Limited Partners vs Unlimited Technologies: How Tech Could Transform Investing in Private Capital Funds

Ludovic Phalippou*

University of Oxford, Said Business School

Abstract

We examine how artificial intelligence and machine learning may alter decision-making in private markets. Unlike public equity, where frequent and standardized disclosures enable rapid validation of predictive models, private markets are defined by sparse, delayed, and strategically framed disclosures. This mismatch is the central tension between *limited partners and unlimited technologies*: investors face opacity and decade-long feedback loops, while algorithms are designed for abundant and high-frequency signals. Recent evidence shows that textual features of fundraising documents and GP reports contain systematic predictive content, that deep learning methods can improve cash flow forecasting, and that benchmarking approaches can realign fund categories with underlying exposures. At the same time, data scarcity, model interpretability, and confidentiality restrictions remain binding. We outline a research agenda that integrates computational methods with economic reasoning, emphasizing oversight, causal inference, and transparency as preconditions for reliable adoption.

* Park End Street, Oxford OX1 1HP, United Kingdom. E-mail: ludovic.phalippou@sbs.ox.ac.uk. I am deeply indebted to Elise Gourier for her excellent comments and suggestions. All errors are my own.

Introduction

Over the past two decades, institutional investors have poured trillions of dollars into private market funds, making private equity, private credit, infrastructure, and real estate central pillars of global portfolios. Once peripheral, these asset classes now sit at the core of how pension funds, sovereign wealth funds, and insurers pursue diversification and return. At the center of this transformation stand limited partners (LPs): the providers of capital responsible for selecting, underwriting, and monitoring fund managers.

In listed markets, investors rely on abundant, standardized, and continuously updated information: quarterly filings, analyst calls, and real-time prices. By contrast, LPs operate in an informational environment that is fragmented, delayed, and strategically curated. Fund metrics are noisy and often shaped to support fundraising narratives. Key disclosures (private placement memoranda, interim reports, due diligence questionnaires) are rich in qualitative detail but unstandardized, making cross-fund comparison cumbersome. For an asset class that is both institutionalized and systemically significant, decision tools remain surprisingly analog.

These frictions matter. LPs must commit billions to long-horizon, illiquid vehicles based on incomplete and inconsistent disclosures. They must forecast cash flows under commit-and-call structures, often during stress episodes such as the global financial crisis, the COVID-19 pandemic, or the recent interest-rate tightening cycle. They must report to boards, principals, and regulators in transparent and auditable ways, even as the underlying data resist standardization. They must also respond to growing ESG mandates, despite fragmented and often unverifiable reporting. At every stage, including fund selection, co-investments, portfolio monitoring, or secondary transactions, the challenge is not the absence of information per se, but the difficulty of converting heterogeneous and strategically framed disclosures into actionable insight.

It is against this backdrop that rapid advances in artificial intelligence (AI), machine learning (ML), and natural language processing (NLP) have captured attention. In public markets, text-as-data methods already demonstrate predictive value: the tone of 10-K filings predicts returns and litigation risk (Loughran & McDonald, 2011); earnings calls reveal political and epidemic exposures (Hassan, Hollander, van Lent, & Tahoun, 2019, 2023); and surveys highlight the rise of computational text analysis as a paradigm (Gentzkow, Kelly, & Taddy, 2019). The arrival of large language models (LLMs) amplifies these arguments. They offer tools to parse complex disclosures with nuance that older approaches cannot match (Kim, Muhn, & Nikolaev, 2023). Yet the structural mismatch between decade-long feedback loops in private markets and algorithms designed for abundant, high-frequency data, remains a challenge.

The structural contrast between private and public markets is stark indeed. Listed firms disclose frequently in standardized formats, enabling researchers to link narratives to market outcomes in real time. In private markets, disclosures occur quarterly at best and are invariably three to six months old, with scope and format left to managerial discretion and lacking comparability. Outcomes materialize only after four to ten years, creating an exceptionally slow feedback loop for model validation. Core documents are long, narrative-heavy, and heterogeneous, complicating systematic analysis. Methods that thrive in public markets therefore cannot simply be transplanted into private ones.

Yet opacity also creates opportunity. Narrative disclosures are dense with information: they articulate strategies, governance mechanisms, and risk frameworks. Recent studies confirm that such narratives contain systematic and forward-looking signals. Fundraising prospectuses correlate with ultimate performance, and the tone of interim GP reports predicts subsequent deal outcomes even after controlling for the reported (interim) asset values (Braun, Fernández-Tamayo, López-de-Silanes, Phalippou, & Sigrist, 2023, 2025). These findings suggest that tools once reserved for public-market analysis can extract meaningful signals from private equity disclosures long regarded as anecdotal.

At the same time, a commercial ecosystem of "LP tech" has emerged. Vendors now market platforms for parsing fund documents, aggregating portfolio data, and assisting with secondary transactions. These products promise to automate labor-intensive processes and deliver structured dashboards. But their effectiveness remains uncertain: few disclose validation procedures, and the line between genuine innovation and rebranded workflow automation is often blurred. The risk is that technological enthusiasm adds complexity without improving outcomes.

This paper contributes by situating these developments within the broader finance and economics literature. It does so in two steps. First, it synthesizes the emerging academic evidence on AI, ML, and NLP in private markets, highlighting both results and limitations. Second, it outlines a forward-looking agenda for research and practice, emphasizing interpretability, governance, and the integration of human oversight. The central argument is that AI will not eliminate the distinctive frictions of private markets, but it can make them more tractable and systematically managed.

The paper proceeds as follows. Section 2 reviews the LP role and decision environment. Section 3 examines the data constraints that shape statistical modeling. Section 4 surveys academic applications of AI and ML in both public and private markets. Section 5 describes the rise of commercial tools. Section 6 concludes with implications for the broader theme of *limited partners versus unlimited technologies*.

2. Overview of the LP Role in Private Markets

Limited Partners (LPs) are the primary capital providers to private market funds. They include pension funds, endowments, sovereign wealth funds, and insurance companies, but also a growing number of fund-of-funds, feeder funds, banks, asset managers, individuals, family offices, and semi-liquid vehicles. Despite this diversity, most LPs share a common set of functions and constraints, which this section describes.

2.1. Capital Allocation and Fund Selection

At the core of the LP role is the task of allocating capital across fund vintages, strategies, geographies, and managers. LPs are typically multi-asset investors balancing private equity, private credit, infrastructure, real estate, and other alternatives within a broader portfolio. Fund selection decisions are informed by a combination of quantitative signals such as historical returns (IRR, TVPI, DPI) and qualitative inputs such as the GP's strategic positioning, team continuity, reputation, and narrative consistency.

