Stanford SOCIAL INNOVATION Review

Features

Toward the Efficient Impact Frontier

By Michael McCreless

Stanford Social Innovation Review
Winter 2017

At Root Capital, leaders are using tools from mainstream financial analysis to calibrate the role that subsidies play in their investing practice.

TOWARD THE EFFICIENT IMPACT FRONTIER

BY MICHAEL McCRELESS

Richard Tugume is a portfolio manager at Root Capital, a nonprofit lender to agricultural enterprises that connect small-scale farmers in Africa and Latin America to markets for their crops. His portfolio includes about 20 loans to businesses in Uganda and the Democratic Republic of Congo (DRC). Those loans, like all loans made by Root Capital, are designed to enable borrowers to create positive social or environmental impact.

In 2015, Tugume's portfolio included large borrowers such as Uganda Cocoa & Commodities (UCC) and Gulu Agricultural Development Company (GADC). UCC connects more than 5,000 cocoa farmers to export markets. GADC sources cotton, sesame, chilies, and sunflower from 60,000 smallholder farmers in Gulu, a district in Uganda that is recovering from 25 years of conflict.

One of the smaller businesses in Tugume's portfolio in 2015 was Furaha, a coffee cooperative in war-torn eastern DRC. Back in 2013, when Tugume first conducted due diligence on the company, he had to plan his visits carefully to avoid local militias that were active in the region. Many Congolese farmers had no option but to smuggle coffee into Rwanda, where they bartered it for food and supplies. Farmers who join Furaha, by contrast, not only gain a route to a safer and more reliable market, but also receive a price premium that Furaha has negotiated with foreign coffee buyers. In addition, Furaha provides clean water and electricity to farmers. Unlike the loans to UCC and GADC, Root Capital's loan to Furaha required a subsidy: The cost of lending to the cooperative was greater than the interest that it would pay to Root Capital.

For Tugume, building a high-impact, financially sustainable loan portfolio requires a delicate balancing act. "Each year, I try to make five or six big loans to large, well-established businesses," he says. "These loans provide revenue to Root Capital, and the businesses meet our social and environmental criteria: They purchase crops from local farmers and often provide services like agronomic training and farm inputs. Then, in the rest of my portfolio, I make much smaller loans to earlier-stage businesses that have a harder time getting loans but show potential for growth."

Tugume, in short, has developed an intuitive approach to creating a portfolio that generates both impact and revenue. His intuition is a powerful tool—but intuition is fallible, and it isn't scalable. To achieve impact on a large scale and to do so efficiently, he and other portfolio managers need analytical tools that will equip them to make lending decisions in a quantitative and holistic way. My colleagues and I at Root Capital are implementing such tools, and we believe that other practitioners in the broad field of impact investing may find them useful.

DEVELOPING A TOOL KIT

At Root Capital, we practice a specific kind of impact investing. We operate in a space between traditional philanthropy (in which donors expect no return on grants that they make) and mainstream financial markets (in which investors expect market-rate returns). We draw on both grant funding and private capital, and our fund is "concessionary," in that our investors forgo the chance to earn the highest possible financial return on their money. Most of our loans fall into one of two categories: Either they yield a negative financial return to Root Capital and thus require a subsidy, or they yield a positive but below-market return. Critical to our investment strategy is our ability to subsidize loans in the first category with revenue from loans in the second category. We also rely on grant funding to support the subsidized loans.

But are we allocating our philanthropic and investment funding in the best possible way? For most of our history, we have not had a way to answer that question with a high degree of rigor. In that regard, we are not alone. Until recently, intuitive methods have had to suffice as impact investors' primary means of balancing financial return with impact. From rural Uganda to Wall Street, investors have lacked tools that quantitatively and comprehensively take into account all of the factors—financial, social, and environmental—that define the performance of both individual investments and investment portfolios. As a result, the process of allocating capital to achieve impact has been inefficient at best and inaccurate at worst.

We set out to create a tool kit that would support decision making at the level of either a single investment or an entire portfolio. To do so, we first had to develop a way to *integrate* data on the financial, social, and environmental (FSE) performance of our loans. We needed a way to view FSE data as part of a single picture—a way to analyze how different FSE goals relate to each other and to identify where trade-offs between impact goals and financial goals might be necessary.

