IMPACT OF PRIVATE EQUITY

Handbook in Economics: Corporate Finance 1 Private Equity and Entrepreneurial Finance, Elsevier, forthcoming

Morten Sorensen
Tuck School of Business

Ayako Yasuda
University of California, Davis

May 2022

Abstract

We survey the academic literature about the impact of private equity investments in the broader economy. Private equity fund managers respond to high-powered incentives and seek to maximize shareholder values via a variety of channels. The literature identifies two broad approaches to value creation taken by private equity funds with sharply divergent outcomes for stakeholders and the aggregate economy. The first approach, associated with public-to-private deals, exploits leverage and interest tax shields, cost reduction, and operating margin improvement. The second approach, associated with private-to-private deals, targets growth-oriented and capital-constrained companies and adds value by relaxing financing constraint, imparting operational and managerial expertise, increasing investment, and inducing top-line revenue growth. Innovation tends to increase with the latter approach (private-to-private deals) while it either declines relatively or becomes more narrowly focused with the former approach (public-to-private deals). For employees, post-buyout high-skilled workers tend to benefit from increased IT investments and upskilling in the jobs, whereas low-skilled workers tend to be hurt from automation and job cuts. For consumers, private-to-private deals imply greater variety and broader geographic availability of products, whereas public-to-private deals imply higher prices and reduced availability. In regulated or subsidized industries, distortion in incentives given by the regulatory framework tends to get magnified when combined with high-powered incentives of private equity. The literature provides evidence of this in healthcare, for-profit education, insurance, and the fracking industry. Collectively, the emerging evidence suggests that welfare outcomes for the broader environment and society depend sharply on the regulatory and competitive structures within which the private equity portfolio companies operate. Thus, regulators need to consider the impact of the high-powered incentives of private equity when assessing the market impact of a given regulatory policy or decision. Finally, impact funds are posited as a mechanism for explicitly aligning the shareholder preferences with the broader public interest. Impact fund investors derive utility from holding impact funds that generate positive impact, and thus are rationally willing to invest in them even though their expected financial return alone may be lower than that from investing in non-impact private equity funds. The result is consistent with the theory of sustainable investing in equilibrium with explicitly pro-ESG investors. Suggestions for future research are discussed.

Keywords: Private equity; leveraged buyout; stakeholders; sustainable investing; impact investing

---

We thank Gordon Phillips (editor) for helpful comments and suggestions. All errors and omissions are our own. Morten Sorensen, Dartmouth College, Tuck School of Business, 100 Tuck Hall, Hanover, NH 03755, email: morten.sorensen@tuck.dartmouth.edu; and Ayako Yasuda, University of California Davis, Graduate School of Management, 3206 Gallagher Hall, Davis CA, 95616-8609, email: asyasuda@ucdavis.edu.

Table of Contents

1. Introduction 3

Traditional Dichotomy: Jensen (1989) versus Shleifer and Summers (1988) 4

2. Incentives in Private Equity 5

2.1 Incentives for maximizing deal profits 6
2.2 Incentives of a fund, deal flow 9
2.3 Incentives of private equity firms, raising future funds 10
2.4 Growth of Environmental, Social, and Governance (ESG) Investing and Private Equity 11

3. Private Equity Management Model 13

3.1 Principal-Agent Problems in Public Corporations 13
3.2 Posited Solutions with Private Equity Management Model 1.0: Traditional View 14
3.3 Posited Solutions with Private Equity Management Model 2.0: Growth, Competency, and non-Financial Outcomes 15
3.4 Ambiguity of the Effect of Private Equity Management Model on Stakeholders and Society at Large 16

4. Impact on Management, Productivity, and Efficiency 17

4.1 Executives, Management, Board of Directors 17
4.2 Productivity and Efficiency 20
4.3 Financial Distress and Bankruptcy 26

5. Impact on Employees and Other Stakeholders 26

5.1 Employment and Wages at Portfolio Companies 27
5.2 Impact on Workers' Welfare 29
5.3 Other Stakeholders 31

6. Impact on Broader Society 31

6.1 Impact on Consumers through Portfolio Company Products and Services 32
6.2 Impact on Governments and Taxpayers 35
6.3 Impact on the Environment 38
6.4 Impact on Innovation and Industry Spillovers 39
6.5 Pro-ESG Investing and Private Equity 41

7. Open Questions and Suggestions for Further Research 43

8. References 45

1. INTRODUCTION

Private equity investments are growing, as documented extensively elsewhere in this volume. A major form of private equity is buyouts, also known as leveraged buyouts (LBOs), where private equity funds, also known as buyout funds, acquire majority equity stakes in portfolio companies using large amounts of debt financing. Buyout funds are active investors, and they are intimately involved with and exert substantial control over the portfolio companies they acquire. In fact, it is reasonable to consider private equity ownership as a separate corporate governance form—distinct from, for example, being a publicly-traded or a family-owned business. There is now a substantial body of academic literature that investigates the impact on and implications of this governance form for (i) the portfolio company itself (including its management), (ii) employees and other stakeholders, and (iii) the broader society, including consumers, governments, the environment, and the industry. This chapter summarizes this literature and its main findings. We further review the nascent literature on impact funds—a type of dual-objective private equity fund—and frame its findings within the rapidly-growing literature on the environmental, social and governance (ESG) and sustainable investing.

Existing Surveys

In addition to the other chapters in this volume, there are a few other academic surveys of private equity. Brown, Harris, Jenkinson, Kaplan, and Robinson (2020) provide a historical perspective of the development of the private equity industry and its performance. Kaplan and Strömberg (2009) describe the organization of private equity firms and their effects on the operations of portfolio companies. Metrick and Yasuda (2011) highlight the importance of private ownership, and the information asymmetry and illiquidity associated with private ownership, as a key explanatory factor of what makes private equity different from other asset classes.

In contrast to these other surveys, this survey focuses on the impact of private equity ownership on the broader society and aims to shed light on critical market, industry, and regulatory factors that determine when the impact of private equity ownership is positive or negative, and for whom. A key insight is that there is an inherent ambiguity in the sign of the relationship between private equity ownership and the societal impact of the firm's activities when the private equity fund managers do not internalize those societal impacts in their objective functions.

Early studies

The modern version of buyouts dates back to the 1980s. The seminal acquisition of RJR Nabisco by Kohlberg, Kravis, and Roberts & Co, vividly described in the book "Barbarians at the Gate" (Burrough and Helyar 1990) took place in 1988. Early studies focused on analyzing the impact of private equity ownership on shareholders. Kaplan (1989a) analyzes data for 76 buyouts (LBOs and management buyouts) and documents that these transactions were followed by significant increases in operating margins and cash flows for the involved portfolio companies, both measured in absolute terms and relative to similar public companies. In another early study, Lichtenberg and Siegel (1990) find that the total factor productivity (TFP) increases for U.S. manufacturing companies that were acquired in buyouts. While recent studies add nuance and depth to our understanding of private equity transactions, the general conclusion that private equity transactions on average create value for, or at least do not harm, shareholders of portfolio companies emerged during this period and remains unchallenged.

Traditional Dichotomy: Jensen (1989) versus Shleifer and Summers (1988)

Discussions of the impact of buyout transactions are often framed as a horse race between Shleifer and Summers (1988) and Jensen (1989). Jensen (1989) emphasizes efficiency gains from private equity transactions. In particular, he argues that public companies suffer from agency problems due to free cash flows, and that private equity ownership with increased leverage can mitigate this problem and both improve social welfare and create gains for shareholders. While he recognizes the possibility of shareholder gains being offset by losses to other financial constituencies (e.g., bondholders) and stakeholders (e.g., employees), he argues that these losses are smaller and short-term compared to the efficiency gains that are long-term, such that takeovers are socially desirable.

In contrast, Shleifer and Summers (1988) emphasize that returns from private equity can often result from transfers from stakeholders, such as employees and suppliers, and not from efficiency gains. Thus, private equity gains need not equate social welfare gains. By the same token, worker layoffs need not equate social welfare loss, though the authors point out that "the redistribution is probably antiegalitarian." (p. 35)

In addition, Shleifer and Summers (1988) note that there are potential negative spillover effects from buyouts, such as the decline of the local economy after factory closures. They propose a particular mechanism of efficiency loss due to breach of implicit contracts. These negative spillover effects could result in overall social welfare loss and make private equity socially undesirable.

While Shleifer and Summers (1988) conjecture that the negative impact of private equity dominates and propose a particular theory of breach of implicit contracts, they do not conduct systematic empirical studies to measure the actual impact. Likewise, although Jensen (1986, 1989) conjectures that the positive impact of private equity dominates and posits that private equity solves a particular agency problem due to free cash flows, he does not conduct systematic empirical analysis to isolate this channel relative to other mechanisms.

In practice, both the positive and negative impacts of private equity can be present in a given transaction. Magnitudes of either effect can vary; as a result, the observed shareholder gains can be consistent with (i) efficiency gains alone, (ii) wealth transfers alone, or (iii) both efficiency gains and wealth transfers. Furthermore, each of these cases may or may not be accompanied by negative externality effects on the broader economy. Empirical findings of efficiency gains – often perceived to be in support of Jensen's arguments – do not necessarily reject Shleifer and Summers' concerns about shareholder gains due to transfers, and likewise, findings of employment loss – often perceived to be in support of Shleifer and Summers' arguments – do not necessarily reject Jensen's theory about efficiency gains. Therefore, instead of testing these hypotheses against each other in a horserace as if they are mutually exclusive, it is more informative to investigate the magnitudes and prevalence of the two effects, while allowing them to be simultaneously present.

One insight that has collectively emerged in the literature in the last three decades is that the pre-deal ownership type matters. Formerly-publicly-traded companies that get acquired by private equity firms (public-to-private deals) undergo systematically different transformations than formerly-privately-owned companies (private-to-private deals), with differential impact on the stakeholders and the broader society. Wealth transfers appear more predominant in public-to-private deals, whereas productivity gains and growth are associated more with private-to-private deals. Interestingly, private-to-private deals are typically less leveraged and more growth-oriented, and thus neither fit the inefficiently-run, cash-cow target architype that Jensen identified, nor are explained by the breach of implicit contracts à la Shleifer and Summers. This heterogeneity in value-generating mechanisms and outcomes underscores the limitation of any sweeping generalization about the social impact of private equity ownership/governance.

2. INCENTIVES IN PRIVATE EQUITY

There are several excellent descriptions of the organization and structure of private equity firms, including in the other chapters of this volume. To establish the terminology and economic relations for our discussion, we review a few central features below, but this is not meant as a comprehensive overview.

A private equity firm manages one or more funds. Each fund raises capital from a group of limited partners (LPs), typically pension funds and other institutional investors. The firm serves as the general partner (GP) of the fund, which is typically organized as a ten-year limited partnership. The fund uses the capital to acquire private assets, and the fund's type is determined by the nature of these assets. Two common types of private equity funds are venture capital funds that acquire equity stakes in young start-ups and buyout funds that acquire equity stakes in more mature companies.

Note that the term "private equity" is sometimes used to refer to either (i) the entire private equity asset class encompassing venture capital, growth, buyout, and distress investing, and (ii) buyout investing alone, which is the largest category of private equity in terms of assets under management. In this chapter we focus mostly on the literature that studies buyout funds and their portfolio companies, except when we discuss impact funds, where most funds raised early in this category (and thus studied in the academic literature) were venture capital funds and buyout impact funds came into existence only in the last few years as of the writing of the chapter.

The day-to-day operations of a fund are managed by its general partner. The general partner decides which companies to acquire, how to manage these acquired portfolio companies, and when and how to sell them again. The general partner also decides when to raise follow-on funds and manages this fundraising process. In return, the general partner receives management fees and carried interest. Management fees are charged by the general partner for managing the limited partners' capital and they are around 2% annually of the fund's total committed capital. Carried interest is the general partner's profit share, which is typically 20% of the fund's overall profits.

2.1 Incentives for maximizing deal profits

Consider, in isolation, a single deal where a buyout fund acquires a portfolio company. The private equity structure provides the general partner with strong incentives for maximizing the profits from this deal. The general partner typically receives carried interest of 20% of the fund's profits, although this does not necessarily mean that the general partner's marginal benefit of increasing the profit from the deal by one dollar is exactly twenty cents, due to complications such as the fund being underwater due to losses in other deals, hurdle rates, and catch-up provisions. These complications are second-order, though, and in most situations the general partner has strong incentives to maximize the profits generated in each deal.

An analytical framework, known as the buyout model, is useful for illustrating the implications of these incentives. In this framework, a buyout fund purchases a portfolio company

with enterprise value EV_0 and excess cash EC_0. The enterprise value, EV_0, is the economic value of the company's ongoing business, which is typically the present value of its future free cash flows. Excess cash, EC_0, is the cash available to the company that exceeds the cash needed to sustain its ongoing operations. The zero subscripts indicate that these amounts are dated at the time just after the private equity fund acquires the company. The sum of the enterprise value and the excess cash is the company's total value, and it typically equals the combined value of the company's equity and debt.

The buyout fund finances the acquisition with debt, D_0, and equity, E_0:

$$EV_0 + EC_0 = D_0 + E_0 \tag{1}$$

The debt, D_0, is the portfolio's company's debt just after the closing of the private equity transaction, which is often substantially larger than the company's debt before the transaction, meaning that the private equity fund pays some of the acquisition price by taking on new debt in the portfolio company. Economically, this is analogous to a home buyer financing part of the purchase of a house by taking out a mortgage. The equity, E_0, is the remaining part of the purchase price, which is paid by the private equity fund, and this is also the initial value of the private equity fund's equity stake in the portfolio company.

A similar relationship holds at the time of the exit when the portfolio company is sold by the private equity fund. Using subscript one to denote the values at the time of this sale, the price of the equity, E_1, reflects the company's updated enterprise value, excess cash, and debt at the time of the exit. It holds that:

$$EV_1 + EC_1 = D_1 + E_1 \tag{2}$$

The fund's profit (or loss) from the deal, denoted P, derives from the change in the value of the fund's equity. This change is:

$$P = E_1 - E_0 = (EV_1 - EV_0) + (EC_1 - EC_0) - (D_1 - D_0) \tag{3}$$

The right-hand side of this expression shows that the profit from a private equity deal consists of three components. It arises from an increase in the enterprise value of the portfolio company, $EV_1 - EV_0$, corresponding to an increase in its economic value. It arises from the excess cash accumulated by the portfolio company during the holding period, $EC_1 - EC_0$. And it arises from the reduction in the portfolio company's debt during the holding period, $D_1 - D_0$.

One implication of this framework is that an immediate way for a private equity fund to increase the profits from a deal is to increase the portfolio company's enterprise value, i.e., increase its economic value as perceived by the acquirer in the future exit. Indeed, a survey of private equity firms by Gompers, Kaplan, and Mukharlyamov (2016) reports that these firms place a heavy emphasis on adding value to their portfolio companies. The sources of added value, in order of importance, are increasing revenue, improving incentives and governance, facilitating a high-value exit or sale, making additional acquisitions, replacing management, and reducing costs. All of these sources directly serve to increase the portfolio company's enterprise value at the time of the exit.

The framework also shows the limits to how private equity funds can create profits in a deal. For example, an increase in the amount of leverage used to finance a deal does not mechanically increase the dollar amount of profits from the deal. An increase in leverage typically increases both D_0 and D_1, by largely similar amounts, so the effect mostly cancels out in the last term of the framework (the increases in the amounts may not exactly offset since, for example, an increase in initial leverage increases interest payments during the holding period, which may reduce either excess cash or limit the portfolio company's debt repayment over the holding period; these effects, however, are largely second order).

However, by making the cost basis of equity investment E_0 smaller for a given dollar amount of profit, an increase in the initial leverage does tend to increase the *return* on the equity investment, while also making it riskier.[1] Interest tax shields generated from higher leverage also contribute to increasing the enterprise value of the firm, ceteris paribus. Gompers, Kaplan, and Mukharlyamov (2016) find that two-thirds of private equity investor survey respondents say they raise as much debt as the market will bear and that limited partners in private equity funds focus more on absolute performance as opposed to risk-adjusted returns.

As another example of the economics of a deal, a portfolio company can pay a special dividend to the fund. Such a dividend would be paid out of the portfolio company's excess cash, and it would thus reduce the price at which the portfolio company is sold in a future exit. In the

[1] Suppose $EV_0 = 100$, $EC_0 = 10$, $EV_1 = 120$, and $EC_1 = 10$. In a low-leverage case, suppose $D_{L0} = D_{L1} = 50$; it implies that $E_{L0} = 100 + 10 - 50 = 60$ and $E_{L1} = 120 + 10 - 50 = 80$, so the return on equity investment is $(E_{L1} / E_{L0}) - 1 =$ is $(80 / 60) - 1 = 33\%$. In contrast, in a high leverage case, suppose $D_{H0} = D_{H1} = 100$; it implies that $E_{H0} = 100 + 10 - 5 = 100 = 10$ and $E_{H1} = 120 + 10 - 100 = 30$, so the return on equity investment is $(E_{H1} / E_{H0}) - 1 =$ is $(30 / 10) - 1 = 200\%$.

framework, a special divided is offset by a corresponding reduction in the excess cash, leaving the fund's profits from the deal largely unchanged (again, the changes may not be exactly similar, and the profits may not be exactly unchanged, due to changes in interest payments and other similar effects).

Other considerations may influence the fund's decision to pay a dividend. For example, a special divided could allow the private equity fund to return capital to its limited partners earlier, which can increase the fund's internal rate of return (IRR) even if the dollar amount of profits remains largely unchanged. This effect would encourage funds to pay special dividends.

Alternatively, the reduction in the portfolio company's excess cash following a special dividend may affect the portfolio company's operations. If the portfolio company suffers from being liquidity constrained it would reduce its enterprise value. The fund only benefits from the dividend, but it suffers both from the reduction in excess cash, which is as large as the dividend, plus the loss arising from the decrease in the portfolio company's enterprise value, and a dividend that negatively affects the portfolio company's operations would therefore normally reduce the fund's overall profits from the deal. Hence, this effect would discourage funds from paying special dividends. In their investigation of corporate tax filings from leveraged buyouts during 2005 to 2009, Cohn, Mills, and Towery (2014) find that portfolio companies only make limited dividend payments, and they find no evidence that private equity firms "strip" value from otherwise healthy portfolio companies.

2.2 Incentives of a fund, deal flow

The above discussion considered the profits from a single deal. However, a fund typically invests in several deals, with 10-15 deals being normal numbers. The general partner's carried interest typically depends on the fund's aggregate profit across all deals net of management fees, where management fees are the total amount paid to the general partner during the lifetime of the fund:

$$\textit{Fund Profits} = \Sigma_d P_d \text{ - } \textit{Management Fees}. \quad (5)$$

Deals are indexed by d, and P_d is the profits from deal d, as defined above. Management fees are specified contractually at the inception of a fund and are independent of the fund's performance. However, since management fees come out of the fund's committed capital, larger management fees will reduce the amount of capital available for investing in portfolio companies. Existing research finds that as much as two thirds of the present value of the general partner's compensation comes from management fees (Metrick and Yasuda 2010).

To maximize carried interest, holding the set of deals constant, the general partner would simply maximize the profits in each deal. However, holding the set of deals constant is not a trivial assumption. Private equity firms are concerned about their ability to identify and invest in attractive deals, but that may require maintaining a certain reputation to be able to access this "deal flow." Hence, even pure profit maximization at the fund level means that the private equity firm must be mindful about how its actions in a given deal affect the firm's reputation and ability to invest in other attractive deals. It is not inconsistent with profit maximization that a private equity firm leaves money on the table in a specific deal to maintain a reputation as a reasonable investor that other potential portfolio companies would want to work with in future deals. This reputational concern, for example, may induce a private equity owner to incur costs to maintain more stringent regulatory compliance standards in sectors where the framework is transparent and well-enforced than non-private equity owners. On the other hand, the same private equity owner may reverse its policy once the regulation is rolled back.

2.3 Incentives of private equity firms, raising future funds

Private equity firms typically manage a series of funds. The general partner's ultimate incentives arise from its concern about its aggregate management fees and carried interest across all its funds. Holding the set of funds constant, a general partner maximizes the aggregate compensation it receives by maximizing the combination of management fees and fund profits, as defined in Equation (5) for each fund. However, holding the set of funds constant is again not a trivial assumption. The aggregate compensation is:

$$\textit{Firm Revenue} = \Sigma_f \left[h_f(\textit{Fund Profits}_f) + \text{Management fees}_f \right] \qquad (6)$$

The function h_f captures the private equity firm's profit share for each fund, indexed by f. This function is weakly increasing in the fund profits, and a simple version of this function is $h_f(p) = \max(0, 20\% \; p)$, although this function can be more complex due to hurdle rates, catch-up provisions, and other features of the fund's "waterfall." Importantly, however, the sum is over both current and future funds, and Chung, Sensoy, Stern, and Weisbach (2012) argue that about half of a private equity firms' incentive pay arises from the effects of its current fund performance on its ability to raise future, typically larger, funds, which generate management fees and carried interest. The other half comes from carried interest earned on the current fund. Metrick and Yasuda (2010) and Chung, Sensoy, Stern, and Weisbach (2012) find that successful buyout fund managers increase their pay by scaling up the size of future funds rather than by increasing their compensation per dollar managed. The capital for these funds is provided by limited partners, which are typically institutional investors, such as pension funds, university endowments, and

sovereign wealth funds. Hence, another concern for a private equity firm is whether the firm's investments and actions are consistent with the preferences and priorities of the limited partners and whether these investments and actions promote the private equity firm's future fundraising.

To summarize the above discussion, the incentives facing private equity firms and the general partners managing their funds are more complex than they may immediately appear. It may be natural to think that the large profit share—in the form of a 20% carried interest—simply means that private equity investors only maximize short-term profits in each deal above all else. However, the private equity organizational form creates more complex incentives for the private equity investors. First, a private equity fund only profits when a portfolio company is sold in an exit, and the sales price depends on the portfolio company's future economic viability. Hence, it is unlikely that extracting short-term profits at the expense of long-term viability is a significant source of profits for private equity funds. Second, a substantial part of a private equity firm's business hinges on its ability to make future investments in companies. Therefore, private equity firms may be concerned about maintaining a reputation for being reasonable investors even if this means leaving money on the table in a specific deal. Third, the private equity firm's ability to raise future and larger funds and earn both management fees and carried interest from those funds, hinges on the firm acting according to the preferences and priorities of limited partners. How these various trade-offs are resolved in practice is an empirical question.

2.4 Growth of Environmental, Social, and Governance (ESG) Investing and Private Equity

In recent years, an increasing share of limited partners who invest in private equity funds have integrated environmental, social, and governance (ESG) issues into their investment analysis. As a sign of increasing demand for ESG-conscious investment practices, as of March 2020, 3,038 organizations representing $103.4 trillion in asset under management have become signatories to the United Nations Principles of Responsible Investment (UNPRI). These principles state that the signatories commit to "[i]ncorporate ESG issues into investment analysis," "[i]ncorporate ESG issues into our ownership policies and practices," and "[s]eek appropriate disclosure on ESG issues by the entities in which we invest."

There are two distinct motivations behind investors' push to incorporate ESG issues into their investment analysis. One is based on materiality, or the idea that the ESG practice of business has material impacts on their future financial performance, and therefore it is within the realm of asset managers' fiduciary duty to the beneficiary of its financial assets to incorporate ESG-related information into its investment analysis. Under this branch of responsible investment, the goal is still wealth maximization, but ESG issues are value-relevant and therefore should be disclosed and

monitored just like other material information. This materiality-based motivation is consistent with U.S. fiduciary investors' recent push to incorporate ESG into their investment practice.

The other motivation is based on impact investing, or the idea that limited partners enjoy non-pecuniary benefits from generating positive externalities via their investments. This type of responsible investment is called dual-objective or double-bottom line because investors explicitly seek to generate both financial returns and positive societal impact at the same time. Indeed, 88% of UNPRI signatories are either investment managers or asset owners, and the UNPRI states that incorporation of ESG issues both fulfills the signatory's fiduciary duty as institutional investors, and also "may align investors with broader objectives of society." Notably, European fiduciary investors are permitted (and even required) to take into consideration broader social and public interests when screening their investments, while U.S. fiduciary investors are under strict regulatory guidelines to consider only the financial impact of investment decisions (or consider non-financial impact only when it is guaranteed not to lower financial return). This explains to a large degree the difference in expressed preferences and priorities of American. vs. European investors towards ESG and sustainable investments.

While European institutional investors have been quicker to integrate ESG issues into their investment process, in the last several years the pace of adoption has accelerated even among the American institutional investors. For example, in his influential annual letter to the CEOs of the largest corporations, the BlackRock CEO Larry Fink wrote: "[A] company cannot achieve long-term profits without embracing purpose and considering the needs of a broad range of stakeholders" (Fink 2020) and "with the world undergoing the largest transfer of wealth in history: $24 trillion from baby boomers to millennials. As wealth shifts and investing preferences change, environmental, social, and governance issues will be increasingly material to corporate valuations." (Fink 2019) Also, climate risk is increasingly seen as an investment risk, endangering sustainability of investment portfolios. In response, major asset managers and asset owners around the world have started planning for transitions to net-zero portfolios, or portfolios with net-zero greenhouse gas emissions.

Pedersen, Fitzgibbons, and Pomorski (2021) build an ESG-adjusted capital asset pricing model in which three types of investors differ in their preferences and information sets with respect to assets' ESG characteristics, and their respective portfolio decisions affect equilibrium asset prices and returns. Pro-ESG investors derive utility from holding high-ESG score assets. ESG-aware investors use firms' ESG scores to update their views on risk and expected returns. ESG-unaware investors are unaware and therefore ignore ESG scores. In the model, ESG-aware investors may get superior outcomes when ESG scores convey value-relevant information. This

insight suggests that materiality of ESG information may drive limited partner demand for ESG disclosures by private equity funds, even if they are not pro-ESG investors.

Given this accelerating shift in investors' preferences for ESG incorporation, general partners in private equity firms are responding by committing to greater ESG disclosures and in some cases raising impact funds that explicitly target generation of positive externalities. Indeed, Preqin reports that "[s]ince 2011, more than 4,400 ESG-committed private capital funds have closed, totaling $3.06tn in combined assets" (Preqin 2020). This is consistent with the incentives of general partners to maximize their ability to raise future funds by catering to the preferences of their limited partners.

3. PRIVATE EQUITY MANAGEMENT MODEL

As described in Section 2, private equity funds use a contractual format of finite-life, closed-end limited partnerships that provide general partners with a particular set of incentives and access to their portfolio companies. What does this imply about the effect of private equity ownership on the shareholders, other companies' stakeholders, and society at large? The literature has mostly focused on analyzing the principal-agent problem between portfolio company shareholders and their managers, and how the private equity model purports to solve this agency problem better than public companies can solve it (Jensen 1989). In contrast, the effect of the private equity management model on portfolio company stakeholders and other parties in the society at large is underexplored. In this section we (i) review the proposed private equity solution to the principal-agent problem, (ii) posit the private equity model as a potential solution to the ESG monitoring problem of responsible investors, and (iii) point out the ambiguity of the relationship between private equity ownership and the impact on stakeholders.

3.1 Principal-Agent Problems in Public Corporations

The literature has identified three ways in which managers of public corporations may behave in value-destroying manners from shareholders' point of view. First, Jensen (1986, 1989) argues that when managers have too much discretion over how to spend any free cash flows of the company they manage, they tend to engage in empire building by investing in projects with negative net present value (NPV) to increase the managers' spheres of influence, and/or to consume private benefits. Second, Bertrand and Mullainathan (2003) find that, when public company managers are shielded from takeovers by state laws, they invest less, pay higher wages to themselves and their peers, and enjoy the quiet life, to the detriment of profitability and productivity of the companies. Finally, while public investors scrutinize hard information such as quarterly earnings and stock

returns, they face free-riding problems and lack incentives to produce soft information about a manager's efforts, ability and quality. Subsequently, this narrow focus on quarterly earnings and share prices by public market investors leads to short-termism by company managers, which in turn destroys shareholder values by distracting the managers from long-term strategic thinking and investment horizons (see Bebchuk and Weisbach 2010 for a survey). Moreover, all three of the agency problems are exacerbated when boards are captured by insiders. Both regulators and firms themselves propose using independent directors as a remedy for these agency problems, and the literature documents some benefits associated with independence of directors (e.g., Chhaochharia and Grinstein, 2006). However, independent directors' effectiveness as monitors depends on their own incentives, and thus can be diminished due to lack of information access, busy-ness, or both (see Bebchuk, Grinstein, and Peyer 2010 and Fich and Shivdasani 2012).