The process typically begins with sourcing and screening. LPs maintain internal lists of active and potential fund managers, often informed by relationships, consultant pipelines, and databases. GPs provide marketing decks, private placement memoranda (PPMs), and in many cases access to a data room. Most GPs provide access to a virtual data room, which contains key documents to support LP analysis: the PPM, the LPA, detailed track record files, audited financial statements, deal case studies, DDQs, and occasionally ESG policies and diversity disclosures. These documents are unstandardized, often lengthy, and vary in structure even within the same GP across fund vintages.

These materials allow LPs to conduct an initial screen. Typical first questions include: Does the GP's strategy align with their mandate? Are returns within expected parameters given the risk profile and market context? Has the team been stable? Is the fund size appropriate relative to past performance and current opportunity set?

Prior research shows that LP commitment decisions are sensitive to interim performance metrics reported during fundraising periods, which may be strategically managed by GPs (Barber & Yasuda, 2017; Brown, Gredil & Kaplan, 2019). More broadly, survey evidence highlights the role of qualitative considerations such as relationships, team impressions, and mandate alignment, in fund selection (Da Rin & Phalippou, 2017).

2.2. Fund Due Diligence

Once a shortlist of fund opportunities is identified, LPs undertake formal due diligence to assess return potential, risk profile, operational integrity, and legal structure. While mandates differ, the process generally follows three pillars.

Investment due diligence evaluates the GP's historical performance, portfolio construction discipline, sourcing model, and sector expertise. A central question is whether past outcomes reflect repeatable skill or idiosyncratic circumstances.

Operational due diligence assesses governance and infrastructure, including compliance systems, valuation practices, and IT platforms, to ensure that the fund can execute its strategy.

Legal due diligence examines Limited Partnership Agreements (LPAs), side letters, and alignment mechanisms such as key-person clauses and waterfalls. For intermediated investors (e.g., fund-of-funds or outsourced CIOs), compliance with end-client restrictions is also verified.

LPs synthesize these inputs into internal memos and investment committee materials, often with input from external consultants and lawyers. Experienced LPs also benchmark current disclosures against prior vintages of the same GP, testing for consistency in terms and track record claims.

2.3. Single-Investment Underwriting

Single-deal underwriting poses a distinct challenge for LPs. Co-investments, GP-led secondaries, and direct transactions promise higher net returns and targeted exposures, but they compress the decision cycle and increase downside risk. LPs must rapidly digest investment memoranda, management presentations, and financial models, and form an independent view on value, structure, and risk. Key questions include: Are base case assumptions credible? How sensitive are

valuations to macro or operational shocks? What exit scenarios are plausible under different market conditions? Is the GP, or lead investor, aligned in both economics and governance?

Unlike fund diligence, which is portfolio-based and relatively standardized, single-deal underwriting is bottom-up, bespoke, and idiosyncratic. It requires benchmarking the target company against peers, identifying potential asymmetries, and negotiating rights—all while managing the relational cost of rejecting opportunities.

These settings are also where technological tools could matter most. Because timelines are compressed, the ability to quickly parse documents, benchmark metrics, or simulate scenarios can significantly improve clarity and consistency. While technology cannot eliminate the trade-offs between speed, rigor, and relationship management, it can make them more transparent and systematically navigable.

2.4. Cash Flow Management and Liquidity Planning

Unlike traditional investments, private market funds operate under a "commit-and-call" structure. LPs make capital commitments upfront, but the timing of capital calls and distributions is at the discretion of the GP. As a result, LPs must manage liquidity across two dimensions: funding sudden capital calls and absorbing lumpy distributions. This dynamic introduces significant complexity into both strategic asset allocation and treasury operations.

Even in normal market conditions, forecasting net cash flows is challenging. GPs rarely provide capital call schedules in advance, and distributions depend on exit markets. In periods of market stress, the difficulty intensifies: capital calls may accelerate, while distributions slow significantly. This "liquidity crunch" dynamic was especially evident during the Global Financial Crisis, the COVID-19 crisis, and more recently after the increase in interest rates (2024-2025).

These constraints are particularly acute for LPs managing retail-facing or semi-liquid products (e.g., defined contribution platforms, feeder funds, interval funds). Such vehicles must respond to inflows and redemption requests in real time, even while their underlying capital is locked into illiquid structures. Liquidity buffers, overdraft facilities, and secondary sales all become tools of necessity rather than design.

At a more strategic level, liquidity planning also affects how LPs size commitments. Many rely on internal pacing models that project capital call patterns, expected distributions, and target private markets exposure. These models are typically rule-based and backward-looking, relying on historical averages rather than forward-looking signals. This creates a tension between staying fully allocated and avoiding overcommitment risk especially when market conditions shift or exit environments deteriorate.

The use of subscription lines of credit adds complexity to forecasting in private markets. By allowing GPs to fund investments without immediately drawing down LP capital, these facilities delay actual cash outflows and mask the true timing of capital deployment. As a result, effective liquidity management becomes a vital yet often overlooked aspect of LP investing. It affects asset allocation, risk management, and portfolio operations. This is an area where improved forecasting, simulation tools, and better data integration could substantially enhance decision-making and outcomes.

2.5. Portfolio Monitoring and Benchmarking

Once capital is committed, LPs shift from due diligence to ongoing monitoring. This phase is both operationally demanding and analytically complex. Most funds report on a quarterly basis, but valuations are lagged, often stale, and subject to GP discretion. Disclosures vary significantly in format, content, and quality making systematic comparison difficult.

Typical quarterly reports include updates on Net Asset Value (NAV), cash flows, and portfolio company developments. At the core of monitoring is the tracking of key metrics such as NAV, distributions, and performance measures like DPI, TVPI, and IRR. To evaluate performance, LPs compare these figures against internal expectations, peer funds, and relevant market benchmarks.

To manage this process, many LPs build internal dashboards or rely on external consultants. These systems aggregate data across funds but often depend on manual inputs, spreadsheet consolidation, and reconciliation across inconsistent formats. As portfolios grow in size and complexity, this resource-intensive approach struggles to keep pace.

Effective monitoring can reduce surprises, enable more timely engagement with GPs, and support better portfolio-level decisions. It also improves transparency and accountability in reporting to investment committees, boards, clients, and regulators. As the scale and scope of private market portfolios expand, the need for more systematic approaches to monitoring and benchmarking continues to increase.

2.6. Secondary Market Participation

Beyond primary fund commitments, many LPs actively engage in the secondary market to manage exposure, enhance liquidity, or rebalance strategies. These transactions fall into two categories: GP-led processes (e.g., continuation vehicles, already discussed above) and LP-led sales of existing fund stakes. The latter has grown into a large market, with volumes exceeding $100 billion annually and portfolios often comprising stakes in dozens of funds. For the buyer, a single transaction may imply exposure to hundreds of underlying companies, each requiring assessment of value, risk, and governance.