A breakthrough came when we plotted our loans on a graph that measures expected return on one axis and expected impact on the other axis. Determining the expected return of each loan was a relatively straightforward task. But to measure expected impact, we had to develop a tool that synthesizes data on the social and environmental performance of each borrower—together with an estimation of the investment impact of each loan—into a single metric. By joining that metric to our financial return metric on a standard graph, we were able to generate an integrated picture of FSE performance.

In working to make sense of those data, we have found inspiration in concepts used by investors in mainstream capital markets. One especially powerful concept is that of an *efficient frontier*. Simply put, a portfolio that lies on the efficient frontier offers the greatest possible return for a given level of risk and for a given set of investment opportunities. We broadened this concept so that it would encompass not just risk and return, but also impact.

A portfolio of investments that lies on what we call the "efficient impact frontier" would offer the highest level of overall impact, relative to the cumulative financial return of those investments. (In the case of Root Capital, the return on our portfolio is usually negative; it takes the form of a required subsidy.) The idea of applying the "efficient frontier" concept to impact investing isn't new. But we are putting this idea into practice.

By adopting and adapting concepts like the efficient frontier, impact investors can make better-informed decisions about individual investments, set more comprehensive goals for their portfolio as a whole, and collaborate more effectively with donors, investors, and public officials. In the absence of such tools, they risk wasting philanthropic dollars on activities that commercial capital can fund, and they risk missing opportunities to subsidize investments that would not otherwise occur and that efficiently create a desired impact.

GATHERING DATA

We start by calculating the expected return—that is, the net profit or loss to Root Capital—and the expected impact of each loan in our portfolio.

Expected return | For the sake of simplicity, we have combined financial risk and financial return into a single metric: risk-adjusted financial return. By analyzing the performance of 1,200 loans that we made in the past, we built a predictive model that estimates a risk premium for each prospective loan. We translate this premium into a dollar value and subtract that sum from the loan's expected revenue.

Many of our loans are unprofitable, and any loan that has an expected loss is economically equivalent to a grant in the amount of that loss. Our borrowers are typically early-stage businesses with limited track records and limited collateral, and they operate in remote areas and in challenging business environments. As a result, the full cost of serving them is high. Yet few of them are able to bear that cost: The interest payments on an unsubsidized loan would cripple their operations. We therefore charge interest rates and fees that are on par with those of local commercial banks, and we seek philanthropic support to cover the rest of our expenses.

MICHAEL McCRELESS is senior director of strategy and impact at Root Capital.

Consider the loan that Root Capital made to Furaha in 2015. It was a one-year loan for $75,000, with interest and fees that came to 11 percent of that sum. The loan would thus generate $8,250 in revenue for us. (I have rounded and simplified some of the numbers in this and other examples.) Given the early stage of this business and its hard-to-reach location, the fully loaded cost to underwrite and monitor the loan was about $23,600. (As Root Capital grows and achieves economies of scale, these operational costs are declining.) The sum that we pay to investors for use of their capital typically comes to 2.5 percent of the loan amount; in the case of the Furaha loan, that cost was $1,875. Add in a risk premium of $6,300, and the total cost of this loan was $31,775. Even if Furaha were to repay the loan in full, in other words, that revenue would cover only about one-third of our costs. The expected "return" on the loan, therefore, was a net loss (or subsidy) of $23,525.

We funded more than half of that loss ($13,025) through a cross-subsidy from larger, more profitable loans, and we drew on philanthropic support to fund the remainder ($10,500). So for this investment in Furaha, the value that Root Capital, our investors, and our donors create lies in transforming a $10,500 "grant" into a $75,000 loan.

Our loans to GADC and UCC, by contrast, had more favorable expected returns. We expected GADC to borrow $1.1 million and to pay $109,800 in interest and fees. We estimated our operational costs for this loan at $13,300. (These costs were much lower than those for the Furaha loan because GADC is in a more accessible location and we were more familiar with that enterprise from previous loans.) We paid investors $28,000 for the use of their capital on this loan, and we assigned a risk premium of $36,800 to it. Total costs came to $78,100, and the loan thus carried an expected profit of $31,700. (If GADC were to repay its loan in full, then the risk premium would no longer apply, and the total realized profit would be $68,500.)

With our loan to UCC, we expected the company to borrow $457,800 and to pay $59,300 in interest and fees. We estimated our operational costs at $19,200, we paid investors $11,400, and we assigned a risk premium of $32,000. All told, we incurred $62,600 in costs, and the loan therefore carried an expected loss of $3,300. (If UCC were to repay its loan in full, then the loan would have a total realized profit of $28,700.)

