Miller Center for Global Impact

THE TRUE COST OF IMPACT-FIRST INVESTING

A Comparative Analysis of the Economics and Subsidy Structure of Impact-First Investing

June 16, 2026

Acumen

Global Partnerships

Halcyon

kiva

mce Social Capital

Santa Clara University Miller Center for Global Impact

Open Road Impact

village capital

AUTHORS

THE TRUE COST OF IMPACT-FIRST INVESTING

Lead Author

Brigit Helms, Executive Director, Miller Center for Global Impact

Co-Authors

Daniel Barker, President & CEO, Halcyon

Caroline Bressan, CEO, Open Road Alliance

Kathy Guis, Vice President, Investments, Kiva

Heather Matranga, Managing Director, Impact Investments, Village Capital

Diana McKeage, Impact Investing, Miller Center for Global Impact

Tara Murphy Forde, Chief Capital & Impact Officer, Global Partnerships

Camilla Nestor, CEO, MCE Social Capital

Daniel Waldron, Director of Insights, Acumen

© 2026 Miller Center for Global Impact. All rights reserved.

TABLE OF CONTENTS

THE TRUE COST OF IMPACT-FIRST INVESTING

EXECUTIVE SUMMARY

Impact-first investing exists to move capital where markets systematically do not: to early-stage, high-impact social enterprises that require smaller, more flexible, and more hands-on financing than traditional capital markets are designed to provide. Drawing on a collaborative cost study of eight leading impact-first investment organizations, representing $197.2M in assets under management across 362 deals, this paper examines the true cost of meaningful impact-first investing and why that cost is structurally underfunded.

The findings suggest that impact-first funds are not inefficient. In this sample, they completed more deals per year than comparable traditional funds and operated at a lower average cost per deal. The real difference emerges when costs are measured per dollar deployed. Because impact-first funds often make smaller, higher-touch investments, they spend more to move each dollar of capital. Under conventional financial logic, avoiding these deals is rational. Under an impact-first logic, however, this additional cost is the cost of inclusion: the cost of carrying capital, support, and risk to enterprises and communities the market otherwise excludes.

This paper argues that reliable, market-making subsidy is not a patch for weak fund performance, but a necessary ingredient for building inclusive markets. When structured through impact-first funds, philanthropic and catalytic capital can multiply its effect by de-risking enterprises, recycling capital, and enabling durable social and environmental outcomes beyond a single grant cycle. Closing the subsidy gap at the fund level is therefore one of the most direct ways to address the enterprise-level "Valley of Death" and to reclaim impact investing for the purpose it was meant to serve.

INTRODUCTION

"Open Road Impact has saved us twice, related to investors that screwed us over and funds that didn't come through at the last minute. We were completely out of cash, and Open Road bridged us. We were already pretty mature and had large operations, but difficult circumstances happened to us: war and regulations."

– Galen Welsch, Jibu

Impact investing cohered as a practice in the mid-2000s based on the principle that capital could be used to generate positive impact for people and the planet. Twenty years later, leaders within the industry are raising important questions about whether the field is meeting its own goals. Studies have shown that sufficient capital is still not reaching those who need it most and that key activities remain underfunded or misunderstood.

At the same time, the impact investing industry largely operates using return expectations inherited from traditional finance. As of 2025, 84% of impact investors target market-rate returns. Yet, it is becoming clear that these inherited assumptions are structurally incompatible with the inclusive, long-term change that impact investing seeks to support. In response, a small community of "impact-first" investors has emerged to re-center the sector on social and environmental impact.

By analyzing the true cost of impact-first investing, we can understand why the most impactful work often remains unfunded and what must shift to change that. Drawing on a cost-analysis exercise that represents the collaborative efforts of eight leading impact-first investment fund managers, this paper offers new data and language to illuminate the true cost of doing impact-first investing properly, with implications for the future of impact investing as a whole.

CONTEXT

The collaboration emerged from shared frustration. After decades of impact investing, each of these organizations still witnesses social enterprises that are successfully accomplishing the highest-impact work struggling to access capital. At the enterprise-level, many of these social enterprises fall into a persistent early-stage funding gap, known as the "Valley of Death," where they require fit-for-purpose capital to move from early-stage to growth-stage, and that capital is structurally unavailable.[1]

Impact-first investing (as opposed to finance-first investing) connotes an investment mandate that prioritizes the achievement of impact in pursuit of financial return.[2] Many impact-first investors—including the authors of this paper—have strong track records of recouping invested capital, coupled with generating outsized social or environmental impact. These funds, however, exist within a financial system organized around the pursuit of market-rate returns, creating a structural misalignment between risk, return, and impact. Impact-first funds have developed a range of bold, innovative, and effective solutions to bridge the financing gap for social enterprises.

While impact-first investing may have the potential to meet the need, the funds that can deliver this capital cannot supply it at the levels required because they provide below market-rate returns.[3] Impact-first funds require reliable subsidy because they absorb uncompensated support costs required to deploy high-touch, affordable capital.[4]

We hypothesize that the enterprise-level "Valley of Death" is driven at least in part by a subsidy gap: impact-first funds lack the reliable subsidy they need because they are capitalized and evaluated against market-rate vehicles. This paper, therefore, focuses on the true cost of impact-first investing in order to illuminate how the fund-level subsidy gap constrains capital for enterprises facing early-stage funding shortfalls. In doing so, we argue that closing this gap is a high-value use of philanthropic capital.

METHODS

Over six months, eight impact-first investment organizations conducted a collaborative cost study representing a combined 170 years of institutional experience.[5] The participating organizations included: Acumen, Global Partnerships, Halcyon, Kiva, MCE Social Capital, Miller Center for Global Impact, Open Road Impact, and Village Capital. The group mutually developed a data-sharing agreement and a data collection template. Each organization shared cost, capitalization, transaction, and impact data regarding one fund.[6] Findings were debated and collectively developed during a series of interactive workshop sessions.

