

The Foundation Review

Volume 9 | Issue 2

6-30-2017

The Philanthropy As One Big Impact Investment: A Framework For Evaluating A Foundation's Blended Performance

Rohit T. Aggarwala
Columbia University

Claudine A. Frasch
Gensler

Follow this and additional works at: <https://scholarworks.gvsu.edu/tfr>



Part of the [Nonprofit Administration and Management Commons](#), and the [Public Affairs, Public Policy and Public Administration Commons](#)

Recommended Citation

Aggarwala, R. T., & Frasch, C. A. (2017). The Philanthropy As One Big Impact Investment: A Framework For Evaluating A Foundation's Blended Performance. *The Foundation Review*, 9(2). <https://doi.org/10.9707/1944-5660.1370>

Copyright © 2017 Dorothy A. Johnson Center for Philanthropy at Grand Valley State University. The Foundation Review is reproduced electronically by ScholarWorks@GVSU. <https://scholarworks.gvsu.edu/tfr>

The Philanthropy As One Big Impact Investment: A Framework For Evaluating A Foundation's Blended Performance

Rohit T. Aggarwala, Ph.D., Columbia University, and Claudine A. Frasch, M.B.A., Gensler

Keywords: Impact investment, philanthropy, foundations, performance evaluation, benchmarking, venture philanthropy, endowment management

Introduction

There are few hotter topics in the philanthropic world these days than impact investing. From the White House to the World Economic Forum to the Giving Pledge, the idea of making investments that also yield social good has received significant attention (Rodin & Brandenburg, 2014; Brest & Born, 2013). And yet, it remains a troubled field for most foundations and philanthropists (Daniels, 2016; Foley, 2015). While some foundations — notably, the F.B. Heron Foundation — have put their entire focus on impact investing, philanthropy still lacks the tools that enable such investments to be made with the same rigor as the best financial investments and philanthropic grants (Miller & Johnson, 2015; Miller & Rogers, 2014). As Antony Bugg-Levine and Jed Emerson (2011) have pointed out, the key challenge is to find a rigorous approach to evaluate the “blended value” of impact investments.

In this article, we propose a framework for evaluating the blended performance of an entire foundation's outlays — both grants and financial investments — by quantifying both impact and financial returns separately, and using them as two axes on a graph. Inspired by Harry Markowitz's work, which underpins modern financial portfolio theory, this approach uses a foundation's existing outlays to chart its overall blended performance by creating an “endowment-grant allocation line” (Markowitz, 1952). This, in turn, makes it possible to evaluate whether the blended value of impact investments could improve the foundation's overall performance, even if those investments generate

Key Points

- While some foundations have put their entire focus on impact investing, philanthropy still lacks the tools that enable such investments to be made with the same rigor as the best financial investments and philanthropic grants. This reveals a more fundamental problem: We do not currently manage foundations as the integrated portfolios that they are.
- This article proposes a framework for evaluating a foundation's blended performance that enables both grantmaking and endowment investing to be evaluated jointly, and thus also allows a complete evaluation of how impact investments could improve — or fail to improve — overall performance.
- The article demonstrates the framework's utility by using it to evaluate a set of actual impact investments in the field of the environment. Using this framework to assess foundations' performance would not only improve fundamental performance, but also potentially unlock vast new areas of social entrepreneurship.

below-market financial returns and smaller impacts than traditional grants. Fundamentally, the framework presented here evaluates the entire foundation as one big impact investment, even if the foundation currently uses only the most traditional tools of grantmaking and an endowment focused solely on financial returns.

We have used this approach to evaluate a set of actual impact-investment opportunities that

were seeking funding in 2012. First, using the long-term average returns of a high-performing nonprofit endowment and a high-performance environmental grant that we had analyzed, we created a theoretical philanthropy and graphed its endowment-grant allocation line. We then analyzed the promised impact and financial returns of a set of 22 potential impact investments related to climate change in a standardized way that also took into account the duration of investment; with this, we were able to rigorously evaluate each in a way that could fairly compare dramatically different types of investment opportunities. Significantly, when compared to the theoretical philanthropy's endowment-grant allocation line, several of these impact investments that seemed attractive individually actually would have reduced the overall performance of the philanthropy, while others that seemed less attractive on their own actually would have improved the portfolio's performance. In this way, the framework eliminates the uncertainty of blended-value analysis and allows the investor's judgment to focus instead on the most important question about any investment, which is execution risk.

The Trouble With Impact Investing

For all the discussion about impact investing, it is a term and a field that seems to raise more questions than answers. Foundations and philanthropists struggle to figure out whether impact-focused investments make sense and what to expect from them. Should these investments avoid financial risk, or seek outsized returns? Should they take first-loss positions in order to catalyze the participation of traditional investors, or would that simply be subsidizing someone else's return? How much of a foundation's assets should go into "impact investments"? And there is also the issue of how to calculate the impact of an investment — especially before you make it, which is when that information is really useful (Brest & Born, 2013; Bugg-Levine & Emerson 2011).

Adding to the confusion, advocates of impact investing take widely differing positions on what level of financial returns should be expected. Some argue that there is an unlimited set of

opportunities that have lots of impact while receiving market-rate returns or better — which has the unsettling implication that we ought just to cancel grantmaking altogether. Others use the term "patient capital" — nicely complimenting those willing to wait a long time to get their money back, but often ignoring how much value even low inflation rates eat up over a decade or two. And there is always the disturbing possibility that the social entrepreneurs pitching to you might really be planning to sell their "social business" for millions, and see you more as a source of low-cost capital than anything else (Rose-Smith, 2016; Milligan & Schöning, 2011).