Expected impact | How do we know whether the impact of a loan will justify the cost of the subsidy embedded in that loan? To answer that question, we developed a tool that we call the *expected impact rating*. This rating synthesizes various kinds of impact data that we collect on each borrower into a single number, thereby enabling us to compare the expected impact of disparate loans and to measure each loan's expected impact against its expected return. (Our

purpose here is not to advocate for impact ratings in general or for our rating in particular. Instead, it is to describe one tool that has allowed us to develop integrated FSE data.)

Our use of this tool aligns with the framework set forth by Paul Brest and Kelly Born in a 2013 article in *Stanford Social Innovation Review*. Brest and Born distinguish between enterprise impact (that is, the impact that an enterprise has on its customers, its suppliers, or the environment) and investment impact (that is, the impact that a particular investment has on that enterprise).[1] Another term for investment impact is *additionality*. According to Brest and Born, additionality reflects the extent to which a given investment provides resources that *add* to what other investors would have provided in its absence.

The expected impact rating takes the form of a number from 0 to 10. To calculate that number, our team first sorts a loan into one of three categories of additionality. The lowest category (0 to 3.0) applies to cases in which a borrower likely could have received a similar loan from a commercial lender. The intermediate category (3.0 to 6.5) includes loans that a borrower likely could have obtained from some other mission-driven organization, but not from a commercial lender. And the highest category (6.5 to 10) applies to cases in which a borrower likely could not have received a similar loan on similar terms from any other source.

Then, within a given category of additionality, we assign to each loan a score for its expected enterprise impact. This score, which ranges from 0 to 3.5, is a composite of the baseline social and environmental need that a borrower aims to address (1 point), its expected performance in addressing that problem (2 points), and its operational scale (0.5 point). For this score, we give equal weight to social and environmental considerations. (In the case of the lowest level of additionality the enterprise impact score goes up only to 3.0.) To quantify enterprise impact, we gather data on the following factors:

- Poverty level in the regions where an enterprise operates
- Expected performance of an enterprise in addressing poverty
- Environmental vulnerability, as measured by water scarcity, soil degradation, threats to biodiversity, and exposure to climate change
- Expected performance of an enterprise in addressing environmental vulnerability
- Scale, as measured by the number of farmers and workers reached by an enterprise

We elevate investment impact above enterprise impact because our aim is to subsidize only those loans that would not happen in a commercial market and because we have prescreened all of our borrowers for expected social and environmental impact. We understand that additionality is one of the most challenging aspects of our expected impact rating to evaluate: It requires our loan officers to make difficult judgments about the alternate lending options that an enterprise may or may not have. But we also understand that even if loan officers misjudge a certain portion of loans, they will make better lending decisions overall if we include additionality in our rating than if we do not. (To ensure that loan officers apply this metric consistently, we are developing a training curriculum on this topic.)

Our loan to Furaha qualified for the highest category of additionality (6.5 points). Because that business operates in an extremely challenging region of DRC, the likelihood that it could have obtained a similar loan from any other lender is very low. Furaha also received 1.5 points for expected enterprise impact. That's an average score for loans in our portfolio, and it reflects several important attributes of the company: the high poverty level of its affiliated farmers, the location of those farmers in a biodiversity hot spot, its commitment to offering a price premiums to farmers, and its provision of electricity to farmers. In sum, we awarded Furaha a score of 8.0 on our expected impact rating—a higher score than 90 percent of all loans in our portfolio.

GADC, meanwhile, has loans from other social lenders but not from commercial lenders. For that reason, our loan to that company qualified for the intermediate category of additionality (3.0 points). But GADC received a higher-than-average score (2.5 points) for expected enterprise impact. It reaches 60,000 farmers, the vast majority of whom live on less than $2.50 per person per day in communities that are recovering from violent conflict. Those farmers also live and work in a climate change hot spot. GADC helps them to obtain

Expected Return and Expected Impact

By plotting expected return on one axis of a graph and expected impact on the other, Root Capital is able to integrate data on its loans. This graph schematically renders 259 loans that were active in 2015.

