

LEARNING FROM EXPERIENCE

A Guide to Social
Impact Bond Investing

Gordon L. Berlin

MARCH
2016

mdrc
BUILDING KNOWLEDGE
TO IMPROVE SOCIAL POLICY
■

Learning from Experience

A Guide to Social Impact Bond Investing

Gordon L. Berlin

March 2016



This report is made possible through funding from Bloomberg Philanthropies.

Dissemination of MDRC publications is supported by the following funders that help finance MDRC's public policy outreach and expanding efforts to communicate the results and implications of our work to policymakers, practitioners, and others: The Annie E. Casey Foundation, Charles and Lynn Schusterman Family Foundation, The Edna McConnell Clark Foundation, Ford Foundation, The George Gund Foundation, Daniel and Corinne Goldman, The Harry and Jeanette Weinberg Foundation, Inc., The JBP Foundation, The Joyce Foundation, The Kresge Foundation, Laura and John Arnold Foundation, Sandler Foundation, and The Starr Foundation.

In addition, earnings from the MDRC Endowment help sustain our dissemination efforts. Contributors to the MDRC Endowment include Alcoa Foundation, The Ambrose Monell Foundation, Anheuser-Busch Foundation, Bristol-Myers Squibb Foundation, Charles Stewart Mott Foundation, Ford Foundation, The George Gund Foundation, The Grable Foundation, The Lizabeth and Frank Newman Charitable Foundation, The New York Times Company Foundation, Jan Nicholson, Paul H. O'Neill Charitable Foundation, John S. Reed, Sandler Foundation, and The Stupski Family Fund, as well as other individual contributors.

The findings and conclusions in this report do not necessarily represent the official positions or policies of the funders.

Acknowledgments

This paper benefited enormously from the advice and experience of David Butler, Tim Rudd, Cindy Redcross, Elisa Nicoletti, and Mariana Veras, the team responsible for the design and operation of the Adolescent Behavioral Learning Experience social impact bond. Dan Bloom and Jean Grossman provided critical review. Mary Jo Bane, Antony Bugg-Levine, Andrea Phillips, Josh McGee, Erica Brown, and Katie Appel Duda read and commented on early versions of the paper. John Hutchins was an essential partner as editor, collaborator, and final arbiter; Joshua Malbin provided final editing and production assistance, and Carolyn Thomas prepared the paper for publication.

For information about MDRC and copies of our publications, see our website: www.mdrc.org.

Copyright © 2016 by MDRC®. All rights reserved.

Preface

Leaders increasingly recognize that confronting the challenges of today and tomorrow means finding new ways to solve problems. The social impact bond is a promising approach to addressing what are often thorny issues.

A great deal of energy surrounds this still relatively experimental approach, which makes sense because the social impact bond promises a lot: more mission funds for nonprofit organizations; the space for government entities to experiment by lowering their risk; new profit opportunities for socially minded investors; and for society, the chance to test and expand more effective interventions to existing challenges.

For all the excitement, though, one would be mistaken to believe that social impact bonds are a well-honed approach. The reality is that this innovative financing mechanism is still relatively new. It was only six years ago that the city of Peterborough in the United Kingdom became the first local government to experiment with a social impact bond. In 2012, Bloomberg Philanthropies and its partners launched the Rikers Island social impact bond — the first in the United States. At the time, there was no blueprint and little information available. That made things harder and less certain.

Since then, more than a dozen other experiments have gotten under way around the world.

Today, as the results of these pioneering projects become known, theories about how to do social impact bonds are giving way to actual lessons from the field. That's why we are pleased to support this effort by MDRC's president, Gordon Berlin, collecting insights about our experience with the Rikers Island social impact bond: what worked, what was a lot harder than expected, and what implications and lessons it offers future efforts.

It is an open question whether social impact bonds will eventually occupy a prominent and sustained position in leaders' problem-solving toolboxes. What is certain, however, is that Berlin's analysis will serve as a helpful guide to problem-solving experimentation. And that is critically important if decision makers are to get smarter at grappling with the challenges of today and tomorrow.

James Anderson
Bloomberg Philanthropies

The social sector’s hottest “impact investing” product — the social impact bond (SIB) — has generated a range of reactions, from excitement to angst. Indeed, SIB enthusiasts often promote an ideal, while critics rail at the prospect of private investors profiting from social ills. In the simplest terms, a SIB uses private funds to pay for a social, educational, or health program, and the government repays investors (plus a return) only if the program achieves prespecified results. Also referred to as pay-for-success contracts, social impact bonds bring together the concepts of lending for a social purpose, return on investment, and payment based on performance.¹

Not surprisingly, as early adopters gain real-world operating experience, reality is turning out to be more nuanced than either proponents or detractors have promised. On balance this is a good development. While much of the early activity in the field was focused on deal making, SIBs’ ultimate success or failure will be determined by operating results, not term sheets.²

In July, the intervention financed by the first social impact bond in the United States — the Adolescent Behavioral Learning Experience (ABLE) program at Rikers Island jail in New York City — was discontinued after three years when results from an independent evaluation demonstrated that it was not meeting its goal of reducing recidivism among 16- to 18-year-olds (that is, reducing the rate at which they committed new crimes or were reincarcerated). As president of the organization that was the “intermediary” at the center of the Rikers deal, I think it is important that the field not learn the wrong lessons from our experience — for example, that SIBs should focus on less disadvantaged populations in less risky settings, or that we should be satisfied with less rigorous methods for measuring the impact of SIB-financed initiatives, or that the Rikers experience was post hoc confirmation that SIBs are inherently misguided.

Instead, emerging lessons from the Rikers deal and others reveal both SIBs’ value to government entities and also the reality that this value will only be realized if the tensions inherent in structuring the terms of a SIB deal can be addressed squarely. These include:

- **The balance of risk and reward.** The returns governments are willing to pay may not be proportional to the risks some lenders are able to take.
- **The focus on government savings.** Many deals still depend on the possibility of government savings. But insisting on government savings can unnecessarily rule out projects that might otherwise offer valuable social returns.
- **The tyranny of SIB metrics.** For SIB-funded programs to meet the expectations of all parties, it is not enough for them to have the desired effect on participants.

¹For an early description of the SIB concept see Overholser (2007). See Social Finance, Inc. (2012) to understand how the concept evolved in the U.S. context; also see Liebman (2011), who makes the case from a federal and state policy perspective and emphasizes the role of evidence.

²See Third Sector Capital Partners (2013); Third Sector Capital Partners (2014).

To meet a deal's savings targets they must also serve a prespecified number of people and do so in a fixed span of time. These expectations are codified in relatively inflexible lending agreements. Yet to be successful, social interventions must be able to adapt to unforeseen developments.

- **The role of evidence.** SIBs reduce risk for government entities by promising that they will have to pay *only* for successful interventions.³ To fulfill that promise, a SIB must include independent, rigorous evaluation of a program's effectiveness over the status quo — a requirement that poses a new form of investor risk. Unfortunately, efforts to mitigate that risk have driven some SIB deals to rely instead on simple outcome measures, which may misleadingly provide only the illusion of benefits and savings to government entities.⁴

In short, at the core of the continuing dispute about the potential role of SIBs in helping to address America's social needs lies a failure to appreciate the essential ways that SIBs differ from more traditional social lending — namely, that repayment depends on a social program's actual effects on participants, a feature that fundamentally alters the risk calculus for investors.