3.2 Posited Solutions with Private Equity Management Model 1.0: Traditional View

Kaplan and Stromberg (2009) identify three value-creating activities associated with the private equity management model: Financial engineering, governance engineering, and operational engineering. Each of these activities is characterized as a potential solution to the aforementioned principal-agent problem in public corporations. First, under financial engineering, buyouts typically result in the elevated leverage ratios for portfolio companies for a sustained period of time. Jensen (1986, 1989) argues that pledging future cash flows to pay down the debt reduces the agency problems and leads managers to engage in less empire building. Moreover, the higher interest payments increase the present value of the interest tax shields, thus enhancing the value of the company, as long as the cost of financial distress does not increase significantly as a result. These are described as the financial engineering aspect of the private equity management model.

Second, under governance engineering, the private equity sponsors typically obtain majority ownership of the portfolio companies they acquire. This eliminates the free-rider problem of monitoring public companies and enables private equity sponsors to closely monitor their portfolio company executives using both hard and soft information about the executives' managerial performance. Combined with the lack of public stock prices and extended holding period, the monitoring focus shifts from short-term earnings management to long-term value realization. Additionally, portfolio company executives are typically given greater equity incentives under private equity ownership than when they are publicly held. Meanwhile, they receive less generous perks (e.g., less frequent use of private jets) than their public company counterparts. Since the equity shares are privately held and lack liquidity, executives are incentivized to realize long-term capital gains for the company, which they can unlock only after the investments are exited, rather than short-term upswings in the stock prices.

Finally, private equity firms accumulate in-house industry and operating expertise by hiring operating partners with operating backgrounds and an industry focus, and then apply this expertise to add value to their portfolio companies. Kaplan and Stromberg (2009) note that, while financial and governance engineering were common by the late 1980s, operating engineering was added in more recent years to private equity firms' repertoire of value-creating tools. Private equity firms use this industry-specific know-how to identify attractive targets, to develop value creation plans, and to implement the plans.

3.3 Posited Solutions with Private Equity Management Model 2.0: Growth, Competency, and non-Financial Outcomes

While the early literature analyzed the private equity management model through the lens of fixing what is broken in the public company governance mechanism, more recently the literature has expanded its focus to the rationale behind private-to-private private equity transactions. After all, most private equity transactions are private-to-private deals, so how can the private equity funds unlock value for companies that already have concentrated ownership and illiquidity? The literature offers two possibilities. First, many private companies face limited access to capital, and this constrains their ability to grow. By unlocking the access to both debt and equity capital, the private equity sponsors may spur growth in the private companies they acquire. Second, private companies may lack the scale, name recognition, and prestige of public companies to attract professionalized managers (e.g., MBA graduates), and thus be less effective at executing their corporate strategy. By becoming associated with prominent private equity firms, which routinely recruit professional executives for their portfolio companies and often have in-house operating partners with senior industry expertise, private companies may improve the managerial skill level of their personnel.

The recent push by institutional investors to incorporate ESG factors into their investment process creates yet another possibility that the private equity management model offers a differential outcome for investors with pro-social or pro-environmental preferences. For example, Hart and Zingales (2017) argue that when (i) shareholders have pro-social preferences, (ii) profit-making and damage-generating activities of companies are non-separable, and (iii) government cannot perfectly internalize negative externalities through laws and regulations, then companies should seek to maximize shareholder welfare rather than market value. Prosocial investors in the model are ethical in plausibly limited ways: If put to a vote, they put positive weight on externalities generated by their decisions, but without such votes they are "willing to hold shares in tobacco or gun or oil companies, and indeed will pay full price for these shares" (p. 267). This and the atomistic weight of individual shareholders creates an "amoral drift," a tendency for public

companies to "underweight social surplus much more than privately held companies" (p. 258). Focusing on public companies, the authors propose allowing broader shareholder proxy voting on corporate policy as a mechanism to implement the welfare-maximizing objective that aggregates the preferences of the shareholders. This proposal is interesting because with the private equity fund structure, limited partners and general partners can write a contract that aggregates the preferences of the limited partners at the inception of the fund. Moreover, the weight of individual investors in the partnership and in the ownership structure of each portfolio company is substantial, and thus prevents the amoral drift. Though Fama (2021) raises concerns about the difficulty of coordinating heterogenous and multi-dimensional ESG preferences of investors in public companies, the 10-year commitment requirement of private equity funds prevents limited partner turnover and temporal shift in aggregate investor preferences. Thus, the debate on ESG incorporation raises a possibility that private equity-owned companies make choices that are more consistent with social preferences of their ultimate owners than either public companies or non-private equity private companies.[2]

3.4 Ambiguity of the Effect of Private Equity Management Model on Stakeholders and Society at Large

While research suggests that the private equity management model has a positive effect on shareholder value, its effect on the company's other stakeholders or society at large is far from clear. For example, suppose a private equity fund takes over a portfolio company and provides its executives with high-powered incentives of significant long-term equity stakes. The company executives, in response, implement heavy investments in high-growth divisions, upgrade the company's IT system, and upskill its high-skilled labor through IT training, while automating and offshoring low-skill jobs and divesting non-performing divisions. On the one hand, this might result in rent extraction from the company's manufacturing plant employees, who face either job losses, or wage and benefit reductions, as posited in Shleifer and Summers (1988). On the other hand, the same strategy implementation may also result in enhanced rent sharing with the so-called knowledge workers of the company, who experience upskilling, productivity increases, and thus wage and benefit increases.

[2] Also see Broccardo, Hart and Zingales (2021), Pedersen, Fitzgibbons, and Pomorski (2021), and Pastor, Stambaugh, and Taylor (2021).

Likewise, the effects of high-powered incentives of private equity portfolio company executives on the company's customers, the government/taxpayers (via tax revenues or subsidy payments), industry-level innovation, or the environment are a priori ambiguous and remain open empirical questions. This is natural in the sense that the logic/rationale of the private equity management model is traditionally shareholder-centric, and the effects on other related parties are side products of the pursuit of a single-objective maximization. One exception is the newly emerging realm of impact investing, which is an explicitly dual-objective investment model where investors intentionally wish to pursue financial as well as non-financial goals simultaneously.

4. IMPACT ON MANAGEMENT, PRODUCTIVITY, AND EFFICIENCY

4.1 Executives, Management, Board of Directors

It is common for private equity investors to replace a portfolio company's management in connection with a private equity transaction. In a survey of (mostly U.S.-based) private equity investors, Gompers, Kaplan, and Mukharlyamov (2016) find that CEOs and CFOs are replaced in 30.6%–42.9% of private equity transactions and that 57.8% of the surveyed private equity firms routinely recruit their own senior management teams. Biesinger, Bircan, and Ljungqvist (2020) study value creation plans for the individual deals of (mostly emerging market) private equity funds; 20% of the plans explicitly mention replacing the CEO, 20% mention replacing the CFO, and 26% of the plans mention replacing other managers.

Replacing managers matters. Looking at buyout transactions, Guo, Hotchkiss, and Song (2011) find that operating cash flows improve more when the private equity firm replaces the CEO at or soon after the buyout. Focusing on the individual CEOs of the portfolio companies and their personalities and traits, Kaplan, Klebanov, and Sorensen (2012) report that a CEO's execution ability is particularly important for a successful outcome. Interestingly, Biesinger, Bircan, and Ljungqvist (2020) reach a similar conclusion from the value creation plans. They find that the action items that are actually implemented are more important than the specific choice of strategy, with diminishing returns to making plans ever more detailed, leading them to conclude that "execution is the key."

Given this focus on the management it is perhaps unsurprising that portfolio companies tend to have better management practices. In an extensive study of management practices, Bloom, Sadun, and Van Reenen (2015) conduct about 15,000 interviews with managers in about 10,000 manufacturing plants in 34 countries and score their management practices. They find that private equity-owned companies have better management practices than most other company types such

as family-run, founder owned, or government owned firms. The only exception is dispersed shareholder firms (e.g., publicly-listed firms), which have similar levels of the management score as private equity-owned firms. They find a significant gap in the quality of management practices between private equity-owned portfolio companies and the practices of family-owned and other private companies. This gap is robust not only in developed countries but also in developing countries – where private ownership of companies is more prevalent and where capital markets are less developed than in developed economies – suggesting that in those countries private equity acquisition is an important mechanism through which better management practices are introduced to private companies. This finding is noteworthy because most companies are private companies, especially in developing countries with less developed public markets, and these private companies are the typical targets of private equity acquisitions.

An important aspect of management practices is management compensation. Gompers, Kaplan, and Mukharlyamov (2016) report that private equity investors use more aggressive compensation packages to incentivize the senior management of portfolio companies. In 61.1%–65.1% of the transactions in their study, the private equity investors specifically mention improved managerial incentives as a source of increased value. Cronqvist and Fahlenbrach (2013) find that private equity investors increase the CEO's base salary and bonus by 25%, with the salary increases concentrated among newly appointed CEOs. About half of the equity grants to portfolio company CEOs only vest at the time of an exit event, such as a sale or IPO of the portfolio company, which aligns the CEOs' incentives with the investors' need for a timely exit. Severance contracts for portfolio company CEOs are also stricter with respect to unvested equity, which is often forfeited. Cronqvist and Fahlenbrach (2013) further report that private equity investors are less likely to tie CEO compensation and bonuses to qualitative, nonfinancial, and earnings-based performance measures. Instead, CEO compensation and bonuses depend on more quantitative targets, such as cash flow-based measures (e.g., EBITDA) with less accounting discretion. Bloom, Sadun, and Van Reenen (2015) find a slightly weaker relation in their global survey of management practices. They report that private equity-owned companies provide significantly stronger managerial incentives with more direct links to the managers' effort and ability. However, only the unconditional difference is statistically significant, and it becomes insignificant when they include country, company, industry, and other controls. Nevertheless, the broader evidence suggests that private equity ownership provides CEOs with steep financial incentives to align the managers' interests with the interests of the private equity investor.

Another important aspect of management practice is managerial turnover. Interestingly, while private equity investors are initially more likely to replace management at the time of the buyout

transaction, Cornelli and Karakaş (2015) document that once the deal is completed, private equity investors are less likely to replace management going forward. They attribute this lower turnover rate to private equity investors having more inside information and being more effective at monitoring the managers, which in turn allows the private equity investors to evaluate the managers' performance over a longer time horizon relative to their publicly-traded counterparts. Consistent with this interpretation, Bloom, Sadun, and Van Reenen (2015) find that portfolio companies have particularly strong monitoring practices (described as "practices around continuous performance measurement, improvement, and feedback") as compared to other types of companies.

Cornelli, Kominek, and Ljungqvist (2013) use monitoring reports from the European Bank for Reconstruction and Development (ERBD) to study determinants of individual turnover decisions. The reports allow them to distinguish between the role of "soft" and "hard" information when the boards of portfolio companies replace the CEO. Interestingly, soft information seems to play a larger role than hard information. One interpretation is that private equity investors have closer and more direct connections to the managers and boards of their portfolio companies, and that these investors are therefore less likely to replace a CEO for bad performance that is due to external factors, e.g., bad luck. Instead, they argue, private equity investors are more deliberate when replacing CEOs, which is in contrast to the boards of publicly-traded companies, which appear to place less weight on the reasons for poor performance. To identify the causal effect, the authors exploit staggered governance reforms that increase boards' personnel authority to dismiss CEOs and find that forced CEO turnover causes improved firm performance.

Acharya, Kehoe, and Reyner (2009) explore the board dynamics of UK-based portfolio companies in more detail. They confirm that CEOs are replaced in connection with private equity transactions, as mentioned above, and they also interview twenty chairpersons and CEOs that have experience with both public and private boards. Contrasting the board structures of private equity-owned portfolio companies with the structures in UK public companies (PLCs) in the FTSE 100, they report that private equity boards are more efficient and have a more concentrated focus on value creation. PLC boards suffer from being larger, being more focused on quarterly profits and market expectations, and having greater concerns about the reactions of external stakeholders than the impact of their decisions on business performance. They identify three main characteristics of private equity boards: they are well-aligned with a focus on value creation, clearly articulate and insist on strategic and performance priorities, and have a greater engagement by the board members. In their sample, three-quarter of the interviewees report that private equity boards add more value than PLC boards. None reported that the public counterparts were better.

4.2 Productivity and Efficiency

In his classical analysis of the agency problems facing different types of companies, especially conglomerates, and the value of the private equity governance model, Jensen (1989) focuses on the "free cash flow problem" in poorly governed companies. This agency problem arises when mature companies with stable cash flows and few profitable investment opportunities make wasteful investments to promote managerial empire building (see also Jensen 1986). According to Jensen (1989), the ability to mitigate this problem is a main benefit of private equity ownership. In support of the free cash flow hypothesis, Opler and Titman (1993) find that high cash flow firms with low Tobin's *q* are more likely to be acquired in a buyout. A more direct test of this hypothesis is the study by Edgerton (2012) of corporate fleets of private jets. Larger jet fleets typically reflect wasteful managerial perks, and consistent with Jensen's hypothesis Edgerton (2012) finds that private equity-owned portfolio companies have significantly smaller fleets than other publicly-traded or private companies, on average, and he finds clear reductions in fleet size after public companies are taken private in buyouts. Many public companies have fleets that appear large by the standards of private equity-owned companies, which he argues is consistent with agency problems in these public companies, although with the caveat that excessive fleets are far from ubiquitous in public companies.

Early Studies Although Jensen (1989) originally focused specifically on the "free cash flow problem" a number of other studies published around the same time found more general improvements in the productivity of portfolio companies, and they generally reported that the increases in productivity were not primarily due to reductions in employment.

In their study of a single prominent transaction, Baker and Wruck (1989) explore the highly levered 1986 acquisition of the O.M. Scott & Sons Company as a divisional buyout from the ITT conglomerate. Baker and Wruck report that "operating performance improved dramatically following the buyout" with sales increasing by 25% and earnings before interest and taxes increasing by 56%. The improvements are attributed to three factors: the constraints imposed by high leverage, changes in managerial compensation, and improvements in monitoring and advising of Scott's top management. Moreover, they find no evidence that the improvements came at the expense of employees, although annual employment declined by 9% through natural attrition over the first two years following the buyout. Additionally, the improvements were not caused by a reduction in spending on R&D, marketing, or capital expenditures, which actually increased by 23% after this buyout.

Other early studies, using larger samples of private equity transactions, largely confirm these findings. Kaplan (1989a) find substantial increases in operating income, net cash flows, and market value in a sample of 76 portfolio companies, and he argues that these increases are due to improved managerial incentives and reduced agency problems, not transfers from employees. Smith (1990) reports similar increases in operating cash flows in a sample of 58 portfolio companies, and she highlights the role of improved working capital management and managerial incentives. She also confirms that the increase in operating returns is not due to layoffs or reductions in advertising, maintenance, or research and development, although she does find a decline in capital expenditures.

Muscarella and Vetsuypens (1990) consider 72 reverse LBOs, i.e., transactions where a portfolio company that was previously public and is then taken private in a private equity transaction is later listed again in an IPO. As part of the IPO process, the company must disclose several years of financial statements, and these reverse LBOs therefore offer a view into the financials and operations of portfolio companies before and after the private equity transaction, although the subsample of portfolio companies that undergo this specific sequence of transactions may not be representative. Consistent with the previous studies, Muscarella and Vetsuypens (1990) find that portfolio companies in their sample show significant improvements in profitability, mainly due to cost reductions. They do not find evidence of reductions in employment.

Finally, Lichtenberg and Siegel (1990) use Census data for 72 portfolio companies to study changes in total factor productivity (TFP). In the three years following a buyout the average TFP increases by 8.3% relative to the industry average. Consistent with the above studies, they do not find any change in the employment of production workers ("blue-collar"), and they report cumulative increases in the compensation of these production workers of 2.3% to 3.6%. The employment of non-production workers ("white-collar"), however, declines by 8.5%.

Later Studies Later studies of private equity transactions that took place during the 1990s and 2000s find more mixed results and generally do not find statistically significant effects on productivity and profitability. In an unpublished working paper, Leslie and Oyer (2008) contrast portfolio companies with publicly-traded comparable companies. Like Muscarella and Vetsuypens (1990), Leslie and Oyer study reverse LBOs—i.e., portfolio companies that were publicly traded before being acquired by a private equity fund and which became publicly traded again when the fund listed them in an IPO. Their sample contains 144 such transactions taking place from 1996 to 2005. They report that top managers of portfolio companies have substantially higher-powered compensation contracts. These managers own more equity, have lower base salaries, and have a larger fraction of variable compensation. In terms of operational improvements, they find that

portfolio companies improve the measure Sales per Employee. However, they find no statistically significant changes in their other performance measures: Return on Assets (ROA), EBITDA / Total Assets, and Employees / Total Assets.

When interpreting accounting measures of operating performance, a concern is whether portfolio companies engage in more aggressive earnings management. Katz (2009) finds that private equity-backed companies have higher earnings quality than those that do not have private equity sponsorship, engage less in earnings management, and report more conservatively both before and after the IPO. These findings are consistent with tighter monitoring and reputational considerations exhibited by buyout investors.

Following Leslie and Oyer (2008), Guo, Hotchkiss, and Song (2011) also revisit the productivity impact of private equity in a sample of 194 public-to-private U.S. buyout transactions from 1990 to 2006, with deal values exceeding $100 million. They find an 11% increase in EBITDA and net cash flows. They also find some improvements in operating performance, although these are not statistically significant.

Cohn, Mills, and Towery (2014) study 317 buyouts of previously public companies, taking place from 1995 to 2007, using earnings and revenue information from corporate tax filings. Consistent with the above studies, in this sample of public-to-private transactions, they find no significant effects on Return on Sales, Return on Assets, and their measure of economic value added (EVA).

Following up on the previous study, however, Cohn, Hotchkiss, and Towery (2022) use the tax filings to study 288 buyouts, taking place from 1995 to 2009, but now focusing on buyouts of private companies. Interestingly, the findings are different in this sample of private-to-private transactions. They find a moderate but significant increase in profitability, both in absolute terms and relative to the industry. Moreover, they find large and rapid increases in revenue after a buyout, which they argue reflects both organic and acquisition-driven growth. Their interpretation is that for the private targets considered in this sample, unlike for public targets, a main source of value creation is unlocking growth opportunities by relaxing financing constraints. The relatively active market for acquisitions of U.S. private companies enables private equity acquirers to use portfolio companies as platforms for acquiring other small companies. While buyouts of public companies may alleviate overinvestment problems, their results suggest that buyouts of private companies solve underinvestment problems. They find no evidence that financial engineering is a significant source of value creation for buyouts of private companies, since these companies already have

relatively high levels of leverage and therefore only see a smaller increase in leverage following the transaction.

Davis, Haltiwanger, Handley, Jarmin, Lerner, and Miranda (2014) use Census data to investigate, among others, the effects of U.S. buyouts on the total factor productivity (TFP) of manufacturing plants. For this analysis, they have TFP data for 286 multi-unit manufacturing companies acquired by private equity firms between 1980 and 2003. The risk of a plant exiting in the two years after the acquisition depends critically on the productivity of the plant. In the bottom tercile of the TFP distribution, the exit probability is significantly higher for plants operated by private equity-owned portfolio companies than for plants operated by companies in the control group. In contrast, there are no significant differences in the exit probabilities in the middle and top terciles. They report a similar pattern for the opening of new plants. Plants opened by private equity-owned companies are substantially more likely to be in the top tercile and significantly less likely to be in the bottom tercile than plants opened by companies in the control group. For continuing plants, they do not find evidence of changes in productivity relative to plants in the control group. Overall, they summarize their findings as evidence that private equity firms reallocate activity to raise TFP, that the large TFP advantage of portfolio companies reflects a concentration of new plants in the upper part of the TFP distribution and exits of plants in the lower end of the TFP distribution, and that their results refute the view that the returns to private equity rest entirely on private gains to financial engineering and wealth transfers from other stakeholders.

In a recent working paper, Davis, Haltiwanger, Handley, Lipsius, Lerner, and Miranda (2021) revisit the analysis using an extended sample and refined empirical methods. Their results, while preliminary, show even larger average productivity gains when both manufacturing and non-manufacturing industries are included the sample (in contrast, the previous study only included manufacturing industries). Moreover, this analysis also suggests that there are significant differences between public-to-private and private-to-private buyouts, consistent with the different findings in Cohn, Mills, and Towery (2014), focusing on public-to-private transactions, and Cohn, Hotchkiss, and Towery (2022), focusing on private-to-private transactions.

Studies of Non-U.S. Transactions The studies mentioned above have primarily studied U.S.-based portfolio companies. It is interesting to supplement these studies with studies of buyouts in other countries. Other countries have different regulatory and governance systems. They typically have relatively more privately held companies, and they may have capital markets that make it more difficult for these companies to raise external capital, leaving them more capital constrained. Other countries also often have better data for portfolio companies than what is available for U.S.

private companies. Interestingly, studies of non-U.S. buyouts find more consistent evidence of operational improvements than the studies of U.S. buyouts during the same periods. This difference may be due to differences in the nature of buyouts in these counties or to the composition of buyouts in these other countries having relatively more private than public targets.

In an early study of non-U.S. buyouts, Bergstrom, Grubb, and Jonsson (2007) consider the entire universe of 73 Swedish buyout exits, exceeding $5m, during the period 1998 to 2006. They find substantial and significant improvements in EBITDA margins and ROIC, and smaller increases in revenue which are less statistically significant. Moreover, employment and wage levels in their portfolio companies have not declined relative to the levels in comparable Swedish companies. A natural concern about this study is selection bias, since the sample is a sample of realized exits. Portfolio companies that have gone bankrupt or are being held for extended time periods may be underrepresented in this sample, although Bergstrom, Grubb, and Jonsson argue that this is rare in Sweden and therefore unlikely to affect their results.

Boucly, Sraer, and Thesmar (2011) find even stronger effects in a sample of 839 French buyouts during 1994 to 2004. They use financial statements from tax filings to track the portfolio companies before and after the buyout, and they find large and statistically significant growth in profitability, employment, sales, and capital expenditures. From four years before to four years after the transaction, employment grows by 18%, assets grow by 12%, and sales grow by 12%, on average, relative to comparable companies. Interestingly, this growth is concentrated in portfolio companies that are privately-held before the buyout, i.e., private-to-private transactions, as opposed to divisional buyouts or buyouts of publicly-traded companies. In France, these privately-held companies are often owned by an individual or a family that is cashing out of its business. Moreover, the improvements are concentrated in portfolio companies in industries that rely more on external capital. Overall, the evidence suggests that private equity investments relax credit constraints for portfolio companies and that this benefit may be particularly important in France due to its relatively underdeveloped capital markets, at least during this sample period. The findings and interpretation appear consistent with Cohn, Hotchkiss, and Towery (2022) and Davis, Haltiwanger, Handley, Lipsius, Lerner, and Miranda (2021), discussed above, who study recent buyouts of U.S.-based private companies, and who also find strong growth and evidence of private equity investments alleviating capital constraints for these portfolio companies.

Focusing on innovation and patenting, Amess, Stiebale, and Wright (2016) study 407 U.K. buyouts between 1998 and 2005. They find a 6% increase in quality-adjusted patent stock three years after the buyout, but improvement in innovation is even stronger for private-to-private transactions and for portfolio companies in financially dependent industries, which is also

consistent with private equity investments relaxing financial constraints in portfolio companies and facilitating their investments in innovation activity. Nikoskelainen and Wright (2007) and Renneboog, Simons, and Wright (2007) provide further UK evidence, and Cumming, Siegel, and Wright (2007) provide a literature review of global evidence related to governance and financial and real returns to private equity.

Studying European acquisitions more generally, and excluding private equity transactions, Erel, Jang, and Weisbach (2015) find significant increases in investments, reduced cash holdings, and lower investment cash flow sensitivities after these acquisitions, which they interpret as evidence that financial constraints are reduced for target companies in general acquisitions, and not just for private equity-driven ones.

Role of General Partners Acharya, Gottschalg, Hahn, and Kehoe (2013) study the role of general partners in a sample of 395 buyout acquisitions of European portfolio companies from 1991 to 2007. They find that the EBITDA to Sales ratio increases by 0.4% annually, on average relative to the sector mean, and that the deal multiple, EBITDA to Enterprise Value, increases by around 1, on average relative to the sector mean. They interpret these improvements as the causal impact of private equity ownership, which creates economic value through operational improvements. Interestingly, they go further and explore the identities and backgrounds on the individual general partners that are responsible for managing the transactions for the private equity firms. Deal partners with an operational background, typically ex-consultants or ex-industry managers, are associated with greater outperformance in "organic" deals the portfolio company improves internally. In contrast, partners with a background in finance, i.e., ex-bankers or ex-accountants, are associated with "inorganic" and M&A driven strategies.

Echoing the heterogeneity of value-creation strategies employed by private equity firms, Davis, Haltiwanger, Handley, Lipsius, Lerner, and Miranda (2021) find that employment effects of private equity ownership is highly persistent over time at the general partner level. Spaenjers and Steiner (2021) study specialist vs. generalist private equity investors in the U.S. hotel industry and find that specialist private equity investors are associated with value creation through operational performance improvement. In contrast, generalist private equity investors are associated with value creation through cheaper source of debt financing. Biesinger, Bircan, and Ljungqvist (2020) find systematic differences across funds in their ability to achieve the objectives set out in the value creation plans for their deals. Funds with focused, homogeneous portfolios of predominantly minority positions are better at implementing these plans than other funds. Bernstein and Sheen (2016) find that private equity firms with specific experience in the restaurant industry do significantly better than firm with more general operational experience.

4.3 Financial Distress and Bankruptcy

A natural concern is whether the high leverage imposed on portfolio companies in connection with a private equity deal makes the company more financially fragile. In a study of U.S. private equity deals, Hotchkiss, Smith, and Strömberg (2021) find that private equity-backed companies have higher leverage and because of this leverage these companies default at higher rates than other companies borrowing in leveraged loan markets. However, conditional on being in default, private equity-backed companies restructure more quickly and more frequently out of court, and private equity owners are less likely to be wiped out in this process.

Wilson and Wright (2013) consider a sample of UK companies, focusing on UK private equity deals during 1995 to 2010. Controlling for size, age, sector and other conditions, they claim that portfolio companies in private equity-backed buyouts are no more likely to become insolvent than other similar companies.

Neither study considers the externality cost of distress and bankruptcy on other stakeholders or the broader economy. Thus, while private equity investors may avoid some of the negative impact of distress and bankruptcy despite the high leverage of portfolio companies, the impact on other stakeholders and the broader economy remains under-studied.

5. IMPACT ON EMPLOYEES AND OTHER STAKEHOLDERS

A common concern about private equity is that it does not actually create value but instead expropriates value from a broader set of the portfolio company's stakeholders. As mentioned earlier, in their seminal work, Shleifer and Summers (1988) study this question in the context of hostile takeovers and conclude that "transfers from stakeholders to shareholders could make for a large part of the takeover premium" (p. 53), including transfers from workers, suppliers, and the government (in the form of tax savings from high leverage). They note, however, that private equity encompasses both hostile takeovers and friendly M&As, and they are "targeted at very different companies" and represent difference "economic processes." This observation echoes the heterogenous outcomes on growth and productivity documented in Section 4 for public-to-private versus private-to-private transactions.

5.1 Employment and Wages at Portfolio Companies

The two most comprehensive studies of the effects of U.S. buyouts on employment and wages are Davis, Haltiwanger, Handley, Jarmin, Lerner, and Miranda (2014) and their follow-on study Davis, Haltiwanger, Handley, Lipsius, Lerner, and Miranda (2021).[3]

Davis, Haltiwanger, Handley, Jarmin, Lerner, and Miranda (2014) use U.S. census data to track employment at 3,200 portfolio companies and their 150,000 establishments that were acquired by private equity funds from 1980 to 2005. They match portfolio companies and establishments to similar controls and follow the targets and matched controls over several years after the private equity transaction. They find that portfolio companies reduce employment by less than 1% relative to control firms in the first two years, but this small net effect masks large differences in the outcomes at the establishment level. Specifically, private equity owners actively re-allocate employment across existing establishments by reducing employment or exiting establishments altogether at some locations, while expanding employment at other locations. Further, portfolio companies create more greenfield jobs at new establishments than control firms. Portfolio companies also acquire and divest more than controls. Finally, earnings per worker in portfolio companies declines by 2.4% relative to controls.