LP-led secondaries are resource-intensive. Pricing typically starts from GP-reported NAVs, which are three to six months old, discretionary, and not comparable across funds. Buyers must adjust for fund-specific terms, residual duration, and GP quality, while also modeling portfolio company prospects under different scenarios. Sellers, meanwhile, weigh bids against book value and opportunity cost, often under liquidity pressure. Negotiations are bilateral, confidential, and frequently involve bespoke structures (e.g., deferred payments, portfolio carve-outs).

These features make LP-led secondaries a natural domain for technological support. Parsing GP reports, normalizing metrics, and aggregating data across hundreds of portfolio companies can materially improve decision speed and accuracy. Machine learning may aid in benchmarking pricing across vintages and strategies, while natural language processing can extract key terms from LPAs or track shifts in GP disclosures. Given the scale of information embedded in secondary portfolios, even partial automation could reshape how LPs and intermediaries operate, highlighting both the challenges and the opportunities for AI-driven tools in private capital investing.

2.7. Reporting Obligations

Reporting has become one of the most complex and resource-intensive functions of the LP role. It now spans ESG and thematic monitoring, internal governance and oversight, client communications, and regulatory compliance. What was once a relatively narrow task of tracking cash flows and portfolio valuations has expanded into a multidimensional reporting ecosystem, where the expectations of boards, clients, regulators, and broader stakeholders all converge.

On the ESG side, many LPs face explicit requirements to monitor climate risk, labor practices, diversity metrics, and other sustainability-related exposures. Public institutions in particular are under pressure to demonstrate alignment with global frameworks such as the UN Sustainable Development Goals (SDGs) or net-zero pledges. Others, especially European investors, are subject to regulatory disclosure mandates under regimes like the Sustainable Finance Disclosure Regulation (SFDR). Beyond these formal obligations, many LPs have made voluntary commitments to responsible investing, stewardship, or thematic initiatives such as social equity or decarbonization. These expectations translate into concrete reporting requirements: portfolio-level carbon footprints, sectoral exposure to “brown” or “green” assets, diversity statistics at the GP and portfolio company level, or qualitative narratives of engagement activity.

The difficulty is that the underlying data are highly uneven. ESG disclosures remain voluntary for most GPs, varying in both scope and quality. Some funds produce annual sustainability reports with detailed carbon and governance metrics, while others provide only cursory commentary. Even when quantitative indicators are provided, they may not be measured consistently across funds or jurisdictions. This makes it hard for LPs to compare results across managers or to aggregate exposures at the portfolio level. For many institutions, the process of preparing ESG reports involves collecting whatever data are available, synthesizing them into internal templates, and carefully caveating the resulting outputs.

Alongside ESG, traditional internal and client reporting obligations have also intensified. LPs must provide regular updates to investment committees, boards of trustees, and other oversight bodies. These reports typically cover performance, pacing and pipeline activity, attribution analysis, risk exposures, liquidity planning, and fee and expense tracking. For governance purposes, they often serve as the basis for decisions about re-ups, portfolio rebalancing, or risk mitigation. In many institutions, this requires coordination across investment teams, finance, compliance, and risk functions, making reporting a cross-departmental effort rather than a narrow investment task.

The burden is even heavier for intermediated investors such as fund-of-funds, multi-manager platforms, or outsourced CIOs. These LPs must deliver tailored dashboards and custom reports for multiple clients, often across jurisdictions. A pension plan may ask for attribution by sector and geography; an endowment may want exposure to climate-related risks; a sovereign wealth fund may request detailed fee transparency; while regulatory bodies demand standardized disclosures of valuation methods or risk factors. Meeting these heterogeneous demands requires both data flexibility and reporting scalability.

Fee and cost transparency is a further pain point. GP disclosures of fees, expenses, and carried interest are often opaque, inconsistent, or buried in footnotes. Yet LPs are increasingly expected

to produce comprehensive reports on the all-in cost of private equity exposure, broken down at both the fund and portfolio level. This requires reconciling data from multiple vintages, legal structures, and currencies, and allocating expenses in ways that are often not straightforward. For many LPs, this remains a highly manual exercise.

Regulatory scrutiny has added yet another layer of complexity. Supervisors in several jurisdictions now require regular disclosure of private market exposures, valuation methodologies, liquidity stress tests, or ESG risks. These demands not only raise the bar for internal reporting quality but also require auditable documentation and robust systems integration. For LPs operating across multiple jurisdictions, the challenge is multiplied by inconsistent regulatory frameworks and overlapping reporting timelines.

Taken together, these obligations highlight a central tension of private market investing. LPs must transform a patchwork of GP disclosures into outputs that are consistent, defensible, and tailored to multiple stakeholders. This is not merely a compliance burden. It is increasingly a test of operational scalability and credibility. Institutions that fail to meet reporting expectations risk reputational damage, regulatory sanction, or client dissatisfaction.

These pressures also create fertile ground for technological solutions. Tools that can extract structured data from GP reports, standardize ESG disclosures, or automate the reconciliation of fees and cash flows could materially reduce reporting costs and errors. Natural language processing may help classify and normalize diverse ESG narratives, while machine learning could assist in mapping disparate fund terms to consistent reporting categories. Even partial automation—such as standardizing capital call notices or flagging inconsistencies in GP disclosures—can deliver meaningful efficiency gains. As portfolios grow in size and complexity, the ability to meet escalating reporting demands will increasingly depend on such tools.

In sum, reporting obligations now sit at the intersection of investment oversight, client communication, ESG stewardship, and regulatory compliance. They illustrate both the breadth of the LP role and the depth of the informational frictions it faces. They also underscore why private markets provide such a natural testing ground for AI and data-driven solutions: the task is not to generate more data, but to transform heterogeneous, lagged, and strategically framed disclosures into consistent and actionable information.

2.8. Summary

The functions of LPs range from capital allocation and due diligence to liquidity management, monitoring, secondaries, and reporting, and each is shaped by incomplete, delayed, and strategically curated disclosures. These frictions make private markets operationally demanding while creating clear pain points where technology could add value such as accelerating document review, standardizing unstructured data, or supporting large-scale analyses in secondary transactions. At the same time, the very features that define LP investing also constrain what AI and machine learning can deliver. This tension sets the stage for Section 3, which examines the data environment in private equity and its implications for statistical modeling.

3. Data Constraints in Private Equity

3.1. The issue

One of the central challenges in using statistical models in private markets is the constraint on data availability. Unlike public markets, where firms are required to disclose quarterly earnings, hold public calls, and issue standardized filings that generate large corpora of text data, private markets operate in a much more opaque informational environment. These informational frictions shape both what is feasible for LPs and how technology can be applied.