This paper draws on lessons from the implementation of the ABLE program and the handful of other SIBs with actual operating experience to provide valuable insights into the inner workings of SIB deals, using ABLE as a vehicle for explaining both the challenges and the potential of SIBs. It begins with an overview of the SIB promise, and then briefly describes the ABLE project. Next it frames critical decisions and identifies central tensions in the design of a deal, including the unique requirements of due diligence in a SIB project. Then it walks through the steps involved in structuring a deal, showing how the tensions inherent in the SIB mechanism must be confronted. It concludes with thoughts about how to move beyond a series of “bespoke” deals to effective, ongoing programs operating on a large scale, whether through sustained government funding or through the creation of functioning capital markets that afford a continuing role for private investors, especially philanthropies.

Social Impact Bonds: The Win/Win/Win Promise?

In its idealized form, the SIB structure promises three benefits: (1) the government gets risk-free flexible funding to test and expand new approaches; (2) investors reap double-bottom-line returns — a social return and repayment of their investments plus interest; and (3) nonprofit service deliverers get access to new sources of capital, committed up front for multiple years, to take programs to a large scale. In contrast to other forms of social lending, wherein rents or rev-

³Liebman (2011).

⁴The trend toward outcomes-based SIBs can be seen in Figure 10, page 20, of Gustafsson-Wright, Gardiner, and Putcha (2015).

enue streams from the delivery of social services are used to repay loans, in a SIB, repayment depends on the findings of an independent third-party evaluation of the program's impacts. This repayment structure guarantees accountability, because the government only "pays for success" after benchmarks are met. In short, SIBs promise to be a win/win/win venture for all — the government, investors, and programs (and, one hopes, the people they serve).

In practice, however, because a SIB investment involves risk, all three parties can also experience failure. As happened at Rikers, a SIB-supported program that turns out not to work can result in losses for investors, no cost savings for the government, and no opportunity for long-term growth for nonprofit service deliverers. Yet while the intervention may *not* have worked, the SIB structure can still succeed: the government does not pay for ineffective services, both the nonprofit service delivery organization and the government get the flexibility to innovate, and learning what does not work can help point the way to what might. And the social sector, including both government entities and nonprofit service deliverers, will have taken an important step toward basing policy and practice on evidence.

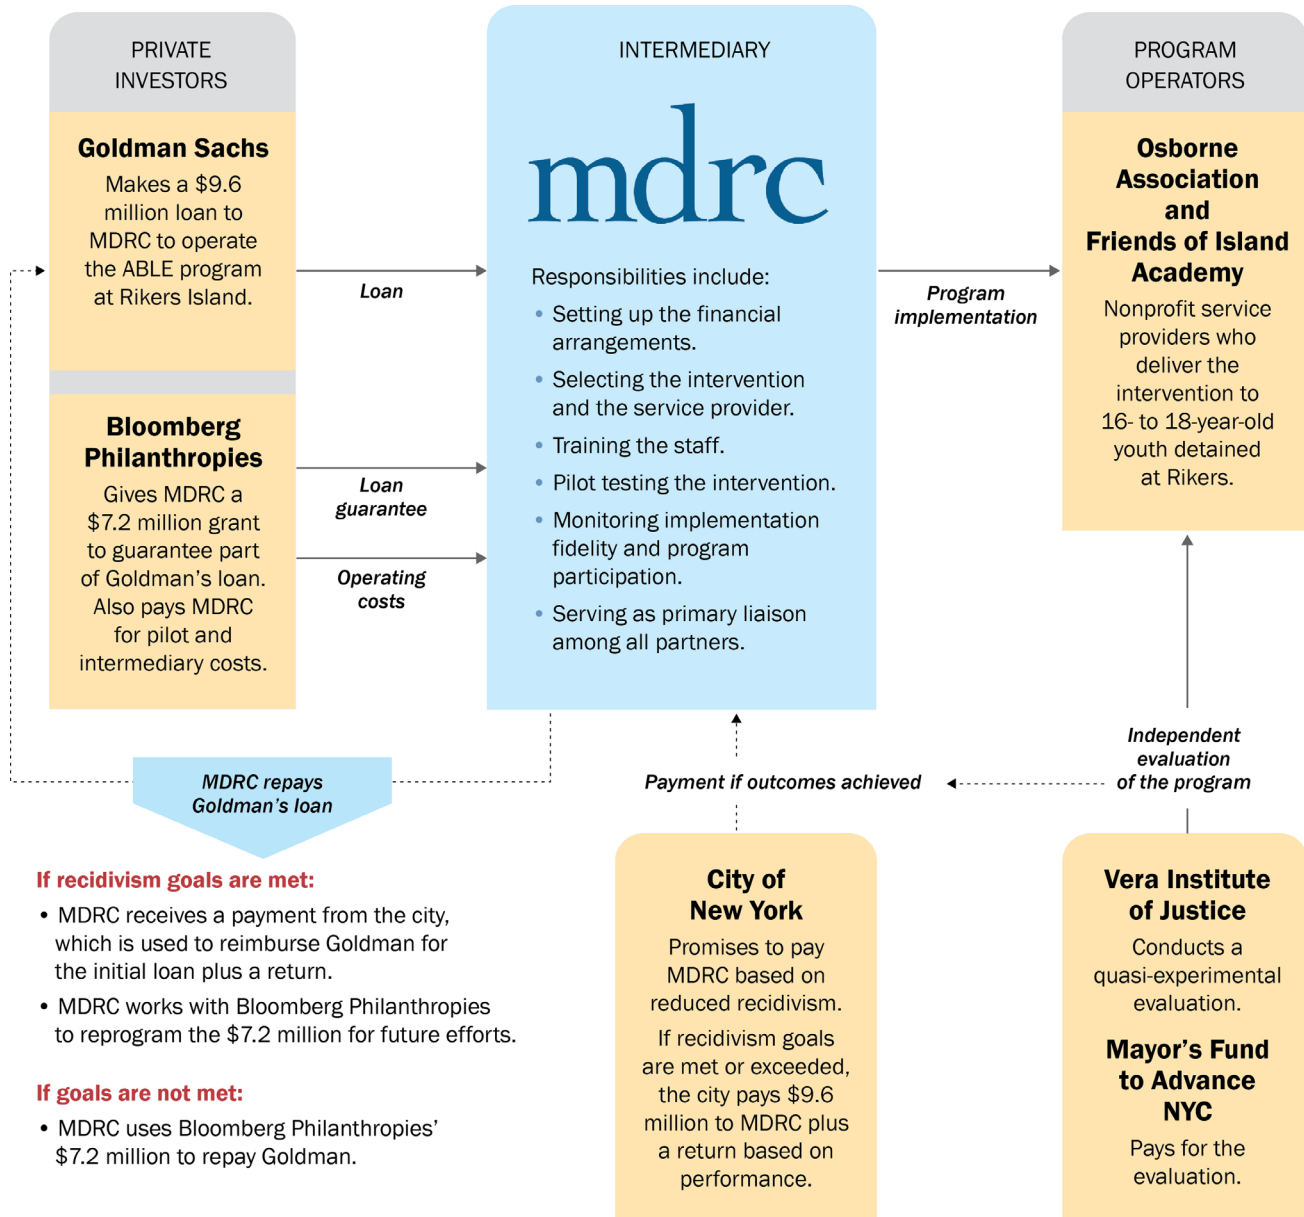
The Rikers Island SIB

At the inception of the Rikers Island SIB, some 3,000 16- to 18-year-old, mostly low-income adolescents of color were spending anywhere from a few days to as much as a year in jail awaiting trial, and nearly one in two returned to jail on a new charge within a year of being released. For the young people themselves and for society at large, this downward spiral was tragic and costly.