In aggregate, the funds represent $197.2M in AUM deployed across 362 deals, with ticket sizes ranging from $7K to $4M (average: $571K). Five were debt funds, and three were equity. Individual fund sizes ranged from $722K to $78.8M AUM, with the number of closed transactions (deals) ranging from 10 to 108 per fund, with an average of 45.

To benchmark against traditional finance, we constructed a comparison dataset using Pitchbook, filtering for private debt and VC funds ($10M–$100M AUM, 2015–2020, fully invested). From several thousand entries, we retained approximately 600 funds and further narrowed them to 79 entries that contained the required data for comparison. We assumed a 2% management fee and full deployment over 4–5 years to estimate cost per deal and cost per dollar deployed.

To incorporate the voice of the entrepreneur, we conducted in-depth qualitative interviews with three high-performing social enterprises representing diverse sectors and geographies.

Traditional Fund Benchmark Assumptions

	All funds	Debt funds	Equity funds
Avg annual transactions (closed deals)	6.4	1.8	6.8
Investment period (years)	4.8	4.0	4.9
Fund life (years)	8.4	-	8.4
Avg AUM (USD millions)	40.3	58.2	38.6
Approx. total deals per fund	~ 31	~ 7	~ 33
Approx. avg ticket size (USD millions)	~ 1.3	~ 8.1	~ 1.2
Number of funds	79	7	72

Impact-first funds are efficient

Is the lower return potential of impact-first funds caused by inefficient fund operations? To explore this question, we compared impact-first funds against conventional finance fund benchmarks on basic measures of operational and cost efficiency. We used cost data to illustrate the challenges of serving customers that are not served by traditional financial markets.

To examine operational efficiency, we asked how much deal activity gets done each year across different fund types. In our benchmark sample of traditional funds with under $100M AUM, the average was 6.4 deals per year.[7] By comparison, the impact-first funds participating in the study averaged over 10 deals per year. The data indicate that the impact-first funds were operationally competitive with traditional funds in terms of deal volume.

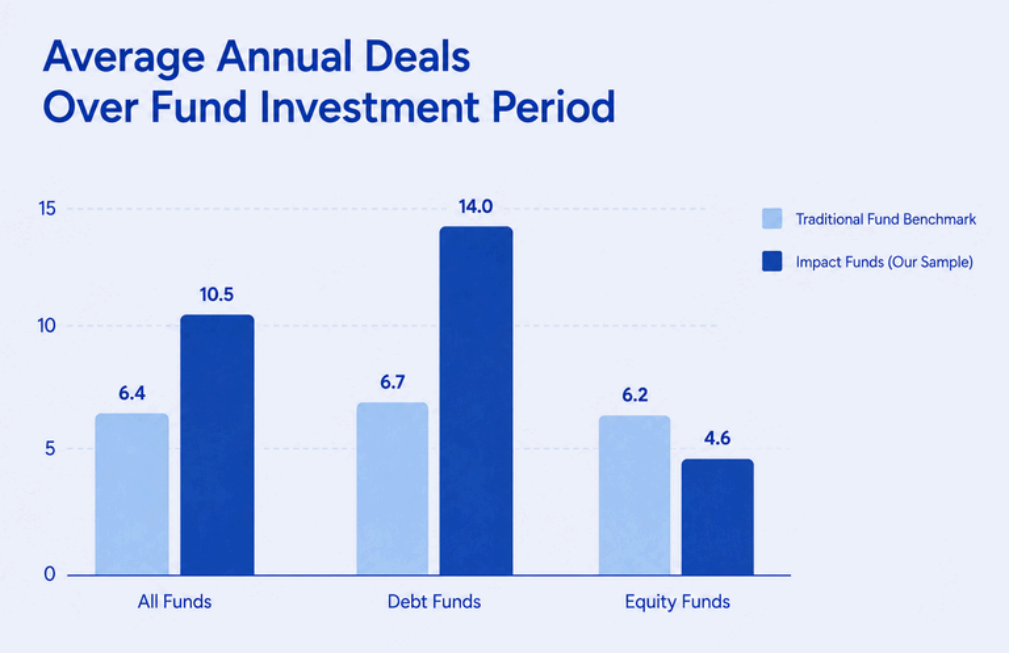

Do impact-first funds cost more to run on a per-deal basis? We next examined cost efficiency by comparing traditional finance with impact-first funds, followed by a subgroup analysis within the impact fund sample. Comparing the average cost to close a single deal across traditional and impact-first funds, we found that funds averaged $125,000 per deal in our benchmark sample. The impact-first funds in our sample, by contrast, averaged approximately $80,000 per deal. We concluded that the impact-first funds performed competitively with traditional funds on cost per deal.

CORE INSIGHTS

Finally, we explored the structural cost dynamics on a per-deal level using a subgroup analysis on the impact-first funds. We wanted to see whether deal size made a difference: specifically, whether larger ticket sizes were more cost-efficient.

Statistically, no relationship was observed between deal size and cost per deal (R^2 = 0.01). This finding confirmed what many of the study participants already understood intuitively: that delivering an impact-first deal usually costs a comparable amount in the tens of thousands of dollars as it does in the hundreds of thousands.