The first constraint is the limited volume of data. Transactions are relatively infrequent, especially compared to the high-frequency time series available in listed markets. Investment horizons are long, and outcomes are only observable years later. A buyout fund raised in 2024 may not be fully realized until the mid-2030s, leaving investors and researchers with a decade-long lag before reliable performance data become available. In addition, the private equity company universe itself has changed substantially over time. Sectoral composition, financing instruments, and reporting norms in venture capital today differ starkly from those two decades ago. This temporal heterogeneity complicates the use of statistical models that rely on stationary processes.

Second, the data that do exist are at low frequency and uneven quality. Fund-level NAVs are reported quarterly and with significant delays. Portfolio company disclosures are sparse and heterogeneous, depending on the GP–LP relationship and on jurisdiction-specific legal requirements. Unlike public markets, there is no mandate for continuous, standardized disclosure. Fee and expense data are often buried, and even when shown they are aggregated at varying levels across funds. Even within the same GP, reporting formats can shift across vintages. These inconsistencies make it difficult to build large, clean datasets suitable for systematic analysis.

Third, while structured numerical data are scarce, private equity generates large volumes of unstructured text. Fund prospectuses, private placement memoranda, LPAs, DDQs, ESG reports, quarterly letters, and investment committee presentations all contain rich narrative information. These documents are not standardized, but they are information-dense: they reveal forward-looking strategies, risk disclosures, contractual arrangements, and governance mechanisms. These are precisely the types of inputs that modern natural language processing tools can parse and structure, transforming what has long been qualitative and anecdotal into something systematic.

Finally, the contrast with public markets highlights both the challenge and the opportunity. Listed firms disclose frequently and in standardized formats, which enables text-as-data research and rapid model validation. By contrast, private markets have far fewer documents, lower frequency, and outcomes that arrive only after long delays. These conditions make the application of statistical tools much harder, but they also underscore why private markets are such a natural testing ground for innovation.

3.2 Machine Learning for Prediction

Machine learning methods in their supervised form link inputs (for example, documents, fund attributes, or cash flows) to outcomes such as returns, distributions, or default rates. In private markets this is inherently difficult because labels often arrive with long delays, are sometimes strategically framed (for example, defaults may be hidden), and reflect heterogeneous structures.

A typical buyout fund takes around twelve years to fully realize. If a researcher trains on one set of vintages, tests on another, and then validates on an out-of-sample period, the implied horizon quickly approaches thirty years. By that time, the macroeconomic cycle, regulatory regimes, fund types, and contractual innovations may all have changed. Models trained on 1990s vintages are poor guides for continuation funds, NAV lending, or semi-liquid vehicles in the 2020s. This slow feedback loop constrains model development and validation.

Two responses follow from these constraints. First, in high-dimensional settings with limited labels, penalized linear models such as lasso and elastic net usually provide robust, interpretable baselines (Gentzkow, Kelly, & Taddy, 2019). Second, researchers shorten the learning cycle by using intermediate performance benchmarks. Recent work shows that fund performance observed at year five contains substantial predictive content for ultimate outcomes; roughly two-thirds of funds that outperform at year five continue to do so at year ten (Braun, Fernández-Tamayo, López-de-Silanes, Phalippou, & Sigrist, 2023). Using such interim labels accelerates feedback and makes supervised learning more feasible.

Another approach is to move from fund-level outcomes to deal-level proxies that arrive earlier. In the venture capital context, early follow-on financing, valuation uplifts, or patenting activity can serve as interim performance signals that are more timely than ultimate fund performance. These types of intermediate indicators are ways to shorten the learning cycle and provide training labels within a few years rather than a decade.

Even with these adjustments, however, the constraints are real. Machine learning models trained on historical vintages may struggle to generalize when the underlying structures of funds and markets evolve. At the same time, the fact that interim performance signals are predictive, and that supervised models such as random forests and gradient boosting can extract structure from limited datasets, suggests that machine learning has a role to play in screening and classification tasks. Moreover, recent work in public markets shows that even small datasets can suffice: Kelly, Malamud, and Zhou (2024) demonstrate that complex ML models can outperform simpler approaches in predicting returns, even when the number of observations is modest. This result underscores that machine learning can be effective even in data-scarce environments, making it directly relevant for private equity. The challenge is to design approaches that balance the promise of high-dimensional models with the realities of delayed outcomes, strategic reporting, and changing market institutions.

3.3 Artificial Intelligence and Large Language Models

Method choice should match the task: when labels are scarce and the target concept is predefined, dictionary indices can suffice; when labels exist and a single attribute must be predicted, supervised text regression is typically preferred (previous subsection); and when the goal is to uncover latent structure or harmonize inconsistent language, generative models and embedding-based representations are appropriate (Gentzkow, Kelly, & Taddy, 2019). In short, supervised ML requires labeled outcomes to learn predictive patterns, whereas LLMs are pretrained on vast external corpora and can interpret new documents without pre-specified labels. Given the limitations discussed above, large language models (LLMs) are particularly attractive in private markets.

In practice, LLMs can already be applied to a wide range of LP workflows. They can standardize private placement memoranda (PPMs), which vary widely in length and terminology across managers. They can extract and compare key clauses in Limited Partnership Agreements (LPAs) and side letters, or process bespoke due diligence questionnaires. They can summarize quarterly GP reports and ESG disclosures into standardized formats, allowing LPs to track exposures and commitments across diverse portfolios. Their main strength lies in workflow acceleration, interpretation, and standardization.

But these strengths are also limitations. Because LLMs are not trained on what "worked" and what "failed" in private equity, they cannot on their own predict future returns. They can organize and highlight patterns in disclosure, but they cannot establish whether a given narrative or contractual feature is associated with superior performance. Their outputs must therefore be complemented by human judgment and economic reasoning.

LLMs also introduce new risks. They reproduce biases embedded in their training corpora, amplifying distortions if certain geographies, sectors, or strategies are underrepresented in the data. When applied to historical PE documents, they are prone to hindsight bias, possibly reading past disclosures in light of subsequent outcomes rather than contemporaneous expectations. They are also known to hallucinate, producing plausible-sounding but inaccurate classifications or summaries, which is particularly problematic in legal contexts.

Recent developments in public markets highlight both the promise and pitfalls of these tools. Cao, Jiang, Yang, and Zhang (2023) show that firms strategically adjust their disclosure language in earnings calls to account for "machine listeners," underscoring that once LLMs are deployed, disclosure itself evolves in response. Kim, Muhn, and Nikolaev (2023) demonstrate that ChatGPT can assist investors in digesting lengthy corporate filings, but their results also highlight that the outputs require validation and correction.