In the hope of reducing recidivism, the New York City government opted to offer a cognitive behavioral therapy program with a strong record of success — Moral Reconnection Therapy (MRT) — to all 16- to 18-year-olds on Rikers Island, starting in early 2013, as part of a larger effort under the New York City Young Men's Initiative. Delivered through the jail's on-site high school, MRT was designed to equip adolescents with the social and decision-making skills to help them take responsibility for their actions, avoid rash and impulsive reactions, make more constructive life choices, and, ultimately, avoid a costly return to Rikers.

Goldman Sachs provided program financing, Bloomberg Philanthropies provided grant support to the intermediary (MDRC) and guaranteed a portion of Goldman's investment, and MDRC negotiated the deal and oversaw project implementation. Working with the New York City Mayor's Office and the Department of Correction (DOC), MDRC engaged The Osborne Association to run the program in collaboration with Friends of Island Academy. The Vera Institute of Justice conducted an independent evaluation (see Figure 1).

Figure 1: Key Partners in New York City’s Social Impact Bond



Goldman Sachs committed to investing \$9.6 million over four years to fund the program, of which \$7.2 million was guaranteed by Bloomberg Philanthropies. If recidivism rates fell by 10 percent relative to a comparison group, the city would have paid back Goldman Sachs in full; if the program reduced recidivism by more than that amount, the city would have paid an additional return on a capped, sliding scale (see Table 1). The deal also included the possibility of a partial payment if recidivism was reduced by 8.5 percent.

Table 1: Payment Terms for the Rikers Social Impact Bond

Two-Year Recidivism Reduction Rate	Department of Correction Success Payment (\$)	Net Projected Taxpayer Savings (\$)
≥20.0%	11,712,000	20,500,000
≥16.0%	10,944,000	11,700,000
≥13.0%	10,368,000	7,200,000
≥12.5%	10,272,000	6,400,000
≥12.0%	10,176,000	5,600,000
≥11.0%	10,080,000	1,700,000
≥10.0%	9,600,000	<1,000,000
≥8.5%	4,800,000	<1,000,000

Source: Rudd, Nicoletti, Misner, and Bonsu (2013).

At the three-year point in the four-year project, the partners had the opportunity to review one-year follow-up data.⁵ If these early results showed that the project reduced recidivism by more than 9 percent, then the program would continue to be funded for an additional year, and the city would make a \$2.4 million success payment to investors. If the results did not meet that target, investors could end the program early and the city would pay nothing. The latter is what happened.⁶ (For a discussion of why ABLE did not meet its goals, see Box 1.)

Three features of how government savings were calculated in the Rikers deal tend to go underappreciated. These features are essential in many SIB deals, and are therefore important to understand and plan for during negotiations:

- **Government savings would always exceed returns to investors.** The deal was structured to insure that the “return” in savings realized by the city exceeded any profit made by Goldman Sachs.
- **Government savings had to be projected beyond the length of the SIB deal.** While investors were slated to be paid back by the end of Year 4, the program’s

⁵“One-year follow-up data” means data collected 12 months after a participant is released from jail. That 12-month period does not begin *until* a participant is released from jail, so each participant could have a different start to those 12 months. It took about 27 months from the time the program was launched until the last sample member reached his 12-month follow-up point.

⁶In July 2015, the Vera Institute of Justice determined that the program did not lead to reductions in recidivism for program participants and the program did not meet the target required for the city to pay the investors. The comparison group (a pre-ABLE group of 16- to 18-year-olds and a group of contemporaneous 19-year-olds) and the study group (16- to 18-year-olds) displayed similar trends in reductions of rates of recidivism over time, indicating that these shifts were the result of factors other than the ABLE program. The program, therefore, did not achieve the predefined threshold of success for the one-year evaluation of a 9 percent reduction in recidivism. See Vera Institute of Justice (2015).

Box 1

Reflections on Why ABLE Did Not Meet Its Recidivism Reduction Targets

The Vera Institute for Justice determined that the ABLE program did not reduce recidivism for 16- to 18-year-olds after one year of follow-up, and the program was discontinued in August 2015. The Vera analysis was not designed to explore the reasons why the program didn't work, but program data and observations by the MDRC intermediary team offer some possible explanations.

What went well? The New York City Department of Correction was able to offer the program daily to adolescents at an unprecedented level: Over the course of the three-year project, ABLE served nearly 2,000 adolescents each year on Rikers Island. The ABLE program was implemented with a high degree of fidelity, as confirmed by regular assessments by Correctional Counseling, Inc. (the curriculum developer) and MDRC. It was able to generate significant support from Rikers managers, nonuniformed and uniformed staff members, and school staff members. The project maintained a well-trained staff that was committed to the project, with very little attrition, which is impressive given the well-documented difficulties of working inside a jail. Finally, ABLE provided some amount of Moral Reconciliation Therapy to nearly all eligible adolescents.

What were the challenges? While MRT was delivered to almost all of the target population, the amount of MRT that adolescents received was less than ideal. Forty percent of participants were in a number of sessions associated in earlier studies with positive outcomes, but only 9 percent completed all 12 stages of the program. The expectation going into the SIB was that 30 percent of participants would complete the program. The number of sessions participants received was directly related to the instability of the jail and the inmates' indeterminate length of stay:

- On a daily basis, Rikers experiences security-related events, like lockdowns and alarms, that create extensive delays and cancellations for scheduled activities, such as school classes, MRT sessions, and recreation. When fights occur, inmates are often moved to different housing areas, which further disrupts the continuity of MRT classes.
- Because the length of stay for young people awaiting trial is unpredictable, ABLE offered MRT in several locations in the community for adolescents who left Rikers without completing the program. Few of those young people attended, however, as adolescents leaving Rikers face a host of competing obligations, including school, court-mandated activities, and sometimes jobs.

Finally, adolescent inmates experience an enormous amount of stress during their time at Rikers, limiting their ability to internalize and practice even the most promising therapeutic program. Most inmates on Rikers Island are detainees awaiting sentencing, so they do not know whether they will get to return home to their families or be sent upstate for a lengthy prison sentence. The well-documented culture of violence and intimidation at Rikers, exacerbated by a strong gang presence, means that adolescents sometimes feel compelled to align themselves with other inmates who engage in negative behaviors; this perceived compulsion conflicts with the main goals of MRT, including promoting personal responsibility and sound decision making.

total projected savings from reductions in recidivism depended on the city continuing to operate the program in Year 5 and Year 6, long after investors would have been paid.