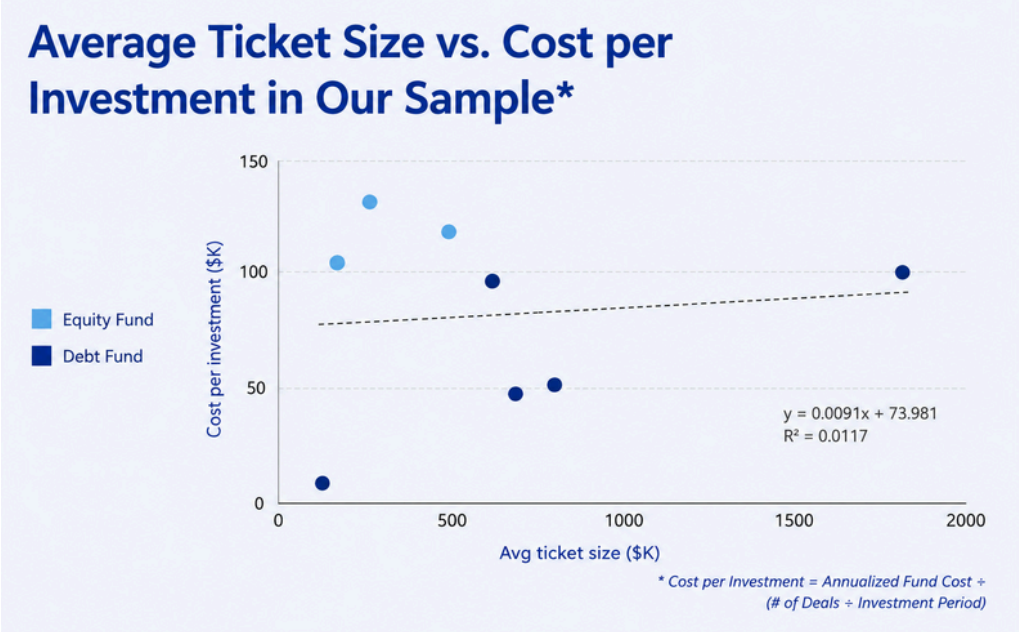

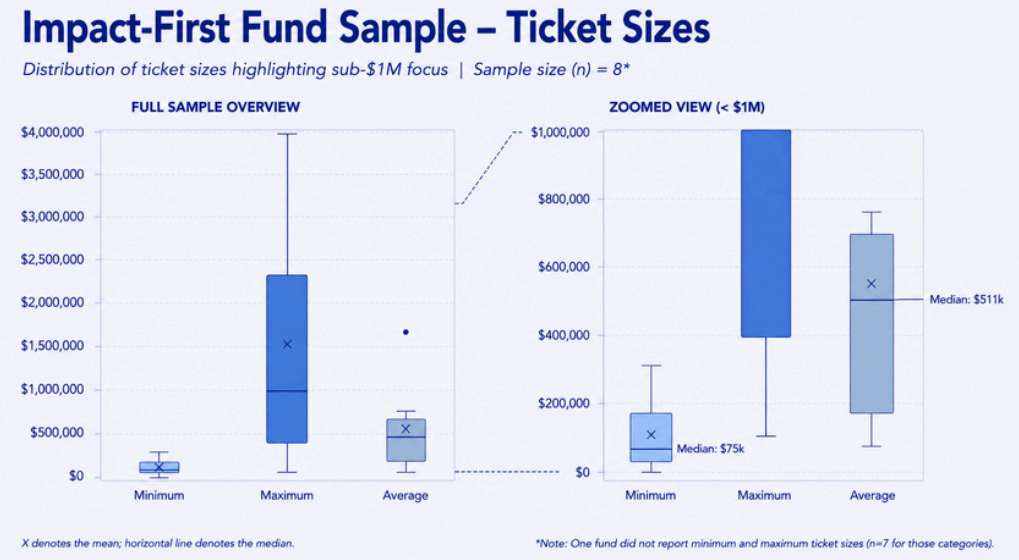

The Valley of Death is a "rational decision"

At this point, we had not uncovered evidence that impact-first funds operated less efficiently than what the traditional market can offer. So we looked at the cost per dollar deployed: how much does it cost for a capital provider to move one dollar to a capital user? In our sample, it costs impact-first funds an average of $0.21 to deploy $1. Traditional funds' cost per dollar came in closer to $0.08, about one-third of the cost for impact-first funds.[8] In this study, the premium for inclusion (including defying traditional scale economics) appears as approximately 13 cents on the dollar.

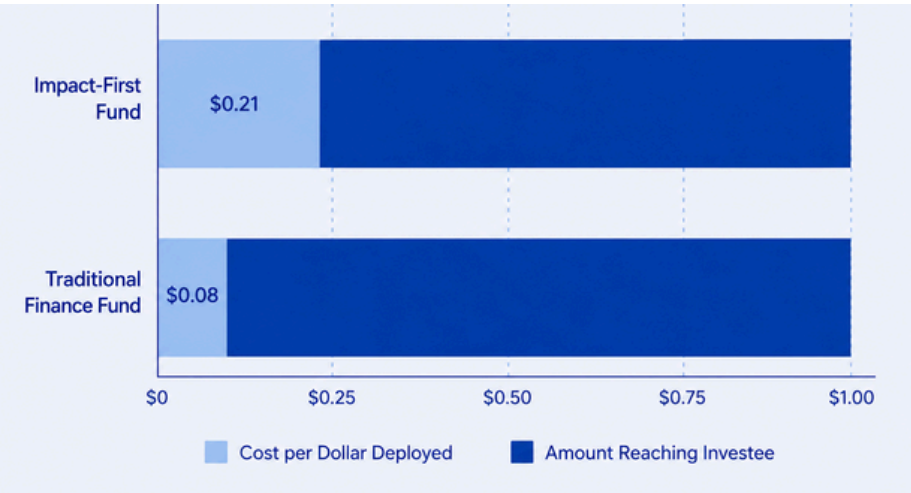

On the surface, this result might be seen to support the claim that impact-first funds are less efficient. Below the surface, we must ask what efficiency even means in this context. Impact-first deals are smaller, more complex, and require more hands-on support than traditional finance deals. Traditional funds appear more cost-efficient per dollar by avoiding the uncompensated value creation that impact funds take on.