Finally, there is the very real issue of how to make, and manage, impact investments. Most program officers have an advocacy or public-sector background; they aren't accustomed to evaluating for-profit business plans. The money managers could do so, but they won't be as focused on impact, and their incentives aren't well aligned since they usually get compensated on the financial performance of the endowment. For the occasional impact investment where the impacts are huge and simply couldn't be achieved through a grant, the grantmakers will likely get excited and be willing to fund it out of their budget; in the same way, impact opportunities that really are financial home runs — beating the risk and return standards of the overall endowment — will get done through the normal investing process. The ones in between will either fall into a no-man's-land or require direct intervention from the CEO to get done (Godeke & Burckart, 2015).

Unfortunately, this no-man's-land is where most impact-investment opportunities lie; even when they have the potential for high returns, they come with greater uncertainty. Navigating the no-man's-land requires a level of integration across functions that few foundations achieve. A potential impact investment must be evaluated on both its impact and its financial return, funded either from the grantmaking budget or the endowment, and managed for both aspects of its performance.

And this is a problem, because it is easy to conclude that such opportunities fail both the impact test and the market test. For philanthropy to embrace impact investing fully, it must find a way to understand these opportunities in terms of their blended value of both impact and financial return. Several efforts are underway to standardize and make transparent the impacts of for-profit entities, which will help evaluate individual impact investments in a systematic way (Godeke & Burckart, 2015; Miller & Johnson, 2015; Miller & Rogers, 2014; Bugg-Levine & Emerson, 2011).

Even with rigorous analysis, however, any investment is good only if it improves the overall performance of the fund or firm making the investment. Thus, for any given impact-investment opportunity, the rigorously evaluated blended value must be compared to the rigorously evaluated blended performance of the foundation considering it. And that is the bigger problem — because, while foundations generally evaluate both their endowment performances and their grants, they rarely consider the two in conjunction with one another (Coffman & Beer, 2016).

Visualizing the Foundation's Blended Performance to Evaluate Investments

The irony here is that the foundation itself is, essentially, one big impact investment. A donor puts an endowment into a foundation; in return, he or she expects to receive both impact on the world and the preservation of financial value to enable future impact. Each year, the investment team works to make financial returns, and the grantmaking team is given a portion of those returns to create impact. At the end of the year, the foundation has two metrics by which it can understand its performance: the net change in the endowment and the impact generated. Taken together, in terms of Bugg-Levine and Emerson's concept of blended value, this is its blended return.

The fact that every foundation has a blended return is clear every time a board decides to spend more than its required distribution to

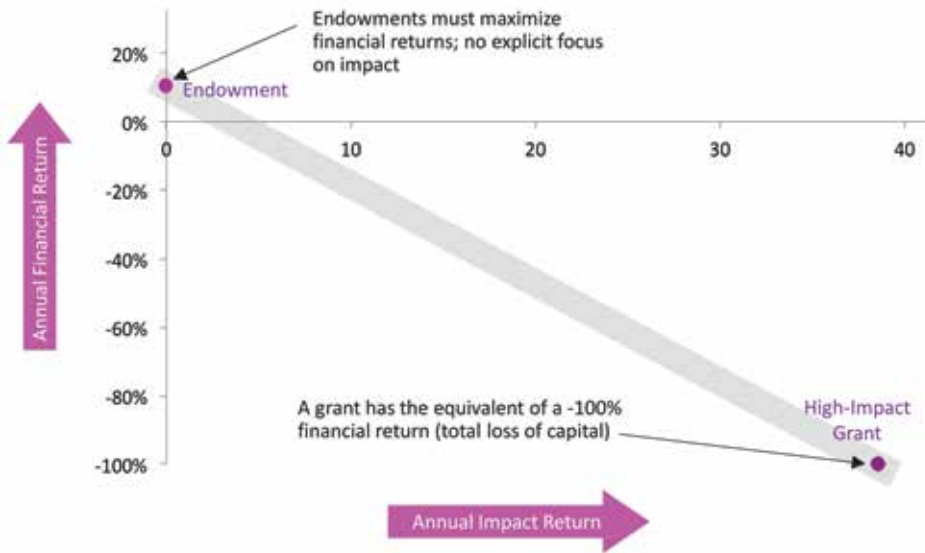
[T]he foundation has two metrics by which it can understand its performance: the net change in the endowment and the impact generated. Taken together, in terms of Bugg-Levine and Emerson's concept of blended value, this is its blended return.

realize some time-sensitive impact opportunity, or to cut grantmaking to preserve capital in times when the endowment shrinks due to poor financial performance. Thus, a good board instinctively manages for both financial and impact performance, even if it sticks with the traditional tools of pure grantmaking on the program side and financial-return-only objectives on the endowment side. In other words, it seeks to optimize blended value — which is to say, it manages itself like an impact investment.

If the foundation is an impact investment, why doesn't the end-of-year board meeting struggle with the same challenges as impact investments do individually? The truth, of course, is that foundations rarely attempt to measure themselves on a blended, quantitative basis. While most grantmakers today present detailed metrics for each grant, only a few foundations really evaluate their overall impact in hard numbers. Those that do struggle to boil impact performance down to even a few metrics (Colby, Fishman, & Pickell, 2011). Even more, the resource-intensity of the impact generated is almost never considered; the endowment managers have usually left the boardroom before the discussion of the grantmaking program has begun.