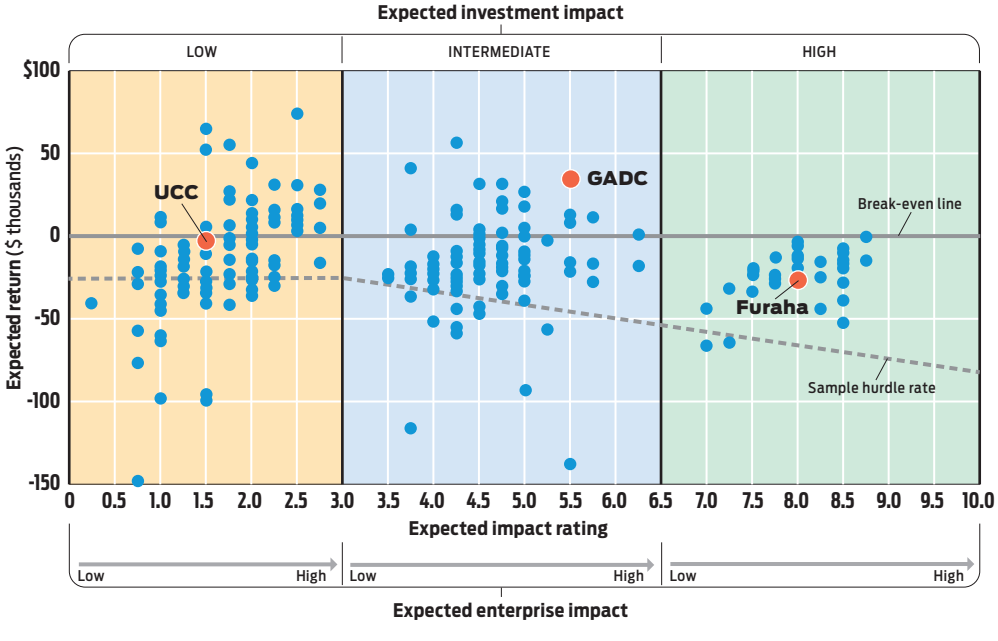

organic certification and provides them with above-market wages and access to affordable housing and health insurance. This loan, in sum, received a total expected impact rating of 5.5.

Our loan to UCC exhibited the lowest level of additionality (0 points) because the company was able to access financing from commercial banks. UCC received an expected enterprise impact score of 1.5. Most of its farmers have very low incomes, and they live in a region that is both a climate change hot spot and a biodiversity hot spot. UCC provides employees with health care and meals, and it recently began offering agronomic training to farmers. The total expected impact rating for this loan was 1.5.

ANALYZING INVESTMENTS

To create a comprehensive picture of how loans in our portfolio align with our FSE goals, we plot each loan on a graph that measures both expected impact and expected return. In this article, I have included a version of this graph that features a schematic presentation of 259 loans that were active in 2015. (See "Expected Return and Expected Impact" on page 51.) These loans are ones for which we have complete data on both expected return and expected impact. The financial data show expected returns for each loan at the time of its approval, and those data incorporate the risk premiums that our team designated at that time. (On the graph, I have highlighted the loans to Furaha, GADC, and UCC.)

Using this graph, we can determine the level of return that we might expect for a given level of expected impact. That information in turn improves our ability to evaluate individual loans. To facilitate this analysis, we have adapted a concept—that of a *hurdle rate*—commonly used by investors and financial managers. A hurdle rate is the minimum rate of return that would justify making a given investment. Similarly, an *impact/return hurdle rate* (as we call it) is the minimum score on our expected impact rating that would warrant making a specific loan. In our model, that score will vary according to the expected return for each loan. The dotted line in the "Expected Return and Expected Impact" graph shows one possible impact/return hurdle rate. If we were to adopt this particular hurdle rate, we would then decline to underwrite any loan that fell below the line.

In practice, the impact/return hurdle rate that we use is not as precise as this example would suggest. We have adapted this tool into a set of rules of thumb for our loan officers. Using these rules, loan officers work to meet targets for impact, revenue, risk, and portfolio size. Within these parameters, they have flexibility to construct a portfolio that reflects the mix of enterprises in their assigned territory and the capital needs of those enterprises. Regional directors have discretion to approve loans that don't adhere to these rules—loans, in other words, that would fall below a given hurdle rate.

At an organization-wide level, meanwhile, we can shift the hurdle rate in response to changing market conditions. If we have limited funding, for instance, we can move the hurdle rate upward and reserve our scarce subsidy funds for higher-impact loans. If more funding becomes available, we can shift the hurdle rate downward and pursue a wider variety of loans.