For illustration, imagine that a traditional fund spends $100,000 to place $1.3 million with a well-connected urban business that has audited accounts and collateral. This transaction costs approximately 8 cents to move each dollar. Now imagine an impact-first fund spending $80,000 to structure a $380,000 investment for a rural water enterprise: field visits, translation, business coaching, meetings to design fit-for-purpose capital. In the financial model, the second loan costs roughly 13 cents more per dollar; on the ground, that difference is carrying capital to people the market does not reach. In this context, delivering deep impact for 13 cents on the dollar could be considered excellent value for money.

CORE INSIGHTS

Having established that impact-first funds pay more per dollar deployed, we wanted to unpack what might be driving that difference. To test whether deal size played a role in the impact-first funds' higher costs, we conducted a subgroup analysis within the impact-first funds, plotting deal size against cost per dollar deployed. Here we found a weak but observable inverse relationship ($R^2 = 0.28$): smaller deals cost more per dollar deployed.[9]

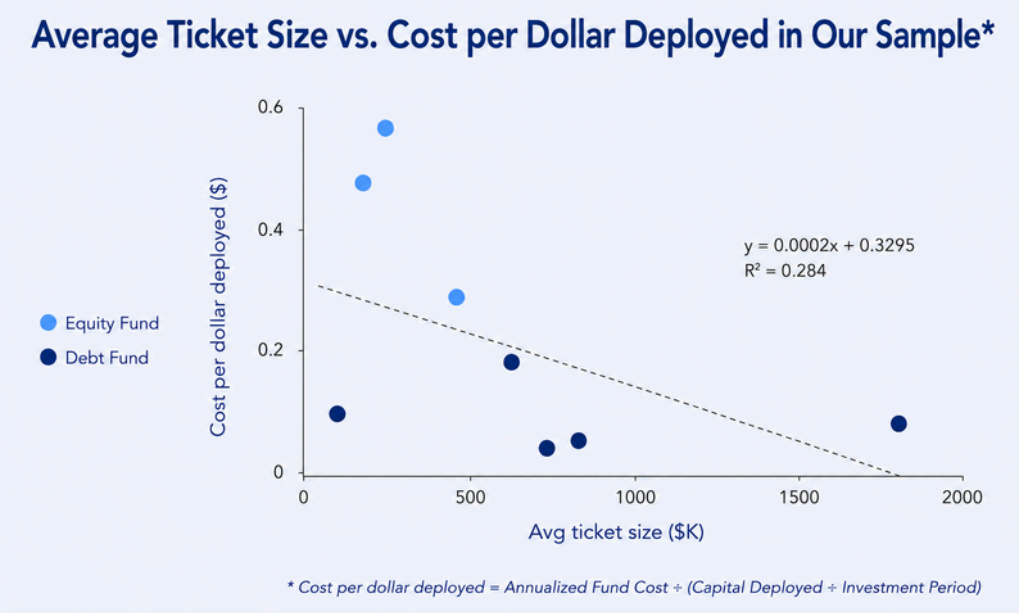

With these fund economics, the rational financial decision for a fund manager would be to avoid smaller deals, where one spends more to move less capital. From a mainstream fiduciary perspective, it is not rational to invest in a fund that deliberately specializes in these smaller, high-touch investments with structurally higher cost per dollar deployed.

This fund-level disincentive is one of the clearest explanations for why the Valley of Death persists at the enterprise level. Early-stage, high-impact enterprises tend to exist in this range of small-ticket sizes, requiring smaller, bespoke capital that is expensive to originate. Using a cost-per-dollar lens and a mainstream interpretation of fiduciary duty, these are usually the first deals to be screened out.

While many factors contribute to early-stage capital scarcity, our findings suggest that under prevailing fund economics, smaller deals are structurally disincentivized. This finding does not prove causality; instead, it makes the underlying mechanism visible. When efficiency is defined as cost per dollar, capital predictably flows toward larger, lower-friction deals and away from the smaller, high-touch transactions usually required by high-impact work. Impact-first funds are among the few models explicitly designed to carry those higher costs while keeping capital affordable for social entrepreneurs.

Market-making subsidy is a force multiplier

Impact-first funds can only carry higher costs while keeping capital affordable for the markets they serve because subsidy is intentionally baked into the fund's capitalization.

The left panel shows the full range of fund capitalizations up to $78.8M. The right panel is a subset of the left panel, with the y-axis maximum at $7M to illuminate the capitalization of smaller funds.

These summary statistics indicate that non-commercial capital represents a substantial share of most impact-first fund capitalizations. In this chart, "non-commercial" refers to capital that carries an intentional subsidy relative to market terms, i.e., grants, first-loss or deeply concessional capital, or other catalytic funding that does not seek a market-rate financial return. This subsidized layer often comprises half or more of the impact funds' AUM, and in some cases, all of it. In our sample, market-making subsidy is therefore not a marginal addition, but rather a defining feature of impact-first fund structures.[10]