And this leaves unanswered perhaps the most important question the board should consider — the foundation's blended performance. In the corporate world, a key metric of performance

FIGURE 1 Visualizing Financial and Impact Returns Together



Sector

is return on invested capital (ROIC), which is to say, “How much money did we make taking into account how much money we are tying up in order to make money?” A foundation’s equivalent to ROIC would be its combined impact and financial return, divided by the total assets in its endowment at the beginning of the year. This would be its blended performance.

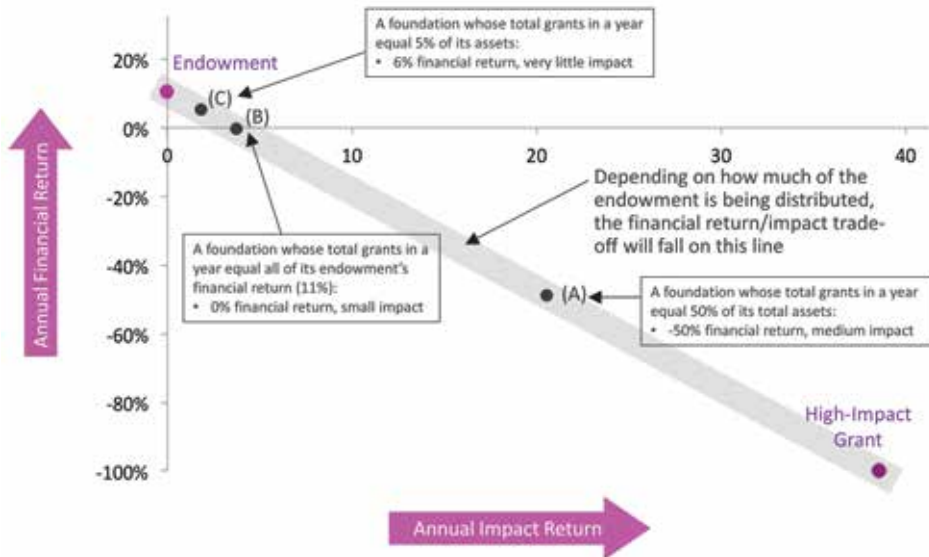
One reason foundations don’t attempt to measure blended performance is that there would be very little utility to the number. In theory, if a bunch of foundations adopted the same approach, it could serve as a comparative metric. And while this would be useful, it wouldn’t necessarily improve performance; just as different industries have different average ROICs, foundations in different fields or focused on different priorities would have structurally different levels of blended performance.

The better use of blended performance would be to evaluate what specific investments, grants, or impact investments would actually improve the overall performance of the whole foundation.

Most companies convert their ROIC into a “hurdle rate” for such decisions — the rate of return specific to that company below which an investment destroys value by reducing its ROIC below an acceptable level.

Doing this for a foundation relies on returning to the source of modern financial-portfolio theory. In 1952, in an article that ultimately won him a Nobel Prize, Markowitz argued that portfolio managers were evaluating their investments in an entirely misguided way by focusing only on the return of an individual asset. Instead, he argued, they should look at the contribution that each investment made to their overall portfolio’s performance, which had not one metric, but two: financial return and risk. By plotting the return and risk of each risky asset (stocks, bonds, real estate, etc.), he argued, an investor could find the “efficient frontier” where any desired increase in return required the acceptance of more risk. Further, he pointed out that by blending this efficient portfolio of risky assets with “risk free” treasury bonds that returned less but had zero functional risk, an

FIGURE 2 The Endowment-Grant Line



investor could create a portfolio with any desired level of risk along a “capital allocation line” that connected the efficient frontier with the return associated with the risk-free asset (Markowitz, 1952; Rubenstein, 2002).

Markowitz’s overall approach allows us to create the same kind of benchmark for a foundation. We start by imagining the entire foundation as a single portfolio that makes capital outlays in pursuit of both financial return and impact. (See Figure 1.) The vertical (y) axis indicates the annualized financial return, and the horizontal (x) axis indicates the impact achieved. Financial returns, as usual, are represented as an annual percentage or, in other words, on a per-dollar basis, because a 10 percent return means 10 cents returned for each dollar in the endowment each year. To be consistent, we would show impact in standardized units of annualized impact per dollar in the grant budget. (See Appendix A.)

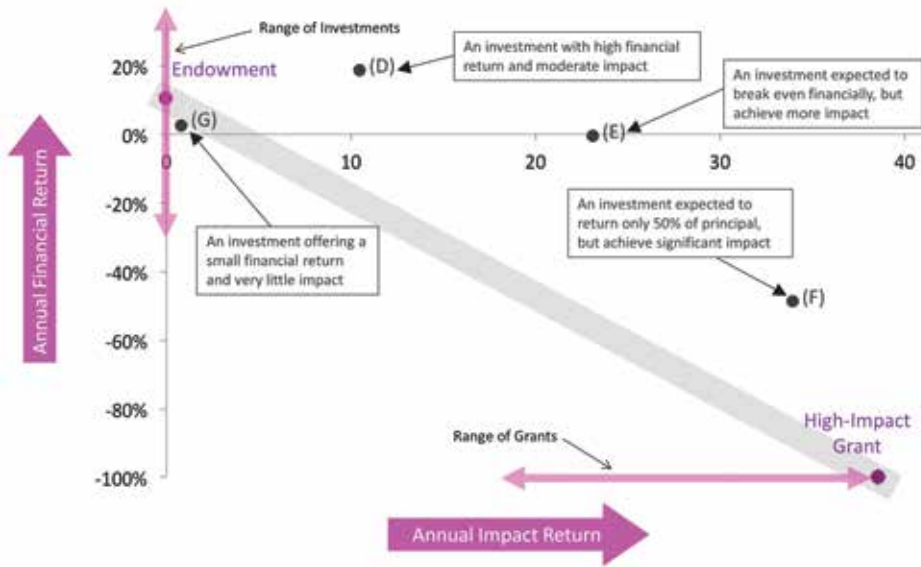
In general, the endowment taken alone would be in the top left of the graph, with (one hopes) a 10 percent to 15 percent financial return, but no

impact yielded. The grant program, considered alone, would be at the lower right corner, yielding impact but losing all of its investment — in other words, a negative 100 percent financial return.