- **Government savings were predicated on the program serving a minimum number of participants.** The contract required the program to serve approximately 3,000 participants each year, because fewer participants would mean fewer opportunities to reduce recidivism. If the program had been unable to serve the intended number of inmates, then success payments were to be prorated.

Negotiating a Deal: Major Decision Points

Selecting a Financing Mechanism

Social impact bonds are a financing mechanism, not a solution to a particular social problem. Start with the problem to be solved and the solution to be tried, then determine whether a social impact bond might be the right funding instrument. In fact, a good first step is to ask: Should the government (or a philanthropy) fund the initiative without resorting to the intricacies of SIB financing?

In the case of ABLE, the Department of Correction and the city’s Office of Management and Budget, which must approve all budget decisions, had little if any discretionary funding for innovation and experimentation. Funding a program that might reduce the likelihood of future incarceration would have been a stretch for an agency consumed by day-to-day crises. Indeed, like many resource-constrained government agencies, the DOC struggled to effectively manage daily life in the jail — from securing, feeding, clothing, and housing inmates to educating them. Further, the city did not know whether cognitive behavioral therapy could be delivered effectively in the jail, nor whether it would actually reduce recidivism among the population of adolescents housed there (even though there was good evidence demonstrating its effects elsewhere).

Outside resources marshaled under the ABLE program made it possible for the DOC to try this new way of preventing recidivism. If the SIB had not provided the necessary “risk capital,” the DOC would have been unlikely to experiment with cognitive behavioral therapy for the adolescents it housed.

Balancing Risk and Reward

A cardinal rule for most lenders is: don’t lose principal. In a conventional social lending project — for example, building low-income housing — deals are structured so that the lenders’

principal is paid back plus interest once tenants move in and rental income begins to flow. To ensure this outcome, due diligence is intense and the lending terms are exacting, designed to protect lenders against possible contingencies. Even when a flawed due-diligence assessment of a housing deal results in lower-than-anticipated rental income or construction delays lead to cost overruns, typically a lender can still expect to receive rental income eventually and thus to recover its principal. And even if a developer goes bankrupt, or unexpectedly fails to qualify for federal rental subsidies, the lender could sell the property to recover some or all of its principal. In short, a loan to build low-income housing involves comparatively knowable and manageable risks, making lenders comfortable with a return of principal plus a rate of interest that is often in the single-digit range.

But the repayment structure anticipated by a SIB differs fundamentally from this traditional loan structure. Unlike the lenders in most loan agreements, a SIB lender typically has no recourse to the assets of the borrower — the terms of the agreement provide for the loan to be forgiven if the benchmarks are not met.⁷ Repayment is triggered by the results of an independent third-party evaluation of the difference a program makes in participants' outcomes — ideally measured as the program's "net impact." In the ABLE case, the hoped-for difference was a reduction in adolescents' recidivism rate. This "trigger-point" structure alters the calculations of nonprofit and commercial lenders, because it is now possible that they might not recover some or all of their principal. Putting principal at risk makes a SIB more like an equity investment for a venture investor than a loan or "bond" repayment for a community development lender. Adding to lender risk, it can be an uphill climb to successfully expand evidence-backed social programs to a larger scale, as the program must now serve new target groups, operate in new settings, and yet strive to maintain its quality and its effects.⁸

Decisions about risk and reward in a SIB have exposed an unresolved tension between two competing visions of the model: One sees SIBs as a way to invest in innovation and test it at a larger scale — a financing vehicle appropriate for investors able to take risks similar to those taken by venture philanthropists. The other envisions SIBs as a way to replicate relatively low-risk, proven interventions capable of providing the "safe" returns sought by community development lenders — the Community Reinvestment Act units in major banks, the program-related-investment arms of foundations, and nonprofit community development lenders. The challenge is that many SIB deals have the risk profile of the first vision but investors primarily drawn from the ranks of institutions reflecting the second vision. This is not surprising. Community development finance institutions have the infrastructure and the know-how needed to make social investing deals, making them the "go-to" organizations for SIBs. But

⁷By the time the ABLE deal was ended early, for example, Goldman Sachs had invested \$7.2 million, \$6 million of which was guaranteed by Bloomberg Philanthropies, leaving Goldman Sachs with a \$1.2 million loss.

⁸See, for example, Garet et al. (2008); Miller et al. (2005).

their “lender” perspective — indeed, the type of funding they have to invest — predisposes them to favor lower-risk projects.

Conceptually, then, a SIB might be used to fund one of three categories of programs: (1) an innovative new approach not previously tried, (2) a program with mixed-to-good evidence about its effectiveness on a limited scale but no reliable information about its effects at a large scale, or (3) a program with strong evidence of effectiveness at some considerable scale.⁹ SIBs were originally conceived as a source of capital for expanding and replicating innovative new approaches to complex social problems.¹⁰ The projects would focus on demonstrating, testing, and building evidence about effectiveness as programs were tested in new settings, with new populations, and at a larger scale — that is, the projects would mostly fall into the second and third categories.

And the early SIB deals have indeed clustered in the second and third categories, many with a focus on expanding innovative approaches to a larger scale. Not surprisingly, the mix of appropriate funders for any given deal changes to fit that project’s risk profile. A private lender may be more willing to participate in a higher-risk deal when its risks are at least partially mitigated by a philanthropy’s guarantee, like the one Bloomberg Philanthropies provided in the ABLE deal.

In short, the risks inherent in a SIB may be more likely to resemble equity risks, in which investors take an ownership interest in a business that is expanding or in a new start-up. These are ventures where one might expect a large number of failures. Equity investors take more risk but have the potential to reap larger rewards. Not so in a SIB, where a government entity may only be willing to pay bond-like returns, altering the risk/reward ratio in ways that can limit a SIB’s viability.

Choosing a Benefit/Cost Perspective

By design, nearly all of the early SIBs were premised on government-budget savings. Indeed, in those deals, payments to investors depended on those savings. But when deals are premised exclusively on savings for government entities, SIBs set a very high bar; few social initiatives have consistently saved governments more than they cost — and those that have done so have often failed to benefit participants substantially.¹¹ However, the government-budget

⁹See Butler, Bloom, and Rudd (2013) for a discussion of the risk continuum and appropriate roles for different kinds of investors.

¹⁰Liebman (2011).

¹¹SIBs typically support programs in which participants volunteer, but some of the most consistent savings from social initiatives have come from *mandatory* programs that were designed to save the government money. For example, welfare-to-work programs, which require welfare recipients to participate as a condition of receiving benefits, reduce welfare payments and yield savings; however, while participants’ earnings from work rise, their overall incomes generally do not. See Greenberg, Deitch, and Hamilton (2010).

perspective is only one of several possible benefit/cost perspectives. In fact, when economists report positive benefit/cost ratios, they typically examine *more* than just the government budget perspective. They also take into account, for example, the participant perspective and the larger societal perspective.

