




September 2012

## Doing Diverse Good Balancing Risk and Return

Working paper



**Redstone Strategy Group** is a leading advisor to private foundations and non-profits worldwide. We help clients identify their highest-return investments, track and learn from results, and continually improve their efforts to solve urgent social problems. Our approach combines substantial experience across all sectors of philanthropy with deep appreciation of our clients' knowledge and expertise. This allows us to collaborate effectively with clients as they improve their ability to achieve social good and learn from their results.

Portfolio theory has much to offer grantmakers and grantees that are becoming increasingly sophisticated in how they seek social return on investment

## Introduction

In 1952, The Journal of Finance published “Portfolio Selection” by Harry Markowitz, who would go on to win the Nobel Prize in Economics for his work. “Modern portfolio theory” sprang from that first paper, providing investors with a mathematical model to underpin the practice of diversification – investing in a range of companies, assets, and industries to optimize risk and return.

Diversification predates financial markets. For centuries it has offered investors ways to minimize downside risk without reducing the average rate of return. Markowitz, in reviewing the history of the formal theory his paper helped establish<sup>1</sup>, cites the opening of Shakespeare’s Merchant of Venice, when the merchant Antonio declares:

My ventures are not in one bottom trusted,  
Nor to one place; nor is my whole estate  
Upon the fortune of this present year;  
Therefore, my merchandise makes me not sad.

*Act 1, Scene 1*

Diversification appeals not only to fund managers and Venetian merchants. It reduces the risks of a wide range of activity. Research suggests, for instance, that Indian farmers marry their children to suitors in distant villages to help diversify their income. When local crops fail, they can fall back on support from their in-laws, who with any luck remain unaffected by the disaster<sup>2</sup>.

Modern portfolio theory builds on that age-old, intuitive course of diversification by pointing the way to metrics and analyses that investors can use to more accurately balance risk and return. Analysis like portfolio optimization, for example, allows young investors to identify a high risk, high return portfolio to grow their wealth over time, while also helping older investors choose a conservative portfolio to preserve their nest eggs.

Today, portfolio theory has much to offer grantmakers and grantees that are becoming increasingly sophisticated in how they seek social return on investment. Indeed in 2007, Randall Ottinger provided a few general examples of how the Prostate Cancer Foundation applied the logic of portfolio theory to its activities<sup>3</sup>. In this paper, we explore in detail how portfolio theory can help philanthropists seeking to balance the risk and return of their grant portfolios – sets of grants with shared goals. By applying these principles to some of the decisions the Hewlett Foundation’s Education Program made when designing its “deeper learning” initiative, we show that portfolio theory offers practical insights about how to balance risk and social return.

## Why diversification matters

Many philanthropies understandably are concerned about the risk that grants and programs will fail. There are pressing needs in the world, and every dollar spent on a failing program is a dollar not invested in a potential success. In financial investing,

If a portfolio concentrates investments in highly related grants, all of those grants are exposed to similar risks and uncertainties, and the risk of the portfolio increases

there is the risk of financial loss with every investment. Stock prices drop, bond issuers default. In social investing, the downside risk is the cost of failure – the persistence of social problems that philanthropy could ameliorate with the right grants.

For several years, philanthropies have been estimating the expected social return on investment of their grants. For example, the Hewlett Foundation developed an “expected return” methodology to estimate the benefits, costs and likelihood of success – or risk – of its grants<sup>4</sup>. Similarly, the ClimateWorks Foundation applies “expected value analysis” to determine the effectiveness of applying its scarce program dollars towards efforts to reduce greenhouse gas emissions around the world. Embedded in these estimates are theories of change grounded in empirical data and expert intuition that describe how grants generate social benefits, and how risky they are<sup>5</sup>.