If we assume that both the investment managers and the program officers are doing their jobs well, then each should be at the outer edge of what is possible for their tools: the endowment team simply can’t get a sustained higher annual return, and the grantmaking team can’t improve its overall impact per dollar in their current programs.

Therefore, a line between these two points on the graph defines the combined financial and impact performance of the foundation’s total portfolio in any giving scenario. Using only these two instruments, a foundation cannot achieve results beyond this line. For example, if a foundation’s grants in a single year equaled 50 percent of its total assets, its financial return would be just above negative 50 percent, but with a much greater impact. (See Figure 2.) If it gave away its endowment’s total earnings, its overall

FIGURE 3 Plotting Impact Investments



performance would be where the diagonal line crosses the x-axis. The average foundation — earning perhaps 11 percent on its endowment and allocating 5 percent of its start-of-year assets to grantmaking — would see an overall financial return of approximately 6 percent and a relatively small impact return.

Thus, for a foundation using only market-rate investments and best-in-class grants, the amount of money allocated to grantmaking is the key determinant of impact achieved. This “endowment-grant allocation line” is the equivalent of Markowitz’s capital allocation line — the investor can achieve any point on the line simply by reallocating assets, but it cannot move above or below the line unless new asset classes emerge.


With our “endowment-grant allocation line” defining the foundation’s overall potential performance, every other investment can be evaluated comparatively: The endowment includes many investments, which range from financial loss to occasional outsized returns, and the grants include some home runs and a few that

didn’t succeed. Impact investments, generating both kinds of return, will appear in the much larger space between the axes. (See Figure 3.)

And this, quite clearly, demonstrates how value-creating impact investments can be identified. Any investment that is above and to the right of the line generates a blend of financial and impact returns that the current endowment-grant structure could not achieve; it will improve the overall performance of the foundation. Any investment that falls below the line destroys value; the money is better kept in the endowment’s financial portfolio and the proceeds used to fund grants. (See Figure 4.)

Can We Really Quantify Impact?

Of course, the real challenge is how to turn this theory into an actual, usable set of numbers that informs an investor prior to an investment decision. The field of impact investing has expended significant effort in ways to quantify and report the impact associated with an investment, such as the Impact Reporting and Investment Standards (IRIS) tool, but these generally are not

FIGURE 4 The Endowment-Grant Line as the Foundation's Hurdle Rate

designed from the outset to help select the next one (Bhouri, 2011).

On the financial side, quantifying value is easy because it is precisely what the concept of “money” exists to do. In the 19th century, economist William Stanley Jevons defined money as offering four functions: a medium of exchange, a common measure of value, a standard of deferred payment, and a store of value. If philanthropy had a unit of impact that could accomplish these four things, our analytical challenge would evaporate. But a metric that could cut across multiple fields of philanthropy eludes us (Jevons, 1875).

However, some program areas do lend themselves to quantitative impact analysis. The Robin Hood Foundation works hard to quantify the impact on poverty alleviation expected from its portfolio of grants. The foundation assigns a monetary value to the expected benefits of the intervention’s outcomes (e.g., one additional person graduating from high school increases his lifetime earnings by X dollars), and then it

calculates a cost/benefit ratio for each potential grant (Weinstein, 2009).

Another area that lends itself to rigorous impact quantification is environmental philanthropy, which today is highly focused on mitigating climate change. Directly or indirectly, climate-change philanthropy is about spending money to keep greenhouse gases (GHGs) out of the atmosphere. Because GHGs are fungible across the planet, are quantifiable in a single unit of carbon dioxide equivalent (CO₂e), and have predictable utility across time, climate-change philanthropy has a metric that is just as analytically powerful as money.

The key, then, is to analyze impacts with the same rigor as one analyzes promised financial results. And here, too, we run into a gap: we lack the conventions that financial analysts take for granted. In climate change, the rules developed for carbon-trading systems focus on precise determination of “additionality,” to ensure that the public is actually getting the GHG reductions it paid for. These are useful rules for their

purpose, but they don't help the ex-ante evaluation of climate-change-focused investments. They are, in fact, more like the accounting standards used in a financial audit. Any honest investment manager will tell you that the numbers you crunch before making an investment would never pass an audit; they are full of uncertainties and estimates because they are trying to predict the future without allowing uncertainty to prevent action. They are kept within reason by a set of conventions about how financial returns should be estimated. Similarly, much of our legal structure exists to ensure that financial flows do not get double-counted: when two people claim the same money, a lawsuit ensues. But every grantmaker knows that many parties often claim the same impact. So, we need not only usable ex-ante estimates of impact, but also an approach to determining how much of the overall impact can be claimed by any one investor.