Consider public education. Individuals enjoy higher earnings and incomes as a result of investments the government makes in their education. And the government benefits from higher tax revenue on those earnings. Beyond government budgets, taxpayers and the population at large benefit from rising productivity and economic growth, and perhaps even from non-economic outcomes that come from living in a safer community with fewer social problems. These societal benefits are not reflected in the government's budget, but they may make a particular social intervention well worth the investment.

Early on, however, most SIBs considered benefits only from the government-budget perspective. For example, deals involving preschool programs have been promoted with the promise that they would reduce the number of children referred for special needs services, a relatively rare but costly event. Similarly, SIB-financed home visiting programs have been premised largely on a reduction in high-cost preterm births — also a fairly rare occurrence. In the ABLE program, government savings were estimated purely as operational savings in the jail system. Yet a successful program might also produce savings in other parts of the safety net, as well as social benefits for participants, families, and communities that are not accounted for in the typical definition of a SIB's success. Thus, a focus on government budgets alone carries a risk: Programs that do not yield savings to government budgets but that benefit participants or society at large could be deemed “failures” from an investment point of view — even if they produce other positive benefits.

Agreeing on a Measure of Success

Having identified a problem and a potential solution, settled on SIB financing, and selected a benefit/cost perspective, the next step is to reach agreement on a measure of success. Multiple parties are more likely to agree on metrics that are simple to understand, easy to measure, and readily available from existing administrative data sources. The metrics must also capture the project's principal goal, be that a reduction in the number of days incarcerated, the number of placements in special education, or the incidence of preterm births. In each case, a dollar value can be attached to the outcome, and the measure represents clear benefits to a government entity, society, or both, which is essential for calculating benefits and costs. Finally, the simpler the outcome measure, the lower the project's evaluation costs and thus the total cost of the SIB.

Simplicity does have a downside, though. As noted above, using a single outcome to define success may miss a range of other benefits that might result from the program — benefits

that also have real value but will go unmeasured. Theoretically, the cognitive behavioral therapy in the ABLE project could have beneficial effects on life *in* the jail, possibly reducing confrontation and the use of punitive segregation, even though it did not lower recidivism. A nurse home visiting program for newborns and toddlers could improve parenting and children's eventual school performance but not reduce preterm births. A preschool program could increase high school graduation but not reduce special education placements. Recognizing these limitations, home visiting projects currently under development in South Carolina and New York are planning to measure a broad range of outcomes, although they are still tying payment to only a few particular outcomes.

Conducting Due Diligence

Constructing and executing a SIB is a complex undertaking requiring different expertise than is routinely maintained by government entities, investors, or nonprofit service deliverers. In a typical public/private social investment, due diligence is focused on gauging the ability of the loan recipient to carry out the deal and the likelihood that future revenues will be sufficient to pay back lenders/investors. Are neighborhood amenities, crime rates, transportation, schools, and income levels likely to attract tenants at rents that make a housing deal attractive? Will the population density and income profile of the neighborhood support the retail stores planned for a strip mall? In short, will a property's rent receipts ensure the payback of principal and interest to a lender? Investors, developers, and government partners understand this calculus and possess the skills to structure finance deals using it.

Due diligence for a social impact bond is different. Here, expertise is required in three areas: knowledge about social service program design and context, skill at assessing evaluation evidence, and acumen in public cost accounting, which is needed to estimate government-budget savings. The Rikers experience is a good case study.

Understanding Program Design and Context

Due diligence for program design requires detailed information on the characteristics and behaviors of the population to be served, an on-the-ground operational understanding of the program to be tested, and intimate knowledge of the inner workings of the system within which it will operate. For the ABLE deal, MDRC spent considerable time up front to learn about Rikers' operation; it had to understand the flow of individuals from intake to housing assignment to release, the intricate workings of existing programs, and the island's procedures for custody management. It was critical to gauge how a cognitive behavioral intervention could fit into the daily rhythms and routines of Rikers — how it would be experienced by inmates and by corrections officers and other staff members.

Assessing Potential Effectiveness

The Rikers SIB had high stakes for its investors, the city government, and the young people themselves. It was important, therefore, that the evidence demonstrate convincingly that cognitive behavioral therapy could have effects on recidivism large enough to yield meaningful savings; those savings needed to be sufficient for investors to risk their capital and for the city government to participate. As summarized and synthesized in a series of carefully conducted meta-analyses (involving some two dozen rigorous studies), the available evidence indicated that cognitive behavioral therapy programs had the potential to reduce recidivism by roughly 25 percent.¹²

But which program model was best suited to the Rikers setting? Most cognitive behavioral therapy models are designed to work with the same group of people from start to finish, but adolescents at Rikers arrive and leave unpredictably. The intervention therefore required a model that could work in an open-entry/open-exit environment. Moral Reconciliation Therapy was selected as the best fit for Rikers because of its flexibility. MRT has open groups, is self-paced, and uses workbooks to document progress. Previous evaluations of MRT found effects even when young people left before completing all 12 stages of the program. But the model's potential was unproven in the Rikers' setting, where it had to serve an adolescent population at an unprecedented scale.

Estimating Costs and Savings

Armed with an understanding of the operational challenges and the evidence, the next step was to estimate likely savings and thus payment terms. First up: selecting an outcome measure (and eventually a research design to measure the program's net impact) that was a good indicator of costs — and, thus, the potential savings that could be used to pay back investors. Incarceration is expensive, and an alarming one in two adolescents housed at Rikers returned in a year. If an intervention could reduce recidivism the savings would be large, and days spent in jail were easily measured using Department of Correction administrative data.

After reaching agreement on an acceptable measure of success, the next challenge was to quantify likely savings — how big would the net reduction in recidivism have to be to make this a worthwhile investment of private and eventually public dollars? Answering this question required mastering the city budget formulas that capture operating costs at Rikers. New York City's relatively sophisticated budget office had this information, but many other cities do not have the ability to do this kind of analysis. If a government entity does not have reliable data on its costs it is impossible to reliably estimate savings, and it therefore becomes impossible to estimate the scale of the program required to make a SIB deal worthwhile.

¹²Lipsey, Landenberger, and Wilson (2007).

In summary, due diligence for a social impact bond bears little resemblance to the due-diligence process with which most social investors are familiar. SIBs are often sold as a means of bringing market discipline to social programs — “market discipline” being defined as due diligence, accountability, and managerial expertise. Yet private investors are unlikely to possess the specialized content expertise needed to master social program design and context; the understanding of evaluation research needed to assess the potential effectiveness of an intervention; or the familiarity with government-budget formulas needed to estimate the costs and potential savings involved. It is important to note that governments could and should develop expertise in these areas — not just to negotiate SIB deals but as the first step in tying procurement and programs more effectively to evidence of effectiveness. After all, government agencies are providing and procuring educational, social, and health services every day without strong evidence of effectiveness. SIB deals can provide the impetus for them to build this essential expertise that they could eventually apply more broadly.

Structuring a Deal

Once one has all the components in place — the intervention, the measures of success, the estimated governmental and societal costs and potential savings — the next challenge is putting those components into a deal whose timing makes sense, that locks in expectations for all its partners, and yet that still allows programs some flexibility to deal with inevitable operational challenges.