A further vulnerability is the susceptibility of LLMs to prompt injection. Security researchers recently documented cases in which academic manuscripts contained hidden instructions (in tiny fonts or white-on-white text) directing the model to "ignore all previous instructions" and produce favorable reviews. These hidden prompts were invisible to human readers but processed by AI systems, effectively tricking them into biased outputs. New tools such as *PhantomLint* attempt to detect and strip such hidden instructions from documents, but the episode illustrates a broader concern: in any setting where LPs rely on AI to parse contractual or narrative documents, adversarial formatting could compromise results without being visible to human reviewers.

The implication is clear: LLMs should be treated as complements, not substitutes, for human expertise. They can drastically reduce the cost of parsing heterogeneous disclosures, but their outputs must be validated and contextualized by experienced LPs. Without oversight, they risk introducing new errors and biases into decision-making; with oversight, they offer the possibility of transforming highly manual processes into scalable and systematic ones, easing one of the central bottlenecks in private market investing.

3.4 Confidentiality issues

A new and potentially transformative development complicates this picture. Some GPs have begun instructing LPs not to submit any of their documents to LLMs, even when these are premium services that promise not to retrain on user inputs. This reflects concerns over confidentiality, intellectual property, and compliance, especially in contexts governed by strict contractual and fiduciary obligations. The implication is profound: the very disclosures that constitute the richest textual data in private markets may no longer be available for systematic analysis if such restrictions are enforced. What had seemed like an abundant source of unstructured information could, in practice, remain locked in analog form.

Whether GPs are contractually or legally entitled to prohibit LPs from processing documents with LLMs is itself an open question. Limited Partnership Agreements often contain confidentiality clauses, but these were drafted long before the advent of AI tools. If LLM use is interpreted as a form of unauthorized disclosure or data transfer, LPs may face contractual liability for feeding documents into automated systems. On the other hand, LPs can reasonably argue that using secure, closed-loop services for internal analysis is no different from hiring external consultants or deploying traditional software. This ambiguity creates a new fault line in GP–LP relations, raising issues around technological adoption that have yet to be resolved.

If LLMs are effectively blocked, LPs are left with supervised ML methods as the fallback. Training models on structured databases does not involve sending documents to external platforms and therefore falls outside the immediate scope of most GP restrictions. Technically, however, one could argue that even building an internal predictive model on GP documents constitutes "feeding" them into a machine. Such a broad interpretation would be implausible. Large LPs have for decades applied regressions, portfolio models, and other statistical tools to their internal archives, and treating those practices as equivalent to external LLM use would represent a dramatic and unprecedented departure from established norms.

For both research and practice, the stakes are high. Much of the promise of AI in private equity rests on the ability to analyze unstructured GP disclosures. If GPs succeed in blocking such applications, LPs and researchers will need to develop alternative strategies: relying on public regulatory filings, constructing synthetic datasets from secondary transactions, or negotiating explicit contractual rights to apply computational tools to confidential materials. The outcome of this debate will shape not only the trajectory of academic work but also the practical feasibility of AI adoption in private markets.

3.5. Summary and further discussions

This section highlights both the opportunities and the limits of advanced statistical methods in private equity. Supervised ML offers predictive power in principle, but delayed outcomes and shifting market structures make validation difficult. Workarounds such as intermediate labels or deal-level proxies shorten the horizon, yet the feedback loop remains slow. LLMs, by contrast,

excel at interpretation and standardization, but cannot connect disclosures to ultimate outcomes and may introduce new risks of bias, hallucination, and manipulation. Confidentiality concerns add a further constraint: if GPs succeed in blocking the use of LLMs on their documents, much of the above discussion is void. At the same time, training models on structured datasets or internal archives is a long-established practice. This means that the benefits of AI will depend not only on technical feasibility but also on the evolving contractual and institutional environment.

4. Academic Literature

4.1 Text-as-Data in Public Markets

The public equity literature provides a clear demonstration of how textual information can enrich financial analysis. Public firms generate frequent and relatively standardized disclosures that can be systematically linked to high-frequency market outcomes such as stock returns, volatility, and analyst forecasts. This rich data environment enables reproducible and scalable contributions that apply natural language processing and machine learning techniques to financial texts.

A foundational contribution is Loughran and McDonald (2011), who construct a finance-specific sentiment lexicon and show that tone in 10-K filings is systematically related to stock returns, volatility, and litigation risk. Their work established that textual disclosures contain latent information that is not captured by conventional numerical variables. Building on this insight, Huang, Teoh, and Zhang (2014) analyze over 18,000 earnings press releases and develop an "abnormal positive tone" measure. They show that firms using unusually promotional language tend to experience subsequent earnings disappointments, higher restatement risk, and stock price reversals. These findings show how textual analysis can uncover patterns of strategic disclosure.

Another influential strand of the literature leverages conference call transcripts. Hassan, Hollander, van Lent, and Tahoun (2019) introduce measures of firm-level political risk based on the frequency and context of risk-related terms in earnings calls, demonstrating predictive power for investment, hiring, and risk premia. Subsequent work extends this approach to other shocks: Hassan, Hollander, van Lent, and Tahoun (2023) analyze firm-level exposure to epidemic diseases, while Hassan, Schreger, Schwedeler, and Tahoun (2025) study the transmission of country-level risks such as Brexit and the diffusion of emerging technologies. These studies show how text can be used not only to measure sentiment but also to quantify exposures, distinguish between bad news and uncertainty, and trace the transmission of shocks across firms, industries, and geographies.

Gentzkow, Kelly, and Taddy (2019) provide a comprehensive survey in the *Journal of Economic Literature* of the "text as data" paradigm, situating financial applications within the broader economics literature. They emphasize how digitization and computational methods allow text to be systematically transformed into measures of sentiment, expectations, and risk that complement traditional datasets.

More recent advances center on large language models. For example, Cao, Jiang, Yang, and Zhang (2023) show that firms strategically adapt their disclosure language in earnings calls to account for algorithmic trading systems that monitor sentiment. Kim, Muhn, and Nikolaev (2023) study whether generative AI can improve the processing of lengthy corporate disclosures, finding that LLMs assist investors in digesting complex information. These contributions highlight both

opportunities (greater nuance and richer thematic analysis) and challenges around interpretability, reproducibility, and bias.

Taken together, this literature establishes that textual disclosures in public markets contain systematic and decision-relevant information beyond what is available in numerical data. Advances in NLP and ML provide tools to extract forward-looking indicators that enhance forecasting and risk management. When labeled data exist, supervised text regression generally performs better than dictionary indices, and careful cross-validation and hold-out testing are essential for honest assessment of model fit (Gentzkow, Kelly, & Taddy, 2019).