So, could a climate-change-focused foundation, using avoided GHGs as its currency of impact, evaluate an impact investment? Easily. It should be no more difficult to model the GHG reductions expected from an investment than it is to estimate its future revenues and costs. It should be feasible to adopt a notional discount rate for GHGs, on the basis that a ton of GHG reductions today is worth more than a ton of GHG reductions 10 years from now. Just as every investment plan has to discount expected returns based on uncertainty in the broader market and the risk of poor execution, it should be feasible to discount our impact estimates to account for the risk that circumstances might change and the risk that our managers might fail to deliver on their plans. Because even a huge foundation doesn't have unlimited resources, we will also need to consider how much capital must be invested, and how long it is tied up, in order to achieve the expected amount of impact — much as an investor already considers these investment characteristics in a financial internal rate of return. And, if a key purpose of the investment is to prove that a new business model works so that others adopt it (the “demonstration effect”), a real-options approach can incorporate the value of those future impacts. Perhaps the most difficult aspect of this would be to determine how much of the

overall GHG reductions are attributable to any particular investment, taking into account other players working on the same goals, other investors in the same project, and external factors that might influence the outcome.

This sounds like a long list, but each analysis requires only a set of reasonable rules and some analytical legwork to get done. And we are not aiming for auditable figures: just as the law uses the “reasonable man” standard, we can use the “reasonable board member” as the person we must convince, rather than an auditor. And our reasonable board member should be impressed; chances are, these are far more rigorous a set of rules than boards are accustomed to seeing used to assess philanthropic initiatives. (See Table 1.)

So, Does It Work in Practice?

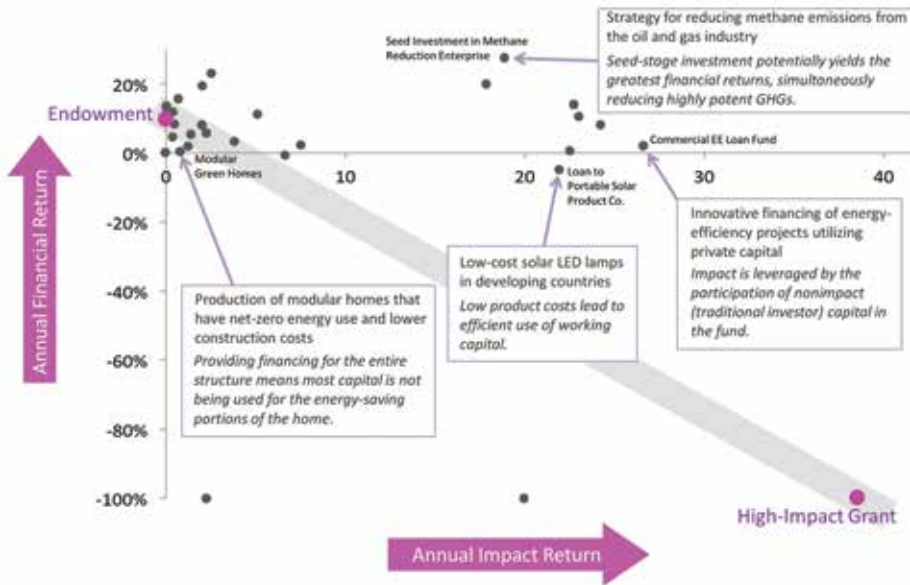
To test this approach, we modeled a set of 22 potential impact investments and three potential grants whose nonfinancial purpose was to reduce GHG emissions, using the conventions described above. Most of the impact investments were active opportunities undertaking fundraising in 2012, when this research started; five were theoretical opportunities for which no business plan had yet emerged. The set was chosen to span a broad range of investment types, including debt and equity, early-stage venture capital to project finance, initiatives with easily quantified impact and those with indirect and/or shared impact, and with terms that range from two to 20 years. As the end points of our “endowment-grant benchmark line,” we used the average 10-year financial returns of the best-performing large university endowment and one large GHG-related grant in which we had been involved and for which we had detailed cost, impact, and allocation data.

The results indicated that there is much insight to be gained from a rigorous approach to analyzing blended value. (See Figure 5.) Each of the potential investments identified had smart, experienced proponents whose intuition led them to think the investments would have great impact. But, even accounting for the fact that they were expected to return money to the foundation, most had such small impacts per dollar invested

TABLE 1 Impact Quantification Guidelines

	Issue	Approach	Method
Establish pro-forma value flows.	<ul style="list-style-type: none"> How much impact do we expect over the duration of the investment term? 	<ul style="list-style-type: none"> Just as a financial analyst models expected periodic cash flows, model the expected impact quantities associated with the same time periods. 	<ul style="list-style-type: none"> Create a pro-forma model of expected impact flows, by time period. Often, but not always, impact flows will follow the same expected growth trajectory as the business itself.
Determine ownership of value flows.	<ul style="list-style-type: none"> How much of the impact generated by this investment can we legitimately claim? 	<ul style="list-style-type: none"> Consider other investors (including owners) and capital providers: Do they also have claim to a portion of the impact? Are other impact-motivated players involved elsewhere in the delivery channel, or is the investment building on other impact-driven work, such as regulation? 	<ul style="list-style-type: none"> Similar to allocating enterprise ownership based on an investor's capital contribution, but expanded to consider external influences contributing to the accumulation of impact, in order to avoid double-counting of impact claims
Discount for the time value of impact.	<ul style="list-style-type: none"> Some impacts, such as carbon reductions, are less important if made tomorrow than the same amount of reductions made today. 	<ul style="list-style-type: none"> "Time value of impact" 	<ul style="list-style-type: none"> Apply an appropriate discount rate to determine the present value of expected impact.
Adjust for the uncertainty of getting your return.	<ul style="list-style-type: none"> Will the business plan be executed? Will the product be used as intended? Will the widgets work, and how long will they last? 	<ul style="list-style-type: none"> Estimate execution risk, permanence risk, and other types of uncertainties for each investment. 	<ul style="list-style-type: none"> Apply cumulative "haircuts" to quantities of estimated impact.
Include the value of potential follow-on impact.	<ul style="list-style-type: none"> How do we value the fact that our purpose is to demonstrate that this business model is feasible? 	<ul style="list-style-type: none"> Perform a real-options analysis using likelihood and scale of follow-on impact. Consider how much of the follow-on impact we can attribute to our initial investment. 	<ul style="list-style-type: none"> Develop informed scenarios of the potential timing and quantities of impact beyond the investment term; assign probabilities.
Determine the "impact IRR."	<ul style="list-style-type: none"> Is this investment better than a grant, or than an investment that yields greater returns but takes longer? 	<ul style="list-style-type: none"> Consider how much capital must be invested and for how long in order to achieve the total estimated impact we can claim. 	<ul style="list-style-type: none"> Calculate the quantity of impact per dollar invested, per year of the investment term.