Payback Periods and Government Reimbursement

A number of complex structural issues have to be resolved before investors and a government entity are likely to sign onto a multiyear SIB arrangement. Chief sources of tension are the inherent conflicts regarding (1) the time needed to deliver the program’s services to the intended population; (2) the time investors must wait for the evaluator to measure effects, so they can begin to recover their principal and receive a return; and (3) the time required before the government entity sees the promised savings from which it has agreed to pay back investors. As noted above, the ABLE program aimed to provide cognitive behavioral therapy to the 3,000 adolescents incarcerated at Rikers Island each year, with the intent of getting as many as possible through the 12-stage MRT process. To know whether the program was having the desired effect on recidivism, it was necessary to collect an additional year or more of follow-up data after the intervention, to learn whether the adolescents returned to jail at a lower rate than would otherwise have been the case. And it would take several more years beyond that for the city government to realize its savings.

Why the lag in government savings? The benefits of a social program may accrue slowly at first and then ramp up exponentially in later years. In ABLE, if the cognitive behav-

ioral therapy program had reduced recidivism, the effect in Year 1 would have been fewer occupied jail beds each day. At the time of the deal, the marginal cost savings of each empty bed equaled \$4,600 per year. The benefit of a given individual's not returning to jail would accumulate over several years as the person remained out of jail not just in Year 1, but also in Year 2 and possibly Year 3 and beyond. The marginal cost savings per bed, per year would leap to \$28,000 if the program reduced recidivism by 100 beds per year, enough to close an entire housing unit of the jail.

Still, the jail would have to close several housing units and accumulate multiple years of savings before the city government would save more than the intervention cost. This financial reality is not uncommon. In many of the early SIB agreements, a government entity has to continue to run the program beyond the agreement's evaluation period in order to accumulate enough savings to pay back investors. Because investors may not be willing to wait more than three or four years for a return on their investment, early deals have frequently assumed that the benefits documented by the evaluation for an early sample of participants will continue to accrue in later years, as the government entity continues to operate the program.

Another risk related to timing: Will a future government honor the payment commitments? By law, governments generally cannot create budget obligations for future legislatures. But SIBs are multiyear agreements that commit a government to pay for success several years *after* services are delivered. Uncertainty about a government's authority to enter into a SIB agreement poses "government counterparty" risk for potential investors. To address this issue, Massachusetts passed legislation setting up a trust fund that authorizes the executive branch to enter into SIB-like arrangements and that appropriates money specifically for SIB projects.¹³ While New York did not have comparable legislation in place, concerns about counterparty risk did not deter investors in the ABLE deal.

Three Numbers and the Tyranny of Inflexible Loan Agreements

In a social impact bond, government-budget savings and investor reimbursement depend on three numbers: the net impact or difference the program must make (for the Rikers SIB, a 10 percent reduction in recidivism relative to a comparison group was required for full repayment), the total number of people to be served, and the period of time over which those services will be delivered. Relative to social services contracts, loan agreements and term sheets are comparatively ironclad commitments. In a low-income-housing development deal, the agreement commits the developer to build a number of units, according to predesigned specifications, at a named cost, in a fixed time period. SIBs follow a similar format in that the agreement codifies the program metrics spelled out in the accompanying Pay for Success contract.

¹³See Goldberg (2012) for a comprehensive treatment of this issue (known as the "sovereign immunity" issue) and possible responses.

But serving disadvantaged families or individuals, particularly within public systems, demands operational flexibility. At Rikers, the plan was to deliver cognitive behavioral therapy through the jail's school, but the pilot period demonstrated that the service provider would not be able to rely on Department of Education teachers to deliver the curriculum. As a result, the provider had to pivot and hire therapy facilitators at the same time it was moving to implement the program at full scale. In addition, lockdowns disrupted the school day and the frequency of program delivery even more than had been anticipated. And more young people were placed in administrative and punitive segregation than expected, making it necessary to find ways to serve all these young people who were not attending the jail's school. Changes had to be made in staffing models and other service-delivery and design parameters.¹⁴

While Rikers Island is an unusually arduous setting in which to deliver a program, recruitment and engagement are challenges in every demonstration project and every effort to expand programs to a larger scale. For example, both the Center for Employment Opportunities (CEO) Pay for Success (PFS) project in New York State and the Massachusetts Roca PFS project encountered early challenges in meeting recruitment goals.¹⁵ These challenges are the norm, not the exception. And troubleshooting, adaptation, and change are commonplace as systems struggle to accommodate innovation.

Unfortunately, the Pay for Success terms embedded in loan agreements can be notoriously inflexible. That inflexibility is the other side of one of a SIB's notable benefits: explicit agreement among multiple parties concerning time periods, benchmarks, and standards of evidence.

When programs struggle to serve the promised number of participants in the agreed-upon time frame, two consequences follow. First, when even an effective program serves fewer individuals than required in the time allotted, government savings fall well short of the amount required to both reduce the government entity's expenditures and pay back investors. Second, fewer people served means a smaller-than-planned sample size for the evaluation design, with the likely result that the research will be less able to reliably detect the program's true effects.

The obvious solution is to extend recruitment for a longer period, say another year. But a simple change like this — common in traditional demonstration projects — wreaks havoc in a SIB. That extra year of operation requires funding. Who will provide that funding? Investors would have to put more money at risk, and the government entity would have to agree to pay

¹⁴Rudd, Nicoletti, Misner, and Bonsu (2013).

¹⁵Moazed (2015). The Massachusetts and New York State pay-for-success projects both seek to reduce recidivism among formerly incarcerated individuals. See Third Sector Capital Partners (2015); Social Finance, Inc. (2014). Earlier research by MDRC found that the CEO transitional jobs model being used in New York reduces recidivism among program participants, especially for recently released and higher-risk subgroups of ex-offenders. See Redcross, Millenky, Rudd, and Levshin (2012).

back that added amount plus interest. Of course, the savings calculation for the government entity would also change, and the deal, so painstakingly negotiated, would need to be restructured.

As these fundamental structural obstacles come into focus, creative responses are emerging. In the Rikers project, the flow of 16- to 18-year-olds into the jail began to fall well below the contractually specified number to be served (possibly due, in part, to changes in the courts or policing). After extended negotiations, the parties reached an agreement to expand the target population to 19- and 20-year-olds and to extend the project's payback period, a concession agreed to by a mayoral administration not in office when the deal was originally negotiated.

Anticipating difficulties of this nature, the New York State CEO PFS deal deliberately spread recruitment over four years: enrollment of 1,000 sample members in the first year, then a one-year cushion period in which the program could make up for any shortfall, followed by another 1,000 recruits in Year 3, with a fourth year to make up for any shortfall. The two phases will be evaluated separately after a year of observation and measurement that starts at the end of each enrollment period, with investors repaid in stages based on the results from each phase. While this process creatively addresses the recruitment challenges it does stretch out the time until final results are known, and stretches out the time until final payment can be made to more than six years.

In short, it can be difficult to negotiate operational changes within the confines of a Pay for Success loan agreement. Every SIB will confront this tension between inflexible contract terms and the need to respond to unexpected operational challenges.