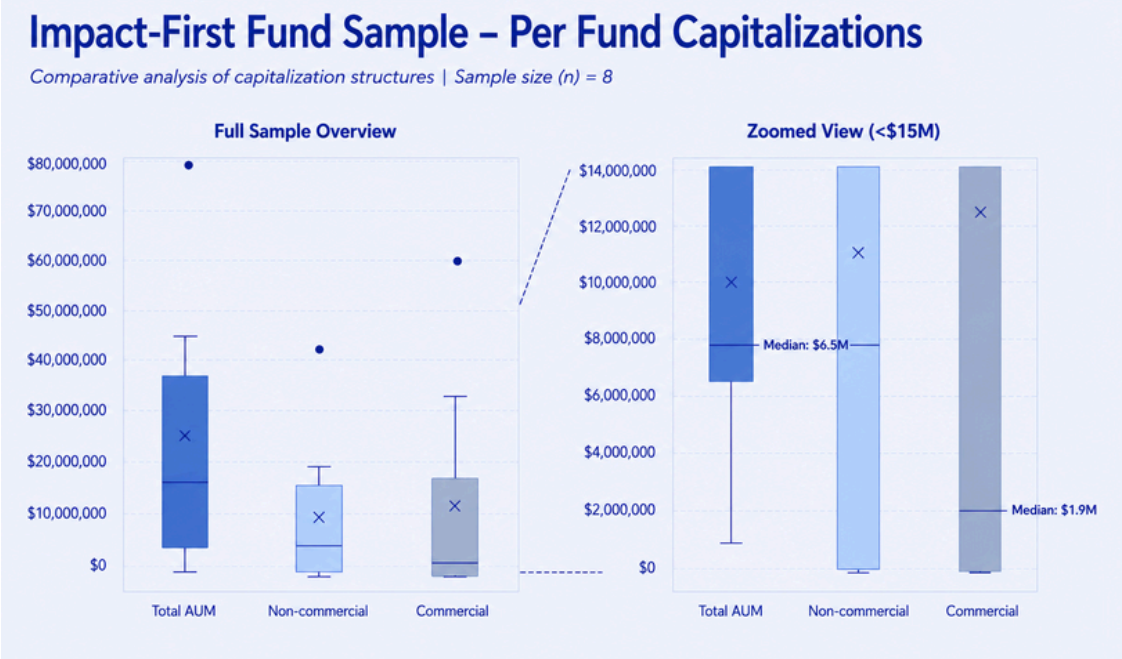

CORE INSIGHTS

	Grant-Funded Project	Impact-First Fund
Capital Efficiency	One-time outflow	Capital is deployed, repaid, and recycled
Use of Subsidy	Directly funds services or outputs	De-risks social enterprises delivering those outcomes
Systems Change	Difficult to sustain impact beyond the grant cycle	Grows enterprises that keep solving the problem after subsidy is gone

From a philanthropic perspective, the more interesting question is how subsidized capital gets allocated. To make the opportunity cost of a philanthropic subsidy visible, we compare how a dollar of philanthropy behaves when funding a traditional project versus an impact-first fund. The table above summarizes these differences in a comparative analysis across three dimensions.

Imagine allocating $100,000 of philanthropic funding for rural water access. As a grant, the dollars pay once for a fixed number of community water points. As a first-loss or deeply concessional capital layer in an impact-first fund, those dollars can unlock a multiple of that amount via repayable investment in an early-stage water enterprise. That enterprise can then go on to deliver sustainable water services into the future.

With increasingly extreme demands on philanthropic dollars, every dollar needs to go farther. Impact-first funds provide an opportunity for capital providers to contribute efficiently and effectively. By deploying catalytic financing strategies, capital providers can support capital users in achieving incredible outcomes that grant financing alone could not.

Case studies illustrating the multiplier effect

Case studies of investees that have borrowed from multiple funds participating in the study bring to life the multiplier effect of catalytic capital.

Jibu

Founded in 2012, Jibu has grown from a handful of struggling franchises into a pan-African safe water network with 11,000 retail points across eight countries. The turning point came when a USAID PACE facility offered a $1 million grant on the condition that Jibu raise three times that amount in private capital, just as the team emerged from Miller Center's accelerator program with a sharpened model and proof that the unit economics could work. When Jibu later pursued larger rounds, a handful of investors put the company in precarious territory: prolonged, extractive due diligence, "drip-fed" funding, and last-minute withdrawals that left the team over-extended on inventory and nearly out of cash. By contrast, fast, flexible catalytic funders like Open Road Impact stepped in to bridge those gaps created by other investors, keeping Jibu alive in war-torn Democratic Republic of Congo and underscoring that the design and reliability of capital can be as important as its amount.