Expanding on this original study, Davis, Haltiwanger, Handley, Lipsius, Lerner, and Miranda (2021) examine U.S. private equity buyouts from 1980 to 2013, increasing the sample to about 3,600 targets and 6.4 million total company-level observations. The main insight of this follow-on study as it relates to employment impact of private equity buyouts is the sharp contrast in outcomes between public-to-private versus private-to-private deals: Employment shrinks 13% over two years after buyouts of publicly-listed firms relative to control firms, but expands 13% after buyouts of privately-held firms. The authors argue that "[f]or targets that trade publicly before the buyout, private equity groups may focus on tackling the agency problems … whether manifested as excess headcounts, wasteful perquisites, or value-destroying "pet projects." They also point out that the results are consistent with the workforce re-contracting hypothesis (through the breach of implicit long-term contracts) of Shleifer and Summers (1988). In contrast, at targets that were privately

[3] Kaplan (1989a) finds that raw median employment growth (excluding divestitures) of large public-to-private buyouts is 4.9% and negative but insignificant when industry-adjusted. Faccio and Hsu (2017) find evidence of higher job creation by targets of politically connected private equity firms than non-connected private equity firms. Consistent with an exchange of favors story, politically connected private equity firms increase employment more during election years and in states with high levels of corruption.

held before buyouts, the main constraints that private equity ownership addresses are the access to capital markets and managerial competence, not agency costs.

Antoni, Maug, and Obernberger (2019) study 511 German private equity buyouts and match establishment-level data with individual data. They find that buyout establishments reduce employment by 8.96% relative to controls, which consists of an increase in the separation rate by 18.75% and an increase in the hiring rate of 9.79%. The active job creation and destruction results resemble those of Davis et al. (2014, 2021). While the large average job losses at the establishment level appears to be at odds with Davis et al. (2014, 2021), note that Antoni et al. (2019) do not include new establishments that are opened after the buyouts, which Davis et al. (2014, 2021) include. Employment at existing establishments in Davis et al. (2014), for example, declined by 3% relative to controls for U.S. buyouts.

Evidence on the effect of private equity ownership on wages is mixed. Part of the empirical challenge is that it is difficult to isolate wage changes for a given job from compositional changes in the overall mix of jobs in the total employment pool. Researchers typically observe the total compensation and the headcount, rather than wages for specific workers. Thus, the average compensation per worker could change after buyouts due to compositional changes, even if the wage for each position remains unchanged. In the U.S., Davis, Haltiwanger, Handley, Lipsius, Lerner, and Miranda (2021) find that compensation per worker rises in divisional targets while it falls for private-to-private deals, while the changes are not statistically significant for public-to-private and secondary deals. The authors suggest that increases at divisional buyouts may reflect "job title upgrading"—i.e., pay increase that comes with new titles (e.g., CEO as opposed to divisional manager) and increased responsibilities, mostly concentrated at top managerial positions. In another earlier U.S. study, Lichtenberg and Siegel (1990) find that wages for white-collar workers decline after buyouts, whereas those for blue-collar workers remain unchanged.

In a study of 1,350 UK buyouts, Amess and Wright (2007) find that LBOs have significantly lower wage growth than non-LBOs. Similarly, Antoni, Maug, and Obernberger (2019) find that the average German buyout target employee loses €980 in annual earnings relatives to employees at control group firms, or 2.8% of median earnings. Note that Antoni et al. (2019) overcome the data issues described above by using individual worker-level data.

Fang, Goldman, and Roulet (2021) study private equity buyouts matched to French administrative data on employee pays and find that the pay gaps between young and old, men and women, and mangers and non-managers decrease after the buyout. Compositional effect drives these results: the companies replace expensive employees with cheaper ones and the remaining employees

receive small pay increases. The results suggest that wage inequality declines as a side effect of private equity pursuing profit maximization. The compositional change is consistent with the overall active re-allocation and shuffling of jobs by private equity reported in Davis et al. (2021).

5.2 Impact on Workers' Welfare

Most existing studies on employment and private equity are focused on jobs (positions) rather than workers (individuals). Several papers shed light on the impact of private equity ownership on career outcomes of workers who were employed at portfolio companies at the time of buyouts, and determinants of variation in their outcomes. Focusing on individual workers as the unit of analysis and tracing them over time after departures from portfolio companies is useful for understanding the distributional impact of private equity ownership on stakeholders and the broader society. It also helps evaluate evidence of wealth transfers as opposed to value creation, as suggested by Shleifer and Summers (1988): "[t]o see whether the parties that lose association with the acquired firm suffer wealth losses, one must trace their subsequent employment."

Reporting from Sweden, using employer-employee linked register data, Olsson and Tåg (2017) study the differential incidence of unemployment among workers performing different job tasks and at different positions within the wage distribution at portfolio companies. They find that workers performing automatable routine tasks at targets that lagged behind peers in productivity pre-deal were 10.2 percentage points more likely to experience unemployment spells, and workers performing offshorable job tasks were similarly 8.6 percentage points more likely to experience unemployment spells. This is despite the fact that there is little evidence of average changes in unemployment after the buyouts; in other words, the unemployment increases are unevenly distributed and is concentrated in workers whose job tasks were replaceable by either internal IT investment in automation or by offshoring. Routine task workers experience 12.7% decrease in labor income after separation from the portfolio company, suggesting that they receive lower wages at a new job. Both findings of unemployment and lower labor income are consistent with wealth transfers from separated workers to acquired companies. At the same time, productivity gain at low-productivity companies from automation and offshoring is consistent with value creation. Finally, layoffs are concentrated among workers in the middle of the wage distribution at portfolio companies, thus supporting the notion that private equity ownership accelerates job polarization within portfolio companies.

The polarization result appears somewhat at odds with the French results in Fang, Goldmanm, and Roulet (2021), who report that expensive (older, male) workers were more likely to separate and are replaced by cheaper (e.g., younger) workers, resulting in lower wage inequality

within firms. However, Olsson and Tåg (2017) note that the elevated unemployment results in their study are concentrated in low-productivity Swedish firms, which may be more similar to the U.S. public-to-private targets studied in Davis et al. (2014), whereas high-productivity Swedish firms may resemble growth-oriented French portfolio companies studied in Boucly, Sraer, and Thesmar (2011). Thus, it is possible that public-to-private and private-to-private deals also have opposite wage distributional effects.

Antoni, Maug, and Obernberger (2019) study German buyouts and find that managers and older workers are not more likely to be fired than other types of workers, but conditional on leaving the portfolio companies they appear to have a more difficult time finding new jobs with equivalent pay. In contrast, low-wage workers are more likely to be fired but will quickly start at another low-wage job. Consequently, managers and older workers are worse off than other workers as a result of buyouts. On the other hand, the authors also find that jobs that require stronger IT skills increase in two years after buyouts, consistent with private equity implementing investments in IT upgrading that increases productivity of skilled workers while replacing less-skilled workers with automation.

Focusing more on the spillover effects of IT upgrading investments on the welfare of skilled workers at portfolio companies, Agrawal and Tambe (2016) study long-term career paths of employees at portfolio companies relative to matched non-private equity employees. First, they find that portfolio companies hire more IT workers relative to controls in years after 2000, suggesting that private equity firms regard IT upgrading as part of their operational improvement strategies, especially after the Internet/e-commerce boom of the late 1990s. Second, they find that employees at portfolio companies experience 6-9 percentage points longer employment spells relative to control groups, both while working at the target firms and after separating from them and moving to other employers. The effect is driven by workers whose jobs are transformed by IT diffusion (including production), consistent with the interpretation that these workers are now expected to acquire IT-complementary skills on the jobs, which then become transferrable skills that help the workers advance their careers both within and outside of the original employers. Further, workers who perform IT-complementary tasks experience shorter unemployment spells after leaving the portfolio company and earn higher long-run wages.

Beyond long-run employability and wage growth at the worker level, Cohn, Nestoriak, and Wardlaw (2021) present evidence of a large, persistent decline in establishment-level workplace injury rates after buyouts of publicly-traded U.S. companies. Annual injuries per employee fall by 0.74 to 1.00 percentage points relative to control groups, or 11.1% to 15.0% of the pre-buyout mean. The injury rates decline more sharply at firms that were under more short-term performance

pressure before buyouts (e.g., more analyst coverage, transitory institutional ownership, and discretionary accruals). The authors argue that the results dovetail with the view that "the private nature of private equity ownership promotes long-term investment by removing a firm from the scrutiny of public markets."

The studies reviewed in this section illustrate the bifurcated outcomes for worker welfare depending on (1) pre-private equity ownership status (public or private) and (2) whether IT complements or substitutes your job. For either workers at private-to-private targets or workers who perform IT-complementary tasks in their jobs, private equity ownership is associated with increased employment and improved long-run wage growth in general. In contrast, for either workers at public-to-private targets or workers who perform automatable or offshorable tasks, private equity ownership is associated heightened risk and longer spells of unemployment, and slower long-run wage growth.

5.3 Other Stakeholders

Brown, Fee, and Thomas (2009) find that suppliers to LBO companies experience significantly negative abnormal returns at the announcements of downstream LBOs. They also find that suppliers who have likely made substantial relationship-specific investments are more negatively affected, both in terms of abnormal stock returns and reduced profit margins, than suppliers of commodity products or transitory suppliers. Interestingly, the results are not present for recapitalizations, suggesting that "increases in leverage *combined* with changes in organizational form result in supplier price concessions." The results are consistent with wealth transfers from suppliers to portfolio companies under the new management who uses the elevated leverage as a commitment device to drive harder bargains with suppliers.

6. IMPACT ON BROADER SOCIETY

In this section, we review the literature that examines the impact of private equity on broader society. As discussed in Section 3, the effect of high-powered incentives of private equity portfolio company executives on the broader society is ambiguous and remains an open empirical question. As shown below, one emerging insight from the extant literature is that the welfare outcomes for the broader environment and society depend on the regulatory and competitive structures within which the private equity portfolio companies operate. In competitive industries and industries that rely little on governments as payers, private equity ownership tends to result in enhanced consumer welfare, whereas in more concentrated industries and industries heavily dependent on governments as payers, private equity ownership tends to lead to pursuit of profit maximizations at the expense

of consumers. Similarly, the impact of private equity ownership on the environment is sensitive to the regulatory regimes under which the private equity portfolio companies operate, and the incentives the regulations give to the company executives.

6.1 Impact on Consumers through Portfolio Company Products and Services

As private equity funds transform their portfolio companies' operations, what impact will they have on consumers who either purchase products and services from the portfolio companies themselves or their competitors? For private equity ownership to have any impact on consumers, the private equity management model needs to include some operational engineering elements, and not just financial engineering elements. Moreover, even when operational engineering is at work, its impact on consumers can be positive or negative. Indeed, the literature finds divergent results depending on the competitive structure of the sectors that companies operate in. In lightly regulated, competitive industries with price-elastic demand, private equity ownership of private targets tends to be associated with enhanced consumer welfare via improved service, flat prices, and greater product variety. In contrast, in more heavily regulated, or government-subsidized industries, especially when operating in low-competition markets, private equity ownership can lead to diminished consumer welfare via higher prices, lower service quality, or both.

6.1.1 Competitive Industries

Bernstein and Sheen (2016) study the impact of private equity ownership on the operations of fast-food chains by studying comprehensive health inspection records at franchise locations in Florida. The authors find that restaurants become cleaner, safer, and better maintained after private equity buyouts, suggesting that private equity owners use their industry expertise to improve the business operations of their portfolio companies in the chain restaurant industry. Importantly, causal inference is made possible by comparing franchise-owned and private equity-owned restaurants. Improved health inspection performance is associated with greater customer satisfaction and restaurant profitability, so these operational changes enable private equity to achieve financial gains while also generating benefits for restaurant customers.

Fracassi, Previtero, and Sheen (2021) use micro-level retail scanner data to compare product varieties and prices of consumer products sold by private equity-owned and non-private equity-owned companies in retail stores. The authors find that private equity-acquired companies increase sales by 50% compared to matched control companies, not by increasing prices but by launching new products and expanding geographically to new store locations. Their sale growth squeezes their competitors' products out of shelf space. Interestingly, these results are driven entirely by

private-to-private deals. The findings are consistent with growth-oriented buyouts of capital-constrained, privately-owned companies. In contrast, public targets raise prices and reduce sales for existing products, echoing earlier results in Chevalier (1995a, 1995b) that find that private equity-acquired supermarket chains (most of which were public targets) tend to increase prices in local markets with other, highly leveraged rivals. Ewens, Gupta and Howell (2022) study private equity investments in local newspapers and find that production of local news content declines after private equity acquisitions, which is consistent with cost reduction.

In summary, private equity acquisition tends to have benign impact on consumers when targets are private and/or operate in competitive markets with low barriers to entry. When the targets are public or operate in less competitive local markets, private equity acquisition can result in higher prices and reduced sales or product availability to consumers.

6.1.2 Regulated or Subsidized Industries

The private equity management model manifests itself differently when it acquires targets in regulated or subsidized industries. In this section we review evidence from the healthcare and education sectors.

Healthcare Sector Evidence The healthcare sector is highly regulated, where revenues of healthcare providers are highly dependent on the reimbursement rates accepted by either the government payers (e.g., Medicare) or private insurers. Since consumers (recipients of healthcare services) typically do not fully pay out of pocket, incentives for providers to compete on price are considerably weaker. The opaqueness of the reimbursement process also makes it difficult for consumers to shop on prices ex ante. Finally, it is difficult to assess ex ante the product or service quality, though pro-competitive policies can mitigate this. These features appear to significantly shape the way in which private equity ownership impacts the healthcare service quality and price that consumers receive.

Gupta, Howell, Yannelis, and Gupta (2020) study the impact of private equity ownership on the quality of care for short-stay Medicare patients at for-profit nursing homes. They find that nursing staffing declines, while bed utilization increases, resulting in improved operational efficiency. However, the efficiency comes at the expense of declining quality, i.e., higher short-term mortality and lower incidence of Five Star ratings at private equity-owned facilities. The authors emphasize that because nursing homes largely rely on Medicare and Medicaid programs (that pay fixed rates per patient per day) for revenue, main levers for increasing profitability are staffing cost reduction, admitting more lucrative Medicare patients, and making them stay longer.

In a complementary study, Huang and Bowblis (2019) study the impact of private equity ownership on patient outcomes for long-stay Medicaid residents by comparing private equity-owned and non-private equity-owned for-profit nursing homes. In contrast to Gupta et al. (2020) that find declining health outcomes for short-stay post-acute elderly patients for private equity-acquired facilities, Huang and Bowblis (2019) find no significant difference between private equity -owned and non-private equity-owned for-profit nursing homes for long-stay Medicaid (non-elderly disabled) patients.

It is interesting that even in the highly regulated nursing home industry, the impact of private equity ownership on consumers is not uniform. What can explain the divergent findings? Gandhi, Song, and Updrashta (2021) suggest (i) heterogenous local market competitiveness and (ii) private equity managers' heightened responsiveness to competitive incentives may be key. They find that private equity acquirers compete on quality in locally competitive markets by increasing high-skilled registered nurse staffing more aggressively, while doing so only modestly in less competitive markets. They also document that, after the introduction of the Five Star System that improves transparency of staffing quality to consumers, private equity-owned facilities in competitive local markets increased staffing expenditure significantly whereas in low-competition markets they decreased staffing expenditure. Thus, even within the same industry, the impact of private equity ownership can differ depending on the competitive incentives the companies receive in the marketplace.

Beyond quality, other researchers have looked at private equity ownership's effect on consumer healthcare spending. Liu (2021) uses insurance claims data of privately insured individuals to study the impact of private equity entry to local hospital markets on healthcare spending by consumers. Counterfactual analyses of structural model estimations suggest that if private equity ownership of hospitals were banned, healthcare spending in local markets where private equity is present would drop by 11%. Higher healthcare spending in private equity-affected markets is driven mostly by higher negotiated prices with insurers, rather than with higher hospital service utilization. Rival hospitals in local markets also raise prices but only if they share a common insurer with private equity-backed hospitals. In addition to superior negotiation expertise of private equity owners, the bargaining model estimation implies that credible bankruptcy threat of highly-leveraged private equity-backed hospitals weakens the bargaining position of insurers, leading to higher prices. Furthermore, consistent with non-pecuniary benefit enjoyed by non-profit hospitals, price increases are larger when non-profit hospitals are acquired by private equity. In sharp contrast to the nursing home industry where the government payout rate is fixed, operational efficiency is barely changed under private equity ownership; instead, the prices charged to insurers

for the same quality service is sharply increased. To the extent that insurers pass on their increased spending to consumers via increased premiums, the study implies that private equity ownership of hospitals is detrimental to patient-consumer surplus via increased prices, while quantity and quality of services rendered are insignificantly changed.

Education Sector Evidence Higher education is another sector where, similar to the healthcare sector, government subsidies (federal student aid for low-income students) are crucial sources of revenues and many customers (students) pay subsidized tuition, thus weakening the providers' incentives to compete on price. The net tuition that students will pay (after grants and loans) is opaque, making it difficult for students to shop on price. There is a reputational incentive not to compete on price by providers, as high price is used as a signal for high quality. Finally, it is difficult to measure the quality of education both ex ante and ex post.

Eaton, Howell, and Yannelis (2020) study the impact of private equity acquisition on for-profit colleges' actions. The authors find that private equity acquisition leads to higher tuition and per-student debt, while education inputs, graduation rates, and earnings among graduates decline. Government aid is exploited more aggressively, and while loan repayment rates are lower, schools are not hurt because of government loan guarantees. The findings suggest that profit maximization through exploitation of government subsidies and loan guarantees may result in worse outcomes for customers (students) both through poorer education quality and higher prices.

For-profit schools, private equity-owned or not, appear to exploit federal aid programs and charge higher tuitions while capturing federal subsidies (Cellini and Goldin 2014). Eaton, Howell, and Yannelis (2020) argue that high-powered incentives of private equity ownership induce private equity-owned schools to more aggressively pursue profit maximization at the expense of students and the federal government.

6.2 Impact on Governments and Taxpayers

The private equity management model can affect government revenues and expenditures through several channels. First, financial engineering reduces tax liabilities of portfolio companies via increased interest tax shields (Kaplan 1989b). While this reduces tax collection from the private equity portfolio companies, governments may collect more taxes from banks who collect the increased interest payments from the portfolio companies, so the net effect is likely moderated. Second, private equity-backed companies may engage in other tax- and cost-avoidance activities as part of a shareholder value-creation strategy, including regulatory arbitrage. Note that with non-

leverage-related tax reductions, government tax revenues are effectively reduced (not recovered from banks).

Third, private equity-backed companies may engage in aggressive subsidy capture or risk-taking behavior that potentially costs taxpayers in events of distress or defaults by either the companies themselves or their stakeholders.

Interest Tax Shields Cohn, Mills, and Towery (2014) study U.S. federal corporate tax return data and find that private equity-acquired companies' leverage remain elevated even several years after the buyout, suggesting that private equity takeovers represent a one-time permanent change in the capital structure of the companies. Elevated leverage creates value via increased present values of interest tax shields and implies reduced corporate tax revenues from private equity portfolio companies. This suggests that as a greater portion of companies in the economy are acquired by private equity, permanently higher leverage of these companies leads to significantly lower corporate tax revenues for the governments, all else equal.

Non-Leveraged-Based Tax Avoidance Badertscher, Katz, and Rego (2013) study private U.S. companies with public debt and find that private equity-owned companies engage in greater tax avoidance than management-owned companies. Extending this study, Olbert and Severin (2021) study European buyouts and find that target companies' effective tax rates decrease by 15% after the private equity buyout. Targets engaging in post-buyout tax avoidance invest less in physical assets and employment and fit the category of buyouts that create value via cost cutting rather than growth. The authors examine industry-wide real effects and find that private equity ownership reduces overall corporate tax revenues and industry-wide effective tax rates without creating positive spillovers for other tax bases (e.g., consumption tax). Together, these findings suggest that some private equity investors impose a negative externality on local domestic governments through increased tax avoidance.

In a study of private equity investments in life insurance companies, Kirti and Sarin (2020) also report that private equity-owned insurers aggressively engage in tax arbitrage by reinsuring their contracts with subsidiaries domiciled in tax havens with 0% corporate tax rates.

Collectively, these papers suggest that minimizing tax payments is one of the value-creation strategies pursued by private equity owners, especially for companies with moderate growth prospects, and this has a potential negative effect on domestic government tax revenues.

Indirect Effects on Governments and Taxpayers Before the Financial Crisis of 2008–09, poorly-rated private-label ABS holdings required a higher capital charge. But in the aftermath of the Crisis, insurance regulators exempted insurers from this capital charge requirements to prevent massive fire sales of downgrade private-label ABS by insurers. Becker, Opp, and Saidi (forthcoming) and Koijen and Yogo (2016) document that U.S. insurance companies, under pressure for reaching-for-yield in a low-interest environment, exploit this regulatory forbearance by holding high-yielding ABS while avoiding capital charge.

With these industry dynamics as a backdrop, Kirti and Sarin (2020) study private equity ownership of life insurance companies and find evidence that private equity ownership exacerbates this regulatory arbitrage incentive. The authors document that acquired insurers aggressively engage in regulatory arbitrage by selling high-rated but low-yielding corporate bonds and buying poorly-rated and high-yielding ABS within days of the buyouts. The compositional changes of their portfolios increase profitability of the private equity-owned insurers but elevates the riskiness of their portfolios and may exacerbate a hidden cost on taxpayers in the event of distress or failures of these insurers.

This finding echoes the findings of Eaton, Howell, and Yannelis (2020), which show that private equity ownership exacerbates the distortions in the regulatory framework exploited by for-profit schools targeting students who receive federal aid. Since student defaults are guaranteed by the federal government, for-profit schools lack incentives to enable students to acquire enough earning power to repay the loans or to keep the loan amount down at a sustainable level. Similarly, since insurance companies are not punished via higher capital requirements to hold risky ABS after the removal of regulatory capital requirements for ABS in 2009, their incentives are distorted to hold high-yield ABS without regard to their credit risk. While all for-profit schools and insurers are given distorted incentives, private equity ownership is shown to induce more aggressive exploitation of the regulatory arbitrage opportunities.

As a counterpoint, Johnston-Ross, Ma, and Puri (2020) study private equity participation in the failed bank resolution process during the 2008-2009 Financial Crisis and find that private equity acquirers helped stabilize the financial system by providing capital to failed bank resolutions and saving taxpayers resolution costs in the process. The authors find that private equity investors acquire banks in poorer health and in need of greater capital injection ex ante, and yet these banks recover better (i.e., keep branches open and re-grow deposits). The private equity owners are repeat bank acquirers and introduce highly experienced and skilled management teams to failed banks, often with turnaround expertise. The authors estimate that private equity acquisitions allowed the FDIC to reduce the resolution costs by $3.63 billion.

It is interesting that even in the highly regulated and subsidized industry of banking, private equity incentives and the government incentives seem to be better aligned in the case of failed bank resolution during the crisis, whereas they appear to be more misaligned in for-profit healthcare and education. What determines the degree of (mis)alignments between private equity and public interests, and what policy interventions, if any, can mitigate them, is an important avenue for future research.

6.3 Impact on the Environment

As discussed in 2.5 and 3.3, a growing share of limited partners in private equity funds require that ESG be incorporated into the funds' investment processes. Whether their motivations are materiality-driven or impact-driven, limited partners who invest across public and private assets increasingly demand to know whether private portfolio companies that private equity funds invest in promote good environmental practice. The empirical evidence on this question is currently quite limited, and there is an acute need for more research. The question on private equity's role on the environment relates to a broader and growing literature on green banking and investor responsibility for monitoring the environmental impact of projects or companies that they fund.

Shive and Forster (2020) study mandated disclosures of greenhouse gas emissions of U.S. companies and find that, while private independent companies pollute less and are less likely to incur EPA violations than their public counterparts, private equity-owned companies do not differ from public companies in their emissions and violation rates. For a subset of utility companies for which reduction in emissions is shown to be costly and for which electricity output is measurable, the results hold after scaling, suggesting that public and private equity-owned companies eschew the costly pro-environmental actions and choose instead to narrowly maximize profits, going close to the legal limits and thus incurring more actions and violations from the EPA. In contrast, independent owners appear to take more costly pro-environmental actions relative to their counterparts. In light of the argument by Hart and Zingales (2017), see Section 3.3, the findings are consistent with the view that both the public company management and private equity-backed company management do not internalize pro-environmental shareholders' welfare, at least in the sample period of 2006-2017. It is possible that the environmental impact of portfolio companies' real activities was not a salient concern for limited partners in most of the sample period, and it would be interesting to see if private equity responds differentially (relative to public companies) to the heightened concerns for ESG performance among institutional investors in the post-Paris Accord era.

Bellon (2020) studies satellite imaging and administrative datasets for fracking wells to study the impact of private equity ownership on pollution decisions at individual well locations. On average, private equity ownership leads to a significant reduction in use of toxic chemicals for extraction and CO2 emissions from flaring. However, this average effect hides significant heterogeneities. Portfolio companies increase pollution in locations and periods where environmental liability risk is low, such as when the environmental regulation on federal land was rolled back. Overall, high-powered incentives to maximize shareholder value may benefit environmental outcomes when the risk of environmental regulation is high.

These studies confirm the insights from other studies of private equity firm behavior in regulated industries: high-powered incentives of private equity owners is a double-edged sword, and can either powerfully aid the policy goals when incentives are well aligned with the policymakers' intended goals (e.g., failed bank resolutions), or significantly exacerbate the distortion in the framework when incentives are misaligned (e.g., federal aid capture, oligopolistic price bargaining between hospitals and insurers, capital regulation forbearance in insurance, and roll-backs of environmental regulation on federal land). Collectively, the emerging evidence suggests that regulators need to consider the impact of the high-powered incentives of private equity when assessing the market impact of a given regulatory policy or decision.

6.4 Impact on Innovation and Industry Spillovers

As discussed in Section 5, the impact of private equity ownership on the company's real outcomes diverges between public-to-private and private-to-private deals. The shareholder value-creation proposition for public-to-private deals tends to be centered around efficiency gains and cost cutting. This is in contrast to the relaxation of capital constraints and top-line growth for private-to-private deals. This raises a question on the aggregate impact of private equity penetration on industry-wide innovation level, and whether the sign of the impact depends on the composition of public and private companies in an industry. A related second questions is whether public targets' pre-deal level of innovation is optimal, either from shareholders' or societal perspective.

Public-to-Private Deals

Lerner, Sorensen, and Stromberg (2011) study U.S. public-to-private buyouts and find that, while the level of patenting is unchanged post-buyouts, the patents in the post-buyout programs are better cited, and more narrowly focused. Ayash and Egan (2019) study U.S. public-to-private buyouts using a difference-in-difference approach and find that, compared to the matched control companies, private equity-owned companies reduce patent flows by one third, driven by both a

decline in new patents (23%) and fewer purchases (7%). The difference in inferences between the two studies stems in part from the fact that patenting level generally increased over time in the sample period. In a contemporaneous study, Cumming, Peter, and Tarsalewska (2020) study international public-to-private buyouts and find that private equity-backed buyouts are associated with a significant reduction in patents, patent citations, and innovator employment.

Private-to-Private Deals

Amess, Stiebale, and Wright (2016) study UK private equity buyouts and find that for private-to-private deals, quality-adjusted patent stocks increase by 14%, accompanied by relaxation of financial constraints. In contrast, quality-adjusted patent stocks weakly decline in the case of public-to-private deals. The findings are consistent with the view that private-to-private deals tend to provide growth equity, employ less leverage, and target smaller companies that hold growth options but are capital constrained. Driver, Kolasinski, and Stanfield (2021) find that private equity-held firms, though equally innovative as other private firms, skew their strategies toward development and away from research; however, their study does not differentiate between public-to-private and private-to-private deals.