4.2 Applications in Private Markets

4.2.1 Textual Analysis of Fund Documents

A first strand of applications adapts text-as-data methods to the private equity context by analyzing the core documents exchanged between general partners and limited partners. Unlike the standardized 10-Ks or earnings calls of public firms, these documents—private placement memoranda (PPMs), limited partnership agreements (LPAs), due diligence questionnaires, quarterly letters, and interim GP reports—are unstandardized, narrative-heavy, and produced with strategic intent. They are nonetheless the richest sources of information available to LPs and have recently been studied with modern computational methods.

Braun, Fernández Tamayo, López-de-Silanes, Phalippou, and Sigrist (2023) assemble a novel dataset of over 1,000 PPMs and examine whether their textual features can predict eventual fund performance. They show that document length, emphasis on themes such as “risk” or “value creation,” and the overall structure of disclosures are correlated with realized outcomes. Using supervised learning methods, including random forests and gradient boosting, they demonstrate meaningful out-of-sample classification accuracy. Importantly, the strongest signals derive not from isolated keywords but from broader disclosure patterns, suggesting that narrative framing encodes latent information about fund quality.

Braun, Fernández Tamayo, López-de-Silanes, Phalippou, and Sigrist (2025) extend this approach to GP interim reports, assembling more than 25,000 reports covering nearly 20,000 deals in Europe and North America. They analyze over 600,000 sentences of qualitative commentary, measuring tone with FinBERT and related large language models. They show that report tone is a robust predictor of final multiples on invested capital (MOIC), even after controlling for interim NAVs and deal characteristics. An increase in tone sentiment translates into significantly higher subsequent MOIC uplift. Tone remains predictive across the life of investments but is especially informative during fundraising periods, when interim NAVs lose predictive power but textual narratives retain it. The authors also document regional differences: interim valuations correlate positively with outcomes in Europe but negatively in North America, whereas tone predicts higher returns in both regions.

They further apply machine learning algorithms directly to these reports and show that Lasso and Random Forest models achieve high predictive accuracy (ROC–AUC around 0.76). Deals predicted to outperform deliver exit MOICs more than double those of predicted underperformers. These findings highlight that interim GP reports, which have often been treated as routine

communication, in fact contain systematically valuable forward-looking information embedded in their qualitative narratives.

Together, these two papers demonstrate that private equity disclosures encode predictive content not only in quantitative metrics but also in narrative choices. PPMs at the fundraising stage and GP reports during the life of a fund both provide textual signals that can be extracted with modern NLP and ML methods. This line of research moves beyond anecdotal impressions of disclosure "style" and offers the first systematic evidence that qualitative GP narratives contain forward-looking information relevant for fund and deal outcomes.

4.2.2 Machine Learning with Structured Fund Characteristics

A second strand of research applies machine learning not to text, but to structured fund and manager attributes. These models use observable characteristics such as fund size, vintage year, sector focus, GP experience, and macroeconomic conditions at inception to classify funds ex ante or to predict outperformance.

Empirical applications in private equity remain limited. In their master's thesis, Pachnanda and Raj apply methods such as random forests and gradient boosting to fund-level variables and report some predictive power for performance classification. While suggestive, this evidence is preliminary and should be interpreted with caution.

In contrast, in public markets, similar challenges of noisy and high-dimensional data have been studied extensively. Kelly, Malamud, and Zhou (2024) show that complex machine learning models can outperform simpler regressions in predicting returns, even when the number of observations is modest. Bryzgalova, Pelger, and Zhu (2025) demonstrate that random forests can uncover economically meaningful clusters in noisy stock return data. Finally, in their survey, Gentzkow, Kelly, & Taddy (2019) highlight that in settings with many covariates relative to observations penalized linear models provide strong baselines, while tree-based approaches benefit from prior dimension reduction and careful cross-validation to control overfitting.

4.2.3 Forecasting Fund Cash Flows and Liquidity

A third set of studies applies machine learning techniques to the forecasting of private equity fund cash flows. This is a problem of first-order importance for LPs, given that private market funds operate under a commit-and-call structure. Capital is committed upfront but drawn down and distributed at the discretion of the GP, creating uncertainty for treasury planning, asset allocation, and risk management. The difficulty is compounded by the fact that commitments are typically large and illiquid, and redemption features are increasingly offered to end-investors through semi-liquid products.

Traditional approaches, such as the Yale pacing model, rely on deterministic curves calibrated from historical averages of calls and distributions. These rule-based models cannot adapt to changing macroeconomic regimes, market shocks, or innovations in fund structures. They therefore perform poorly during stress episodes, such as the Global Financial Crisis, the COVID-19 pandemic, or the recent rise in interest rates.

Karatas, Klinkert, and Hirsa (2021) introduce deep learning methods to this domain, training recurrent neural networks, long short-term memory networks, and gated recurrent unit models on

irregular fund-level cash flow series. They augment these with macroeconomic variables such as interest rates, volatility indices, and credit spreads. Their models deliver improved out-of-sample accuracy compared to rule-based pacing models, particularly when forecasting the timing of capital calls and distributions. The authors employ a sliding-window technique to mitigate the slow feedback problem inherent in PE data. Nevertheless, the predictive horizon remains limited.

While deep learning methods show potential, their performance is highly sensitive to architecture, hyper-parameter tuning, and data volume. In public equity markets, evidence is mixed: in some cases complex models add predictive power, but in others simpler approaches perform comparably (Kelly, Malamud & Zhou, 2024). Non-deep methods such as random forests or boosting have proven effective at extracting structure from noisy datasets (Bryzgalova, Pelger & Zhu, 2025), suggesting that they may be better aligned with the sparse and heterogeneous data of private markets. Whether deep architectures can ultimately deliver an edge in this environment remains an open question.

In private markets, the case for embedding economic structure into predictive models is especially compelling. Purely statistical approaches may capture correlations, but they risk missing the mechanisms that drive it. For example, Franzoni, Nowak, and Phalippou (2012) show that buyout returns are shaped by debt market conditions, in particular refinancing costs. Not only the statistical link is strong but there is a natural economic reason for it. A data-driven model without economic context would probably struggle to pick up that exact factor. Ignoring these structural drivers reduces explanatory power and can lead to spurious inferences, particularly in settings with long horizons and shifting market regimes. A more promising path is therefore hybrid: using machine learning to detect patterns and improve short-term forecasts, while anchoring models in economic theories of leverage, incentives, and contractual design. Such integration not only enhances interpretability but also increases the likelihood that results remain relevant across vintages and evolving market structures.

For LPs, the practical implication is that forecasting tools are not replacements for domain expertise but as complements. Machine learning models may improve the accuracy of short-term cash flow projections, which can aid liquidity planning and stress testing. Yet their output must be interpreted in light of the economic mechanisms that drive capital calls and distributions. Without this grounding, ML risks becoming "good at prediction but bad at explanation," offering statistical lift without addressing the causal drivers that LPs and regulators need to understand.