FIGURE 5 Results of Analysis, With Characteristics of Selected Opportunities



Sector


that they were really best considered only on their financial merits. Looking more closely, some specific opportunities surprised us. One investment sounded great but proved to destroy value; in retrospect, it was so capital-intensive per unit of impact that it could not compete with the baseline endowment-grant combination. Another was unimpressive at first glance, but it achieved so much financial leverage by bringing in nonimpact investors that the large quantity of impact easily justified reduced financial returns to the impact investor. Another took so long to realize impact that the discount rate ate up its value. Above all, though, our results also demonstrated this heartening fact: a number of the opportunities we saw do have the potential to improve the overall performance of the foundation. And some clearly did trade off financial returns for impact, giving the philanthropic impact investor a key role to play. (See Appendix B.)

These results, therefore, demonstrate also the limitations of the after-the-fact reporting formats many impact investors use to monitor the level of impact their investments have had. Tools such as IRIS are useful for the auditing function, to

ensure over the long term that the field is not selling snake oil. But they are not always useful to predict the impact of potential individual investments. If impact investing is to compete for capital with the traditional philanthropic approach of earning money and then making donations, it must adopt far more precise ways of selecting those investments that will create value for the foundation as a whole. Making impact investments imprecisely and then evaluating them rigorously will — inevitably — lead to lackluster performance. Only when quantitative analysis drives investments will the field’s actual performance realize its potential.

Our results also demonstrated one additional benefit to this kind of rigorous analysis: that there is a huge, unexplored white space that philanthropy can and should consider. Most of the impact-investment opportunities we analyzed fell into two clusters: some were truly value-creating, while most clustered closely around the profitable end of the endowment-grant benchmark line, indicating, essentially, “business plans with some positive social benefit.” (See Figure 6). This makes sense; while lots of

FIGURE 6 Expected Financial and Impact Returns for a Set of Impact Investments, Plotted Against the Endowment-Grant Line



people seek proposals for grants and profitable business plans, there is not the comparable demand for business plans that lose some, but not all, of their money while creating social good. Of course, there are lots of organizations that do good work while generating revenues insufficient to cover their costs; think of most arts and educational institutions that augment revenues from tuition and ticket fees with grants from donors. The field of social entrepreneurship would clearly benefit if philanthropist-investors were to find a way to value equity stakes in ventures that might never make a profit but aren't truly charities.

Putting This Into Practice

The approach here represents a first-cut attempt at developing a methodology that would consider an impact investment in its true context. We are encouraged by the initial results and what we have learned through building this analytic model and data set. It will be necessary to refine the methodology through the analysis of additional climate-focused investments — to more specifically address the riskiness of financial returns (through standardized assumptions

by asset class and business stage), as well as more consistently consider and account for the impact upside of the demonstration affect, a key rationale behind many impact investments (which can be done through a real-options-style analysis). Further, just as every investment firm develops its own models to reflect its beliefs, preferences, and risk tolerance, any impact investor will need to tailor an approach such as this to its own situation and purpose in the selection of opportunities.

Even when the framework is refined, acting on the opportunities it identifies will still require management finesse, because the traditional separation of a foundation's investment and program teams provides no obvious place from which to analyze or manage investments seeking both financial and impact returns. Asking these two groups to work together under existing structures seems destined for failure — especially if the investment managers are compensated on financial performance alone and the program officers lack financial experience. Alternatively, a separate impact-investing team would need not only to attract the best

There will be challenges inherent in a transition from traditional foundation operations to this unified analytic and investment approach. But this framework and our results are evidence that it can and should be done, and with worthwhile result. The difficulty, and messiness, involved in identifying and quantifying a foundation's impact is the necessary price of determining its effectiveness.

from both the program and financial sides of the house, but also ensure that they are fully working together and can compete fairly for investment dollars with both the grantmaking program and the endowment. For example, the Heron Foundation has merged its investment and program teams into one “capital deployment” team as it shifts its strategy to one fully focused on impact investing (Wallace, 2013).

Our framework provides a basis on which to solve these management challenges. A quantitative, dual return metric offers program and financial staff a neutral, common ground that requires each to think deeply in the terms of the other side of the house. It provides a way for a board of directors to determine smart allocations of money away from grants or endowment funds and into impact investments at a scale that matters. And it offers a basis for results-based compensation structures that can attract money managers into the impact-investing space, with the right incentives.

There will be challenges inherent in a transition from traditional foundation operations to this unified analytic and investment approach. But this framework and our results are evidence that it can and should be done, and with worthwhile result. The difficulty, and messiness, involved in identifying and quantifying a foundation's impact is the necessary price of determining its effectiveness.