A Guarantee That the Government Pays Only for Success

In the early conceptual work on social impact bonds, building evidence was an important goal alongside innovation and program expansion.¹⁶ No party has a bigger stake in knowing whether a given SIB investment yielded positive returns than the government and the taxpayers it serves. Indeed, a SIB's most sacred pledge is the promise that the government will pay *only* for success. If SIBs rely on weak evidence — for example, the simple output and outcome measures that government agencies, nonprofit service providers, and private investors typically use to measure performance — government entities are likely to end up paying for programs that neither benefit the people they serve nor save money.

Too often even sophisticated actors equate program *outcomes* (for example, the percentage of participants placed in jobs for 90 days) with *net impacts* — the change in outcomes caused by the program above and beyond what would have happened otherwise. Consider a

¹⁶Liebman (2011).

program where 60 percent of participants retain their jobs 90 days after job placement. Is that number good or bad? One way to judge would be to ask whether it is higher than the percentage in the previous year. But what if the previous year was 2009, the nadir of the Great Recession? Or what if the comparison year was in the late 1990s, when the national unemployment rate was 4 percent? Forced to rely on outcomes alone, one would have no way of knowing whether any change in participant performance was caused by the program or by the changing economy.

The outcomes-versus-impacts distinction can be vexing for policymakers and program managers alike. If recidivism rates among adolescents have held steady for five years and then fall after the introduction of a cognitive behavioral therapy program, isn't that sufficient evidence of the program's success? Unfortunately, no. In the Rikers case, the parties realized that recidivism could decline for a range of "external" reasons unrelated to the program, including changes in policing or court practices. Interventions to ameliorate other social problems would face similar challenges.¹⁷

To judge fairly, one needs a measure of the program's net impact: the difference over time that the program makes in the outcomes of program enrollees relative to the outcomes of a like group of people not enrolled. The independent evaluator in a SIB has the task of identifying just such a reliable comparison group, preferably using an experimental research design or, when that is not possible, a rigorous quasi-experimental evaluation design, as the Vera Institute of Justice did in the ABLE deal.

Not surprisingly, the various parties to a SIB often balk at including evaluation expenses in the deal. Yet if a government entity ends up paying unwittingly for programs that do not work, it will not just perpetuate the lack of accountability SIBs were designed to counter; it will end up actually costing the government more than if the program had never been tried.

Indeed, this is the very rationale that early proponents of SIBs offered for relying on an independent evaluator to build reliable evidence of effectiveness.¹⁸ When programs are found to work, the government and taxpayers can pay with confidence. When they do not, the field can benefit from learning what hasn't made a difference, laying the groundwork for learning what will work.

Recent SIB deals have attempted to find a neutral ground in the outcomes-versus-impact tug-of-war by basing payment on outcomes while also supporting an experimental re-

¹⁷The risks of incorrectly attributing causality to observational studies are well known. For example, the 1985 Nurses Health Study suggested that hormone replacement therapy reduced heart attacks; these findings were overturned in 2002 when a rigorous random assignment impact study showed that the therapy actually caused an increase in the risk of heart disease. See Ioannidis (2005); Taubes (2007); Tai, Grey, and Bolland (2014).

¹⁸Liebman (2011).

search design to learn what difference the program actually makes. For example, under a Pay for Success contract in Santa Clara County, California, the county will base payment on the number of months of housing stability a supportive housing program helps homeless individuals achieve.¹⁹ At the same time, though, the project will include a randomized controlled trial to evaluate whether permanent supportive housing generates efficiencies and economic benefit for the county.²⁰ Although chronically homeless people in Santa Clara may turn out to be well served by the county’s focus on housing stability outcomes, in other settings this approach sets up the risk that investors will be paid on the basis of outcomes, even though a subsequent, reliable evaluation finds that the program did not make a net difference after all. In the end, neither government entities nor philanthropic organizations can fulfill their objectives when a SIB does not insist on strong evidence, as the ABLE project did.

From “Bespoke” Deals to Sustained Government Investment in What Works

As has been the case with other new financial instruments, social impact bonds will probably evolve as investors gain experience. But whether they will become an established investment tool or a forgotten fad will depend on two developments: (1) whether social impact bonds can eventually attain the characteristics of other capital markets, with established risk profiles and expected rates of return; and (2) whether SIB-financed programs that prove successful are sustained after their deals end. One way to promote the first of these developments would be to focus on deals at the “proven” end of the risk continuum described above. Unfortunately, as explained below, there are currently very few programs that can truly be considered “proven.” Given this reality, SIB’s long-run future probably requires philanthropies to play a central role, not unlike the role they played in the early days of financing for low-income housing. Philanthropies have the tools necessary to manage the risks associated with SIB deals that will inevitably be located toward the higher-risk end of the deal continuum, and they have a clear interest in getting governments to adopt and sustain effective programs.

Creating Capital Markets

While the SIB deals so far negotiated by foundations, banks, and other community development lenders and investors are encouraging, they nonetheless have a “bespoke” quality. Like any custom-tailored product, each SIB has required specialized due diligence to investigate the problem addressed, the population served, and the program model’s effectiveness; each has involved mastering a different government agency’s cost-accounting system; each has had to identify acceptable evidence-based net-outcome measures and payment-benchmark thresholds;

¹⁹County of Santa Clara, California, and Abode Services (2015).

²⁰County of Santa Clara, California (2015).

and each has entailed negotiating its own loan agreement. This process is time-consuming and expensive. For banks and other community development lenders to consider SIBs a viable option in the long run, deals would have to become standardized, like the housing lending deals they now do routinely.

The J.B. and M.K. Pritzker Family Foundation, the Goldman Sachs Urban Investment Group, and the United Way of Salt Lake are attempting to build precisely this kind of capital market in Utah with a social impact bond that funds preschool. The Utah SIB aims to enroll children identified as being at risk of eventually needing costly special education services into preschool programs that might prevent those placements; the idea is that the city and state will realize savings as fewer children require special education when they enter elementary school. Research (much of it nonexperimental) suggests that participation in early childhood education programs can reduce the need for special education services, with the better programs reducing the number of children assigned to special education by 12 percentage points or more.²¹

These circumstances — a high-cost problem (special education referrals) and encouraging but not definitive evidence of a solution (high-quality preschool) — exemplify one kind of problem SIBs could be used to address. As principals at Goldman Sachs and its partners in the early childhood deal explained:

As investors, we will only get a return of principal and interest on our loan if the preschool program is successful in preparing children to start kindergarten and reducing the number of children that need special and remedial education.... We believe this model holds promise because it is scalable, replicable and sustainable. It provides a new framework for thinking about how the public and private sectors can work together to address pressing social needs in a way that results in better outcomes for children, alleviates some of the financial burden on taxpayers and generates savings for governments.²²

The project started with a prototype serving roughly 600 children, 109 of whom were assessed to be at high risk of needing special education placement in kindergarten. The deal was structured in stages: an initial pilot group first, then if the pilot group was successful, four additional groups. Investors put up \$7 million to fund the preschool programs, with payback contingent on special education savings through grade 6.²³ Early results from the pilot year appeared encouraging: All but one of the 109 children judged to be at risk were promoted to kindergarten without placement in special education, and a success payment was made to investors on the assumption that *all* of the at-risk children would have been placed in special education otherwise. That assumption has been criticized, however, as requiring a leap of faith, especially for chil-

²¹Barnett (1995); Muschkin, Ladd, and Dodge (2015); McCoy et al. (2015).