Aggregate Impact

The divergent findings for public-to-private and private-to-private deals on innovation suggest that the aggregate impact of private equity penetration on a given industry depends on the composition of public and private companies in an industry that become private equity targets. An industry where private equity targets are primarily public (private), all else equal, is expected to experience an aggregate decline (increase) in innovation activities after an increase in private equity ownership. Since this composition varies from sector to sector and country to country, the average impact across industries or countries is a priori ambiguous. The impact can also change over time, as the composition changes.

There is a related older literature that documents a negative relationship between indebtedness and innovation for publicly-traded companies and for public-to-private transactions (Baysinger and Hoskisson 1989; Hall 1990; Long and Ravenscraft 1993). According to these studies, it is the elevated leverage rather than the private equity ownership that negatively impacts the company's propensity to invest in long-term and hard-to-assess investments in innovation. It is important to note that the debate on this question during the 1990's was tempered by the fact that private equity targets were mostly old-economy sector companies whose level of innovation was low even before the buyout. Since most innovative companies in the economy, e.g., Silicon

Valley tech companies, were not suitable LBO targets, the impact of private equity acquisition on the aggregate economy was thought to be limited. For example, Hall (1989) argued that "[e]ven if all of this R&D spending went away after going private, this would make a very small dent in overall industrial R&D."

In contrast, today's private equity invests in every sector of the economy, including IT and healthcare, and many of the buyout deals in these high R&D-intensity sectors employ leverage. For example, in 2021 about 40% of U.S. private equity deal values were in IT and healthcare sectors (PitchBook 2022). Enterprise software companies, for example, have become prime buyout targets as the SaaS, or Software-as-a-Service business model has become the industry norm and given management more tools to smooth out cash flows across periods. This is a recent phenomenon, and the literature has not closely examined the effect of either private equity ownership or higher leverage on innovation in high-tech industries. As the share of these high R&D-intensity companies and sectors in the economy grows rapidly, the aggregate effect of private equity penetration in these sectors on the innovation of the portfolio companies themselves and spillover effects on competitors remains an open question.

6.5 Pro-ESG Investing and Private Equity

One of the recurring findings on the impact of private equity on the broader society and the environment is the importance of aligning shareholder preferences with the broader public interest. Any misalignment due to distortion in the regulatory framework tends to be magnified under private equity ownership because of the high-powered incentives of the private equity management model to maximize shareholder value.

What if investors in private equity funds have explicitly prosocial incentives and impose their preferences on the private equity management model? Impact funds are private equity and VC funds that explicitly pursue dual objectives of both financial return and generation of positive (either social or environmental) externality. In essence, much like how nonprofit hospitals may internalize the positive externality of provision of quality healthcare in the local community as nonpecuniary benefit, impact funds aim to internalize the positive externality that the fund's portfolio companies generate as nonpecuniary benefit to fund investors. As investors pressure for-profit companies to adopt prosocial practice, e.g., net-zero pledges to fight climate change, a debate arises as to how for-profit companies can credibly commit to such pledges that are costly and impact profits negatively. Will the private equity fund mechanism work more or less effectively than the public company governance mechanism to induce prosocial behavior in the portfolio companies? What will be the return implications of imposing such preferences on fund activities?

Broccardo, Hart, and Zingales (2021) study voting and exiting as two strategies employed by prosocial investors and consumers to pressure companies to choose clean (vs. dirty) technology.[4] The authors build a model that is meant for public companies and advocate for proxy voting as a more effective mechanism to induce clean technology adoption than divestment. However, in practice it may be hard to implement a proxy voting solution to gather the collective preference of public company investors who are free to trade the stocks at a moment's notice. In contrast, the closed-end fund structure, a small and fixed set of limited partners, and illiquidity of fund interests create a more stable structure in which to articulate the prosocial goal of the fund and execute a long-term strategy to adhere to the goal. Limited partners will then vote with their feet when it is time to decide whether to re-up for the next fund the general partners raise, based on the combined financial and non-financial (environmental and/or social) performance of the first fund.

Geczy, Jeffers, Musto, and Tucker (2021) study limited partnership contracts of impact funds and find that impact funds give limited partners advisory roles that enable them to perform substantial oversight over deal selection, due diligence, conflict of interest, and other material fund activity. At the same time, the study finds that impact funds typically do not tie manager compensation explicitly to impact outcomes. The use of informal governance rather than explicit contracting to monitor impact performance raises several questions: Is this contracting form optimal or a reflection of a still nascent and rapidly evolving industry? Does it indicate the inherent difficulty of impact measurement, or uncertainty about the relationship between impact and financial performance and investors' ambiguity towards the trade-off? Clearly, more research is needed to answer these questions.

Pastor, Stambaugh, and Taylor (2021) analyze financial and real effects of sustainable investing in an equilibrium where (i) companies can either create positive (green) or negative (brown) externality and (ii) investors derive utility (disutility) for holding green (brown) assets, care about companies' aggregate social impact, and care about climate risk. In the model, pro-ESG investors' willingness to forgo return in exchange for investing in green-tilted portfolio lowers green companies' cost of capital. Climate risk also increases brown companies' expected return. Pro-ESG investors enjoy "investor surplus" despite earning negative alpha. This equilibrium framework is useful in understanding expected financial returns of impact funds.

[4] Also see Pursiainen and Tykvova (2021) for a study of how customers "vote with their feet" in response to announcements of buyouts of retail brands.

Impact fund investors derive utility from holding impact funds that generate positive impact, and thus are rationally willing to invest in them even though their expected financial return alone may be lower than that from investing in non-impact private equity funds.

Taking this insight to the fund-level financial performance data, Barber, Morse, and Yasuda (2021) estimate random-utility/willing-to-pay models and find that limited partners accept 2.5-3.7 ppts lower IRRs ex ante for impact funds, compared to comparable non-impact funds. The result is consistent with the view that investors derive nonpecuniary utility from investing in impact funds, thus sacrificing financial return. Development organizations, foundations, financial institutions, public pensions, Europeans, and United Nations Principles of Responsible Investment signatories have high willingness-to-pay for impact. Unpacking the channels behind this heterogeneity across investor types, the authors find that, on one hand, investors with mission objectives and/or facing political pressure have high willingness-to-pay; on the other hand, those subject to fiduciary duty-related restrictions against dual-objective investments are reluctant to invest in impact funds, likely for fear of running afoul of the regulation. These results are consistent with the predictions for pro-ESG investors in Pastor, Stambaugh, and Taylor (2021): pro-ESG investors earn negative alpha in expectation but are rationally willing to do so because of nonpecuniary utility they derive from holding impact funds in their portfolios.

Do impact funds actually generate positive externalities? Does the externality they generate correlate positively or negatively with the fund's financial returns? And within a given fund, does the externality each portfolio company investment generates correlate positively or negatively with the fund's financial returns from the investment? How does a fund measure the externality generated at each portfolio company, and how does the company attribute the externality generated to the investment by the impact fund vs. other investments it receives? These are just examples of a myriad of questions that remain open and are fruitful areas of future research.

7. OPEN QUESTIONS AND SUGGESTIONS FOR FURTHER RESEARCH

As discussed throughout the chapter, many questions remain open for future research and are summarized below.

Why do private equity-backed companies tend to exploit regulatory arbitrage and tax avoidance more aggressively than non-private equity-backed peers? Faccio and Hsu (2017) suggest that some private equity firms may benefit from political connections. Do such connections also enable them to pursue either regulatory capture or tax avoidance more successfully or at lower cost?

In regulated or subsidized industries, distortion in incentives given by the regulatory framework tends to get magnified when combined with high-powered incentives of private equity. What policy interventions, if any, can mitigate the misalignment of incentives between private equity and public interests?

While in the past private equity deals are concentrated in low-tech, consumer or industrial sectors, today more than a third of private equity investments are in IT and healthcare, the most innovation-driven segments of the economy. What is the aggregate impact of private equity's greater presence in the tech industry on innovation? Does it depend on the composition of public-to-private vs. private-to-private deals? What else matters?

How should impact funds govern and provide incentives for impact generation at the portfolio companies? If implicit rather than explicit contracting is optimal, what is the underlying mechanism? This is a broader question for governing ESG practice at public and non-private equity-backed private companies, too, and there is a potential for innovating on contracting that may have broader applicability for aligning shareholder preferences with the broader public interest.

8. REFERENCES

Acharya, Viral, Conor Kehoe, and Michael Reyner (2009) "Private Equity vs. PLC Boards In the U.K.: A Comparison of Practices and Effectiveness" Journal of Applied Corporate Finance 21(1), 45–56

Acharya, Viral, Oliver Gottschalg, Moritz Hahn, and Conor Kehoe (2013) "Corporate Governance and Value Creation: Evidence from Private Equity" Review of Financial Studies, 26(2), 368–402

Agrawal, Ashwini, Prasanna Tambe (2016) "Private Equity and Workers' Career Paths: The Role of Technological Change" Review of Financial Studies, 29(9), 2455–2489

Amess, Kevin and Mike Wright (2007) "The Wage and Employment Effects of Leveraged Buyouts in the UK" International Journal of the Economics of Business, 14(2), 179–195

Amess, Kevin, Joel Stiebale, and Mike Wright (2016) "The Impact of Private Equity on Firms' Innovation Activity" European Economics Review, 86, 147–160

Antoni, Manfred, Ernst Maug, and Stefan Obernberger (2019) "Private Equity and Human Capital Risk" Journal of Financial Economics 133, 634-657

Ayash, Brian and Edward J. Egan (2019) "The Effect of U.S. Public-to-Private Leveraged Buyouts on Innovation" working paper

Badertscher, Brad, Sharon Katz, and Sonja Rego (2013) "The Separation of Ownership and Control and Corporate Tax Avoidance" Journal of Accounting and Economics, 56, 228–250

Barber, Brad, Adair Morse, and Ayako Yasuda (2021) "Impact Investing" Journal of Financial Economics 139, 162-185

Baysinger, B. and R. E. Hoskisson (1989) "Diversification Strategy and R&D intensity in multiproduct firms" Academy of Management Journal 32(2), 310–332.

Bebchuk, L. A., and Weisbach, M. S. (2010) "The State of Corporate Governance Research" The Review of Financial Studies, 23(3), 939-961

Bebchuk, L. A., Grinstein, Y., and Peyer, U. (2010) "Lucky CEOs and Lucky Directors" The Journal of Finance, 65(6), 2363-2401

Bellon, Aymeric (2020) "Does Private Equity Ownership Make Firms Cleaner? The Role Of Environmental Liability Risks" working paper

Bernstein, Shai, Josh Lerner, Filippo Mezzanotti (2018) "Private Equity and Financial Fragility during the Crisis" Review of Financial Studies, 32(4), 1309–1373

Bernstein, Shai, and Albert Sheen (2016) "The Operational Consequences of Private Equity Buyouts: Evidence from the Restaurant Industry" Review of Financial Studies 29(9) 2387–2418

Bertrand, Marianne, and Sendhil Mullainathan (2003) "Enjoying the Quiet Life? Corporate Governance and Managerial Preferences" Journal of Political Economy, 111(5), 1043–1075

Biesinger, Markus, Cagatay Bircan, and Alexander Ljungqvist (2020) "Value Creation in Private Equity" working paper

Bloom, Nicholas, Rafaella Sadun, and John Van Reenen (2015) "Do Private Equity Owned Firms Have Better Management Practices?" American Economic Review: Papers & Proceedings 105(5) 442–446

Bouchly, Quentin, David Sraer, David Thesmar (2011) "Growth LBOs" Journal of Financial Economics 102, 432–453

Braun, Robert Tyler, Hyunkyung Yun, Lawrence P. Casalino, Zachary Myslinski, Farai Kuwonza, Hye-Young Jung, Mark Aaruon Unruh (2020) "Comparative Performance of Private Equity-Owned US Nursing Homes During the COVID-19 Pandemic" JAMA Network Open, 3(10)

Broccardo, Eleonora, Oliver D. Hart, and Luigi Zingales (2021) "Exit vs. Voice" working paper

Brown, David T., C. Edward Fee, and Shawn E. Thomas (2009) "Financial Leverage and Bargaining Power with Suppliers: Evidence from Leveraged Buyouts" Journal of Corporate Finance, 12, 196–211

Brown, Greg, Bob Harris, Tim Jenkinson, Steven Kaplan, and David Robinson (2020) "Private Equity: Accomplishments and Challenges" Journal of Applied Corporate Finance (2020), 32(3), 8–20

Burrough, Bryan and John Helyar (1990) Barbarians at the Gate, HarperBusiness

Chhaochharia, V., and Y. Grinstein (2006) "Corporate Governance and Firm Value: The Impact of the 2002 Governance Rules" Journal of Finance 62:1789–825.

Chevalier, Judith (1995a) "Do LBO Supermarkets Charge More? An Empirical Analysis of the Effects of LBOs on Supermarket Pricing" Journal of Finance, 50(4), 1095–1112

Chevalier, Judith (1995b) "Capital Structure and Product-Market Competition: Empirical Evidence from the Supermarket Industry," American Economic Review, 85(3), 415–435

Chung, Ji-Woong, Berk Sensoy, Lea Stern, and Michael Weisbach (2012) "Pay for Performance from Future Fund Flows: The Case of Private Equity" Review of Financial Studies 25(11) 3259–3304

Cohn, Jonathan, Edith Hotchkiss, Erin Towery (2022) "Sources of Value Creation in Private Equity Buyouts of Private Firms" Review of Finance, 26(2) 257–285

Cohn, Jonathan, Lillian Mills, and Erin Towery (2014) "The Evolution of Capital Structure and Operating Performance after Leveraged Buyouts: Evidence from U.S. Corporate Tax Returns" Journal of Financial Economics, 111(2) 469–494

Cohn, Jonathan, Nicole Nestoriak, and Malcolm Wardlaw (2021) "Private Equity Buyouts and Workplace Safety" The Review of Financial Studies, *34*(10), 4832-4875.

Cornelli, Francesco, Ohuzhan Karakas (2015) "CEO Turnover in LBOs: The Role of Boards" working paper

Cornelli, Francesca, Zbigniew Kominek, Alexander Ljungqvist (2013) "Monitoring Managers: Does it Matter?" Journal of Finance, 68(2), 431–481

Cronqvist, Henrik, and Rüdiger Fahlenbrach (2013) "CEO Contract Design: How do Strong Principals do it?" Journal of Financial Economics, 108, 659–674

Davis, Steven, John C. Haltiwanger, Kyle Handley, Ron Jarmin, Josh Lerner, and Javier Miranda (2014) "Private Equity, Jobs, and Productivity" American Economic Review 104(12): 3956-3990

Davis, Steven, John C. Haltiwanger, Kyle Handley, Ben Lipsius, Josh Lerner, and Javier Miranda (2021) "The (Heterogenous) Economic Effects of Private Equity Buyouts" working paper

Driver, Kolasinski, Stanfield (2020) "R&D or R vs. D? Firm Innovation Strategy and Equity Ownership" working paper

Eaton, Charlie, Sabrina T. Howell, and Constantine Yannelis (2020) "When Investor Incentives and Consumer Interests Diverge: Private Equity in Higher Education" Review of Financial Studies 33, 4024–4060

Edgerton, Jesse (2012) "Agency Problems in Public Firms: Evidence from Corporate Jets in Leveraged Buyouts" Journal of Finance, 67(6), 2187–2213

Ewens, M., A. Gupta, and S. Howell (2022) "Local Journalism under Private Equity Ownership" Working paper

Faccio, Mara, Hung-Chia Hsu (2017) "Politically Connected Private Equity and Employment" Journal of Finance, 72(2), 539–573

Fama, Eugene (2021) "Contract Costs, Stakeholder Capitalism, and ESG" European Financial Management 27, 189–195

Fich, E. M., and Shivdasani, A. (2012) "Are Busy Boards Effective Monitors?" In Corporate Governance (pp. 221-258). Springer, Berlin, Heidelberg

Fink, Larry (2019) "Profit & Purpose", BlackRock website (https://www.blackrock.com/americas-offshore/en/2019-larry-fink-ceo-letter).

Fink, Larry (2020) "A Fundamental Reshaping of Finance", BlackRock website (https://www.blackrock.com/corporate/investor-relations/2020-larry-fink-ceo-letter).

Fracassi, Cesare, Alessandro Previtoro, and Albert Sheen (2021) "Barbarians at the Store? Private Equity, Products, and Consumers" forthcoming in Journal of Finance

Gandhi, Ashvin, YoungJun Song, and Prabhava Upadrashta (2021) "Have Private Equity Owned Nursing Homes Fare Worse Under COVID-19?" working paper

Geczy, Christopher, Jessica Jeffers, David Musto, and Anne Tucker (2021) "Contracts with (Social) Benefits: The Implementation of Impact Investing" Journal of Financial Economics, *142*(2), 697-718.

Gompers, Paul, Steven N. Kaplan, and Vladimir Mukharlyamov (2016) "What do Private Equity Firms Say They Do?" Journal of Financial Economics, 121, 449–476

Guo, Shorun, Edith S. Hotchkiss, Weihong Song (2011) "Do Buyouts (Still) Create Value?" Journal of Finance, 66(2), 479–517

Gupta, Atul, Sabrina T. Howell, Constantine Yannelis, Abhinav Gupta (2020) "Does Private Equity Investment in Healthcare Benefit Patients? Evidence from Nursing Homes" working paper

Hall, B. (1989). Corporate Restructuring and R&D: Hearing before the Subcommittee on Science, Research and Technology of the Committee on Science, Space and Technology, U.S. House of Representatives, 101st Cong. 43 (Testimony of Bronwyn Hall).

Hall, B. (1990) "The Impact of Corporate Restructuring on Industrial Research and Development," The Brookings Institution 21, 85–135.

Hart, Oliver and Luigi Zingales (2017) "Companies Should Maximize Shareholder Welfare Not Market Value" Journal of Law, Finance, and Accounting, 2: 247–274

Hotchkiss, Edith, David Smith, and Per Strömberg (2021) “Private Equity and the Resolution of Financial Distress” The Review of Corporate Finance Studies 10(4) 694–747

Jensen, Michael (1989) “The Eclipse of the Public Corporation” Harvard Business Review 67(5) 61–74

Johnston-Ross, Emily, Song Ma, and Manju Puri (2020) “Private Equity and Financial Stability: Evidence from Failed Bank Resolution in the Crisis” working paper

Kaplan, Steven (1989a) “The Effects of Management Buyouts on Operating Performance and Value,” Journal of Financial Economics, 24(2), 217–254

Kaplan, Steven (1989b) “Management Buyouts: Evidence of Taxes as a Source of Value,” Journal of Finance, 44(3), 611–632

Kaplan, Steven, Mark M. Klebanov, Morten Sorensen (2012) “Which CEO Characteristics and Abilities Matter?” Journal of Finance, 67(3), 973–1007

Kaplan, Steven and Per Strömberg (2009) "Leveraged Buyouts and Private Equity," Journal of Economic Perspectives, 23(1), 121–146

Katz, Sharon (2009) “Earnings Quality and Ownership Structure: The Role of Private Equity Sponsors” The Accounting Review, 84(3), 623–658

Kirti, Divya and Natasha Sarin (2020) “What Private Equity Does Differently: Evidence from Life Insurers” working paper

Koijen, R. S. J., and Yogo M. (2016) Risks of Life Insurers: Recent Trends and Transmission Mechanisms. In Felix Hufeld R. S. K. and Thimann C. (eds.), The Economics, Regulation, and Systemic Risk of Insurance Markets, chap. 4. Oxford: Oxford University Press.

Lerner, Josh, Morten Sorensen, and Per Strömberg (2011) “Private Equity and Long-Run Investment: The Case of Innovation” Journal of Finance, 66(2), 445–477

Leslie, Philip, and Paul Oyer (2008) “Managerial Incentives and Value Creation: Evidence from Private Equity” working paper

Lichtenberg, Frank and Donald Siegel (1990) “The Effects of Leveraged Buyouts on Productivity and Relate Aspects of Firm Behavior” Journal of Financial Economics, 27, 165–194

Liu, Tong (2021) “Bargaining with Private Equity: Implications for Hospital Prices and Patient Welfare” working paper

Long, W. F. and D. J. Ravenscraft (1993) “LBOs, debt and R&D intensity” Strategic Management Journal 14, 119–135.

Metrick, A. and A. Yasuda (2010) “The Economics of Private Equity Funds” The Review of Financial Studies, 23(6), 2303-2341

Metrick, A. and A. Yasuda (2011) “Venture Capital and Other Private Equity: A Survey” European Financial Management 17, 619-654.

Muscarella, Chris and Michael Vetsuypens (1990) “Efficiency and Organizational Structure: A Study of Reverse LBOs” Journal of Finance 45(5) 1389–1413

Olsson, Martin and Joacim Tåg (2017) “Private Equity, Layoffs, and Job Polarization” Journal of Labor Economics, 35(3), 697–754

Pástor, L, Stambaugh, R. F., and Taylor, L. A. (2021) “Sustainable Investing in Equilibrium” Journal of Financial Economics, 142(2), 550-571

Pedersen, Lasse Heje, Shaun Fitzgibbons, and Lukasz Pomorski (2021) “Responsible Investing: The ESG-efficient Frontier” Journal of Financial Economics 142(2) 572–597

PitchBook (2022) “2021 Annual US private equity Breakdown” PitchBook website (https://pitchbook.com/news/reports/2021-annual-us-pe-breakdown)

Preqin (2020) “Preqin Impact Report: The Rise of ESG in Alternative Assets” Preqin website (https://www.preqin.com/insights/research/reports/preqin-impact-report-the-rise-of-esg-in-alternative-assets)

Pursiainen, Vesa and Tereza Tykvova (2021) “Retail Customer Reactions to Private Equity Acquisitions” working paper

Renneboog, Luc, and Cara Vansteenkiste (2017) “Leveraged Buyout: A Survey of the Literature” working paper

Shive, Sophie A. and Margaret M. Forster (2020) “Corporate Governance and Pollution Externalities of Public and Private Firms” Review of Financial Studies 33(3) 1296–1330

Shleifer, Andrei, and Larry Summers (1988) “Breach of Trust in Hostile Takeovers” In Corporate Takeovers: Causes and Consequences: University of Chicago Press

Spaenjers, Christophe and Eva Steiner (2021) “The Value of Specialization in Private Equity: Evidence from the Hotel Industry” working paper

Wilson, Nick and Mike Wright (2013) “Private Equity, Buy-outs and Insolvency Risk” Journal of Business Finance & Accounting 40(7–8) 949–990

Journal of Financial Economics 139 (2021) 162–185

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Impact investing ☆

Brad M. Barber [a], Adair Morse [b,c], Ayako Yasuda [a,*]

[a] Graduate School of Management, UC Davis, One Shields Avenue, Davis, CA 95616, United States
[b] Haas School of Business, UC Berkeley, 2220 Piedmont Avenue, Berkeley, CA 94720, United States
[c] National Bureau of Economic Research, 1050 Massachusetts Avenue, Cambridge, MA 02138, United States

ARTICLE INFO

Article history:
Received 6 August 2018
Revised 13 December 2019
Accepted 10 January 2020
Available online 17 July 2020

JEL classification:
G11
G21
G22
G23
G24
G28
H41
M14

Keywords:
Impact investing
Venture capital
Private equity
Socially responsible investment
United nations principles of responsible investment (UNPRI)
Sustainable investing
Public pension funds
Willingness to pay
Random utility discrete choice models

ABSTRACT

We show that investors derive nonpecuniary utility from investing in dual-objective Venture Capital (VC) funds, thus sacrificing returns. Impact funds earn 4.7 percentage points (ppts) lower internal rates of return (IRRs) ex-post than traditional VC funds. In random utility/willingness-to-pay (WTP) models investors accept 2.5–3.7 ppts lower IRRs ex ante for impact funds. The positive WTP result is robust to fund access rationing and investor heterogeneity in fund expected returns. Development organizations, foundations, financial institutions, public pensions, Europeans, and United Nations Principles of Responsible Investment signatories have high WTP. Investors with mission objectives and/or facing political pressure exhibit high WTP; those subject to legal restrictions (e.g., Employee Retirement Income Security Act) exhibit low WTP.

☆ We are grateful for the comments of Philip Bond, Magnus Dalquist, Lisa Goldberg, Will Gornall, Tim Jenkinson, Oğuzhan Karakaş, Hideo Owan, Ludovic Phalippou, Elena Pikulina, Josh Rauh, Berk Sensoy, Laura Starks, Annette Vissing-Jorgensen, and seminar and conference participants at AFA Annual Meetings (Chicago), Caltech/USC 2016 PE Conference, European Commission Promoting Sustainable Investment Conference, the Fifth International Conference on Financial Frictions in Copenhagen, LBS 9th Private Equity Symposium, 2017 LSE-Chicago Booth Economics of Social Sector Organizations Conference, Masahiko Aoki memorial academic conference, 2016 PERC Conference (Chapel Hill), Sustainable Finance Conference at Swedish House of Finance/ Stockholm School of Economics, 2016 WFA Annual Meetings (Park City), UW 2016 Summer Finance Conference, Cambridge University, Lancaster University, Maastricht University, New York Fed, Oxford University Saïd Business School, Tuck (Dartmouth), UC Berkeley (Law), University of Massachusetts, University of Warwick, University of Porto, University of Tilburg, and Yale School of Management. Radin Ahmadian, Christina Chew, and Derek Lou provided valuable research assistance. All errors are our own.

* Corresponding author.
E-mail address: asyasuda@ucdavis.edu (A. Yasuda).

https://doi.org/10.1016/j.jfineco.2020.07.008
0304-405X/

1. Introduction

Do investors knowingly accept lower expected financial returns in exchange for nonpecuniary benefits from investing in assets with both social and financial objectives? Classic asset pricing models generally define an investor's objective function using utility over wealth or consumption. While there have been innovations in the form of these utility functions (Epstein and Zin, 1989; Laibson, 1997), wealth generation is the common goal of investors. Economists are now taking seriously the possibility that investors might value positive societal externalities in utility in addition to wealth. Theoretical models consider the implications of these nonpecuniary preferences in a variety of settings (e.g., Andreoni, 1989, 1990; Fama and French, 2007; Hart and Zingales, 2017; Niehaus, 2014), yet these models start from a relatively untested assumption that nonpecuniary motives affect the allocation of capital in a way that reflects an intentional willingness to pay for impact.

A natural starting point is to look for indications of demand for nonpecuniary benefits by the sources of capital themselves. As of April 2019, 2372 organizations representing $86 trillion in asset under management have become signatories to the United Nations Principles of Responsible Investment (UNPRI). Virtually all major consulting groups have implemented a social impact practice, and all major investment banks have an impact division to meet corporate, institutional, and private wealth demands for impact considerations in investment. These indications of demand for investing with a social conscience do not imply that investors readily accept a tradeoff between financial returns and nonpecuniary benefits. For instance, the signing of the UNPRI accords does not imply that a holder of capital necessarily must tilt investment toward impact. Rather, UNPRI investors can comply by adhering to principles of governance within their investing entity.

An important, recent empirical literature on socially responsible investment (SRI) mutual funds shows that the demand for responsibility is growing rapidly (Bialkowski and Starks, 2016), reflecting both preferences and social signaling (Riedl and Smeets, 2017). However, performance in public market SRI has not been statistically different from other mutual funds in this period (see the amalgamation of evidence in Bialkowski and Starks, 2016). Hence, the tilt toward SRI need not reflect a willingness to pay in wealth for nonpecuniary benefits.

Thus, we study a different asset market—impact investing—to ask whether the theoretical assumption that investors are willing to pay for impact holds. Two primary instrument types that receive the largest capital allocation among impact investors are private debt and private equity.[1] While private debt is the largest category, we are not aware of any data sources for private debt impact investments. Instead, we focus on impact funds, which are predominantly Venture Capital (VC) and growth equity funds that are structured as traditional private equity funds but with the intentionality that is the hallmark of impact investing. The Global Impact Investing Network (GIIN) defines impact investing as "investments made with the intention to generate positive, measurable social and environmental impact alongside a financial return."[2] Thus, an impact investor exhibits an intention to generate both positive social or environmental returns and positive financial returns. Green washing investments, which is branding for an appearance of impact intentionality (Starks et al., 2017) and purely for-profit investment in sectors that associate with positive externalities (e.g., health, education, clean energy) do not meet the intentionality criteria. In our data collection, we ensure that we only choose impact VC funds that explicitly market a dual agenda.