4.2.4 Benchmarking

A fourth strand of applications focuses on benchmarking. Two examples illustrate the potential: (i) clustering funds into robust "asset classes" and (ii) constructing liquid proxies that allow LPs to manage allocations in real time.

For clustering, in public equity Bryzgalova, Pelger, and Zhu (2025) use random forests to group stocks into economically meaningful categories, showing how machine learning can uncover structure in noisy return data. Translating these approaches to private markets could yield classifications that better capture true economic exposures. One attempt by Goetzmann, Gourier, and Phalippou (2022) applies standard econometric methods to illiquid assets and finds that the resulting groupings often diverge from traditional classifications.

Large language models can also complement this process by standardizing inputs. Fund labels are often inconsistent. For example, one vehicle may be described as "growth equity" and another as "late-stage venture," despite overlapping strategies. LLMs can detect such semantic similarity and assign consistent categories, improving the accuracy and interpretability of clustering. Yet this process is not risk-free: without clear rules, classifications remain a black box, raising concerns that clusters reflect arbitrary distinctions rather than economically meaningful ones.

For liquid proxies, Phalippou (2023) constructs a listed private equity index by using NLP on global news to identify and weight companies most exposed to private equity activity. Despite not being optimized for correlation, the resulting index tracks private-fund benchmarks closely, though with higher volatility. Such tools offer LPs a hedge against denominator effects, a tactical pacing instrument, and a transparent benchmark for reporting.

4.3. Research Agenda

Textual analysis of prospectuses and GP reports reveals that narrative choices encode forward-looking information even in the absence of standardized reporting. ML models applied to fund attributes can deliver statistical lift over traditional regressions, while forecasting exercises suggest that deep learning architectures improve liquidity and pacing projections. Benchmarking applications can expose when GP marketing categories diverge from actual economic exposures. These contributions are substantive, yet this is just infancy.

The most fundamental limitation is the problem of delayed feedback. Unlike public markets, where outcomes are continuously observable, the ultimate performance of a private equity fund is known only after a decade or more. This slows model refinement and increases the risk of overfitting to historical regimes that may not generalize to current market structures. Closely related is the issue of data quality. The main repositories are proprietary, and reporting conventions vary widely across funds and geographies.

Interpretability represents a further challenge. LPs operate under fiduciary and reputational obligations that make reliance on black-box models difficult. For AI/ML methods to become actionable, they must either be interpretable in themselves or embedded in governance frameworks that provide human oversight and accountability. In practice, two guardrails are essential: out-of-sample validation (hold-out testing and cross-validation) and manual audits of small text subsamples to verify that machine-generated labels match the economic concept. For cases that require inference after variable or model selection, sample-splitting or bootstrap procedures help quantify uncertainty in a transparent way (Gentzkow, Kelly, & Taddy, 2019).

The research agenda is also about scope. A deep contribution would be to identify causal mechanisms: for example, whether contractual provisions mitigate risk, or how narratives affect fundraising dynamics. Embedding causal inference within AI frameworks would broaden their relevance.

Finally, important aspects of LP practice remain underexplored. Side letters, which define bespoke investor rights, are central to governance yet absent from empirical work. Secondary transactions, now a large and complex market, offer opportunities to develop valuation models that integrate GP reporting with market pricing. ESG disclosures, which are increasingly material to investors, remain fragmented and voluntary, making them an ideal domain for systematic text-based analysis.

5. The Rise of Commercial Tools for LPs: Mapping the Emerging Ecosystem

Paradoxically, it is in practice rather than in academia that the most concrete efforts to apply artificial intelligence (AI) and machine learning (ML) to LP decision-making have emerged. This may reflect the relative ease of announcing a technology product compared to the time-intensive process of producing peer-reviewed research. In commercial settings, a product can be launched with a press release and a website; academic contributions, by contrast, must meet evidentiary standards and withstand scrutiny. Moreover, there is no formal validation framework for these tools; no equivalent of randomized trials or reproducibility protocols. As a result, while many offerings appear sophisticated and are marketed as such, it remains difficult to evaluate their actual capabilities or reliability. This may help explain why academic contributions often seem cautious by comparison.

Several firms target document processing and workflow automation. Chronograph, for instance, offers a platform for portfolio monitoring and data aggregation. In 2024, it introduced a module labeled "Chrono AI," which applies natural language processing (NLP) techniques to fund documents such as quarterly reports, AGM materials, and capital calls. According to public descriptions, the system extracts recurring metrics and summary information, with a user-facing interface that enables human validation. The underlying architecture is not disclosed but appears to combine large language models (LLMs) with structured data infrastructure. The company claims adoption among institutional clients who use the tool to extract millions of data points from GP documents.

Covenant, another vendor, specializes in the analysis of Limited Partnership Agreements (LPAs). Its platform identifies legal and economic terms—such as fee structures, key person clauses, and "most favored nation" provisions—from typically unstructured legal texts. The firm indicates that it uses a combination of LLMs and custom parsing algorithms, with legal experts involved in validating outputs. One stated use case is the facilitation of side letter negotiations and MFN elections during fund closings. However, no information is available on classification performance, accuracy rates, or the extent of human post-processing required. As with many vendors, it is unclear whether the tool relies on statistically trained models or deterministic extraction logic.

Several newer entrants focus on the secondary market. Clipway, LiquidLP, and AltConvey position themselves as analytics platforms supporting the pricing and execution of LP-led and GP-led secondary transactions. These tools aim to synthesize data from GP reports and market sources to support valuation processes. While they are frequently described as offering "AI-enhanced underwriting" or "real-time analytics," the actual mechanisms remain opaque. It is not clear whether these systems implement predictive modeling based on historical pricing data or simply facilitate standardized presentation of GP-reported NAVs.

Two startups (GovernGPT.ai and iAltA) reportedly apply LLMs to GP communications, including quarterly letters and regulatory filings. Their goal is to identify shifts in tone, highlight governance-relevant language, and flag inconsistencies over time. The conceptual premise aligns with academic studies that relate narrative sentiment to investment risk. However, these firms have not disclosed model architectures, training data, or performance benchmarks. It remains

unknown whether the tools are based on pre-trained general-purpose LLMs or have been fine-tuned on financial documents from private markets.

Related functionality is offered by Ontra and Tap, both of which focus on automating legal workflows such as NDAs and side letters. These tools are said to use machine learning to classify clauses and extract key obligations. In theory, such tools could reduce the burden of legal review, especially in high-volume workflows. But again, without clarity on classification accuracy, false positives, or model robustness, it is difficult to assess their practical reliability especially in edge cases or non-standard documents.

Some platforms focus on transforming unstructured GP reporting into structured dashboards. Canoe Intelligence, for example, processes capital calls, valuation updates, and distribution notices to produce structured portfolio data. It is not publicly documented whether Canoe uses adaptive learning techniques or rule-based parsing. Allocator offers a similar interface for multi-manager portfolio aggregation, but its technical implementation is not disclosed.