This framework also provides a way to continuously evaluate and improve the performance of the entire foundation. One could imagine compensating all staff based on the foundation's combined performance, for example, helping to break down the silos between the financial and program staff. Recognizing the fact that each program area will probably always have a different unit of impact, one could break the foundation's endowment up into a separate account for each program, and evaluate each program using this framework — just as many corporations evaluate performance at the business-unit level. In such a case, fields in which impact investing proves more effective would do more of it; fields in which traditional grants prove more effective would stay where they are. But in both cases, ongoing evaluation would ensure that a foundation did not miss a change in circumstance or a good opportunity.

Either way, the real challenge of impact investing is a challenge to the foundation itself: Can we think about our overall performance in a rigorous, quantitative way that incorporates both our impact and our financial objectives? And can we do so in a way that informs and improves our decision-making process? If we can — and if we are willing to act on it — getting impact investing right offers a way not only to use new tools, but to improve the effectiveness of the foundation as a whole.

Acknowledgments

The authors would like to acknowledge the contributions of John Goldstein, Eric Hallstein, and Peter Kelly, who worked on the analysis of the impact investments as part of a team from Imprint Capital Advisors, now part of Goldman Sachs Asset Management.

References

- BOURI, A. (2011, Summer). How standards emerge: The role of investor leadership in realizing the potential of IRIS. *Innovations: Technology, Governance, Globalization*, 6(3), 117–131.
- BREST, P., & BORN, K. (2013, Fall). When can impact investing create real impact? [with responses]. *Stanford Social Innovation Review*, 11(4), 22–31.
- BUGG-LEVINE, A., & EMERSON, J. (2011). *Impact investing: Transforming how we make money while making a difference*. San Francisco: Jossey-Bass.
- COFFMAN, J., & BEER, T. (2016). How do you measure up? Finding fit between foundations and their evaluation functions. *The Foundation Review* 8(4), 27–43. <https://doi.org/10.9707/1944-5660.1325>
- COLBY, D., FISHMAN, N., & PICKELL, S. (2011). Achieving foundation accountability and transparency: Lessons from the Robert Wood Johnson Foundation's scorecard. *The Foundation Review*, 3(1), 70–80. <https://doi.org/10.4087/FOUNDATIONREVIEW-D-10-00031>
- DANIELS, A. (2016, May 12). Foundation CEOs struggle to boost mission investing. *Chronicle of Philanthropy*. Available at <https://www.philanthropy.com/article/Foundation-CEOs-Struggle-to/236462>
- FOLEY, S. (2015, October 30). Gates cautious on “impact investing.” *Financial Times*. Retrieved from <https://www.ft.com/content/c7f4efa2-7e7d-11e5-a1fe-567b37f80b64>
- GODEKE, S., & BURCKART, W. (2015, March 18). Impact investing can help foundations avoid obsolescence. *Chronicle of Philanthropy*. Retrieved from <https://www.philanthropy.com/article/Opinion-Impact-Investing-Can/228569>
- JEVONS, W. (1896). *Money and the Mechanism of Exchange*. New York: D. Appleton.
- MARKOWITZ, H. (1952, March). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- MILLER, C., & JOHNSON, T. (2015, June 4). Mission-aligned investing: more complex than it seems. *Stanford Social Innovation Review*. Retrieved from https://ssir.org/articles/entry/mission_aligned_investing_more_complex_than_it_seems
- MILLER, C., & ROGERS, J. (2014, July 18). Taking impact investing and accounting full tilt. *Stanford Social Innovation Review*. Retrieved from https://ssir.org/articles/entry/taking_impact_investing_and_accounting_full_tilt
- MILLIGAN, K., & SCHÖNING, M. (2011, Summer). Taking a realistic approach to impact investing. *Innovations: Technology, Governance, Globalization* 6(3), 155–166.
- RODIN, J., & BRANDENBURG, M. (2014). *The power of impact investing: Putting markets to work for profit and global good*. Philadelphia: Wharton Digital Press.
- ROSE-SMITH, I. (2016, August 11). Impact investment: Not kidding around anymore. *Institutional Investor*. Retrieved from <http://www.institutionalinvestor.com/inside-edge/3577813/impact-investment-not-kidding-around-anymore.html#.WOaWcVdYXHg>
- RUBENSTEIN, M. (2002, June). Markowitz's “Portfolio Selection”: A fifty-year retrospective. *Journal of Finance*, 57(3), 1041–1045.
- WALLACE, N. (2013, May 19). A foundation risks all of its endowment on creating jobs. *Chronicle of Philanthropy*. Retrieved from <https://www.philanthropy.com/article/A-Foundation-Risks-All-of-Its/154829>
- WEINSTEIN, M., WITH LAMY, C. (2009). *Measuring success: How Robin Hood estimates the impact of grants*. New York: Robin Hood Foundation.

Rohit T. Aggarwala, Ph.D., is co-head of labs at Sidewalk Labs and an adjunct professor of international and public affairs at Columbia University. From 2010 to 2013, he founded and led the environmental grantmaking program at Bloomberg Philanthropies. Correspondence concerning this article should be addressed to Rohit T. Aggarwala (email: rta4@columbia.edu).

Claudine A. Frasch, M.B.A., is a director in the strategic consulting practice of Gensler, a global architecture, design, and planning firm. From 2012 to 2015 she advised Bloomberg Philanthropies on impact investing.

APPENDIX A Calculating an Annual Impact Return

Financial investment decisions are usually made based on two types of return calculations: net present value (NPV) and internal rate of return (IRR). The NPV factors out the cost of the capital deployed, so has the simplicity of yielding a single number that incorporates the duration of the investment; with NPV, an investment that takes 10 years to yield a return can be compared to an investment that returns cash to the investor in one year.