Besides the data availability, the VC institutional setting brings an additional advantage. Because VC funds only fundraise at the inception of the fund and investors contractually commit their capital for the duration of the fund (typically ten years), the timing of capital flowing in and out of funds is not a concern in our setting. This institutional feature allows us to focus on the investors' discrete choice to invest in traditional VC funds versus impact VC funds among the observable choice set at a given point in time.

These advantages in the impact-versus-traditional VC market provide us with an ideal setting to identify any ex-ante willingness to pay for impact that investors may exhibit. We ask: (a) whether investors are intentionally willing to forego expected financial returns in exchange for expectation of impact, (b) whether this willingness to pay depends on the source of the capital (e.g., pension fund, bank, or development organization), and (c) whether the evidence points to any attributes (e.g., mission objectives, household versus institutional ownership, the legal or regulatory framework governing the allocation of capital) that explain heterogeneity in investor willing to pay for impact.

Using Preqin data, we construct a sample of 24,000 VC and growth equity (to which we refer together as VC for simplicity) investments by about 3500 investors over the period 1995–2014. These investments reflect 4659 funds—the combination of traditional VC and impact VC funds. We manually isolate 159 of these funds as being impact funds using a strict criterion that the fund must state dual objectives in its motivation. Investors are not all alike in their portfolio choice decisions; thus, we also manually look up the ultimate source of capital for each of the 3500 investors and code them into ten investor types. Our final piece of data coding is to codify the impact agenda themselves in more detail. The impact agenda of impact VCs are quite broad, including funds that seek to reduce greenhouse gas emissions, encourage the development of women and minority-owned firms, alleviate poverty in developing countries, or develop local business communities.

Our primary analysis estimates the willingness to pay (WTP) for impact across investor types and attributes. To set the stage for this analysis, we estimate reduced-form regressions of impact fund performance compared to that of traditional VC funds. We show that the annualized internal rate of return (IRR) on impact funds is 4.7 percentage

[1] GIIN annual impact investor survey 2017.

[2] https://thegiin.org/impact-investing/need-to-know/#what-is-impact-investing.

points (ppts) lower than traditional VC funds, after controlling for industry, vintage year, fund sequence, and geography.

Reduced-form estimations suggest investors may be willing to forego returns, but this evidence is not sufficient. Selection in observability of VC fund returns may affect this analysis, and, more fundamentally, ex-post performance estimations do not necessarily reveal ex-ante decisions to invest as a function of expected returns.

To investigate whether investors willingly forego expected return at the time of their investment decision, our primary empirics employ a discrete choice methodology using investors' observed choices of investments (yes/no decisions in a random utility framework) among a large set of VC funds fundraising in a year as the dependent variable. This approach builds on a large literature on hedonic pricing techniques, which provide tools for estimating implicit prices of attributes that a good possesses (e.g., Court, 1939; Griliches, 1961; Rosen, 1974; McFadden, 1974, 1986). Cameron and James (1987) introduce the idea that WTP can be estimated in discrete choices over alternatives. In discrete choice models, the choices made by agents over alternatives can be used to infer the sensitivity of the choice probability to price and other attributes (McFadden, 1974). Cameron and James (1987) note that if one reparameterizes the sensitivity of choice to an attribute by scaling it relative to the sensitivity of choice to price, the result is an estimate of the individual's WTP for that attribute.

A relevant example of the method is Huber and Train (2001), who study households that choose among a set of electricity providers. They are interested in the tradeoffs in price households make when choosing characteristics of the provider (e.g., local utility versus conglomerate), making inference as to people's WTP to do business with a more expensive local provider. Analogously, we study the choice of alternatives of funds and ask whether investors exhibit a WTP for the impact characteristic of a fund.

Our empirical analysis relies on two key independent variables: an impact fund dummy variable (the hedonic variable) and an ex-ante estimate of expected return for each fund (the price variable in a hedonic model), which we model using historic data on a fund's characteristics that investors would observe at the time of fundraising. From investors' choices, we find that both the ex-ante expected returns and the impact fund designation positively relate to the probability of investing in a fund. We estimate a logit model over the choice of funds fundraising in a given vintage, including investor fixed effects (i.e., a conditional logit model) or similar-investor dynamic groupings (to capture time-varying investor demand for the asset class). Our specifications include a rich array of fund and investor characteristics to model dimensions of portfolio choice preference. Measuring how sensitive the investment rate is to a fund's expected return allows us then to convert the desirability of impact into a WTP for impact via standard hedonic methodologies.

We address two main methodological concerns with respect to the estimation and inference in our VC setting. First, unlike traditional hedonic models, our price variable is an estimate—the forecast expected returns—and thus has measurement error. This likely induces overdispersion in expected return forecasts and attenuation bias in the expected return coefficient in the logit model. Since WTP has the expected return coefficient in the denominator, attenuation does not affect the sign of the estimated WTP but increases its magnitude. To address the magnitude issue, we apply a shrinkage estimator, which provides an asymptotic correction for the attenuation bias in the expected return coefficient.

Second, investors may have differential exposure or access to opportunity sets of funds to invest in, thus inducing them to have different expected return forecasts for the same fund. Mis-specifying this heterogeneity may induce a bias in the expected return coefficient, thus affecting the magnitude of our WTP estimates. Heterogeneities in expected returns is plausible in our private investment setting, but the exact mechanism is difficult to pin down with precision given the limitation in our knowledge of the actual expected return model or heuristic used by investors. As empiricists, we are agnostic as to whether parsimony versus specificity in the expected return model brings us closer to the true expected return used by each investor. Thus, we use both a parsimonious homogenous expected return model and a heterogenous expected return model to estimate expected returns and report WTP estimates based on both models to generate a range of plausible WTP estimates. Furthermore, we estimate the model under both rationed and expanded opportunity set assumptions and find that our impact coefficient and WTP estimates are consistently positive and stable. Overall, we report that the aggregate WTP for impact is between 2.5%–3.7% in expected IRR.

WTP for impact is not in equal magnitude across investor types. Five noteworthy investor groups exhibit a positive WTP for impact. (i) Development organizations have a high WTP for impact, presumably reflecting their direct impact mission. (ii) Foundations also have a small but positive WTP for impact in some specifications, again reflecting their mission orientation. (iii) Financial institutions—banks and insurance companies—have high WTPs, likely reflecting their incentives to invest in local communities either to comply with the Community Reinvestment Act (CRA) and/or to garner goodwill from the community or politicians/regulators. (iv) Public pension funds have a high WTP for impact, in line with the tendency for state pensions in the US to prefer investments within their home state (Hochberg and Rauh, 2013) to bring spillover economic benefits, nonpecuniary political benefits, and direct social objective benefits. (v) Investors in Europe, Latin America, and Africa have a higher WTP.

We then explore six investor attributes that might capture differential utility from investing in impact across investors; namely, whether the capital is (1) held by households (as opposed to an organization), (2) intermediated by an asset manager, (3) held by an organization with a mission objective, (4) held by an organization facing regulatory or political pressure to invest in impact, (5) held by an organization subject to laws restricting investments in impact, or (6) held by an organization (e.g., corporation) with charters that restrict investments in impact.

We find that mission focus (i.e., development organizations and foundations) is associated with a positive WTP of 3.4 to 6.2 ppts in expected excess IRR. This result is robust to including the Limited Partner (LP) geography fixed effects interacted with impact (i.e., within each geography, investor mission orientation is positively related to WTP for impact). Likewise, organizations expressing their mission by signing the UNPRI have a similarly higher WTP, especially after their signing. These UNPRI results are robust to including either LP geography fixed effects interacted with impact or LP type fixed effects interacted with impact.

Next, we find that political or regulatory pressure is associated with a positive WTP. In our most conservative models, WTP for impact associated with pressure is 2.3–3.3 ppts in expected excess IRR. Legal restrictions against investments for nonfinancial motives (e.g., the Employee Retirement Income Security Act (ERISA) and the Uniform Prudent Management of Institutional Funds Act (UPMIFA)) are associated with a lower WTP for impact. In contrast, we find no evidence that organizational charters that require a focus on financial returns (e.g., corporate charters that require shareholder wealth maximization) lower the WTP for impact. In addition to the LP geography-impact interaction, these estimates for attributes (4)-(6) are further robust to including the LP type fixed effects interacted with impact, thus exploiting international (e.g., US versus non-US) differences in laws for a given LP type governing attributes (4)-(6).

Finally, we provide evidence on whether investors' WTP varies across the different types of impact, though we characterize this evidence as preliminary given the small sample sizes in each type. Impact funds focused on environmental impact, poverty alleviation, and women or minorities generate the highest WTP estimates. In contrast, impact funds focused on small- and medium-sized enterprises (SMEs) and social infrastructure (e.g. health, education, and mainstream infrastructure) funds do not generate investment rates that reliably differ from those of traditional VC funds. These preliminary findings, which we hope provides fodder for future research, suggest that the internalization of utility from public good investing depends on how much the good is viewed as a public good versus an endeavor that could be profitable.

There is little prior academic work on impact investing by private investment vehicles. Kovner and Lerner (2015) study 28 community development venture capital funds in the US, finding that these funds tend to invest in companies at an earlier stage and in industries outside the VC mainstream and with fewer successful exits. Geczy et al. (2018) analyze contracts of impact funds and show that these contracts provide specific impact goals, indicating that investors intentionally seek impact when investing in these funds.

Our work relates to the broader literature on SRI that dates back as far as Milton Friedman's 1970 doctrine on responsible investing.[3] A survey by Renneboog et al. (2008) highlights the tension of SRI investing, concluding that investors in SRI funds may (but not with certainty) be willing to knowingly forego some expected financial returns for social or moral considerations. Consistent with the idea that investors in SRI funds value attributes other than performance, Benson and Humphrey (2008), Renneboog et al. (2011), and Bialkowski and Starks (2016) show that SRI fund flows are less sensitive to performance than non-SRI flows, while Bollen (2007) shows SRI funds have less volatile flows. Hartzmark and Sussman (2019) show Morningstar sustainability ratings introduced in 2016 resulted in large reallocations of capital toward funds with high sustainability ratings. Similarly, one strand of the SRI literature argues the nonpecuniary interests of investors affect the expected returns of investors; stocks preferred for nonfinancial reasons earn lower returns than spurned stocks. Building on this idea, Hong and Kacpercyzk (2009) find that stocks subject to widespread negative investment screens earn strong returns (also see Chava 2014). In other work, Dimson et al. (2015) provide evidence that investor engagement with the management of publicly traded firms on a collection of environmental, social, and governance issues is associated with positive abnormal returns. The above studies highlight the potential importance of nonpecuniary motives when investing, which dovetails with our analysis of the performance of impact funds and investors' WTP for impact.

Our paper also relates to a strand of the private equity literature that focuses on understanding demand. For example, Lerner et al. (2007) and Sensoy et al. (2014) compare returns earned by different types of LPs. Our findings complement those of Lerner et al. (2007), Hochberg et al. (2014), and Hochberg and Rauh (2013) in finding the importance of relationship and geography in understanding investment patterns in private equity.

2. Data and statistics

2.1. Data and impact funds designation

Our data on funds, investors, and performance come from Preqin's Investor Intelligence and Performance Analyst data sets. We initially search all private equity funds (which include buyout, balanced, and various types of funds of funds) for impact funds. However, the majority of impact funds we identify are venture or growth oriented. Impact buyout funds are a relatively recent phenomenon and were quite rare during much of our sample period. For example, Bain Capital raised its first "Double Impact Fund" only in 2017, and KKR did not set up its impact-investing unit until 2018. Thus, we restrict our study to one of VC and growth equity, which we loosely refer to as VC.

Our first task is to designate funds as being impact or traditional VC, using the criterion that an impact fund must state the dual objectives of generating a positive externality in addition to earning financial returns. To identify such funds, we proceed in the following three steps.

1. We form an impact potentials list, combining (i) text search of articles in Factiva using a list of impact-

[3] "The social responsibility of business is to increase its profits," The New York Times Magazine, September 13, 1970. Also see Geczy, Stambaugh, and Levin (2003).

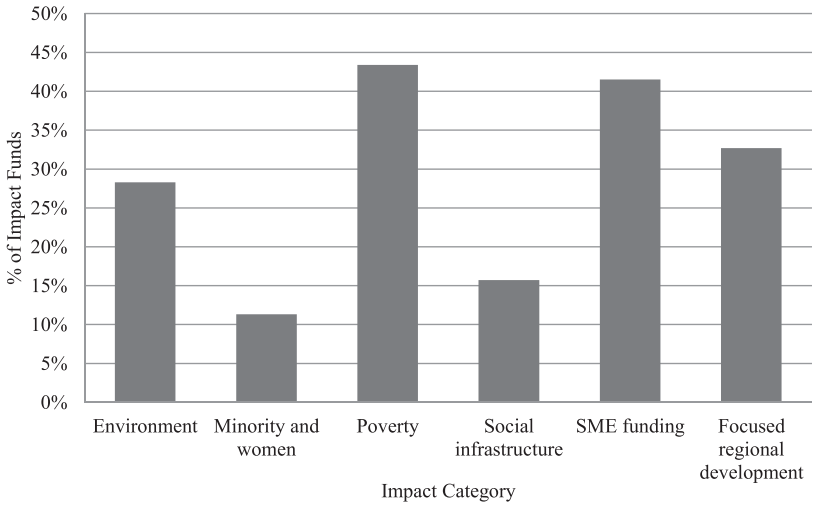

Fig. 1. Distribution of impact categories that impact funds target.
For the sample of 159 impact funds, we identify the impact categories targeted by each impact fund. The figure presents the percentage of sample funds that target each category. Funds can have multiple impact categories. The categories are as follows:
Environment - delivers positive environmental impact (e.g., agriculture, energy, water, and forestry).
Minorities and women - funds firms run by minorities or women.
Poverty - funds firms in impoverished areas.
Social infrastructure - develops infrastructure for societal benefit (e.g., microfinance, health care, schools, and housing).
SME funding - provides capital to SMEs and undercapitalized markets.
Focused regional development - imposes a material geographic constraint on investment.

related keywords[4] to describe funds; (ii) four third-party lists of impact funds and managers (ImpactBase (www.impactbase.org), Impact50 list in ImpactAssets (www.impactassets.org), ethos funds in Preqin (www.preqin.com), and MRI Manager Database in Cambridge Associates (www.cambridgeassociates.com)); and (iii) list of funds with majority geographic focus on countries with GDP per capita less than $1400.[5] Our potentials list consists of 323 VC and growth impact funds once we impose the further restriction that the fund appears in the Preqin Performance Analyst database.

2. We manually read descriptions and online resources about funds and fund families and screen out funds that do not explicitly aim to be double bottom line or state a dual objective. This results in the elimination of 146 funds from the sample (e.g., some large traditional global fund of funds like one managed by HarbourVest Partners that do not bear any resemblance to impact investing).
3. We further restrict the sample to funds with vintage years between 1995 and 2014, and investor information exists for at least one LP per fund in Preqin's Investor Intelligence database. This results in the elimination of 18 funds.

The screening process above results in our final sample of 159 impact funds. Note that we likely fail to designate some funds as impact (false negatives) due to a lack of detailed information, but our approach yields a clean sample of impact funds (i.e., false positives are unlikely).

Impact funds have diverse goals, so it is useful to consider specific examples of impact funds in our final sample. Bridges Ventures is a London-based family of funds "…dedicated to sustainable and impact investment…" that uses an "…impact-driven approach to create returns for both investors and society at-large." Bridges has several funds in our sample including the CarePlaces Fund, which builds care homes for the elderly. Its limited partners include university endowments, banks, pension funds, and high net worth investors. NGEN Partners is a Manhattan-based family of funds that "…invests in companies that positively improve the environment and human wellness" and manages three funds in our impact data set (NGEN Partners I and II and NextGen Enabling Technologies Fund). The North Texas Opportunity Fund "…seeks to invest in companies located in or willing to expand operations to underserved North Texas region markets, with a special emphasis on the southern sector of Dallas. The firm invests in minority or women owned or managed companies located anywhere in North Texas."

To parsimoniously categorize these diverse impact goals, we construct six impact categories: environmental impact, minority and women funding, poverty alleviation, social infrastructure development (e.g., health, education, and mainstream infrastructure), SME funding, and focused regional development (jobs creation and economic development funds in a specific region). For each impact fund, we read fund descriptions in three databases (Preqin, Capital IQ, and ThomsonOne), as well as in the fund's own marketing materials on their websites, and code the impact objectives of the fund using these six categories, allowing funds to have multiple objectives.

[4] See Table A1 for the list of keywords.
[5] See Table A2 for the list of low GDP per capita countries.

Fig. 1 depicts the percentage of the 159 impact funds for each of stated impact goals. The smallest impact categories are minority and women funding (11% of funds) and social infrastructure development, which includes health and education as well as other social or physical infrastructure (16%). The remaining impact categories are more common with the most prevalent being poverty alleviation (43%) and SME funding (42%), followed by focused regional development (33%) and environmental impact (28%).

We augment our Preqin data with the list of UNPRI signatories and signing years, which we obtain from UNPRI. As of November 16, 2015, there were 1422 signatories (297 asset owners, 931 investment managers, and 194 professional service managers) who collectively manage $59 trillion. We match UNPRI signatories to our data set using investor names. Investors that are subsidiaries of a UNPRI signatory are also coded as signatories but not investors that are parents of UNPRI signatory subsidiaries.

2.2. Fund statistics

Our analysis focuses on 4659 funds with vintage years from 1995 to 2014. Table 1 presents descriptive statistics for the 4500 traditional VC funds on the left and the 159 impact funds on the right.

Although traditional VC funds are larger than impact funds ($204.6 million versus $129.6 million at the mean and $102 million versus $83 million at the median), the mean commitment size does not differ by impact-versus-traditional VC. When we average across investors in a fund and then calculate the mean of this average across funds, we find that the mean commitment size for impact funds is $27.1 million, which is not significantly different from the mean commitment size of $22.2 million for traditional VC funds. One might wonder if the difference arises because we are more likely to observe commitment size for traditional funds and thus are more likely to observe smaller capital commitments. This does not appear to be the case, as we observe proportionately more investment amounts for impact investments (38.0%) than for traditional funds (32.6%). Note that the motivation for our decision to use discrete choice of investments rather than commitments in dollars as the outcome variable is transparent in Table 1; of the 23,986 investments, we observe that only 7867 (32.8%) have data on commitment size.

In terms of realized performance, traditional funds have a mean (median) IRR of 11.6% (7.4%), while impact funds have mean (median) IRR of 3.7% (6.4%). The same pattern emerges for value multiples (VMs) and imputed public market equivalents (PMEs). (Note that we do not observe fund cash flows for our sample funds; thus, the imputed PMEs were calculated using regression coefficients from Table IA.IV of Harris et al. (2014), which use the S&P 500 as a benchmark, and observed IRRs and VMs for our sample funds.) The imputed PMEs for impact funds suggest that they do not beat the public market, on average, while traditional VC funds do (albeit with significant time variation).

Our preferred measure of performance, percentile rank, is based on a fund's performance ranking (either IRR or VM, based on data availability) relative to cohort funds of the same vintage and geography (five regions). Percentile ranks adjust for the large temporal and geographic variation that trouble any inference using the other performance measures, which are notoriously difficult to risk adjust (Korteweg and Sorensen, 2010; Sorensen et al., 2014; Korteweg and Nagel, 2016). In particular, VC funds of vintages from the mid to late 1990s realized very right-skewed IRRs, when impact funds were relatively rare (only 14 of our sample funds have vintage years between 1995 and 1999).

Table 1 reports that traditional funds have a mean (median) percentile rank of 0.49 (0.50), while impact funds have a mean (median) rank of 0.34 (0.28). The difference of 0.15 (0.22) in percentile rank translates to a difference of about 3.0% (4.7%) in excess IRR centered at the median in historical returns. Appendix Table A3 provides the mapping of percentile ranks to excess IRRs, which are calculated as a fund IRR less the median IRR for the fund's vintage year and geography cohort. Although this mapping includes the strong VC return years of the late 1990s, the post-2000 sample yields estimates that are within 1.2 ppts of the full sample mapping for percentile rank differences 0.40 or less.

Table 1 reports a large difference in the standard deviation of IRRs for traditional funds versus impact funds (32.06% versus 15.17%). This difference is not statistically significant, but the magnitude of the difference triggers concern about risk differences. When we look into the source of these standard deviation differences, we find that the difference in return dispersion is again due to the lack of impact funds during the dot.com boom in the 1990s together with the highly right-skewed performance of traditional VCs during this period. Among traditional funds of late 1990s vintage, 22 earned IRRs greater than 100%. From 2000 onwards, the standard deviation of IRRs for traditional and impact funds are similar (16.8% versus 14.7%). Likewise, even in the full sample, the downside risk (measured as the standard deviation of funds with IRRs less than 50%) is statistically and economically the same across traditional (14.3%) and impact (15.2%) funds. Furthermore, our results are quantitatively similar if we restrict our sample to funds from 2000 onwards.

Another potential data concern is the observability of returns among VC funds. Preqin data are similar to other databases in return statistics (Harris et al., 2014), but the observability of returns may vary depending on a fund's impact status. In unreported tests, we use the fact that public pensions are often required to disclose their holdings and returns (Metrick and Yasuda, 2010) to ensure robustness of our results to a setting unlikely to be affected by selection in observability.

Panel B of Table 1 reports the geography of impact-versus-traditional funds. We collapse Preqin codes of the geographic focus of fund investments to eight regions and designate a fund to have a geographic focus if more than a third of all geographic descriptors are concentrated in a given region. Most funds (84%) focus on only one of the eight global regions and a small percentage have no geographic focus (3.5%). Impact funds tilt more toward

Table 1
Fund descriptive statistics, 1995 to 2014.
This table presents fund summary statistics for traditional funds (left columns) and impact funds (right columns). Capital commitment is the average capital commitment across investors within a fund. IRR is the final or last observed internal rate of return for the fund. VM is the fund's value multiple. Imputed PME (public market equivalent) is the fund's PME imputed using regression coefficients in Table IA.IV in Harris et al. (2014) and the fund's available IRR and VM. Percentile rank is the fund's percentile rank relative to similar cohort funds (year, region, and fund type). In Panel B, we present the geography focus of fund investments. In Panel C, we present the industry focus of fund investments. Funds can have multiple geography and industry focuses.

	Traditional VC funds				Impact funds			
	N	Mean	Median	Std. dev.	N	Mean	Median	Std. dev.
Panel A: Descriptive statistics								
Vintage year	4500	2005.4	2006.0	5.26	159	2006.7	2008.0	4.44
Fund size ($mil)	4000	204.6	102.0	300.2	147	129.6	83.00	147.3
Capital commitment ($mil)	2717	22.21	14.60	33.85	125	27.09	15.00	32.88
IRR (%)	1207	11.59	7.40	32.06	76	3.70	6.35	15.17
VM - value multiple	1484	1.51	1.22	1.94	91	1.17	1.10	0.56
Imputed PME	1147	1.29	1.09	1.29	65	1.00	0.97	0.42
Percentile rank	1530	0.49	0.50	0.30	94	0.34	0.28	0.30
Fund sequence number	4500	3.95	2.00	5.63	159	3.88	2.00	5.91
Panel B: Geography focus of fund investments								
North America	4500	0.50			159	0.33		
Developed Europe	4500	0.23			159	0.18		
Emerging Europe	4500	0.06			159	0.09		
Africa	4500	0.02			159	0.23		
Central and South America	4500	0.03			159	0.12		
Developed Asia-Pacific	4500	0.07			159	0.01		
Emerging Asia-Pacific	4500	0.17			159	0.14		
Middle East	4500	0.03			159	0.00		
All regions	4500	1.10			159	1.09		
Panel C: Industry Focus of Fund Investments								
Business services	4500	0.03			159	0.03		
Energy	4500	0.06			159	0.19		
Consumer discretionary	4500	0.05			159	0.03		
Diversified	4500	0.27			159	0.48		
Industrials	4500	0.04			159	0.06		
Information technology	4500	0.45			159	0.06		
Health care	4500	0.22			159	0.06		
Infrastructure	4500	0.01			159	0.05		
Food and agriculture	4500	0.01			159	0.04		
Materials	4500	0.01			159	0.04		
Real estate	4500	0.00			159	0.04		
Media and communications	4500	0.12			159	0.03		
All industries	4500	1.27			159	1.12		

developing countries including Africa, Latin America, and Emerging Europe than traditional funds.

Panel C of Table 1 reports the industry foci of impact-versus-traditional funds. We collapse the Preqin codes to 11 different industries (business services, energy, consumer, diversified, industrials, information technology, health care, infrastructure, food and agriculture, real estate, and media/communications) and code a fund as having an industry focus if more than a third of industry sector descriptors are concentrated in a given industry. Both self-described diversified funds and funds that lack any focus on particular industries (according to our coding method) are categorized as "diversified." Impact funds are more likely to be energy or diversified funds and are less likely to be IT, health care, or media and communication funds than traditional VC funds.

2.3. Investor (LP) statistics

We categorize investors into nine types by doing manual web searches for each investor in our sample. We refer to these groupings as LP types, reflecting the limited partner designation of investors in private equity. Our goal is to attribute the investing to the source of capital (rather than the intermediary). Thus, for asset managers, we search each manager to uncover whether the asset manager specializes in serving a particular constituent (e.g., public pensions).

Development organizations include multinational, national, and regional organizations that invest with development purposes in mind (e.g., International Finance Corporation, Ireland Strategic Investment Fund, and New Mexico State Investment Council). *Financial institutions* include

banks and insurance companies. (When we separately analyze banks and insurance companies, we obtain similar results for each group.) *Corporation & government portfolios* include corporations who invest in VC (e.g., Cisco and Siemens), state-owned corporations (e.g., China Steel and China Oceanwide Holdings), and sovereign wealth funds that are not development-oriented (e.g., Abu Dhabi Investment Authority).[6] *Wealth managers* include family offices (e.g., Merrion Family Trust) and advisers who serve retail or high net worth clients (e.g., BNY Mellon Wealth Management). *Private pensions* are primarily corporate pensions but also include multiemployer retirement funds (e.g., Carpenters' Pension Fund of Illinois).[7] *Foundations, Endowments*, and *Public pensions* are self-explanatory. Finally, *Institutional asset managers*, a residual category, include LPs that manage money for a diverse institutional client base (e.g., Adams Street Partners), where the capital appears to be primarily institutional capital with a mixture of constituents.

In Table 2, Panel A, we provide descriptive statistics on LPs. The smallest categories in terms of LP counts are endowments and wealth managers, but even these have close to 200 distinct LPs participating in the market. The total number of investments by LP type generally mirrors the patterns of LP numbers. The average LP makes 6.9 fund investments. The most active investors are public pensions (15.4 funds per investor), private pensions (8.9 funds), and development organizations (8.3 funds). The average LP has 4.3 years of experience as an LP, though this number is positively skewed. Public pensions, private pensions, and endowments are the most experienced LPs. Overall, 9% of LPs are UNPRI signatories. Institutional asset managers are the most likely to sign the UNPRI (17.9%), followed by wealth manager (14.4%), and public pensions (13.4%). Foundations, corporations, and endowments are extremely unlikely to be UNPRI signatories.

The last two rows of Panel A present statistics across the 23,986 investments made by the 3460 LPs. The penultimate row of Panel A, last column, reports that for 33.4% of investments, there is a prior investment relationship between the LP and fund family. The last row of Panel A, last column, reports the home bias rate, which is strikingly large with 75.8% of investments made into funds focusing on the home region of the LP headquarters.

In Table 2, Panel B, we present the regional distribution of LP headquarters. Focusing on all LPs (last column of Table 2), nearly half of all LPs are in North America, while another 29% are in Developed Europe. However, the regional distribution of LPs varies by LP type. For example, 82% of endowment LPs are in North America, while only 34% of financial institution LPs are in North America. Relative to other LPs, development organization LPs have greater presence in Emerging Europe, Africa, Central and South America, and Emerging Asia-Pacific.

[6] We sort sovereign wealth funds into development organization and government portfolios following Dyck and Morse (2011).

[7] There are 81 multiemployer pension funds, and the majority are union-backed. Our results by LP type and LP attributes are qualitatively similar if we group these multiemployer pension funds with public pensions.