In analyzing our active loans from 2015, we also examined the relationships between the expected financial return of those loans and various types of expected impact. Our analysis revealed that (at least in our case) there are trade-offs between financial return and some types of impact. On average, loans with higher levels of additionality and loans to enterprises that serve relatively poor communities require a larger subsidy. Loans to businesses that reach more farmers and are doing more to help them improve their livelihoods or adapt to climate change, by contrast, tend to require a smaller subsidy or to generate a profit. (Of course, these findings are specific to Root Capital, and organizations that operate in other sectors, asset classes, and geographies may arrive at different results.)

CHARTING A FRONTIER

Building on tools like the expected impact rating and the impact/return hurdle rate, we have also developed another tool—the efficient impact frontier—to analyze and optimize the performance of our portfolio as a whole.

To see how this tool works, consider the graph titled "Expected Return and Expected Impact." Toward the bottom left corner of the graph are a handful of loans that have limited impact, that are unprofitable, and that commercial banks likely would have undertaken in our absence. These are the lowest-quality loans in our portfolio, and we should not have made them in the first place. But we could not identify them until we had integrated our FSE data and developed a way to analyze those data. To apply the efficient impact frontier concept, we model the effect on our portfolio of replacing those low-quality loans with higher-quality loans. In particular, we test what would happen if we removed the lowest-quality 10 percent of our loans and replaced them with simulated loans that have greater expected impact, a higher expected return, or both. The decision to swap out 10 percent of loans

The Efficient Impact Frontier

To optimize its performance, Root Capital conducts an analysis that involves swapping out a portion of loans in its current portfolio to create potential alternative portfolios.

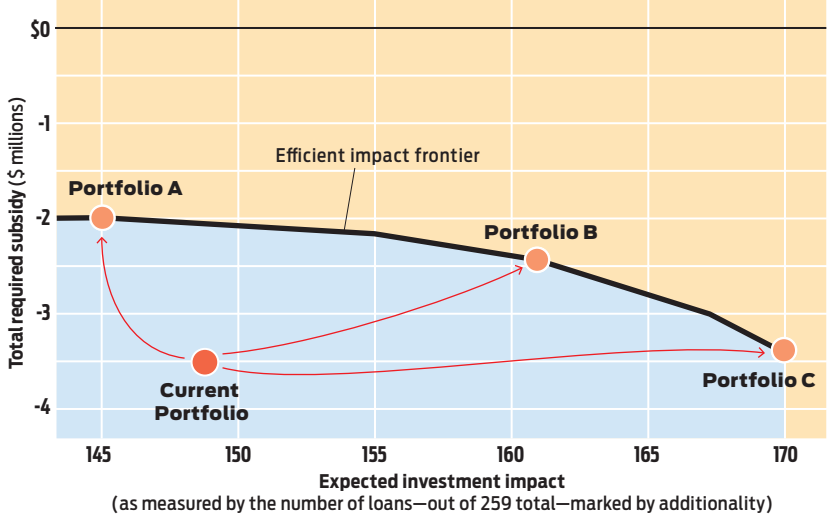

in the portfolio is somewhat arbitrary. Given our experience, we think that swapping out that number of loans annually is a plausible aspiration. But in practice, the proportion might turn out to be higher or lower.

In this article, I have included a graph that illustrates this exercise. (See "The Efficient Impact Frontier" on page 52.) On that graph, each point inside the shaded area represents a potential portfolio, and the line that marks the upper boundary of this area represents the efficient impact frontier. Each point on that line matches a simulated portfolio in which we have eliminated the 26 lowest-quality loans in our current portfolio—in other words, the worst 10 percent of that portfolio's 259 loans. Portfolio A replaces those 26 loans with duplicates of the 26 highest-*return* loans in the current portfolio. Portfolio C replaces them with duplicates of the 26 highest-*impact* loans in the current portfolio. Portfolio B replaces them with duplicates of 26 loans from the current portfolio that combine high impact with a high return. These simulated portfolios define the outer boundary of what we could achieve in light of our current opportunities.

For the sake of this exercise, we are using additionality as the measure of impact. But we have found that the efficient impact frontier assumes essentially the same shape when we use other impact metrics. In each case, if we start with our current, suboptimal portfolio, we can simultaneously increase the expected impact and the expected return of the portfolio simply by making better, more data-driven decisions about which loans to make. Once we reach a portfolio that lies on the efficient impact frontier, we face a trade-off: Achieving greater impact would require a higher level of subsidy. And conversely, achieving greater financial return would require a reduction of impact.