Some philanthropies pursue a “risk averse” approach to grantmaking and make grants with extremely high likelihoods of succeeding, such as grants to local soup kitchens or homeless shelters. These grants are very likely to successfully deliver food and shelter to the dispossessed, and thus carry limited risk. Other philanthropies maximize potential impact by pursuing revolutionary and large-scale solutions, such as promoting a global deal to address climate change, even if these efforts carry a much higher risk of failure – this is considered a “risk neutral” approach because it maximizes return in spite of the risk. Neither of these approaches is inherently better; the challenge lies in deciding on a tolerance for risk and then making the best bets that fit that approach.

To this end, some philanthropies are pushing to better understand how their grants may complement one another to achieve maximum social impact. Sean Stannard-Stockton appropriately terms efforts to develop such integrated grant portfolios as “strategic philanthropy”<sup>6</sup>. For example, the Hewlett Foundation’s grants to develop new tools for teaching “deeper learning” can increase the impact of school systems that receive grants to pilot these approaches. So grants to individual schools or to experts developing tools may be more attractive in combination than they are individually.

However, every grant is exposed to risk that events will conspire against it, and uncertainty about whether the theory of change is accurate. If a philanthropist’s portfolio concentrates investments in highly related grants, all of those grants are exposed to similar risks and uncertainties, and the risk of the portfolio increases. It is like a Venetian merchant who sends three ships on the same route to the same port to sell the same goods. If the weather is good, the port is safe, and the market is strong, then the merchant is a wealthy man. If the merchant is not so lucky he is ruined. And so Antonio diversifies.

Portfolio theory can help philanthropies create sets, or portfolios, of grants that strike the right balance between risk and return – that diversify investment to reduce risk, but carefully concentrate investments for maximum return.

## Rules of thumb for smart diversification

Portfolio theory is helpful when a program is considering various grants with the same social goal in mind – reducing unmet need for reproductive health in francophone west Africa, or reducing greenhouse gas emissions in the US. Before embarking on an analytical approach to portfolio theory, it is worth considering what simple rules of thumb portfolio theory suggests to program officers seeking to avoid unnecessary risks in these situations.

- **All else equal, choose the independent grant.** Just as in financial investing, diversification is a meaningful way for philanthropies to manage the variability of their grant portfolios. It is possible to generate higher returns at lower risk through diversification, so it is worth stopping to consider whether grants that are otherwise equally appealing concentrate resources in one region, one opportunity, one theory of change; or, whether the likelihood they succeed is less correlated with existing grants.
- **Beware the engineer's solution.** A neat, logical, and internally consistent theory of change is often a boon to grantmakers when designing a strategy. But investing in only one elegant theory of change can involve unnecessary risk because every grant rests, like the theory itself, on a few foundational premises. In the world of social change those premises are often highly uncertain. Thus, before finalizing a strategy, identify the premises upon which most of the grants rely, prepare to track them carefully, and seek alternative theories and grants that would not rely on them.
- **“A little bit of everything” is not the right answer either.** If the engineer's solution can be too aggressive, granting to “a little bit of everything” can be too conservative. Program officers often find it difficult to turn away from all of the good options they identify in the early stages of planning. But by choosing to concentrate in a few viable theories of change, a philanthropy can create a rigorous framework for diversifying its portfolios. This approach is likely to work far better than under-funding many potentially attractive opportunities for the sake of diversification.
- **Avoid the riskiest and safest portfolios.** For risk neutral institutions committed to pursuing the highest risk, highest return solutions, a little diversification can go a long way to increasing the likelihood that the portfolio produces meaningful social impact. And for institutions committed to consistently delivering impact year in and year out, making a few risky bets can increase the potential return of the portfolio substantially without significantly increasing risk. The costs of social returns and safety – in added risk or foregone benefits – can rise exponentially.