3. Realized performance results

Our starting point, and the topic of this section, is reduced-form regressions of fund performance. An economic conjecture is that impact funds will earn below average returns because they impose a constraint (the generation of positive externalities) on the investment opportunity set, which hurts performance. Alternatively, it is possible that the market fails to fully price the opportunities in the sectors that impact funds target (e.g., natural resources, infrastructure development), thus resulting in above-market opportunities for impact funds (though this argument requires a friction in pricing). We consider both possibilities and test for performance differences between impact and traditional VC funds.

We analyze the realized (or last reported) performance of funds in our sample: internal rate of return (*IRR*), value multiple (*VM*), and the average percentile rank of a fund relative to its vintage year and region cohort (*Rank*). We include funds with vintage years 1995 through 2012 in this analysis and use last reported performance for funds with later vintage years that are not yet completely liquidated. We regress a fund's *IRR* on a key impact dummy variable (IMP_j) that equals one for impact funds and step in control variables (denoted by the matrix X) in estimating six variations of the following regression:

$$IRR_j = \alpha + \beta IMP_j + X\Gamma + \varepsilon_j \quad (1)$$

In model (1), we estimate a univariate regression with only the key impact dummy, which recovers the average difference in IRR between traditional VC funds and impact funds from Table 1. In model (2), we add controls for fund size, fund sequence number, and vintage year. In model (3), we add controls for fund industry and fund geography.

In the remaining model variations, we introduce time-varying controls for fund industry and geography. Ideally, we would like to include vintage-geography-industry fixed effects, but we lack degrees of freedom to do so (since some geographies and industries have few funds). As a compromise, we consider models with vintage-geography and static industry fixed effects (Model 4), vintage-industry and static geography fixed effects (Model 5), and fixed effects for 60 clusters of vintage year, industry, and geography (Model 6). In the last model, we cluster funds into six three-year vintage groups (1995–1997 to 2010–2012), two geographies (North America/Europe versus the rest of world), and five industry groups (information technology and business services, diversified and consumer discretionary, health care, media and communications, and other industries in Table 1, Panel C).

In each regression, we estimate robust standard errors clustered by vintage year and geography. The six regressions are also estimated using either a fund's VM as the dependent variable or a fund's percentile rank as the dependent variable.

Table 3 reports the coefficient estimates on the key impact dummy variable. We find that impact funds reliably underperform traditional VC funds. Focusing first on IRR results in columns (1) to (3) of Panel A, the univariate regression of column (1) reveals that impact funds underperform traditional VC funds by 7.89 ppts ($p<0.01$). When we

Table 2
Limited partner (LP) descriptive statistics.
For each of the LP types and all LPs, we present descriptive statistics by first averaging all observations for a unique LP and then calculating the mean (standard deviation) for each variable across N LPs. Funds per LP are the total number of unique fund investments by an LP. Vintage year is the average vintage year of fund investments. Years of experience is the number of years since the LPs' first fund commitment (measured at the time of each investment and averaged across all investments for a given LP). The% prior relationship is the percent of capital commitments where the LP and fund's general partner (GP) had a prior investment relationship. The% home bias is the percent of capital commitments by the LP type where the region of the LP and fund are the same (using the eight major global regions of Panel B). In Panel B, we present the regional distribution of LPs by LP type. For development organizations, we manually coded geographic foci of their missions and used them instead of the actual headquarters location. For example, the Inter-American Development Bank is headquartered in the US, but its mission is focused on South and Central America. Standard deviations are in parentheses.

	Dev. org.	Foundation	Financial institutions	Endowment	Corp. & gov't	Institutional	Wealth manager	Private pension	Public pension	Total
Panel A: LP descriptive statistics										
# of LPs	258	453	572	196	404	591	174	440	372	3460
% of total	7.5	13.1	16.5	5.7	11.7	17.1	5.0	12.7	10.8	100.0
# of capital commitments	2147	2770	2473	1287	1513	3541	635	3893	5727	23,986
% of total	9.0	11.5	10.3	5.4	6.3	14.8	2.6	16.2	23.9	100.0
Funds per LP	8.32	6.11	4.32	6.57	3.75	5.99	3.65	8.85	15.40	6.93
	(16.70)	(13.61)	(8.85)	(15.72)	(16.58)	(14.82)	(6.54)	(19.80)	(30.43)	(17.43)
Vintage year	2007.2	2005.8	2005.7	2004.8	2006.6	2005.7	2005.9	2004.7	2005.5	2005.7
	(3.79)	(3.66)	(4.33)	(4.15)	(5.09)	(4.48)	(4.36)	(4.09)	(3.67)	(4.28)
Years of experience	4.39	4.13	3.77	4.64	2.70	3.54	3.67	5.08	7.76	4.34
	(4.47)	(4.69)	(4.38)	(5.39)	(3.46)	(4.30)	(4.58)	(5.23)	(7.52)	(5.11)
% UNPRI signatories	5.4	2.2	11.0	1.5	1.0	17.9	14.4	8.4	13.4	9.0
% Prior relationship	23.7	41.7	22.7	38.8	23.3	25.0	24.6	38.3	41.1	33.4
% Home bias	59.4	78.2	82.4	82.0	72.1	61.5	68.5	78.3	84.5	75.8
Panel B: Regional Distribution of LPs by LP Type (%)										
North America	19	83	34	82	21	32	34	72	62	48
Developed Europe	28	15	36	16	28	40	39	20	29	29
Emerging Europe	5	0	2	0	2	1	2	1	0	1
Africa	5	0	3	1	1	2	1	1	2	2
Central and South America	6	0	1	1	2	1	0	3	2	2
Developed Asia-Pacific	8	1	10	0	20	9	20	2	3	8
Emerging Asia-Pacific	25	0	10	1	24	11	4	0	1	9
Middle East	4	1	4	0	3	5	2	1	1	3

Table 3
The performance of impact funds, vintage years 1995–2012.
Fund performance (Panel A, IRR; Panel B, VM; Panel C, percentile rank) is regressed on a dummy variable for impact funds and controls. Controls include vintage year, log of fund size, log of fund sequence number, fund geography, and fund industry. Models (1) to (3) step in controls without interactions using 5 geographies and 12 industries. Model (4) creates fund group dummy variables based on 6 three-year vintage groups (1995–97 through 2010–12) and 5 fund geographies in place of vintage year and geography FEs of Model (3). Model (5) creates fund group dummy variables based on 6 three-year vintage groups and 12 fund industries in place of vintage year and industry FEs of Model (3). Model (6) creates fund group dummy variables based on 6 three-year vintage groups, 5 fund industries, and North America/Europe v. other funds. The 5 fund industries include (1) information technology and business services, (2) diversified and consumer discretionary, (3) health care, (4) media and communications, and (5) others (energy, industrials, infrastructure, food and ag., materials, real estate). Models that include fund size in the regression lose observations of traditional VC funds with missing fund size. Robust standard errors (in brackets) are calculated by clustering on vintage years and fund geography. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: IRR						
Impact	−7.890***	−9.937***	−4.701**	−4.898**	−4.652*	−5.359**
	[2.548]	[2.386]	[2.282]	[2.440]	[2.555]	[2.520]
Observations	1283	1252	1252	1252	1252	1252
R-squared	0.004	0.146	0.166	0.288	0.19	0.274
Panel B: Value multiple						
Impact	−0.403***	−0.465***	−0.361***	−0.265*	−0.228*	−0.194*
	[0.124]	[0.107]	[0.137]	[0.141]	[0.122]	[0.103]
Observations	1456	1417	1417	1417	1417	1417
R-squared	0.002	0.117	0.125	0.184	0.122	0.204
Panel C: Percentile rank						
Impact	−0.149***	−0.158***	−0.089**	−0.093**	−0.083**	−0.078*
	[0.037]	[0.037]	[0.040]	[0.045]	[0.040]	[0.040]
Observations	1505	1465	1465	1465	1465	1465
R-squared	0.014	0.027	0.068	0.17	0.121	0.164
Controls for all panels in column						
Vintage year FE	NO	YES	YES	NO	NO	NO
Log(fund size)	NO	YES	YES	YES	YES	YES
Log(fund sequence)	NO	YES	YES	YES	YES	YES
Fund geography FE	NO	NO	YES	NO	YES	NO
Fund industry FE	NO	NO	YES	YES	NO	NO
Vintage group*Geography	NO	NO	NO	YES	NO	NO
Vintage group*Industry	NO	NO	NO	NO	YES	NO
Vintage grp.*Industry*Geography	NO	NO	NO	NO	NO	YES

add controls for fund size, sequence number, and vintage year in column (2), the performance spread grows to 9.94 ppts ($p<0.01$). Finally, in column (3) we add controls for fund geography and industry. While fund geography and industry explain some of the performance variation, the performance spread of 4.70 ppts remains reliably negative. Models (4) to (6) introduce finer controls on industry and geography as discussed above and yield results similar to those of model (3). Thus, industry and geography explain some of the variation in performance between impact and other funds. However, models with industry and geography controls may underestimate the differences between impact and traditional funds if we misclassify some impact funds as traditional funds and if impact funds are more prevalent in some industries and geographies.

The analysis of VMs (Panel B) and percentile ranks (Panel C) are qualitatively similar to the analysis of IRRs. VMs for impact funds are reliably less than those of traditional VC funds, ranging from 0.194 to 0.465 depending on model specification. Percentile ranks for impact fund are also reliably less than those of traditional VC funds, ranging from 7.8 to 15.8 ppts depending on model specification.

These performance results represent one contribution of our analysis, as we show impact funds underperform traditional VC funds. However, this fund-level analysis of realized returns is not immune from concerns of selection in observability of VC fund returns. More fundamentally, ex-post performance estimations do not necessarily reveal ex-ante decisions to invest as a function of expected returns. As motivated in the introduction, we are interested in an intentional WTP for impact if any and its variation across different investors. Thus, we now turn to a WTP model of this ex-ante choice, which builds on the hedonic pricing and resource choice literatures.

4. Willingness-to-pay methodology

This section presents the discrete choice hedonic model for estimating investors' WTP for impact funds, closely following Cameron and James (1987) and Huber and Train (2001). In the original applications of these models WTP might be, for example, a homebuyer's WTP for a porch, estimated from homebuyers' purchase choices among the houses for sale at a cross-section in time. Analogously, WTP in our context is the hedonic value of investing in impact, estimated from investors' choices of investments from investment opportunities available at a cross-section in time.[8] Our WTP model is different from the housing example in that the price variable in our setting

[8] VC fund structures only allow investments at fund formation. Thus, the choice of a VC investment is considered relative to other funds of the same vintage.

is an unobserved expected return, requiring an additional layer of estimation that we discuss in detail below.

4.1. Random utility model of willingness to pay

Consider investor i facing a binary choice of whether to invest in fund j. A random utility model of latent utility U_{ij}^{*} from such an investment is given by

$$U_{ij}^{*} = \beta \mathbb{E}\left[\mathrm{r}_j\right] + \Gamma_1' X_{1,j} + \Gamma_2' X_{2,ij} + \mu_i + \delta_i \mathrm{IMPACT}_j + \mathrm{e}_{ij} \tag{2}$$

The first terms $\{\mathbb{E}[\mathrm{r}_j],\ X_{1,j},\ X_{2,ij}\}$ and the four parameters $\{\beta,\ \Gamma_1,\ \Gamma_2,\ \mu_i\}$ govern the creation of utility from an investor's portfolio choice. $\mathbb{E}[\mathrm{r}_j]$ is the expected return for fund j. $X_{1,j}$ is a matrix of nonprice fund characteristics that enter the portfolio choice preference for the investment (e.g., geography, sector, fund size). $X_{2,ij}$ is a matrix of investor characteristics governing investor preferences, including the investors' recent intensity of investing in the asset class, the proximity to the investment, and the prior relationship with the VC firm. Beyond these investor-specific variables, investors may differentially value exposure to the asset class within the larger portfolio choice of all of their capital, which we cannot see. Thus, we allow for investor fixed effects, μ_i, as a heterogeneous baseline utility of investing in VC. Investors' VC investment portfolio size may grow or shrink over time, and their baseline utility may fluctuate. Therefore, we also introduce time-varying investment intensity bins, described more fully in the results Section 5.1.

The final term in Eq. (2), IMPACT_j, is a dummy variable equal to one if fund j is an impact fund (and zero otherwise). Investors may have specific utility for impact; therefore, we index δ_i by i. To make this investor heterogeneity operational, we cluster investors by investor types.[9]

Omitted from Eq. (2) are explicit risk variables that might enter into a standard portfolio choice decision. Differences in liquidity, which might generally carry return implications, are not relevant in our context since investing in the VC asset class—whether traditional or impact—involves liquidity lock-up and no trading. Differences in fund-specific risk might be a concern. We control for the portfolio choice variables $X_{1,j}$ to absorb risk differences as they relate to the industry sectors, geography, and size of funds chosen. Yet, residual risk differences could be correlated with a fund being IMPACT. If so, we would expect the ex-post standard deviation of impact VC performance to be different from traditional VC. As discussed in Section 2.2 in conjunction with Table 1, Panel A, we find no such difference except in the dot.com period when some traditional VC funds had outsized IRRs of 100% or more and drove up the skewness of return distribution for traditional funds. In contrast, downside risk is not statistically and economically different between traditional and impact funds. Finally, our results are robust to restricting our sample to 2000 onwards.

4.2. Logit specification and willingness to pay

Random utility U_{ij}^{*} is not directly observable to the econometrician, who instead only observes the investor's choice to invest or not. The observable, discretized investment decision U_{ij} corresponds to the latent utility U_{ij}^{*} as follows:

$$\begin{aligned} U_{ij} &= 1 \; if f \; U_{ij}^{*} > 0 \\ U_{ij} &= 0 \; if f \; U_{ij}^{*} \leq 0 \end{aligned} \tag{3}$$

Under the assumption that the error term ε_{ij} is distributed *iid* extreme value, this form of random utility can map to a logistic distribution with a mean 0 and variance $\pi^{2/3}$; thus, a logit estimation can uncover the parameters of Eq. (2):

$$\begin{aligned} Logit\left(invest_{ij}\right) = \beta \mathbb{E}\left[\mathrm{r}_j\right] + \Gamma_1' X_{1,j} + \Gamma_2' X_{2,ij} + \mu_i \\ + \delta_i \mathrm{IMPACT}_j + \mathrm{e}_{ij}. \end{aligned} \tag{4}$$

Following Cameron and James (1987) and Huber and Train (2001), the WTP for impact (*WTP*) for investor i is[10]

$$WTP_i = = \frac{\partial \mathbb{E}\left[\mathrm{r}_j\right]}{\left(\partial \mathrm{IMPACT}_j\right)_i} = \frac{\left(\partial U_{ij} / \partial \mathrm{IMPACT}_j\right)_i}{\partial U_{.j} / \partial \mathbb{E}\left[\mathrm{r}_j\right]} = \frac{\delta_i}{\beta}. \tag{5}$$

4.3. Expected returns formation

Estimation of a discrete choice hedonic model requires a price variable, which in our context is the expected return for each fund, $\mathbb{E}[\mathrm{r}_j]$. As motivated by the literature on the determinants of fund performance (Kaplan and Schoar, 2005; Sorensen, 2007), we estimate expected returns, with estimates denoted $\hat{\mathbb{E}}[\mathrm{r}_j]$, based on fund characteristics observable at the time of investment. We start with the assumption that all investors have the same model for predicting fund returns and later relax this assumption. We begin by considering an investor who is making decisions about VC investments offered in the market in a particular vintage year, say 1995, as an example. The investor forms return expectations based on the information set available for the VC asset class at that point in time. The information set consists of the average asset class return observed recently and a fund-specific skill adjustment. Skill in the VC asset class shows up to the econometrician as persistence in fund series performance. To estimate the strength of persistence and the average observed asset class return, 1995 investors would use data covering vintage funds 1983 to 1990 (because of the time lag in realizing returns in VC).[11] Denoting these 1983 to

[9] An alternative is to estimate a random effects logit (mixed logit) model of investor choice (Revelt and Train 1998); however, we found the computation to be prohibitive costly given the very extreme choice (1 fund chosen out of about 100) in VC selection. Also, given that many investors only invest in a few funds, the model was not precisely estimable.

[10] Technically, IMPACT is a discrete choice variable; thus the correct form is $WTP_i = \frac{U_i(.IMPACT_j=1) - U_i(.IMPACT_j=0)}{\partial u / \partial \mathbb{E}[\mathrm{r}_j]}$. The continuous time version is provided above for readability.

[11] The typical private equity fund invests in companies during years 1 to 5 of the fund's life and liquidates those investments after year 5. It is during this liquidation phase that the fund's performance becomes clear

1990 vintage years as being in set v, we assume the 1995 investors use the following simple linear model to gage parameters:

$$r_{j_v} = \alpha_0^{1995} + \alpha_1^{1995} r_{j_v}^{prior} + \varepsilon_{j_v}. \quad (6)$$

The return r_{j_v} of fund j in these look-back vintage years v is a function of the performance of the prior funds managed by the same VC firm ($r_{j_v}^{prior}$) and the overall asset class performance for funds in vintage pool v (the constant).

Using the coefficients from the estimation of Eq. (6), we apply them to funds that are raising capital in 1995 to forecast expected returns for any 1995 fund j as

$$\hat{\mathbb{E}}\left[r_{j\in 1995}\right] = \hat{\alpha}_0^{1995} + \hat{\alpha}_1^{1995} r_{j\in 1995}^{prior}. \quad (7)$$

We roll forward this process to the remaining vintage years, until we have an estimate of expected returns for each fund j with vintage years from 1995 to 2014.

By definition, these forecast expected returns have measurement error since we do not observe the actual expected returns.[12] In our context, this measurement error is a common problem of overdispersion in expected return forecasts, given by the simple relation:

$$\hat{\mathbb{E}}\left[\mathrm{r}_j\right] = \mathbb{E}\left[\mathrm{r}_j\right] + u, \quad (8)$$

where $\mathbb{E}[\mathrm{r}_j]$ is the true but unobservable expected return and u is measurement error that is uncorrelated with $\mathbb{E}[\mathrm{r}_j]$. The importance of this overdispersion comes when we turn to estimating *WTP* in the logit formation. Overdispersion in $\hat{\mathbb{E}}[\mathrm{r}_j]$ may cause attenuation bias in the $\hat{\beta}$ coefficient on $\hat{\mathbb{E}}[\mathrm{r}_j]$ when we estimate the logit Eq. (4), relative to the true β if we had the precise $\mathbb{E}[\mathrm{r}_j]$. Because $WTP_i = \frac{\delta_i}{\beta}$, attenuation in $\hat{\beta}$ implies an overestimate of *WTP*.

We take two steps to correct the bias. First, we seek to remedy a source of error, which is our inability to observe the soft information entering the assessment of skill. We augment Eq. (6) to include indicator variables as to whether the fund is missing prior fund performance information ($Miss_j^{prior}$), is a first-time fund ($First_j$), and/or is an impact fund ($IMPACT_j$). This augmented model, dropping the vintage subscripts to reduce equation clutter, is given by

$$r_j = a_0 + a_1 r_j^{prior} + a_2 Miss_j^{prior} + a_3 First_j + a_4 IMPACT_j + a_5\left(Miss_j^{prior} * IMPACT_j\right) + a_6\left(First_j * IMPACT_j\right) + \varepsilon_j. \quad (9)$$

In estimating these regressions, we use percentile ranks as the performance measure. Table 4 reports a summary of estimates from the 20 rolling regressions estimating Eq. (9), corresponding to fund expectations formed from 1995 to 2014. In Panel A, we summarize the coefficient estimates and associated t-statistics on the model's independent variables across the set of rolling estimations. Consistent with the literature, prior fund performance carries the vast majority of the explained variation. The only other reliable relationship is that first-time funds tend to have subpar performance.[13]

Second, we employ a correction for the logit estimates attenuation by applying a shrinkage procedure used in practice. Because overdispersion is a common issue in portfolio choice, investors knowingly shrink extreme forecasts toward a global mean, as in the seminal portfolio optimization models of Jorion (1986) and as applied in expected returns or cost of capital estimations in Fama and French (1997).

The shrinkage procedure begins with regressing realized fund return (r_j) on the estimated expected returns $\hat{\mathbb{E}}[\mathrm{r}_j]$:

$$r_j = \gamma_0 + \gamma_1 \hat{\mathbb{E}}\left[\mathrm{r}_j\right] + e_j. \quad (10)$$

Our estimates are $\gamma_0 = \underbrace{0.25}_{p<0.001}$ and $\gamma_1 = \underbrace{0.50}_{p<0.001}$. The γ estimates imply that our $\hat{\mathbb{E}}[\mathrm{r}_j]$ has some information about future returns ($\gamma_1 \neq 0$) but that $\hat{\mathbb{E}}[\mathrm{r}_j]$ is imprecise ($\gamma_1 \neq 1$ and $\gamma_0 \neq 0$). Then, following standard shrinkage procedure, we calculate the shrinkage estimate of expected returns $\hat{\mathbb{E}}_{shrink}[\mathrm{r}_j]$ as the prediction from Eq. (10):

$$\hat{\mathbb{E}}_{shrink}\left[\mathrm{r}_j\right] = 0.25 + 0.50\ \hat{\mathbb{E}}\left[\mathrm{r}_j\right]. \quad (11)$$

Whereas $\hat{\mathbb{E}}[\mathrm{r}_j]$ has a ranking range of 0.16 to 0.72 (on a natural percentile rankings range of 0 to 1), $\hat{\mathbb{E}}_{shrink}[\mathrm{r}_j]$ has a range of only 0.28 to 0.61, reflecting the shrinkage to address imprecision. We use $\hat{\mathbb{E}}_{shrink}[\mathrm{r}_j]$ to estimate Eq. (4).

Importantly, using the true realized returns to shrink the dispersion in estimated expected returns asymptotically eliminates attenuation bias in the logit *WTP* estimation when the following two key assumptions are added to a classic errors-in-variable analysis:

(i) The ex-post residual of realized fund returns relative to true expected returns is uncorrelated

to investors. Thus, a 1995 investor would have a good indication regarding the performance of funds with vintage years 1983 to 1990 because these funds would be 6 to 13 years old in 1995. In contrast, funds with vintage years 1991 to 1994 would still be in their investment phase with no or limited liquidations.

[12] Our forecast model uses fully realized fund percentile rank as dependent variables in Eq. (6), while as of 1995, some of the 1983-1990 vintage funds are yet fully realized (assuming ten-year fund life, 1983-1985 funds are at least ten years old, whereas 1986-1990 funds are still less than ten years old). This may introduce measurement errors to our model in one or more ways. For example, if investors apply interim-to-final rank transition in their true expected return formation using their soft information, our model approximates that with some measurement errors. Alternatively, investors may attempt to isolate the component of performance persistence that is due to skill (and investable) from spurious correlation due to contemporaneous exposures (Korteweg and Sorensen, 2017), in which case our naive model also produces forecast expected return with measurement errors. Furthermore, investors may be heterogeneous in their soft information possessed to form their expected returns. We present our expected return model that incorporates investor heterogeneity in the next section.

[13] We consider a number of robustness checks to ensure our results are not driven by the specific model that we use to predict expected returns. First, our results are robust to alternative specifications of the expected return model of Eq. (5). For example, we add additional lags of past fund performance, fund industry fixed effects, and fund geography fixed effects. These additional variables are not consistently related to fund performance nor does their inclusion materially affect the WTP estimates. Consequently, we opt for the more parsimonious model.

Table 4
Summary of expected return regression models.
In each of 20 forecast years, 1995 to 2014, we estimate a regression of realized fund performance (using percentile ranks) on fund attributes as described in the main text. For example, in the 1995 forecast year we estimate relations between fund attributes and performance using data on 1983–1990 vintage-year funds since the performance and attributes of these funds would be observed by an investor looking to invest in 1995. Panel A summarizes the distribution of the 20 coefficient estimates and associated *t*-statistics across the 20 regressions; Panel B summarizes the number of observations and *R*-squareds across regressions. The interaction terms are only estimated for the last 12 of the 20-year rolling window regressions because there are a small number of impact funds in the early part of the sample.

	Mean [*t*-stat.]	% of |*t*-stats| > 1.96	% of coef. > 0	25th percentile [*t*-stat]	Median [*t*-stat.]	75th percentile [*t*-stat]
Panel A: Coefficients						
$R_j{}^{prior}$	0.217	85.0%	100.0%	0.167	0.222	0.280
	[3.12]			[2.93]	[3.23]	[3.95]
$Miss_j{}^{prior}$	-0.052	15.0%	20.0%	-0.079	-0.051	-0.014
	[-0.94]			[-1.49]	[-1.00]	[-0.42]
$First_j{}^{prior}$	-0.076	45.0%	10.0%	-0.146	-0.054	-0.019
	[-1.97]			[-3.63]	[-1.67]	[-0.64]
$Impact_j$	-0.203	0.0%	0.0%	-0.292	-0.186	-0.139
	[-1.15]			[-1.35]	[-1.08]	[-0.88]
$Impact_j * Miss_j{}^{prior}$	-0.033	0.0%	33.3%	-0.160	-0.069	0.067
	[-0.38]			[-0.92]	[-0.31]	[0.24]
$Impact_j * First_j{}^{prior}$	0.068	0.0%	75.0%	0.004	0.062	0.147
	[0.31]			[0.04]	[0.31]	[0.79]
Panel B: Regression statistics						
Observations	459.4	n.a	n.a	240.0	446.0	649.0
R-squared	6.50%	n.a.	n.a.	3.00%	4.50%	10.20%

with $\mathbb{E}[r_j]$. Specifically, $\text{cov}(\mathbb{E}[r_j], \xi) = 0$, where this residual is given by

$$r_j = \mathbb{E}\left[r_j\right] + \xi. \tag{12}$$

This condition assumes that the unexpected part of realized returns is not systematically higher [or lower] for funds with high expected returns and is the implicit assumption made in the asset pricing literature, which uses realized returns to understand temporal and cross-sectional variation in unobserved expected returns.

In venture capital, returns are highly nonlinear and positively skewed. To deal with this issue, our empirical analysis uses percentile ranks relative to cohort funds, which are bounded between zero and one. In theory, these bounds could generate $\text{cov}(\mathbb{E}[r_j], \xi) < 0$ since an expected percentile rank near one (zero) will have negative (positive) estimation error because of the boundary. In practice, we do not believe these boundary conditions are binding on errors since true expectations of percentile ranks would not approach the boundary; stated differently, investors do not expect to be able to pick the top and bottom performing funds from a cohort. The actual range of percentile ranks that we estimate after shrinkage is 0.28 and 0.61, with a standard deviation of 0.035. This range is a reasonable estimate of the range of ex-ante expected percentile ranks, and they are far from the boundary conditions.

(ii) Measurement error u in forecast expected returns, $\hat{\mathbb{E}}[r_j] = \mathbb{E}[r_j] + u$ from Eq. (8), is uncorrelated with the residual in the realized return relative to the true expected return, $\text{cov}(u, \xi) = 0$. We can think of no reason why errors in our forecast returns would be correlated with errors in realized returns.

In Appendix B, we show that the bias in our WTP estimates is positively related to the two covariance terms and is positive if $(cov(\mathbb{E}[r_j], \xi) + cov(u, \xi) > 0)$. To simplify the analysis below, we also assume the measurement error u in forecast expected returns, $\hat{\mathbb{E}}[r_j] = \mathbb{E}[r_j] + u$ from Eq. (8), is uncorrelated to the true expected returns, $\text{cov}(\mathbb{E}[r_j], u) = 0$; this assumption only affects the absolute magnitude of the attenuation bias if the two assumptions above are violated.

To show how shrinkage with realized returns removes attenuation under (*i*) and (*ii*) from above, imagine a simple model of an outcome variable y (investment in a private equity fund in our case) such that $y = a + b\mathbb{E}[r_j] + e$, where the econometrician must estimate with $\hat{\mathbb{E}}[r_j]$ instead of $\mathbb{E}[r_j]$. The standard estimate of the slope coefficient in the

classic errors-in-variable analysis is

$$\text{plim}\left(\hat{b}\right) = \frac{b\sigma^2_{\mathbb{E}[r_j]}}{\sigma^2_{\mathbb{E}[r_j]} + \sigma^2_u} = \lambda b, \tag{13}$$

where $\lambda < 1$ is the attenuation bias. Note that the slope parameter (γ_1) of the shrinkage regression of Eq. (10) yields an estimate of this attenuation bias:

$$\gamma_1 = \frac{cov\left(\widehat{\mathbb{E}}[r_j],\ r_j\right)}{\sigma^2_{\widehat{\mathbb{E}}[r_j]}} = \frac{cov\left(\mathbb{E}[r_j] + u,\ \mathbb{E}[r_j] + \xi\right)}{\sigma^2_{\mathbb{E}[r_j]} + \sigma^2_u}$$
$$= \frac{\sigma^2_{\mathbb{E}[r_j]}}{\sigma^2_{\mathbb{E}[r_j]} + \sigma^2_u} = \lambda. \tag{14}$$

Thus, the shrinkage regression provides a valid correction for the attenuation bias.