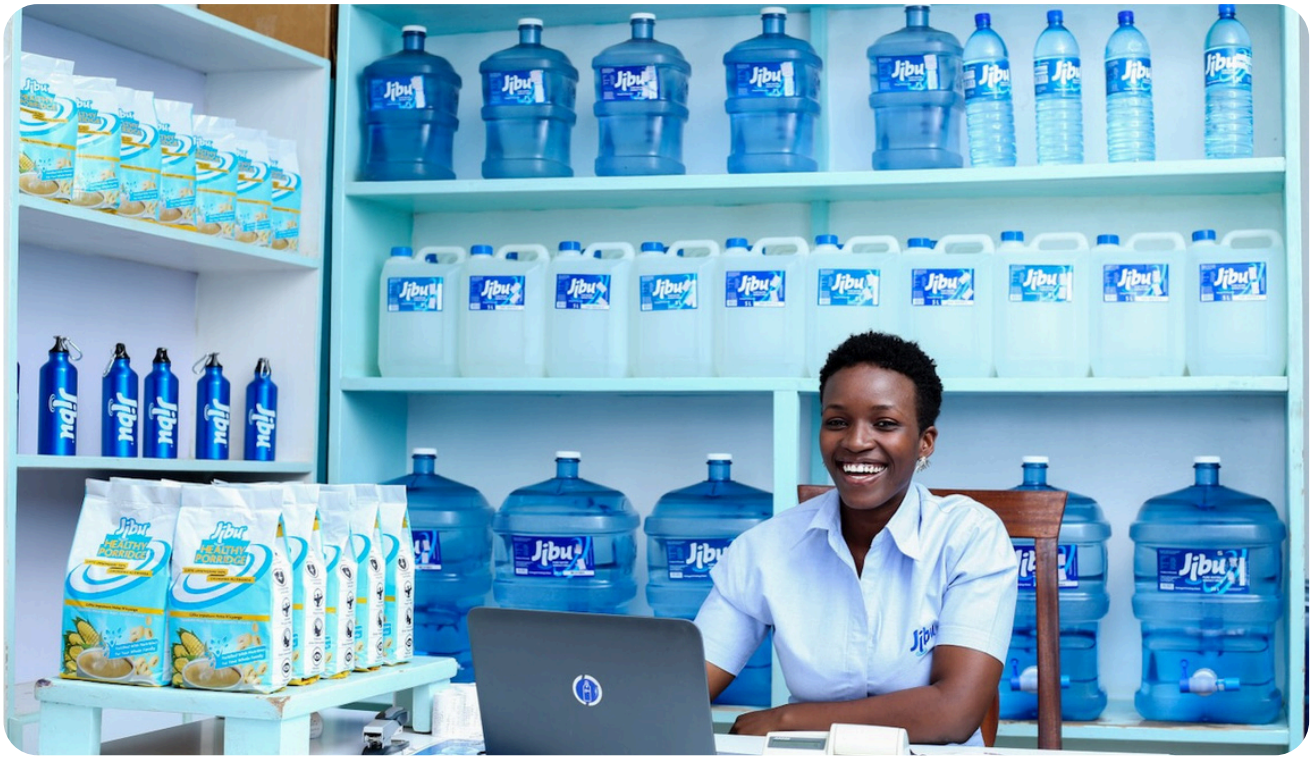

CORE INSIGHTS

Clínicas del Azúcar

Founded in 2011, Clínicas del Azúcar spent its first years proving that a low-cost, one-stop diabetes clinic could serve low-income patients. The catalytic moment came when an impact investor backed the company's national expansion, launching them from a regional network to over 50 clinics across Mexico and a first site in the United States. Clínicas has turned down private equity offers tied to growth at all costs, including raising prices on bottom-of-the-pyramid patients. In the words of Miguel Garza, the enterprise's business development officer, "When things go south in a company, the masks come off," which is why, for Clínicas, mission alignment with funders is a non-negotiable element.

Husk Power Systems

Husk Power Systems started in 2008 as a tiny biomass minigrid experiment in Bihar, long before "minigrids" were a recognized asset class. Building on Shell Foundation's early grant support, Acumen's first-round equity gave Husk room to grow from a handful of pilots into a real company, absorbing early failures while the team rebuilt the model.

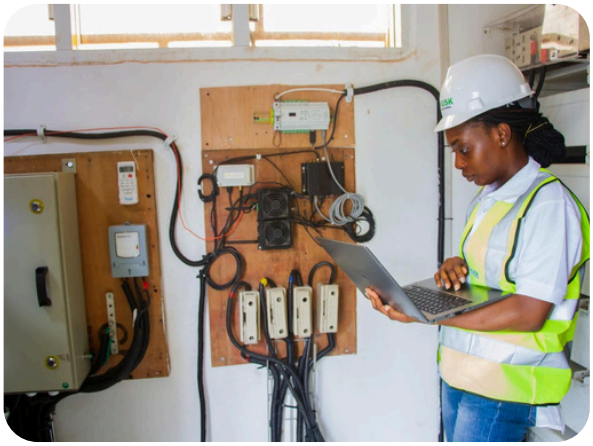

This work was sharpened by Miller Center's accelerator program in 2013, which resulted in a pivot toward scalable, productive-use power rather than household lighting alone. Years later, a USAID PACE grant provided the high-risk capital the business needed to enter Nigeria, turning Husk from a single-country pioneer into a cross-continental utility in the making.

CORE INSIGHTS

	Jibu	Clínicas del Azúcar	HUSK.
	Water | East Africa Founded 2012	Healthcare | Mexico Founded 2011	Energy | India & Sub-Saharan Africa Founded 2008
Where they are now?	**700,000+ lives improved** 11,000 retail points across 8 countries	**500,000+ lives improved** 50+ clinics across Mexico and 1 in US	**2.2M lives improved** 400 mini-grids across 2 continents, defining a new asset class
Catalytic capital amount	**$1M**	**$900K**	**$5M**
Counterfactual estimate without catalytic capital	Failure before proving viability **15,000 lives improved**	Stayed regional, approx. 15 clinics **75,000 lives improved**	Failure before proving viability **125,000 lives improved**
Impact Multiplier	**47x**	**7x**	**18x**

With relatively modest early-stage, high-risk investments ($1M in Jibu, $900K in Clínicas del Azúcar, and $5M in Husk Power Systems), catalytic investors enabled these enterprises to cross thresholds that traditional finance would almost certainly have left unfunded: proving new models, absorbing early failures, and entering new markets considered high risk. Through a conventional lens, these deals either would not have cleared an investment committee or would have been priced far beyond what entrepreneurs and their customers could bear.

Across the three cases, the initial catalytic investment has been recycled into durable enterprises serving hundreds of thousands of people, with "impact multipliers" ranging from 7X to 47X. Instead of funding a single round of services or infrastructure, philanthropic capital was used to build and de-risk enterprises that now continue to deliver impact well beyond the life of the original subsidy. This is why market-making subsidy is not a patch for inefficiency of capital allocators, but rather a force multiplier for capital users. Market-making subsidy allows impact-first funds to deploy capital seeking system-wide outcomes which rational market-rate investors, left to their own logic, would never finance.