Our work with the Hewlett Foundation's Education Program marks a first attempt at bolstering rules of thumb by applying portfolio theory to grantmaking analytically. The next two sections describe two applications of the approach: in a simple case

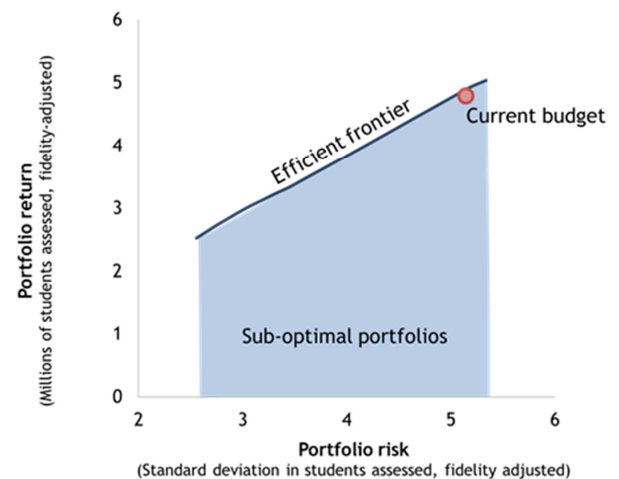
where intuition can be confirmed; and, in a complicated case where intuition can be improved.

## Charting the efficient frontier

In 2009-10, the Hewlett Foundation’s Education Program planned a portfolio of grants to support “deeper learning” in American schools, helping K-12 and community college students develop thinking skills and knowledge that contribute to economic success and civic engagement. As it planned its portfolio, the program realized that success depended on valid assessments of deeper learning being proven as practical and affordable in American classrooms. In 2011, the program identified six grants to support assessments and prove testing at scale. We applied portfolio theory to help the program analyze those grants.

When we began the analysis, the program had a plan. The program recognized that investing in six assessments was not equivalent to a Venetian merchant sending six different ships to six different ports; some of the assessments relied on similar principles, and would be implemented in overlapping sets of schools. They had considered these potential correlations qualitatively, and hoped to use portfolio theory to check their intuition and identify ways to improve.

**Figure 1. The efficient frontier with six grants**



By assuming quantitative correlations between the six grants based on qualitative factors, an analysis based on portfolio theory identified a set of hypothetical portfolios with the highest return at any level of risk. These hypothetical portfolios lie along an “efficient frontier” (Figure 1), and offer the best possible tradeoffs between risk and return.

For these six grants, portfolio return is estimated as the expected number of students assessed by deeper learning as a result of their grants, adjusted by tests’ coverage of deeper learning skills. Portfolio risk is estimated as the variability, or standard deviation, of portfolio return. Riskier portfolios will tend to include more grants that are both individually risky and highly correlated; they will tend to succeed and fail together, producing big swings in outcomes.

When designing a strategy with dozens of grants, quantitative analysis guided by portfolio theory can be a good check on intuition

As the red “current budget” point demonstrates, the program’s plan was a good one – less than five percent from the efficient frontier along either risk or return, and well within the margin of error for this analysis. But the analysis did point to one opportunity for possible improvement: Shift some investment to one of the assessment grants the program was considering, but had not yet approved.

This portfolio of six grants offers a good test case for portfolio theory. The quantitative analysis confirmed the program officers’ rules-of-thumb reasoning, but offered more. First, it offered an option to improve the strategy. And, second, it identified an efficient frontier so that if the program wanted to reduce risk in the future – say because it decided that a small proof of concept was more important than achieving scale – the efficient frontier instantaneously offered a range of alternatives for discussion.

## Navigating complexity

With six grants, rules of thumb work well. But when designing a strategy with dozens of grants, quantitative analysis guided by portfolio theory can be a good check on intuition. Before they faced the question of how to demonstrate affordable assessments, the Hewlett Education Program outlined a multi-year strategy that considered nearly 40 clusters of grants. With that many investments under consideration, rules of thumb can be useful, but even the most disciplined decision-makers will struggle to consistently apply their intuition; a model can help.