4.4. Heterogeneity in expected returns forecast and use

The logit model described in Eq. (4) assumes that investors are homogenous in their forecast of expected returns. In practice, investors may exhibit heterogeneity in their forecast of expected returns because of different forecast mechanisms or because of interest in only a subset of funds (Hochberg et al., 2014; Cavagnaro et al., 2019). The dimension of concern to us is bias at the LP investor type level. Thus, we estimate $\hat{\mathbb{E}}_{type,shrink}[r_j]$ uniquely for each investor type as a robustness check on our results. To implement these expected return forecasts, we limit the set of funds in which an investor has an interest to those funds with investment by investors of the same type (e.g., financial institutions, development organizations, foundations, etc.).[14] Estimating by LP type also allows us to incorporate prior relationships as part of the forecast. Different investor types have different propensities to have a prior relationship with a fund in question, which will provide heterogeneity in the use of soft information.[15] These heterogeneous expected return estimates are correlated with the homogeneous expected returns (with a correlation coefficient of 0.76) but, as anticipated, have more variation (a standard deviation of 0.042 versus 0.035).

Having estimated LP type-specific expected returns for each fund, we then estimate

$$Logit\left(Invest_i\right) = \beta\ \hat{\mathbb{E}}_{type,shrink}[r_j] + \Gamma'_1[X_{1,j}] + \Gamma'_2 X_{2,ij} + \mu_i + \ \delta_{type}\text{IMPACT}_j + \varepsilon_{ij}. \tag{15}$$

Note that we are not after the best or optimal model for predicting future returns—we are after the actual model investors use to form their expectations. On the one hand, some investors might rely on hard information about all funds in the market and use this broad information set in forming expected returns. On the other hand, some investors might examine a narrow set of funds that they are more familiar with (or have access to) and use soft information in forming expected returns. We as econometricians do not know which model is closer to the true model that investors use. We are agnostic about which estimates are superior representations of investor behavior and report estimates from both the homogenous and heterogenous expected return models to generate a range of reasonable WTP estimates.

As a final robustness check in our analysis of WTP patterns across LP types, we allow for the possibility that investors are heterogeneous in their use of expected returns because they face different portfolio choice considerations. The model described in Eq. (15) assumes that investors are homogenous in their use of forecast expected returns, though forecasts vary across LP types. Yet, investors may exhibit heterogeneity in their portfolio diversification model (e.g., preferring investments in a particular industry or geography). Thus, in our analysis of WTP across LP types, we allow heterogeneous expected returns forecasts to interact with the industry and geography.

We amend Eq. (15) to allow for these heterogeneities by investor type:

$$Logit\left(Invest_i\right) = \beta\ \hat{\mathbb{E}}_{type,shrink}[r_j] + B'_1\left[X_{1,j} \cdot \hat{\mathbb{E}}_{type,shrink}[r_j]'\right] + \Gamma'_1[X_{1,j}] + \Gamma'_2 X_{2,ij} + \mu_i + \delta_{type}\text{IMPACT}_j + \varepsilon_{ij}. \tag{16}$$

The resulting WTP for impact embeds a richer investor type-level application of the role of returns in the portfolio choice model while maintaining a baseline property of hedonic discrete choice models that heterogeneities in the magnitude of the coefficient on the hedonic variable (impact designation in our case) be calculated relative to common coefficients on the price variable (representing the change in utility per unit of price) that are fixed across the choice agents. This WTP calculation is

$$WTP_{type} = \hat{\delta}_{type} / \left(\hat{\beta} + \hat{B}'_1\left[\bar{X}_{1,j}\right]_{type}\right), \tag{17}$$

where $[\bar{X}_{1,j}]_{type}$ is the average of exposures to the industry, geography, and discretized fund size buckets by investor type.

5. Willingness to pay results

5.1. Aggregate WTP results

Table 5 reports coefficient estimates from the logit model of investment choice, akin to those used in the discrete choice implementations in venture capital (Ljungqvist et al., 2006; Bottazzi et al., 2016). Panel A presents the results using homogenous expected returns (corresponding to Eq. (4)); Panel B uses heterogenous expected returns (Eq. (15) but with a single impact coefficient). The dependent variable is an investment indicator variable. The set of observations are all potential investments into the funds that close in a given year by all of

[14] The forecast $\hat{\mathbb{E}}_{type,shrink}[r_j]$ will only be defined over a subset of funds. Because we are interested in how our estimate of WTP varies when allowing $\hat{\mathbb{E}}_{shrink}[r_j]$ to be instead unique to an investor type, $\hat{\mathbb{E}}_{type,shrink}[r_j]$, we want to reproduce the full data set of funds as the opportunity set for investment. Thus, we fill in $\hat{\mathbb{E}}_{type,shrink}[r_j] = \hat{\mathbb{E}}_{shrink}[r_j]$ for the funds j not estimated by a particular investor type.

[15] For each fund in a given LP type sample, we calculate the average of the relationship dummy variable across all LPs that invested in that fund and add this relationship propensity as an additional independent variable to Eq. (9) when estimating LP type-specific expected returns.

Table 5
The willingness to pay for impact.
The dependent variable is a dummy variable that equals one if an LP invests in a fund. Observations are determined by crossing all vintage year funds with LPs that make an investment in that year. All columns except column (2) are a logit model with LP investment group controls. LPs are dynamically placed in one of 368 groups according to how many prior three-year investments they make in VC by LP type. Column (2) is a conditional logit model (LP fixed effects). Columns (3) and (4) drop ex-ante top-quartile VC funds and top 15 VCs, respectively, investment opportunities for LPs that have no prior relation with the VC fund families. Column (5) creates an opportunity set assuming that a GP was fundraising in year t (and thus is included in the fund opportunity set for LPs investing in year t) if it closed a fund in year $t+1$ and its predecessor fund was raised in t-5 or older. Column (5) creates an opportunity set assuming that an LP considers investments in year t but does not realize investments until $t+1$. *Impact* equals one for impact funds. *Expected returns* are expressed as percentile ranks relative to vintage year cohort funds and are modeled based on known fund characteristics at the time of investment and are adjusted for shrinkage. In Panel A, a fund's expected return forecast is homogenous across all investors. Panel B allows heterogeneous forecast for each fund by LP type. The WTP estimate is the ratio of the *Impact* coefficient divided by the *Expected returns* coefficient. Standard controls included in all columns are LP experience (log of years since first fund investment plus one), LP-GP relationship (we analyze five regions rather than eight by combining Emerging Europe, Africa, and Central and South America into "Rest of the World" and Emerging Asia-Pacific and Middle East into "Emerging Asia-Pacific"; however, to establish an LP-fund geography match, we continue to employ the eight-region code first and then combine the eight home bias dummies into five), fund-LP geography match (five dummy variables for five regions that equal one if the fund and LP are in the same region), expected fund size, and fixed effects for fund geography (five regions), industry (12 industries), and vintage year. Standard errors in brackets are clustered at the LP level, except for the conditional logit. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Homogeneous expected returns forecast						
Expected returns	3.354***	3.426***	3.248***	2.833***	3.146***	3.307***
	[0.276]	[0.210]	[0.363]	[0.354]	[0.270]	[0.275]
Impact	0.591***	0.585***	0.599***	0.567***	0.590***	0.580***
	[0.0599]	[0.0443]	[0.0645]	[0.0643]	[0.0599]	[0.0595]
WTP estimate	0.176	0.171	0.184	0.200	0.188	0.175
Pseudo *R*-squared	0.261	0.237	0.264	0.269	0.258	0.263
Observations	3047,430	3047,430	2780,390	2944,643	3301,101	3873,720
Panel B: Heterogeneous expected returns forecast						
Expected Returns	4.655***	4.725***	5.072***	5.022***	4.622***	4.655***
	[0.225]	[0.140]	[0.253]	[0.262]	[0.227]	[0.225]
Impact	0.613***	0.602***	0.650***	0.645***	0.618***	0.613***
	[0.0577]	[0.0422]	[0.0589]	[0.0588]	[0.0579]	[0.0577]
WTP Estimate	0.132	0.127	0.128	0.128	0.134	0.132
Pseudo *R*-squared	0.263	0.240	0.267	0.272	0.259	0.263
Observations	3047,430	3047,430	2780,390	2704,939	3159,087	3047,430
Model:						
Logit with eynamic LP invest. groups	Yes	–	Yes	Yes	Yes	Yes
Conditional logit model	–	Yes	–	–	–	–
# F.E. (LP or dynamic LP groups)	368	3460	368	368	368	368
Sample restrictions:						
Drop top quartile unless prior relation	–	–	Yes	–	–	–
Drop top 15 VCs unless prior relation	–	–	–	Yes	–	–
Expanded fundraising years	–	–	–	–	Yes	–
Expanded LP investor set	–	–	–	–	–	Yes

the active LPs with at least one fund investment in that vintage year. This crossing of all LPs and all VC funds active in each vintage year yields over three million fund-LP observations, pooled across years. Our main independent variables of interest are *Expected returns* (forecasted and shrunk per the methodology section) and *Impact*. To prevent the impact coefficient from picking up LPs' portfolio choice demand for particular investment characteristics, we include fixed effects for fund vintage, geography, and industry. We also include two variables capturing paired characteristics between the investor and the fund. First, following Hochberg and Rauh (2013), we include a home bias variable, defined as whether fund j focuses its investments on the home region of investor i, where we consider eight major regions globally. Second, because the prior relationship between an investor and a particular VC fund manager matters (Lerner et al., 2007; Hochberg et al., 2014), we include an indicator variable for a prior investment relationship between investor i and any prior fund managed by fund j's fund manager. We measure expected fund size as the three-year prior average of the median fund size in the vintage and market (US or non-US).

We first show our aggregate WTP result with two models of investor heterogeneity in terms of their preferences for the VC asset class. In column (1), we absorb investment rate heterogeneities with 368 dynamic (i.e., time-varying) buckets of LP type crossed with the discrete number of prior investments in the previous three years. Each investor group consists of investors of same LP type (e.g., development organization, foundation, pension, etc.) and the same average number of investments per year made in the prior three years. In column (2), the model is conditional logit at the individual LP investor level. The conditional logit levels LPs according to their average likelihood of investing in a VC fund. We prefer the dynamic LP investment groups of column (1). The benefit is in allowing

for dynamic appetite for the VC asset class since we cluster together, for example, all foundations have five total investments in the prior three years or public pension funds with 20 investments in the VC asset class in the prior three years, etc.[16]

In Panel A, column (1), the coefficient on impact is 0.591, and the coefficient on expected returns is 3.354 ($p<0.01$ for both coefficients). The WTP estimate is reported as the ratio of these estimates. We find that investors are willing to pay 18 percentile ranks ($0.18 = 0.591/3.354$) for impact, where a percentile rank runs from 0 (0th percentile) to 100 (100th percentile). In column (2), estimates from the conditional logit model imply a similar WTP of 17 percentile ranks. A WTP of 18 percentile implies that the average investor is indifferent (obtains identical utility) between investing in an impact fund at the 41 percentile rank of its vintage-geography cohort and investing in a traditional VC fund at the 59 percentile rank. In terms of the expected excess IRR of the fund, this suggests that investors are willing to give up 3.7 ppts in expected excess IRR to invest in an impact fund (see Appendix Table A3 for the mapping of percentiles to excess IRRs). This 3.7 ppts is 11% of a cross-sectional standard deviation of IRRs (0.32).

In Panel B, column (1), the estimated impact coefficient and expected return coefficient are 0.613 and 4.655, which yields a $\text{WTP} = 0.13 = 0.613/4.655$. A WTP of 13 percentile rank suggests that investors are willing to give up 2.5 ppts in expected excess IRR, which is lower than the estimate in Panel A. In Panel B, column (2), conditional logit estimates yield a similar WTP estimate. Overall, our WTP framework suggests that an investor WTP lies between 2.5% to 3.7% in IRR. These WTP estimates are smaller but within one standard error of the performance shortfall that we estimate in the reduced-form regressions of Table 3. Alternatively, a WTP of 13–18 percentile rank suggests that investors are willing to give up 0.13–0.17 in excess PME (see Table A3).

In columns (1) and (2), we assume that the investment opportunity set for all LPs in a given year is the set of funds completing fundraising that year. In practice, opportunity sets may be either more restrictive or more expansive.

One story is that some LPs are more likely to invest in high-performing funds, either due to LP skill differential or to assortative matching between elite VC firms and elite LPs (Cavagnaro et al., 2019; Lerner et al., 2019). Sensoy et al. (2014) show that access disparities between LP types (e.g., endowments versus pensions), as well as returns to such access disparities, largely dissipated in the 1999–2006 period. Yet the possibility remains that some specific LPs continue to enjoy exclusive access to top-performing VC firms that is denied to the rest and that this in turn makes investors that invest in (less exclusive) impact funds appear to accept lower financial returns.

Another possibility is that fundraising campaigns may last longer than a year for some funds. In such cases, funds with vintage year y were effectively fundraising in year y-1, and LPs that were in the market in year y-1 had the opportunity to invest in that fund. Yet another possibility is that for some LPs the fund selection/due diligence may take more than a year. In those cases, LPs that did not pull the trigger until year $y+1$ were effectively looking to invest in year y and had the opportunity to invest in vintage y funds.

Columns (3)-(6) report results of tests where LPs are designed to have either more restricted or expanded opportunity sets. In column (3) and (4), we present estimates using the same empirical model as column (1), except we restrict the possibility of investing in funds managed by elite VC firms only to a subset of LPs that have already invested in the VC firms' previous funds. The two columns differ in the way we define an elite VC firm. In column (3), elite VC firms are those with at least one fund with top-quartile performance (ranked against its vintage cohorts) among the three previous funds. In column (4), we define elite VCs as the 15 VC firms that are chosen as "Best VCs" in Metrick and Yasuda (2010). We find that our impact coefficient and WTP estimates are quite robust to these rationed opportunities sets, with only slight variation across columns.

In columns (5) and (6), we turn to considering an expansion of opportunity sets rather than to a rationing of fund access. In column (5), we assume that the fund raised in year y was also fundraising in year y-1 if more than five years had lapsed between the vintage years of the current fund and the previous fund. Operationally, we treat these funds as being in the market in both y-1 and y. In column (6), we assume that LPs that made no investments in year y but invested in $y+1$ were in fact already looking to invest in year y but passed. Alternatively, some of those LPs with investments in year $y+1$ actually made the commitments into the funds in year y, but the funds did not close until $y+1$. Either way, operationally we treat these LPs as being in the market in both y and $y+1$. In both columns (5) and (6), the sample size expands because either the set of funds (column (5)) or the set of LPs (column (6)) in a given year is greater than in our baseline sample. Again we find that our aggregate WTP estimates are very robust to the expanded opportunity sets, as the results in columns (5) and (6) are very similar to those in column (1).

Before proceeding to the analysis of WTP across different LP types, we estimate the WTP across five geographic regions using the model of column (1) but interacting impact with five regions. We summarize the WTP by region in Fig. 2. Circumstantial evidence suggests that demand for impact should be higher for investors domiciled in Europe. In their 2014 report, the Global Sustainable Investment Alliance (GSIA) reported that 59% of total managed assets in Europe are in SRI strategies compared to only 18% of assets in the US, 17% of assets in Australia, and 1% of assets in Asia. This suggests that Europeans value positive externalities more than others.[17] Our results strongly confirm the

[16] In earlier drafts of the paper, we estimated linear probability models and obtained similar results.

[17] See Liang and Renneborg (2016) and Dyck et al. (2019) for related evidence.

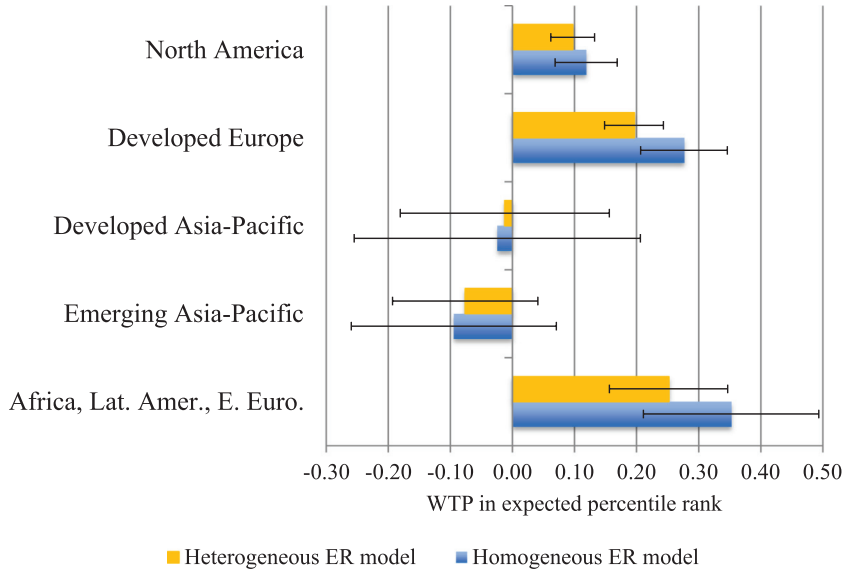

Fig. 2. Willingness to pay (WTP) for impact by geography.
The figure presents estimates of the willingness to pay for impact based on the logit model including the impact coefficient and the impact coefficient interacted with geography using either homogeneous or heterogeneous expected return models. (Geography is not mutually exclusive.) The WTP is the sum of the impact coefficient plus the impact*geography coefficient, all divided by the expected returns coefficient expressed in percentile ranks. Percentiles are based on performance relative to cohort funds. Cohorts are defined by fund vintage year and region. Black bands represent 95% confidence intervals on WTP estimates.

circumstantial evidence. North Americans have a positive and significant WTP for impact, but it is smaller than the baseline estimate (10–12 percentile ranks or 1.6–2.2 ppts in expected excess IRR). In contrast, investors from Developed Europe and from Africa, Latin America, and Eastern Europe have much higher WTP of 20–28 percentile ranks and 25–35 percentile ranks, respectively, corresponding to an expected excess IRR WTP of 4.2 to 8.7 ppts.

5.2. WTP by LP type

In this section, we estimate variation in WTP across LP types. The WTP estimate of 13–18 percentile ranks as reported in Table 5, column (1) is an average effect among all investors in our sample. It does not imply that all investors exhibit the same WTP. In practice, investors are likely to be heterogeneous in their taste for impact with some investors valuing the attribute more than others for social, institutional, legal, or regulatory reasons.

Table 6 presents the results. In all columns we allow investors' taste for impact to vary across nine LP types and five geographic regions. In column (1), we estimate the logit model using homogenous expected returns corresponding to Eq. (4). In column (2), we use heterogenous expected returns corresponding to Eq. (15).

Fig. 3 summarizes the WTP results by LP type. We find that development organizations, financial institutions, and public pensions have large positive WTP for impact with estimates ranging from 13 to 27 percentile ranks (2.5–6.2 ppts in excess IRR). In contrast, endowments, corporations, institutional managers, wealth managers, and private pensions have negligible WTP for impact, as their impact coefficients in the logit model are not significantly different from zero. Foundations have a small positive and statistically significant WTP (6 percentile ranks) in column (2). Although the WTP magnitudes fluctuate across models, the patterns across LP types are very robust.

In addition to testing the null that the individual WTPs are equal to zero, we also test the null hypothesis that LP types have equal WTP; we can easily reject the null hypothesis that the WTP is equal across LP types ($p<0.001$). In pairwise tests of the null hypothesis of equal WTP, we cannot reject the null in pairwise tests for *Development Organizations, Financial Institutions, and Public Pensions*. These are the same LP types that exhibit a robustly positive WTP for impact throughout our analysis. We always reject the null that these three LP types' WTPs are equal pairwise to WTPs of the other six LP types.

One explanation for our results might be that investors look as if they are willing to pay for impact, but in reality they erroneously expect returns on impact funds to be comparable to those earned on other VC funds. Because this story would be applicable to investors new to impact investing but not for investors repeatedly choosing impact VC funds, we can test this prediction. In untabulated results, we reestimate the specification of Table 5 column (1) modified to include the impact dummy interacted with an indicator for an LP having prior impact investing experience. We find that both the impact indicator and the interaction term are positive and significant, with the interaction term coefficient being twice as large as the impact coefficient. Put simply, investors with prior investments in impact funds are much more likely to invest in impact. This result combined with the PME result that impact funds, on average, do not beat the market ex post (Table 1) suggests that our main results

Table 6
Willingness to pay for impact by investor type.
Presented are coefficients and willingness-to-pay estimates from investment choice logit models. The columns vary in their estimation or use of the expected return forecast. Column (1) implements a homogenous model to forecast expected returns, where we estimate a single estimate of the forecast expected returns by fund using all funds in the dataset. Columns (2) and (3) use LP-type specific expected return forecasts but forecast with a smaller set of information (only funds invested by the investor type). Column (3) uses the same forecast as column (2) and also interacts these forecasts with fund characteristics (industry, geography, and size). Column (4) drops the expected return forecast variable altogether. Note that WTP is reported only for columns (1)-(3). Standard errors clusted at the LP level are in brackets. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

	(1)		(2)		(3)		(4)
Expected returns forecast:	Homogenous ER forecast		Heterogenous ER forecast by LP type				No forecast
Reported from logit:	Estimates	WTP	Estimates	WTP	Estimates	WTP	Estimates
Expected return	3.364***		4.591***		5.568***	Note: Not comparable to columns (1) and (2).	
	[0.275]		[0.223]		[1.584]		
Impact estimates by LP type:							
Development org.	0.906***	0.27***	0.738***	0.16***	0.980***	0.14***	0.595***
	[0.180]		[0.183]		[0.155]		[0.180]
Foundation	0.267	–	0.299*	0.06*	0.469***	0.07***	0.00261
	[0.179]		[0.179]		[0.178]		[0.178]
Financial institution	0.765***	0.23***	0.710***	0.15***	0.852***	0.13***	0.483***
	[0.144]		[0.144]		[0.122]		[0.141]
Endowment	−0.518	–	−0.443	–	−0.300	–	−0.802**
	[0.346]		[0.346]		[0.360]		[0.343]
Corporation	−0.0188	–	0.0655	–	0.238	–	−0.316
	[0.233]		[0.224]		[0.194]		[0.232]
Institutional	0.0872	–	0.233	–	0.501***	0.08***	−0.187
	[0.182]		[0.182]		[0.157]		[0.181]
Wealth manager	0.121	–	0.23	–	0.449	–	−0.142
	[0.329]		[0.332]		[0.335]		[0.325]
Private pension	−0.153	–	−0.0746	–	0.0834	–	−0.440***
	[0.168]		[0.168]		[0.174]		[0.165]
Public pension	0.730***	0.22***	0.832***	0.18***	1.028***	0.16***	0.430***
	[0.121]		[0.119]		[0.107]		[0.121]
Region*Impact F.E.	YES		YES		YES		YES
ER interacts with portfolio choice variables	NO		NO		YES		NO
Standard controls	YES		YES		YES		YES
Pseudo *R*-squared	0.261		0.264		0.276		0.260
Observations	3047,430		3047,430		3047,430		3047,430

are capturing investors' preferences rather than inaccurate beliefs.

As a robustness check of observed variation in WTP across LP types, we allow for variation in the portfolio choice considerations of LPs as described in Eq. (16) with results summarized in column (3). The general pattern of WTP across LP types is quite similar with the exception of institutional LPs that have a greater WTP when we consider portfolio choice considerations.

Finally, column (4) of Table 6 reports a model where we exclude the forecast expected returns. Forecasted expected returns, from our main specification, are lower for impact funds. Thus, we expect that, by excluding these forecasts, the coefficients on impact interacted with the investor types should be lower, as this variable is picking up a lower desirability for the fund associated with missing variable of expected returns. Indeed, this is what we find. The coefficients on all the impact interactions with investor type shift negatively, while the patterns of relative magnitudes of impact coefficients across investor types align exactly with our previous specifications.

5.3. Attributes

5.3.1. Discussion of investor attributes

In this section, we analyze the origins of varying utility over impact by studying attributes of investors that could motivate WTP. Table 7 presents six investor attributes (across columns) and their mapping to the ten LP types (rows). The first three attributes characterize inherent LP features. *Household* categorizes investors based on the constituents of the capital (organizations or households). *Intermediated* classifies the LP types based on whether the capital is intermediated through an asset manager, with an observation that intermediation creates distance between the ultimate owner of capital and those who facilitate capital allocations. *Mission* identifies investors (development organizations and foundations) that have an impact mission as a primary goal.

The last three attributes (*Pressure, Laws,* and *Charters*) characterize the implicit or explicit rules governing different investors' ability and desire to invest in impact funds. For these attributes, we exploit the fact that laws govern-

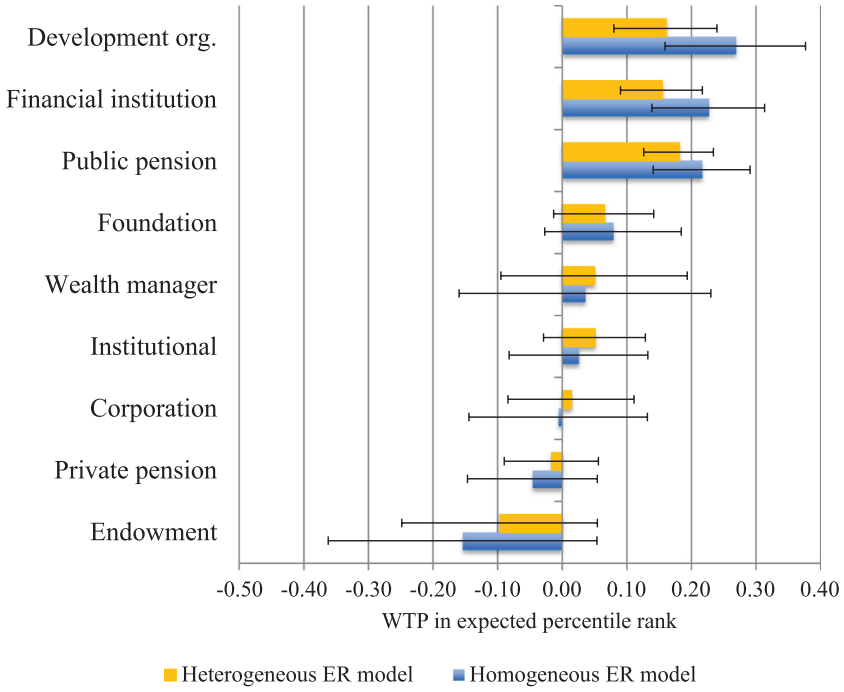

Fig. 3. Willingness to pay (WTP) for impact by investor type.
The figure presents estimates of the willingness to pay for impact derived from the logit estimation of Table 6, column (1) (homogeneous expected return model) and column (3) (heterogeneous expected return model). The WTP magnitude is the ratio of the impact coefficient for the LP type divided by the expected returns coefficient expressed in percentile ranks. Percentiles are based on performance relative to cohort funds, where cohorts are defined by fund vintage year and region. Black bands represent 95% confidence intervals on WTP estimates.

ing these rules vary by geography, thereby allowing us to estimate the WTP associated with these rules within an LP type fixed effect model. Although we do not claim causal identification, this within-LP-type estimation strategy offers suggestive evidence that rules may directly affect WTP for impact.

Pressure (column (4)) identifies regulatory or political pressures that encourage impact investment. Worldwide, public pensions, despite commonly being subject to a fiduciary duty standard, may face political pressure to increase the (perceived or real) welfare of voting populations.[18] Likewise, financial institutions worldwide may have incentives to invest in impact funds that serve low- to moderate-income communities if such investments garner goodwill from customers or regulators. However, in the US, additional regulations (or the threat to regulate) are imposed on financial institutions in a way not operative in other countries. Specifically, US commercial banks are subject to investment obligations to serve their local low- and moderate-income communities under the CRA (CRA Investment Handbook, 2010, p.24). Likewise, insurance companies in some of the large US states (e.g., Texas, New York, and California) must comply with state-level insurance regulations akin to the CRA that require them to invest in local communities. Even outside of those states, insurance companies in the US may face pressure to invest in impact locally to preempt passage of a federal CRA-like regulation for insurance (Gainer, 2009). We exploit this geographic variation by coding *Pressure* equal to one for US financial institutions (banks and insurance companies), as well as for public pensions worldwide, and zero for others.