The NPV, however, is appropriate only in instances where the investor has unlimited access to capital. By incorporating into its analysis the cost of that capital, it focuses only on the returns to the investor deploying the funds, not the fundamental owner of the funds. Thus, its use is most often appropriate in a corporate or project-type setting.

It is tempting to use a similar approach to evaluating impact, by assigning a dollar value to impact a priori (e.g., \$1 per ton of carbon dioxide equivalent) and calculating the value of the total impact achieved in dollar terms. However, this fails to be useful unless the philanthropist is really willing to buy an unlimited amount of impact at that price, which is rarely the case, especially if market mechanisms (such as carbon-trading markets) offer lower prices. In this way, it is similar to the NPV, which implicitly assumes that access to capital is unlimited.

Investors use the IRR because it does not incorporate the cost of capital; rather, the return measured is the total return to the holder of capital, including the rent of the capital deployed. This is appropriate both for the owner of the capital and in instances where the pool of capital is constrained, as in an endowment or investment fund. Thus, the IRR is a more appropriate metric for foundations because their pool of capital is limited by their endowments.

Similarly, because philanthropists are not usually open to buying unlimited amounts of impact at a given price, we believe the IRR is a better inspiration for the quantification of impact than the NPV. This raises a problem, however: the simplicity of the IRR as a metric is due to the fact that the numerator (cash earned) is the same as the denominator (cash invested). When considering impact generated per dollar invested, the numerator and denominator are, of course, different. Thus, we reinterpret the IRR to be “annual return on cash invested,” which can come in the form either of cash or of GHGs, measured in carbon dioxide equivalents. (See Figure 7.)

FIGURE 7 Explanation of the approach used to develop a metric for impact return comparable to the IRR for financial return.

	Financial	Impact	Situation
NPV	$NPV = \sum_{n=0}^N \frac{C_n}{(1+r)^n}$ <p><i>N</i> = total number of periods <i>n</i> = period <i>C</i> = annual CASH flow <i>r</i> = discount rate, CASH</p>	$PV = \sum_{n=0}^N \frac{C_n}{(1+r)^n}$ <p><i>N</i> = total number of periods <i>n</i> = period <i>C</i> = annual CARBON flow <i>r</i> = discount rate, CARBON</p>	When you can borrow freely at the cost of capital
Annual Return	<p>CASH earned per dollar invested, per year</p> <p>IRR = discount rate <i>r</i>, where NPV = 0</p>	<p>CARBON earned per dollar invested, per year</p> <p>A close proxy = Total CARBON (PV, kg) / Investment (\$, PV) /years (years = return-weighted investment term)</p>	When capital is constrained

APPENDIX B Results of Analysis

Project Name:	Type	Actual or Hypothetical	Investment Term (years)	Impact Duration	Investment Size	Expected Annual Financial Return (IRR)	Expected Annual Impact Return (CO ₂ e, kg)
Solar Product Company A	WC Loan	Actual	10	13	5,000,000	-1.9%	20.91
Preinvestment Facility (infrastructure)	Loan	Actual	6	16	5,000,000	1.0%	22.81
Modular Green Homes	Loan Guarantee	Actual	2	15	1,500,000	1.61%	0.80
Clean Cookstoves (with credit revenue)	WC Loan	Actual	5	6	1,000,000	1.88%	-
Solar Product Company B	WC Loan	Actual	2	5	2,000,000	2.45%	22.53
EE Finance Fund*	Loan Fund	Actual	5	15	2,500,000	4.10%	26.62
Truck Retrofit Fund	WC Loan	Actual	7	25	3,000,000	4.34%	7.57
Ranchland Restoration Fund A	Fund	Actual	11	11	2,000,000	8.14%	0.38
REDD Fund B (with partial credit revenue)	Fund	Hypothetical	20	20	75,000,000	8.63%	26.41
Tax Equity Wind	Tax Equity	Actual	15	20	50,000,000	8.68%	1.46
Diesel Replacement	Project Equity	Actual	10	15	1,000,000	9.36%	0.22
EE Project Equity*	Project Equity	Hypothetical	10	15	345,000	10.22%	2.01
Tax Equity Distributed Solar	Tax Equity	Actual	6	20	50,000,000	12.06%	2.23
REDD Fund A (with credit revenue)	Fund	Actual	10	10	10,000,000	12.80%	-
Biomass Power Plant	Equity	Actual	16	16	950,000	13.14%	5.03
EE Project Company*	Equity	Hypothetical	7	14	750,000	13.94%	0.73
Ranchland Restoration Fund B	Project Equity	Actual	6	6	7,745,472	15.27%	0.01
Solar Product Company B	Mezzanine Loan	Actual	2	5	2,000,000	15.40%	22.71
Shipping Technology Series B	Equity	Actual	8	18	5,000,000	17.24%	18.07
NYC Taxi Conversion	Loan Fund	Hypothetical	12	12	26,785,200	18.92%	1.96
LA Taxi Conversion	Loan Fund	Hypothetical	14	14	12,183,730	21.83%	2.58
Energy Emissions Reduction Company	Equity	Actual	10	10	1,500,000	28.08%	18.71
Conservation Tillage Project	Grant	Actual	5	10	400,000	-100.0%	19.98
REDD Fund A (with no credit revenue)	Grant	Hypothetical	5	28	10,000,000	-100.0%	2.23
Climate Advocacy Grant	Grant	Actual	4	20	50,000,000	-100.0%	38.56

*EE = energy efficiency.