As we have discussed, portfolio risk is not merely determined by the risk of individual grants, but also by the degree to which grants in a portfolio are interrelated. Financial analysts use the “Beta” coefficient to track the correlation between an asset and the broader market. Stocks with Betas over one tend to not only move with the index, but exaggerate market swings, and so are relatively risky. For any individual stock, it is easy to calculate a Beta based on historical data on the price of the stock and the price of the market.

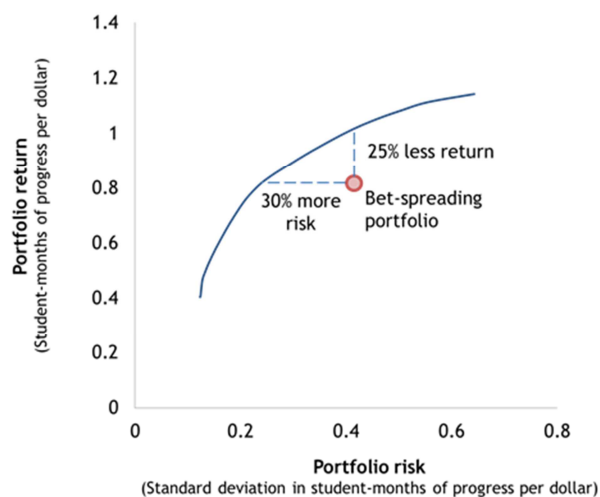
Philanthropic investors, on the other hand, do not have years of historical price data to work with. As a result, they must base their estimates of the relationships between investments on the hypotheses embedded in their theories of change, informed by empirical data and expert opinion. It is nearly impossible to track interactions between 40 investments while making intuitive judgments.

The program wanted to diversify, but did not know when to stop; at what point does diversification become “a little bit of everything”? Early on, the program considered an extreme bet-spreading portfolio, in which they would invest something in nearly 90 percent of the possible investments they identified. These included policy-related grants, the development of new educational tools, the creation of model school networks, and educational research. As a result of this bet spreading, only two investments would be fully funded. Co-funders would need to be involved in the other cases or the activities would be scaled down, reducing their reach and chances of success. By not concentrating funds, the portfolio seemed to minimize risks.

Philanthropic portfolio analysis cannot draw on reams of empirical data; it requires documenting intuitive beliefs and assumptions in quantitative terms

Figure 2 illustrates that this extreme bet-spreading portfolio (the red point) had strayed dangerously into “a little bit of everything” territory, and would be well below the efficient frontier (the blue line). The program could generate a substantially higher return at precisely the same level of risk. Indeed, the return of the bet-spreading portfolio was 25 percent lower than a portfolio with equivalent risk on the efficient frontier, and 30 percent higher than a portfolio with equivalent return.

**Figure 2. Bet-spreading and the efficient frontier**



This is in part due to the sheer number of grants funded by each portfolio. An equivalently risky, higher return portfolio on the efficient frontier would fund half the number of grantmaking activities, with five times as many (11) fully funded by the program. The quantitative estimates involved here are rough, and so the results should be taken with a scoop or two of salt. But large gaps between the proposed portfolio and the estimated efficient frontier should prompt a reevaluation.

Bet spreading is valuable for risk reduction, but portfolio theory can help identify when investments are spread too thin. As in financial investing, exposure to more investments does not necessarily mean increased diversification. Since every retail company is exposed to changes in consumers’ disposable income, most are likely to be quite vulnerable in times of recession. Buying stock in every retailer does not reduce an investor’s exposure to this risk. Instead, cautious financial investors may try to invest in only the healthiest retailers, and then invest in strong companies in a variety of other sectors as well.

## How we performed the analysis

There is no denying it: Philanthropic portfolio analysis cannot draw on reams of empirical data. Instead, it requires documenting intuitive beliefs and assumptions in quantitative terms. As we just saw, a model adds value by revealing implications that a qualitative understanding might miss, not by producing results to be pursued blindly without careful investigation.