Laws (column (5)) identifies investors facing fiduciary duty legal restrictions against impact investing. While most public pensions worldwide face formal (legal or regulatory) restrictions to act solely for the benefit of pension recipients by achieving target investment returns and providing liquidity while minimizing risks and costs, regulations concerning investing principles of other entities such as endowments, foundations, and private pensions are typically less restrictive. However, in the US, foundations, endowments, and private pensions face more restrictive fiduciary standards than their non-US counterparts. US private pensions are subject to the 1974 ERISA, which states that a pension plan fiduciary could consider nonfinancial factors (such as environmental or social impact) only if doing so would result in the same level of return at the same level of risk as comparable investment alternatives.[19]

[18] Public pensions may also face pressure to serve the political interests of their boards, which are often pro-labor and consider local job creation as an important policy goal. Consistent with this idea, Dyck et al. (2016) and Andonov, Hochberg, and Rauh (2018) both show that the investments of public pensions are affected by the degree to which the boards governing the pensions are appointed by government officials.

[19] U.S. National Advisory Board (NAB), 2014. Private capital, public good. The ERISA guideline issued in 2008 and in effect until 2015 went even further, stating that pensions "… may never subordinate the economic interests of the plan to unrelated objectives, and may not select investments

Table 7
Limited partner (LP) types and attributes related to impact motives.
The table summarizes investor attributes by LP type (column (1)) and region. Column (2) indicates whether the primary constituents of the capital are households (versus organization). Column (3) indicates whether the constituent capital is intermediated as opposed to directly invested by the constituent or an administrator (e.g., foundations and pensions). Column (4) indicates whether impact is a primary goal of the constituent. Column (5) identifies legal and political pressure to invest with impact. The last two columns identify laws (e.g., ERISA) and charters (e.g., corporate charters) that restrict impact investment.

Limited Partner	Household	Intermediated	Mission	Pressure toward impact	Laws restricting impact	Charters restricting impact
Development organizations	–	–	yes	–	–	–
Foundations	–	–	yes	–	yes UPMIFA and PRI (US)	–
Financial institutions	–	–	–	yes Community Reinvestment Act & state regulation modeled after CRA (US)	–	yes
Endowments	–	–	–	–	yes UPMIFA (US)	–
Corporate & government portfolios	–	–	–	–	–	yes
Institutional asset managers	–	yes	–	–	–	yes
Wealth managers	yes	yes	–	–	–	–
Private pensions	yes	–	–	–	yes ERISA (US)	yes (non-US)
Public pensions	yes	–	–	yes Political pressure	yes State & national laws	–

Likewise, the UPMIFA, which governs the management of US foundations and university endowments, imposes fiduciary standards similar to those of ERISA (see Geczy et al., 2015). However, unlike the ERISA, the UPMIFA provides an additional duty of obedience to the unique charitable mission of the organization. Furthermore, tax laws in the US create an additional hurdle on foundations. The US tax authority requires foundations to maintain a 5% annual payout rate to keep their tax-exempt status; impact investments in the form of program-related investments (PRIs) can count if certain eligibility tests are met.[20] While the policy may have been intended to encourage PRIs, the ambiguity around the test outcome and the perceived threat of tax-exempt status loss may subdue foundations' WTP for impact in their investment portfolio.

Charters (column (6)) identifies restrictions against impact investment in the form of organizational charters, excluding investors already covered by legal restrictions (column (5)) under the assumption that legal restrictions are more binding. Charters require organizations to maximize value for shareholders, which may constrain investments into impact funds. The list of organizations bound by charters includes financial institutions, corporations, non-US private pensions (subject to fiduciary responsibility via their parent corporate charters), and institutional asset managers (subject to fiduciary standards of the institutional sources of capital).

Finally, as we noted in the introduction, both the number of and the dollar amount of assets managed by organizations that are UNPRI signatories have sharply increased in the recent years. Since investors signing the UNPRI are

on the basis of any factor outside the economic interest of the plan" (p.12 of Johnson, K., 2014, "Introduction to Institutional Financial Duties," International Institution for Sustainable Development research report) and that those who consider noneconomic factors could be challenged later for noncompliance with ERISA absent a written record demonstrating no financial sacrifice was made. The new ERISA guideline issued in 2015 withdraws this language and reverts to the original ERISA restrictions. See: https://www.dol.gov/opa/media/press/ebsa/ebsa20152045.htm.

[20] Specifically, the PRIs must further the foundation's organization mission, and the financial returns cannot be a primary purpose of the investment. In practice, PRI investors are required to demonstrate that conventional investors maximizing returns would not invest at the same term as their investment terms. This is simple if the financial instrument used is a below-market return debt security. Precisely for this reason, below-market-return loans are popular vehicles for PRIs. In contrast, equity vehicles are relatively rare, possibly because of the perceived risk of violating the PRI eligibility requirement if it makes too much profit ex post.

Table 8
The willingness to pay for impact by investor attribute.
This table presents logit model estimates (Panel A) and willingness-to-pay estimates (Panel B) including variables to test the incremental willingness to pay for investor attributes. In columns (1) to (3), a fund's expected return forecast is homogenous across all investors. Columns (4) to (6) allows heterogeneous forecast for each fund by LP type. All columns include the interaction of the impact variable with the six LP attribute dummies, a UNPRI signatory dummy variable (that is one for LPs that signed the UNPRI), and a UNPRI post-signing dummy variable. Columns (2) and (5) add in the interaction of the impact variable with the LP geography. Columns (3) and (6) further add the ten LP types and impact interactions. All models include standard controls (see text and Table 5 for details). Standard errors clustered at the LP level are in brackets. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

	(1)	(2)	(3)	(4)	(5)	(6)
	Homogenous ER forecast			Heterogenous ER forecast		
Panel A: Model estimates						
Expected returns	3.393***	3.381***	3.386***	4.609***	4.607***	4.608***
	[0.276]	[0.276]	[0.276]	[0.223]	[0.223]	[0.222]
Impact estimates by investor attribute						
UNPRI signatory	0.411***	0.317**	0.357***	0.377***	0.284**	0.328**
	[0.132]	[0.140]	[0.134]	[0.135]	[0.142]	[0.136]
UNPRI post-signing	0.737***	0.702***	0.754***	0.791***	0.764***	0.802***
	[0.211]	[0.211]	[0.211]	[0.219]	[0.219]	[0.218]
Mission	0.916***	0.884***		0.866***	0.764**	
	[0.322]	[0.313]		[0.332]	[0.318]	
Household	0.370	0.319		0.422*	0.277	
	[0.234]	[0.219]		[0.240]	[0.228]	
Intermediated	−0.206	−0.224		−0.0528	−0.052	
	[0.178]	[0.178]		[0.180]	[0.179]	
Pressure	0.987***	1.005***	0.553**	0.957***	0.996***	0.569**
	[0.138]	[0.145]	[0.229]	[0.139]	[0.147]	[0.234]
Charter	0.14	0.196	0.404	0.238	0.203	0.382
	[0.305]	[0.293]	[0.515]	[0.315]	[0.305]	[0.517]
Laws	−0.835***	−0.711***	−0.942***	−0.652***	−0.526**	−0.935***
	[0.211]	[0.222]	[0.353]	[0.216]	[0.226]	[0.353]
Impact	0.0668	n/a	n/a	−0.0472	n/a	n/a
	[0.336]			[0.347]		
Panel B: Incremental willingness to pay (WTP)						
UNPRI signatory	0.12***	0.09**	0.11***	0.08***	0.06**	0.07**
UNPRI post-signing	0.22***	0.21***	0.22***	0.17***	0.17***	0.17***
Mission	0.27***	0.26***		0.19***	0.17**	
Household	–	–		0.09*	–	
Intermediated	–	–		–	–	
Pressure	0.29***	0.30***	0.16**	0.21***	0.22***	0.12**
Restrictions by charter	–	–	–	–	–	–
Restictions by laws	−0.24***	−0.21***	−0.28***	−0.14***	−0.11**	−0.2 ***
Standard controls	YES	YES	YES	YES	YES	YES
LP attributes	YES	YES	YES	YES	YES	YES
Impact*LP geo	NO	YES	NO	NO	YES	NO
Impact*LP type	NO	NO	YES	NO	NO	YES
Pseudo *R*-squared	0.262	0.262	0.262	0.264	0.264	0.264
Observations	3047,430	3047,430	3047,430	3047,430	3047,430	3047,430

doing so with a cost of compliance, it is plausible that they also have higher WTP for impact compared to nonsignatories because of a mission objective. This mission objective may be a fixed attribute for the investor or may reflect some time-varying interest in generation of nonpecuniary benefits from their portfolios. Thus, we introduce two final variables, an indicator variable that takes a value of one if the investor is a UNPRI signatory and an indicator variable that takes a value of one for UNPRI signatories in the years after signing.

5.3.2. WTP results by attribute

Table 8 reports the role of investor attributes in generating WTP for impact. The specification is again the logit estimation with dynamic LP investment groups. Columns (1)-(3) present the results using homogenous expected returns; columns (4)-(6) are for heterogenous expected returns. The columns differ as follows. Column (1) and (4) provides the baseline logits. Column (2) and (5) includes fixed effects for LP geography interacted with impact. This forces the estimation to identify attributes' effects beyond regional preferences for impact. Column (3) and (6) includes LP type interacted with impact fixed effects. The inherent LP type attributes—*Mission, Household*, and *Intermediated*—do not vary by geography and thus drop in this specification.

We report three main attributes results that inform our understanding of investors' WTP for impact. First, having a mission objective increases investors' WTP for impact.

Investors with *Mission* objectives have a WTP for impact of 17–27 percentile ranks (3.4–6.2 ppts in expected excess IRR). We also find evidence supporting a mission objective in considering the coefficients on UNPRI variables. Both the *UNPRI signatory* and the *UNPRI post-signing* variables have positive and significant coefficient across columns. The WTP of being a *UNPRI Signatory* is 6–12 percentile ranks, while that for *UNPRI post-signing* is an additional 17–22 percentile ranks. UNPRI signing captures both temporal and cross-sectional differences in investors' WTP for impact.

Second, investors facing *Pressure* from political or regulatory institutions exhibit a high WTP. In the saturated model of columns (3) and (6), the estimated WTPs for *Pressure* are 12–16 percentile ranks (2.3–3.3 ppts in excess IRR). This evidence is consistent with the interpretation that investors facing *Pressure* returns to satisfy the pressure they face from constituents or to comply with regulators to allocate capital to investments that generate positive externalities.

In auxiliary analyses, we find the effect of *Pressure* can be linked to a preference for local investments by investors that face pressure to invest with impact. We previously noted that financial institutions within the US are subject to regulatory pressure to invest locally, while public pensions funds worldwide are subject to political pressure to do so. In both of these scenarios, the mechanism of pressure acts locally. We test whether pressure is a local concept and find that indeed this is the case: investors subject to pressure are much more likely to invest in impact funds that are focused on generating externalities at home than abroad or in another unrelated region (see Appendix Table A4).

Third, we find that *Laws* of fiduciary duty against dual-agenda impact investing have a significantly negative effect on decisions to invest in impact. In the saturated model of columns (3) and (6), the estimated WTPs for *Laws* are −20 to −28 percentile ranks (−4.2 to −6.7 ppts in excess IRR). Laws like the ERISA and UPMIFA matter. In contrast, we find that having *Charter* restrictions against impact alone does not materially affect their demand for impact, on average; shareholders' recourses (e.g., lawsuits and management turnover) do not seem to bind against impact investing in a way that we can identify.

Note that our results are a mixture of utility from regulatory compliance, social signaling, and preferences as underlying investor motivations. We do not attempt to disentangle the sources of utility across different investors, but our results have some counterparts in the literature. Akin to our *Mission* result, Riedl and Smeets (2017) show that both signaling and preferences explain investors' SRI decisions. Yet, in our data, investors that are subject to regulatory *Pressure* may be more driven by the signaling benefits or regulatory compliance. For example, financial institutions may be interested in impact investing as a method of complying with regulator or fostering local goodwill (a form of social signaling). Likewise, pension managers may have signaling incentives over the distilling perception of local job creation that drives portfolio decisions (Dyck et al., 2016; Andonov et al., 2018). As a counterpoint, Bauer et al. (2019) find in an experiment using a Dutch pension program that retirees themselves support allocating more of their retirement portfolios to sustainable investments even when they expect financial returns to be lower. Disentangling between these underlying mechanisms is an important question that we leave for future research.

5.4. WTP by impact category

In this section we examine whether investors' WTP for impact varies by the impact category (e.g., the environment, women and minority businesses, poverty). Fig. 4 presents the results of the logit model estimation of Table 5, column (1) but interacts the impact dummy with each of the six impact categories we describe in Section 2.1. Note that these categories are not mutually exclusive, as a given fund can meet the criteria of more than one impact category.

The results indicate that investors exhibit a positive WTP when considering investing in impact funds focusing on environment, poverty, and women or minority issues. Investors are willing to forego 15–22 percentile ranks (3.0–4.7 ppts in excess IRR) in performance when investing in these impact categories. Notably, these are all arguably categories with high public good or externality content. In contrast, investors do not exhibit significantly higher WTP when considering investing in impact funds focusing on SME funding relative to nonimpact funds. However, 57% of impact funds in the SME category also have a poverty focus and are thus captured by the poverty category. SME funds without a poverty focus often target particular geographic areas (e.g., Oregon Investment Fund) and are unlikely to attract interest from investors other than local financial institutions and pensions.

6. Conclusion

Our goal has been to understand whether investors are willing to accept lower financial returns for nonpecuniary benefits of intentional impact investing. We show that ex-post financial returns earned by impact funds are 4.7 ppts lower than those earned by traditional VC funds, even after controlling for a host of fund characteristics. To examine whether investors in impact funds willingly trade off expected financial returns at the time of investment decisions, we use a hedonic pricing framework of WTP for impact. We find that impact investors are, on average, willing to forego 13 to 18 percentile ranks of vintage-geography benchmarked performance or about 2.5 to 3.7 ppts in expected excess IRR.

From the perspective of portfolio companies that are financed by impact funds, investors' willingness to accept lower return implies lower cost of capital for the portfolio companies. Assuming 20% carried interest and 2% management fees, back-of-the-envelope calculations for mature funds in the sample suggest that companies that get funded by impact funds generate an excess gross (i.e., gross-of-fees) VM that is 0.29 to 0.43 lower than those funded by traditional VC funds (see Appendix C for details). The mean (median) fund-level gross value multiple in the mature fund sample is 2.3 (1.5). Thus, the WTP for

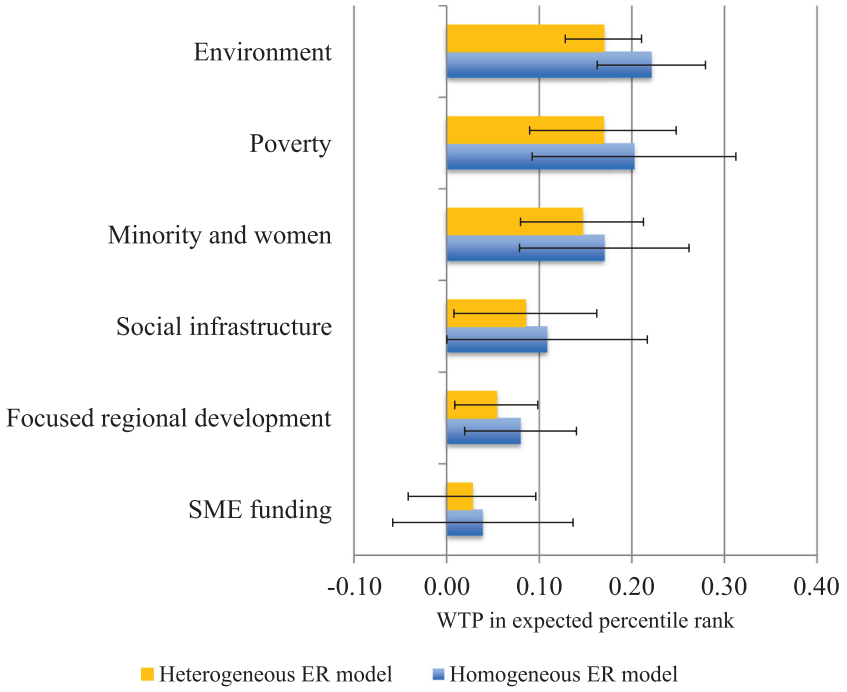

Fig. 4. Willingness to pay (WTP) by impact category.
The figure presents estimates of the willingness to pay for impact by impact category using either homogeneous or heterogeneous expected return models. Estimates are based on a variation of Table 5, column (1) that includes an interaction of impact with the impact type category. The WTP is calculated as the sum of the coefficient on impact and the coefficient on impact*category, all divided by the coefficient on expected returns (expressed in percentile ranks). Percentiles are based on performance relative to cohort funds, defined by fund vintage year and region. Black bands represent 95% confidence intervals on WTP estimates.

impact funds suggests an economically meaningful reduction (e.g., 0.3 reduction is 20% of the median firm's multiple) in the cost of capital for the portfolio companies that they finance.

WTP varies considerably over who controls the capital. To unpack the heterogeneity across investors, we categorize investors into nine broad categories. Investors in three of the nine categories—development organizations, financial institutions, and public pensions—exhibit reliably positive WTP for impact. We then delve into what attributes of investors affect investors' WTP for impact. Not surprisingly, investors with organizational missions and investors that are PRI signatories (especially post signing) have high WTP. In addition, we find that investors facing political and/or regulatory pressure (e.g., banks and insurance companies in the US that face CRA and other equivalent requirements) and those benefiting from political or local goodwill exhibit a higher WTP for impact. In contrast, laws that discourage the sacrifice of financial returns for impact (e.g., ERISA and UPMIFA in the US) may reduce the WTP for impact. Since the number of high-fiduciary LPs affected by such legal restrictions is large (1258 out of 3504 in our sample), this finding has important implications for how subtle shifts in legal interpretations of institutions' fiduciary duty may affect investors' WTP for impact. For example, in the US, the IRS and Treasury issued guidance on mission-related investments in September 2015, assuring that it is possible for private foundations to make a prudent investment using the foundation's assets that advances the foundation's charitable purpose, even if the investment offers a lower rate of return, higher risk, or lower liquidity than alternative investments that do not further charitable purposes. To the extent that (either real or perceived) risk of a tax penalty from making impact investments had a negative effect on their WTP prior to this ruling, this regulatory shift may affect foundations' WTP in the future years. Moreover, recent growth in fundraising by impact buyout and impact infrastructure funds by mainstream General Partners (GPs) like KKR and Bain Capital is consistent with asset managers meeting investors' demand for dual-bottom-line funds. Since a positive loading on the impact implies lower performance sensitivity in fundraising, these GPs may find the impact designation valuable for their objective of maximizing net present value of future fee revenues (Chung et al., 2012).

In combination, our results provide compelling evidence that investors are willing to pay for nonpecuniary characteristics of investments. This result indicates that the capital allocation decisions, though certainly governed by the linchpin risk-return tradeoff of wealth maximization in standard utility models, are also shaped by the real-world consequences of the investments that people make. The WTP for impact varies considerably across legal and regulatory environments, investor geography, and time. This variation opens up a number of avenues for future research that explores the factors that govern the variation that we show.

Appendix

Supplementary material associated with this article can be found, in the online version, at http://jfe.rochester.edu/appendix.htm.

References

Andonov, A., Hochberg, Y.V., Rauh, J.D., 2018. Political representation and governance: evidence from the investment decisions of public pension funds. J. Finance 73, 2041–2086.

Andreoni, J., 1989. Giving with impure altruism: applications to charity and Ricardian equivalence. J. Pol. Econ. 97, 1447–1458.

Andreoni, J., 1990. Impure altruism and donations to public goods: a theory of warm-glow giving. Econ. J. 100, 464–477.

Bauer, R., Ruof, T., Smeets, P., 2019. Get Real! Individuals Prefer More Sustainable Investments. Maastricht University.

Bialkowski, J., Starks, L., 2016. SRI Funds: Investor Demand, Exogenous Shocks and ESG Profiles. Unpublished Working Paper. University of Texas at Austin and University of Canterbury.

Benson, K., Humphrey, J., 2008. Socially responsible investment funds: investor reaction to current and past returns. J. Bank. Finance 32, 1850–1859.

Bollen, N.P.B., 2007. Mutual fund attributes and investor behavior. J. Financ. Quant. Anal. 42, 683–708.

Bottazzi, L., Da Rin, M., Hellmann, T., 2016. The importance of trust for investment: evidence from venture capital. Rev. Financ. Stud. 29, 2283–2318.

Cameron, T.A., James, M.D., 1987. Efficient estimation methods for closed ended contingent valuation survey data. Rev. Econ. Stat. 69, 269–276.

Cavagnaro, D.R., Sensoy, B.A., Wang, Y., Weisbach, M.S., 2019. Measuring institutional investors' skill at making private equity investments. J. Finance 74, 3089–3134.

Chava, S., 2014. Environmental externalities and cost of capital. Manag. Sci. 60, 2223–2247.

Chung, J., Sensoy, B., Stern, L., Weisbach, M., 2012. Pay for performance from future fund flows: the case of private equity. Rev. Financ. Stud. 25, 3259–3304.

Court, A.T., 1939. Hedonic price indexes with automotive examples. In: The Dynamics of Automobile Demand. General Motors, New York, pp. 98–119.

Dimson, E., Karakas, O., Li, X., 2015. Active ownership. Rev. Financ. Stud. 28, 3225–3268.

Dyck, A., Lins, K.V., Roth, L., Wagner, H.F., 2019. Do institutional investors drive corporate social responsibility? International evidence. J. Financ. Econ. 131 (3), 693–714.

Dyck, A., Manoel, P., Morse, A., Pomorski, L., 2016. Bureaucrats, Politicians and Risking-up in Public Pension Funds. University of Toronto and University of California, Berkeley.

Dyck, A., Morse, A., 2011. Sovereign Wealth Fund Portfolios. University of Toronto and University of California, Berkeley.

Epstein, L.G., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. Econ. J. Econ. Soc. 7, 937–969.

Fama, E., French, K., 2007. Disagreement, tastes, and asset prices. J. Financ. Econ. 83, 667–689.

Fama, E., French, K., 1997. Industry costs of equity. J. Financ. Econ. 43, 153–193.

Gainer, B., 2009. *Re*: what lessons does the CRA offer the insurance industry? In: Chakrabarti, P., Erickson, D., Essene, R., Galloway, I., Olson, J. (Eds.), Revisiting the CRA: Perspectives on the Future of the Community Reinvestment Act. Federal Reserve Banks of Boston and San Francisco [includes #s].

Geczy, C., Jeffers, J.S., Musto, D.K., Tucker, A.M., 2018. Contracts with Benefits: the Implementation of Impact Investing. University of Pennsylvania, University of Chicago, and Georgia State University.

Geczy, C., Jeffers, J.S., Musto, D.K., Tucker, A.M., 2015. Institutional Investing When Shareholders are not Supreme. University of Pennsylvania, University of Chicago, and Georgia State University.

Geczy, C., Stambaugh, R.F., Levin, D., 2003. Investing in Socially Responsible Mutual Funds. Wharton School of the University of Pennsylvania.

Global Impact Investing Network (GIIN), 2017. Annual impact investor survey.

Griliches, Z., 1961. Hedonic price indexes for automobiles: an econometric of quality change. In: The Price Statistics of the Federal Government. NBER, Cambridge, pp. 173–196.

Harris, R., Jenkinson, T., Kaplan, S., 2014. Private equity performance: what do we know? J. Finance 69, 1851–1882.

Hart, O., Zingales, L., 2017. Companies Should Maximize Shareholder Welfare not Market Value. Harvard University and University of Chicago.

Hartzmark, S.M., Sussman, A.B., 2019. Do investors value sustainability? A natural experiment examining ranking and fund flows. J. Finance 74, 2789–2837.

Hochberg, Y., Ljungqvist, A., Vissing-Jørgensen, A., 2014. Informational hold-up and performance persistence in venture capital. Rev. Financ. Stud. 27, 102–152.

Hochberg, Y., Rauh, J., 2013. Local overweighting and underperformance: evidence from limited partner private equity investments. Rev. Financ. Stud. 26, 403–451.

Hong, H., Kacperczyk, M., 2009. The price of sin: the effects of social norms on markets. J. Financ. Econ. 93, 15–36.

Huber, J., Train, K.E., 2001. On the similarity of classical and bayesian estimates of individual mean partworths. Mark. Lett. 12, 259–269.

Kaplan, S., Schoar, A., 2005. Private equity performance: returns, persistence and capital flows. J. Finance 60, 1791–1823.

Korteweg, A., Nagel, S., 2016. Risk-adjusting the returns to venture capital. J. Finance 71, 1437–1470.

Korteweg, A., Sorensen, M., 2010. Risk and return characteristics of venture capital-backed entrepreneurial companies. Rev. Financ. Stud. 23, 3738–3772.

Korteweg, A., Sorensen, M., 2017. Skill and luck in private equity performance. J. Financ. Econ. 124, 535–562.

Kovner, A., Lerner, J., 2015. Doing well by doing good? Community development venture capital. J. Econ. Manag. Strat. 24, 643–663.

Laibson, D., 1997. Golden eggs and hyperbolic discounting. Quart. J. Econ. 112, 443–478.

Lerner, J., Schoar, A., Mao, J., Zhang, N.R., 2019. Investing Outside the Box: Evidence from Alternative Vehicles in Private Capital. NBER.

Lerner, J., Schoar, A., Wongsunwai, W., 2007. Smart institutions, foolish choices: the limited partner performance puzzle. J. Finance 62, 731–764.

Liang, H., Renneboog, L., 2016. On the foundations of corporate social responsibility. J. Finance 72, 853–910.

Ljungqvist, A., Marston, F., Wilhelm Jr., W., 2006. Competing for securities underwriting mandates: banking relationships and analyst recommendations. J. Finance 61, 301–340.

McFadden, D.L., 1974. Conditional logit analysis of qualitative choice behavior. In: Zarembka, P. (Ed.), Frontiers in Econometrics. Academic Press, New York [includes #s].

McFadden, D.L., 1986. The choice theory approach to market research. Market. Sci. 5, 275–297.

Metrick, A., Yasuda, A., 2010. Venture Capital and the Finance of Innovation, 2nd Edition John Wiley and Son, Hoboken [includes page #s].

Niehaus, P., 2014. A Theory of Good Intentions. UC San Diego and NBER.

Riedl, A., Smeets, P., 2017. Why do investors hold socially responsible mutual funds? J. Finance 72, 2505–2550.

Renneboog, L., Ter Horst, J., Zhang, C., 2008. Socially responsible investments: institutional aspects, performance, and investor behavior. J. Bank. Finance 32, 1723–1742.

Renneboog, L., Ter Horst, J., Zhang, C., 2011. Is ethical money financially smart? Nonfinancial attributes and money flows of socially responsible investment funds. J. Financ. Intermediation 20, 562–588.

Revelt, D., Train, K.E., 1998. Mixed logit with repeated choices: households' choices of appliance efficiency level. Rev. Econ. Stat. 80, 647–657.

Rosen, S., 1974. Hedonic prices and implicit markets: product differentiation in pure competition. J. Pol. Econ. 82, 34–55.

Sensoy, B., Wang, Y., Weisbach, M., 2014. Limited partner performance and the maturing of the private equity industry. J. Financ. Econ. 112, 320–343.

Sørensen, M., 2007. How smart is smart money? A two-sided matching model of venture capital. J. Finance 62, 2725–2762.

Sorensen, M., Wang, N., Yang, J., 2014. Valuing private equity. Rev. Financ. Stud. 27, 1977–2021.

Starks, L., Venkat, P., Zhu, Q., 2017. Green wash or green walk: the ESG responsiveness of institutional investors. Unpublished Working Paper. University of Texas, Austin.