73 Strings is a more recent entrant focused on valuation and document analysis for tail-end positions, including continuation vehicles. The company advertises AI-based valuation support, though it has not provided public information on model calibration, input assumptions, or validation accuracy. Given the increasing reliance on internal NAVs in secondary pricing, tools in this category may eventually play a critical role, but current opacity limits external evaluation.

Collectively, these tools illustrate growing interest in computational approaches to LP challenges. Across vendors, a common theme is the move from manual processes to partially automated workflows—especially in document parsing, compliance tracking, and portfolio monitoring. However, it is often difficult to separate substantive ML applications from marketing rhetoric. Some vendors rely on rule-based extraction systems, while others appear to use LLMs primarily as summarization or interface tools. Few, if any, have published performance data, academic validations, or third-party audits.

Unlike academic models, which are typically evaluated on holdout samples or out-of-sample vintages, commercial systems remain proprietary and are assessed internally or via client feedback. As a result, their effectiveness remains difficult to evaluate objectively. Nonetheless, their growing adoption, particularly among large-scale LPs with complex portfolios, suggests demand for more scalable and structured decision-support tools.

Whether these tools improve investment decisions, reduce errors, or simply repackage existing workflows with more sophisticated interfaces remains an open question. What is clear is that “LP tech” has emerged as a commercial category. Dedicated industry events, such as the LP Tech Summit (Toronto, October 2025), point to increasing institutional interest. But whether these developments reflect genuine innovation or a temporary wave of technology-driven enthusiasm will depend on future evidence. For now, in the absence of standardized performance metrics or external validation, the practical value of these tools remains a matter of belief rather than proof.

6. Conclusion

This paper examines the potential for artificial intelligence and machine learning to reshape LP decision-making. The core tension is clear: LPs operate with opaque disclosures, delayed feedback, and heterogeneous data, while modern algorithms are designed for abundant, high-frequency signals. This tension defines both the promise and the limits of technological adoption in private markets.

Recent research highlights genuine advances. Textual analysis of GP disclosures shows that narratives contain systematic and forward-looking signals. Machine learning methods can improve performance prediction and liquidity forecasting. Benchmarking applications illustrate how computational tools can align classifications with true economic exposures.

Yet structural constraints remain binding. Outcomes arrive only after a decade, reporting conventions are inconsistent, and models trained on past vintages risk overfitting. Large language models excel at standardization but cannot on their own link disclosures to ultimate outcomes, and their classifications often remain black boxes. Confidentiality restrictions further complicate the use of unstructured GP documents, potentially curtailing some of the most promising applications.

The lesson for academics is that the frontier lies less in chasing predictive accuracy than in embedding causal structure, interpretability, and governance into model design. For practitioners, the most effective use of technology is as a complement rather than a substitute for human expertise—tools that reduce the cost of parsing disclosures, highlight anomalies, and improve comparability, but always with oversight and validation.

Ultimately, the juxtaposition of limited partners and unlimited technologies underscores a broader theme: innovation in private markets is shaped as much by institutional constraints as by technical feasibility. The future of LP investing depends on partnerships—between investors and managers, between algorithms and human judgment, and between academic inquiry and practice.

References

Barber, Brad M., and Ayako Yasuda, 2017, *Interim fund performance and fundraising in private equity*, Journal of Financial Economics 124, 172–194.

Braun, Reiner, Borja Fernández-Tamayo, Florencio López-de-Silanes, Ludovic Phalippou, and Natalia Sigrist, 2023, *Limited Partners versus Unlimited Machines: Artificial Intelligence and the Performance of Private Equity Funds*, SSRN Working Paper No. 4490991.

Braun, Reiner, Borja Fernández-Tamayo, Florencio López-de-Silanes, Ludovic Phalippou, and Natalia Sigrist, 2025, *Would I Lie to You? On Private Equity Intermediary Reports*, SSRN Working Paper No. 5322805.

Brown, Gregory W., Oleg R. Gredil, and Steven N. Kaplan, 2019, *Do private equity funds manipulate reported returns?*, Journal of Financial Economics 132, 267–297.

Bryzgalova, Svetlana, Markus Pelger, and Jason Zhu, 2025, *Forest through the Trees: Building Cross-Sections of Stock Returns*, Journal of Finance, forthcoming.

Cao, Shiyang, Wei Jiang, Baozhong Yang, and Alan L. Zhang, 2023, *How to Talk When Machines Are Listening: Corporate Disclosure in the Age of AI*, Review of Financial Studies 36, 3603–3642.

Da Rin, Marco, and Ludovic Phalippou, 2017, *Investor Experience and Fundraising in Private Equity*, Journal of Financial Economics 124, 543–564.

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

Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019, *Text as Data*, Journal of Economic Literature 57, 535–574.

Goetzmann, William N., Elise Gourier, and Ludovic Phalippou, 2022, *How Alternative Are Private Markets?*, Working Paper.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, 2019, *Firm-level Political Risk: Measurement and Effects*, Quarterly Journal of Economics 134, 2135–2202.

Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, 2023, *Firm-level Exposure to Epidemic Diseases: COVID-19, SARS, and H1N1*, Review of Financial Studies 36, 2409–2449.

Hassan, Tarek A., Jesse Schreger, Markus Schwedeler, and Ahmed Tahoun, 2025, *Sources and Transmission of Country Risk*, Review of Economic Studies 91, 2307–2346.

Hoberg, Gerard, and Gordon M. Phillips, 2010, *Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis*, Review of Financial Studies 23, 3773–3811.

Huang, Alan S., Siew Hong Teoh, and Yinglei Zhang, 2014, *Tone Management*, Accounting Review 89, 1083–1113.

Karatas, İhsan, Pascal Klinkert, and Ali Hirsa, 2021, *Deep Learning for Cash Flow Prediction in Private Equity*.

Kelly, Bryan T., Semyon Malamud, and Kangying Zhou, 2024, *The Virtue of Complexity in Return Prediction*, Journal of Finance 79, 459–503.

Kelly, Bryan T., Petr Kuznetsov, Semyon Malamud, and Hao Xu, 2025, *Artificial Intelligence Asset Pricing Models*, NBER Working Paper No. 33351.

Kim, Andrew G., Maximilian Muhn, and Valeri V. Nikolaev, 2023, *Bloated Disclosures: Can ChatGPT Help Investors Process Information?*, BFI Working Paper No. 23-07.

Loughran, Tim, and Bill McDonald, 2011, *When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-K Filings*, Journal of Finance 66, 35–65.

Phalippou, Ludovic, 2023, Thematic investing with big data: The case of private equity, *Financial Analysts Journal* 79, 30–40.