To estimate portfolio risk and return for the Education Program, we began with the four factors that many philanthropies already consider as they design grant portfolios: the benefit, cost, likelihood of success, and contribution of the potential grants they



identified. Following the Hewlett Foundation's "expected return" approach, we collected research and interviewed experts to estimate these factors.

We then estimated a fifth factor: the correlation between each pair of investments. With a theory of change in hand, there is a relatively intuitive method to turn qualitative insight into a series of quantitative estimates. The key question is this: How many times more likely is "Grant X" to deliver social benefits if "Grant Y" succeeds than if it fails? Even when there are hundreds of pairs of investments with correlated outcomes, intuitively answering this simple question typically takes hours, not days.

Take the decisions at the heart of the Education Program's strategy. It is a truism in American education that what is tested is taught, and so the program believed that valid and affordable tests of deeper learning skills would be essential to widespread adoption. Therefore, if the Program's grants succeeded in producing high-quality tests, we estimated that advocacy to support national adoption of deeper learning educational standards had a reasonable chance of success. If, on the other hand, strong tests did not emerge, national adoption of deeper learning standards would be dead in the water.

For the sake of this analysis, we estimated that national standards adoption were 20 times more likely to succeed if strong tests emerged than if they did not. This does not mean that new standards are a certainty with new tests in hand, only that they are highly unlikely without them. Of course, this estimate is imprecise, but the Program's theory of change posits that the two grants are very strongly related and our estimate reflects this hypothesis. The analysis can be easily updated with new information, and the sensitivity of the analysis to this input can be tested.

With the joint probabilities in hand, we then estimated portfolio risk and return, and identified a portfolio with the highest rate of return at any given level of risk. While these estimates were rough, they provided a richer understanding of grantmaking portfolios than would be possible without them, and produced an efficient frontier for the program to consider as it went about its grant making.

## Conclusion

No one embarks on a complicated analysis because it is easy, and no modeled result deserves to be held up as incontrovertible truth. But portfolio theory strongly suggests that there are powerful forces at work within grantmaking portfolios. Spotting them, describing them, and quantitatively estimating them are often the only steps that will allow a program to effectively navigate them.

For decisions involving a handful of investments – six assessments to prove deeper learning at scale or a merchant's fleet – rules of thumb that have held true since Shakespeare's day are a useful guide. For decisions that involve dozens of investments, a model can be a useful tool.

---

Cover photo by Rafael Matsunaga: <http://www.flickr.com/photos/rednuht/479370088/>

<sup>1</sup> Markowitz, Harry. M. (1999). The Early History of Portfolio Theory: 1600-1960. *Financial Analysts Journal*. 55 (4), pp. 5-16.

<sup>2</sup> Rosenzweig, Mark. R. & Oded, Stark (1989). Consumption Smoothing, Migration and Marriage: Evidence from Rural India. *Economic Journal of Political Economy*. 97 (4), pp. 905-926.

<sup>3</sup> Ottinger, Randall (2007). Portfolio Philanthropy: How philanthropists can apply portfolio theory to make wiser social investments. *Stanford Social Innovation Review*, Fall.

<sup>4</sup> Redstone Strategy Group (2008). Making Every Dollar Count: How Expected Return Can Transform Philanthropy. <http://www.redstonestrategy.com/reports.php?action=detail&publicationID=18>

<sup>5</sup> Brest, Paul (2010). The Power of Theories of Change. *Stanford Social Innovation Review*, Spring..

<sup>6</sup> Stannard-Stockton, Sean. The Three Core Approaches to Effective Philanthropy. *Stanford Social Innovation Review Blog*. [http://www.ssireview.org/blog/entry/the\\_three\\_core\\_approaches\\_to\\_effective\\_philanthropy](http://www.ssireview.org/blog/entry/the_three_core_approaches_to_effective_philanthropy)