Return concessions are the cost of inclusion

Although impact-first investing often targets outcomes and communities more traditionally served by philanthropy, it is still routinely evaluated against traditional finance benchmarks, including in this study. Therefore, we wanted to investigate what that difference from the market-rate actually represents and why it may be better understood as the intentional cost of impact-first finance.

In our study, we conducted a preliminary analysis of the investor return profiles reported by the debt sub-group of impact funds, comparing their financial return profiles to those of market-rate debt benchmarks.[11]

Metric	Impact Debt Funds	Traditional Private Credit Fund Benchmark	Representative Data
Return Expectations	Most target 2–5% (senior) Some target capital preservation	8–10% senior return target	~4–6% investor return differential
Actual Returns	Range: -12% to 6.25% Most clustered between 2–3%	~9% net return to investors	~6% median actual return gap
Borrower Rates	2–11% (often adjusted for risk, hedging and affordability)	18–24% (risk-adjusted pricing)	~10–15% affordability gap

In our sample, impact-first debt funds priced loans approximately 4–20 percentage points below prevailing market lending rates for small and medium enterprises (SMEs). This pricing gap is not evidence of weak performance; rather, it is the financial expression of doing something the market does not pay for. The delta should be considered as the cost of enabling access to capital: in other words, of inclusion.[12]

It is worth underscoring that all of the impact-first debt funds represented had met, exceeded, or were on track to meet their investor return expectations. By definition, these funds were high performing. The concessionary aspect is not a shortfall against expected return, but instead, a deliberate step away from a market-rate counterfactual.

The question then becomes what that foregone return is buying. Our group developed an iceberg analogy to visually represent what is going on "beneath the surface" at impact-first funds.

CORE INSIGHTS

What is visible is the count and amount of transactions. What remains invisible is the depth of support that surrounds every transaction: the intensive engagement required both before and after impact-first capital is deployed. This invisible enabling work comprises the true cost of making impact-first deals happen, which traditional capital markets incentivize rational investors to avoid.[13] Impact-first funds absorb these costs into their operations; traditional finance offloads them onto someone else, or more often, the deals do not get done at all. At the same time, loans are intentionally priced below prevailing rates so that borrowers and their customers can afford them. The combined "below the surface" work and affordable borrower pricing produce lower financial returns for investors.

Applying a systems lens, we should understand every investment as part of a larger system. We cannot look at one fund's financial statements in isolation. When these true costs are pushed off one investor's ledger, they do not disappear. They are simply shifted elsewhere. Rather than "overhead," to impact-first investors, these activities are integral to the product itself. They comprise the core inputs required to deliver outcomes. A higher cost per dollar deployed implies that impact-first funds are paying for what the market leaves out.

CONCLUSION

Impact-first funds are doing difficult work. At the enterprise level, they help early-stage social enterprises bridge funding gaps left open by traditional finance. At the fund level, they operate inside a traditional financial system organized around market-rate benchmarks. They correct for market failure, convert overlooked opportunities into investable ones, and fund work that the market systematically undervalues or treats as an externality to be addressed by someone else. What looks like inefficiency is actually the cost of delivering impact where markets fail.

If we want systems change, we must be willing to fund the full cost of this work, not just what the market can price. For funders, that means resourcing the work beneath the surface: diligence, structuring, accompaniment, and ecosystem support. For the field, it means evaluating impact-first funds on the system-level value they enable rather than the individual returns they concede. Closing the subsidy gap at the fund level is one of the most direct ways to address the capital gaps that enterprises experience as the "Valley of Death."

Impact-first funds are incredibly efficient engines for converting dollars into impact. Our analysis suggests that impact-first investing can be as rigorous as conventional fund management while explicitly prioritizing inclusion. It is the belief of the participants in this study that there is a tradeoff between impact and financial return. Philanthropic capital is not being used to subsidize the failure of funds to make ends meet, but rather, to support self-sustaining solutions that grow over time at an affordable price. At the same time, much of the value created by impact-first funds—including improved resilience, demonstration effects, and ecosystem strengthening)—remains unpriced and hard to quantify. This "impact dark matter" sits outside of standard metrics even as it shapes markets and lives.

This paper adds new data to a growing body of evidence that catalytic capital is one of the most efficient uses of philanthropy and a necessary ingredient in building inclusive markets, and that impact-first investing is a strategy that puts it to work.[14]