The impact/return hurdle rate plays an important part in this process, because it determines the number and nature of the loans that we screen out of the portfolio. Customer development (or deal origination) plays an equally important part, because it determines the kinds of loans that we can bring into the portfolio. The amount of philanthropic funding that we have available plays a crucial part as well, because it determines how far rightward on the efficient impact frontier—that is, in the direction of impact—we can aim: On the "Efficient Impact Frontier" graph, we can mark the amount of available grant funding on the vertical axis, and we can then work to build the portfolio that lies at the corresponding point on the frontier line.

Weighing these factors requires us to make a variety of educated assumptions. We are currently implementing the efficient impact frontier tool, and one lesson to emerge from this work is that high-quality data and well-designed tools do not obviate the need for human judgment. They simply provide better inputs for that judgment.

At Root Capital, we have customized the efficient impact frontier and other tools to suit our needs as a nonprofit agricultural lender. Other organizations that pursue impact investing can adapt the tools for their own use. They can start by plotting their investments on a graph that measures social performance on one axis and financial performance on the other. For any organization, the impact metrics on the horizontal axis will vary according to its mission, its theory of change, and the resources that it has available for measurement efforts. On the vertical axis, investors that seek market-rate returns might replace "required subsidy" with "risk-adjusted financial returns." This approach does not assume any given relationship between expected impact and expected return. Rather, it provides a framework for exploring that relationship and for setting realistic goals to improve impact, return, or both.

SSIR ONLINE

Visit **ssir.org** to learn more about the "efficient impact frontier" tool developed by Root Capital.

- "The Relationships Between Expected Impact and Expected Return" document
- "Root Capital's Expected Impact Rating" document

SCALING UP THE FIELD

Applying tools such as the efficient impact frontier to integrated FSE data helps Root Capital to create greater impact with a given amount of grant and investment funding. These tools, we believe, can provide the transparency and the accountability that are necessary to scale up impact investing efforts—in particular, efforts that involve subsidized (or "concessionary") investments.

There are many socially valuable investment vehicles that generate more revenue than philanthropy but less revenue than commercial investing. In the United States, for example, the median community development loan fund covers 63 percent of its costs, and even the largest such funds cover only 90 percent of their costs.[2] Many microfinance institutions, particularly those that reach highly vulnerable populations, require ongoing grant support. Or consider One Acre Fund, a nonprofit organization that supplies smallholder farmers in East Africa with asset-based financing and agriculture training services. One Acre generates revenue that offsets about 80 percent of its field operating costs and relies on grants to cover the remainder.

For many years, efforts to scale up these kinds of investment vehicles have met with a big obstacle: the lack of tools for analyzing return and impact in tandem and in a rigorous way. Measuring social and environmental impact has always been difficult. Combining grants with investments makes it difficult to evaluate financial performance as well. In the case of Root Capital, the fact that our lending generates earned revenue raises questions for our donors: "Why, given that you are generating revenue, do you need grant funding from us?" "Is my grant just subsidizing returns for your investors?" One donor asked outright, "Am I the dumbest money in the room?" (Investors in Root Capital, after all, get their money back and also earn interest.) Donors, in short, need a way to distinguish between a subsidy that is truly necessary to generate impact and a subsidy that provides cover for operating inefficiencies or poor investment decisions. So Root Capital—along with other financial institutions that blend grant funding and investment capital—must be able to clarify the roles that different kinds of capital play in their work.

The tools that we have developed can help stakeholders on both the supply side and the demand side of the capital market to make better decisions. They provide information that donors can use as they allocate grants and that investors can use as they allocate private capital. And they provide users of philanthropic and investment capital with the insight that they need to set—and achieve—integrated FSE goals. ■

NOTES

1 Paul Brest and Kelly Born, "When Can Impact Investing Create Real Impact?" *Stanford Social Innovation Review*, Fall 2013.

2 Bethany E. Chaney, "Community Development Financial Institutions: A Study on Growth and Sustainability," Mary Reynolds Babcock Foundation, June 2011; and Michael Swack, Jack Northrup, and Eric Hangen, "CDFI Industry Analysis Summary Report," US Department of the Treasury CDFI Fund and the Carsey Institute, Spring 2012.