRR to measure performance: The case of private equity. *Journal of Performance Measurement*, 12, 55-66.

PHALIPPOU, L. 2009. Beware of venturing into private equity. *Journal of Economic Perspectives*, 23, 147-166.

PHALIPPOU, L. & GOTTSCHALG, O. F. 2009. The performance of private equity funds. *Review of Financial Studies*, 22, 1747-1776.

PICARD, P. 1987. On the design of incentive schemes under moral hazard and adverse selection. *Journal of Public Economics*, 33, 305-331.

PICOT, A., DIETL, H. & FRANCK, E. 2008. *Organisation*, Schäffer-Poeschel Verlag.

PRIVATE EQUITY GROWTH CAPITAL COUNCIL. 2013. *Private Equity Trends Q2 2013*. Washington.

ROBINSON, D. T. & SENSOY, B. A. 2013a. Cyclical, performance measurement, and cash flow liquidity in private equity. Working paper, NBER.

ROBINSON, D. T. & SENSOY, B. A. 2013b. Do private equity fund managers earn their fees? Compensation, ownership, and cash flow performance. *Review of Financial Studies*, forthcoming.

ROSS, S. A. 1973. The economic theory of agency: The principal's problem. *American Economic Review*, 63, 134-139.

SCHAEFER, D. 2011. Private equity trapped in zombie funds. *The Financial Times*.

SCHAEFER, D. 2012. Investors push for private equity fee cuts. *The Financial Times*.

SENSOY, B. A., WANG, Y. & WEISBACH, M. S. 2013. Limited partner performance and the maturing of the private equity industry. Working paper, NBER.

SPENCE, M. 1973. Job market signaling. *The Quarterly Journal of Economics*, 87, 355-374.

STOFF, I. & BRAUN, R. 2014. Evolution of private equity fund terms beyond 2 and 20. *Journal of Applied Corporate Finance*, 26, 65-75.

STUCKE, R. 2011. Updating history. Working paper, University of Oxford.

TALMOR, E. & VASVARI, F. 2011. *International Private Equity*. John Wiley & Sons.