APPENDIX

DEFINING IMPACT-FIRST

THE CAPITAL SPECTRUM

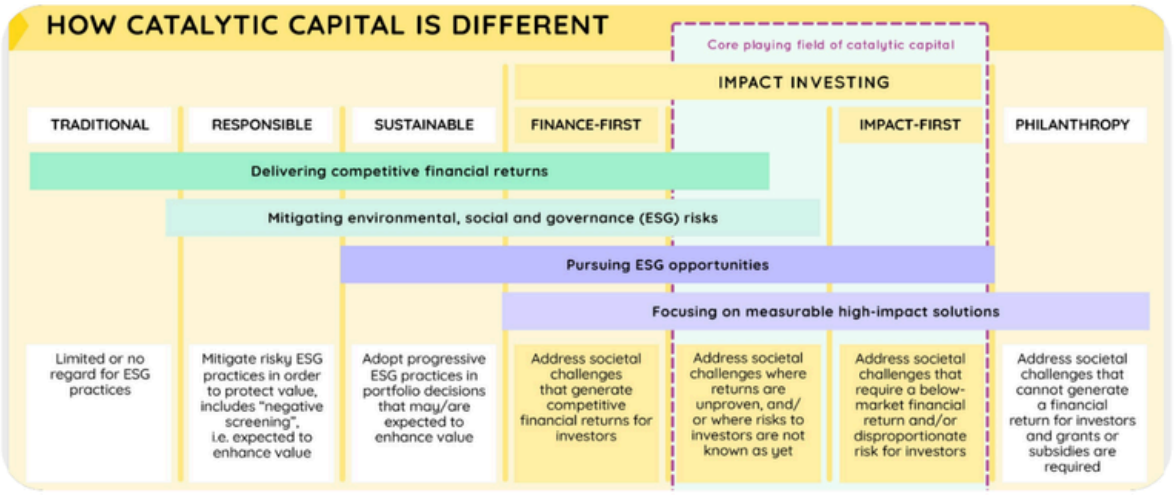

For the purposes of this paper, we use "impact-first" as defined by C3, the Catalytic Capital Consortium, in the above illustration. Beyond this definition, being able to assign financial metrics to this spectrum is one aim behind the drive to pursue this research and collaboration. Along with greater transparency of financial metrics, we aim to provide more attributes to refine the definition of "impact-first" on a qualitative basis.

GENERATIVE AI

Responsible Use of Generative AI on this Project

Over the course of the six-month research study and the subsequent documentation of findings, the project team used ChatGPT-4o and 5.1 as a collaborative tool to support the synthesis of workshop transcripts, data interpretation, structural outlining, and argument and language refinement. Gemini Pro was used to generate charts.

All data analysis, narrative framing, and key insights communicated in this article originated from the authors and their appointed colleagues at the represented funds. All prompts provided to the LLM were detailed, iterative, and grounded in original thinking produced by the humans on the project team.

The use of Generative AI provided a rigorous analytical interlocutor for the project's lead architect, enabling the orchestration of a field-wide effort that would otherwise have remained unpublished.

NOTES

[1] Emerging markets cannot afford the government intervention that seeks to cover some of this gap in developed countries (e.g., in the U.S., SBIR, NSF/I-CORPS, DARPA). This form of subsidy has led to a number of the innovations that our global economy currently relies on.

[2] Impact-first investing is closely related to catalytic capital. For the purposes of this study, we use the definition of impact-first investing and catalytic capital as defined by the Catalytic Capital Consortium.

[3] In many jurisdictions, fiduciary duty has a narrow interpretation, such that investors feel legally bound to pursue market-rate returns. Therefore, investing in funds that prioritize non-financial outcomes is often considered a non-compliant use of funds for asset allocators, even while the United Nations urges countries and institutions to update their fiduciary regulations to recognize the reality of the 21st century.

[4] Subsidy, as used in this study, is the difference between the true cost of deploying capital to underserved markets and what borrowers in those markets can afford to pay.

[5] Because this was an open-ended, collaborative research effort, participation was limited to eight organizations. This n = 8 sample is therefore illustrative, intended to surface mechanisms and benchmarking logic rather than claim population-level representation.

[6] In this study, we used transaction (closed deal) as the unit of analysis for investment activity. A transaction is a discrete financing event that requires underwriting/diligence and a closing process (e.g., a loan origination or equity round). We did not treat routine disbursements/ drawdowns under an existing agreement as separate transactions. In our impact-first sample, most transactions corresponded to unique enterprises. For the traditional finance benchmark, we approximate transaction counts using PitchBook-reported financing rounds/events during the investment period. PitchBook may bundle or classify follow-ons inconsistently so we treated them as the best-available proxy for closed financing events during the investment period.

[7] When we looked closer by subgroup, we found that impact-first debt funds showed even higher throughput relative to traditional benchmarks, especially those lending to financial service providers. Equity funds were somewhat lower, likely due in part to more bespoke and hands-on deal structuring.

[8] Most funds in this sample are US-based and invest in emerging markets, which is a common configuration in catalytic capital, where risk-bearing capital flows toward contexts underserved by commercial finance. Cross-border cost drivers likely influence the observed cost profile, but were not isolated in this phase. Future research may disaggregate cost dynamics by fund and portfolio geography.

[9] Cost per deal and cost per dollar deployed measure distinct efficiency dimensions. While the former shows little sensitivity to deal size, the latter increases as ticket size decreases due to largely fixed transaction and support costs being amortized over smaller capital amounts.

[10] We coin the term “market-making subsidy” to distinguish subsidy that is structurally necessary to enable a market to exist, from subsidy that improves or incentivizes activity within an existing market.

[11] A study of equity returns was out of scope for this phase of the project and presents an interesting opportunity moving forward.

[12] It is worth noting that, impact debt funds are often lending to institutions or in sectors where market-rate lenders will not lend at all, so there is no true "market" rate for these parts of the market.

[13] Penalizing economic activities that need to take place somewhere is a strong indication that externality-denying capitalism is at work. It should be noted that social entrepreneurs often personally bear costs that investors will not bear, for example, contributing their own efforts uncompensated at early stages to make their enterprise financially viable.

[14] This analysis is intended as a first step in a broader, collective effort to reclaim impact investing for the purpose it was meant to serve. Future work is needed to build on this foundation. To make the existing findings more robust, we seek to increase the number of funds participating in the study. We are also aware of a wealth of directly related topics, including and not limited to: developing fit-for-purpose ways to quantify the “dark matter” value created by impact-first funds beyond financial returns alone; examining equity funds in greater depth, including how return expectations shape their ability to remain impact-first; blended finance structures that reinforce impact-first principles; disentangling risk, operating costs, and subsidy in cost attribution; examining salary scales and staff geography as additional cost drivers; benchmarking finance-first impact investing against impact-first funds.