²²Glen, Pritzker, and Bayle (2013).

²³Goldman Sachs (2015); Wood (2015).

dren for whom English was a second language.²⁴ While conceptually the use of a test to identify at-risk children had the potential to significantly strengthen the program’s design and targeting, no test is likely to predict special education placement perfectly. In fact, the record of preschool programs’ effects on special education placements varies widely, from positive to negative.²⁵ To reliably estimate the state’s savings as future children enroll, it would be necessary to evaluate the net difference the program is actually making.

But even when a SIB-funded program has demonstrated its effectiveness, is private capital the right funding structure for sustaining an effective program over the long run? Once an initiative has been expanded to a large scale and proven effective using a SIB financing structure, a strong case can be made that the government should fund such a “proven” program outright. Indeed, it would cost the government more *not* to fund a program that saves money, and regardless of whether it funds the program directly or through a SIB structure, the government is the ultimate payer. This argument would suggest that ongoing funding of programs that have already proven themselves to be effective is not a viable way to establish the kind of capital market mainstream investors seek.

A promising alternative would be to use a SIB financing structure to try a model proven in one place in another location or with a different population. Packaging similar deals in new locations would eliminate much of the bespoke quality of the early SIBs. In addition, negotiations with a local government to strike a SIB deal can reveal the strengths and weaknesses in local data and budgeting systems. As SIBs spread, efforts to strengthen local systems follow. Finally, taking SIB deals to new states and localities could also spread discipline concerning evidence-based assessments and programs, knowledge and skills that have broad applicability in the public sector.

Shoring Up a Weak Evidence Base

The premise underlying the push to establish capital markets for SIB deals is the notion that there are an adequate number of “proven” programs to expand and sustain. But that is not yet the case.

The evidence base on effective social programs operating on a large scale is inadequate. As a result, there simply are not many proven programs to ramp up. As former federal budget director Peter Orszag (Obama administration) and former director of domestic policy John Bridgeland (Bush administration) explain in an op-ed:

Right now, we know too little about the policy solutions that will improve outcomes, and we measure “success” by the numbers of people served or a handful

²⁴Popper (2015).

²⁵Barnett (1995); Muschkin, Ladd, and Dodge (2015); Lipsey et al. (2013); McCoy et al. (2015).

of moving stories.... But rigorous evaluations of programs are still exceedingly rare, and Congress has not provided the resources — or the political will — to invest in building the kind of evidence that can be used to improve programs, make better funding choices and get better results.²⁶

In sum, commercial and nonprofit community development finance lenders — vital sources of hoped-for funding for social programs — may be best suited for deals that would expand to a large scale the handful of programs with the very strongest evidence.²⁷ Yet there are few deals likely to be done at this strong-evidence end of the continuum. Where will the risk capital needed for the innovation and expansion of promising programs come from?

Establishing a Principal Role for Philanthropies

Philanthropies and donor-advised funds may well be the best candidates to finance innovation. Paradoxically, philanthropies have focused for the most part on attracting private, for-profit investors into SIBs, rather than building an ongoing role for themselves.

Yet the philanthropic sector has by far the most to gain. SIBs provide philanthropies with the opportunity to get something they desperately want: the commitment of governments. Foundations spend grant dollars to improve the human condition, often with the ultimate goal of demonstrating what works, then advocate for government action and support. While foundation money is often patient, it is almost never permanent. Obtaining government commitment is often a futile quest, however. SIBs offer an underappreciated advantage: a contractual commitment from a government entity to pay for and sustain successful initiatives, at least initially.

Moreover, providing risk capital for innovation, experimentation, and evidence building is a core competency of philanthropic organizations. Equally important, philanthropies have a broader range of tools with which to work than private investors, including both grants and program-related investments, allowing them to make deals that involve greater risk than private investors might be able to tolerate.²⁸ For example, grant dollars might be used to pay for a pilot period designed to identify start-up and operating problems; to fund an intermediary organization to oversee a deal, troubleshoot problems, and negotiate and implement any contingency plans that turn out to be necessary; or to backstop commercial lenders. Bloomberg Philanthro-

²⁶Bridgeland and Orszag (2013).

²⁷Godeke and Resner (2012).

²⁸As enacted by Congress in the Tax Reform Act of 1969 (see Section 4944[c] of the Internal Revenue Code of 1954, as amended), foundations may use program-related investments to invest or lend their assets for socially beneficial endeavors that promise a below-market rate of return. Program-related investments are often used in conjunction with grants to finance the building of low-income housing, to provide capital for microenterprises, and more.

pies provided funding in all three of these areas in the ABLE deal, and some philanthropic organization has played these roles in the majority of SIBs in the United States.²⁹

In short, SIBs offer foundations an outcome they covet but seldom attain — government commitment — and philanthropies regularly take the kinds of social program risks contemplated in SIBs. Perhaps SIBs are a better investment for philanthropically oriented donors, therefore, than they are for many commercial lenders.³⁰

Last, isn't the ultimate measure of a SIB's success the identification of an effective (and cost-saving) program that can be sustained by government funding? Yet the issue of sustainability is rarely addressed explicitly in SIB agreements. It should be.³¹ If a SIB actually saves the government money, *not* continuing to fund and operate the program would cost the government more, because savings would no longer accrue. Moreover, if the savings are substantial, philanthropies and government entities could multiply those savings by plowing a portion of them back into further improvement efforts — for example, by providing more intensive follow-up services for adolescents when they return to the community from Rikers.

Much of the enthusiasm for SIBs is premised on their promise to attract new forms of capital for testing and expanding innovative approaches to some of the nation's most difficult and costly social problems. But that promise won't be fulfilled if success starts and ends with government entities paying back investors for successfully executed SIBs. To make a lasting difference, governments will have to continue to fund proven programs long after successful SIB agreements end.

Conclusion

The number and complexity of social needs facing the United States are daunting. Government entities need to innovate and experiment, but often lack the wherewithal to do so. And the government sector as a whole faces deep skepticism about the effectiveness of existing educational and social programs. At the same time, promising educational and social programs are starved for capital both to build evidence about their effectiveness and to expand to a scale where they can reach the large numbers needing help. Private wealth that could be directed at these problems is at an all-time high — banks are seeking to invest in their communities, donor-advised funds house more than a billion dollars in wealth, and a new generation of philanthropists has pledged billions more to social purposes.

²⁹Gustafsson-Wright, Gardiner, and Putcha (2015).

³⁰Overholser and Whistler (2013); Rangan and Chase (2015); Overholser (2015); Palandjian and Shumway (2015).

³¹Schaeffer, Shumway, and Brumme (2015).

Social impact bonds and pay-for-success contracts are a vehicle for connecting problem, solution, and capital. They promise to raise new forms of capital to finance risk taking and evidence building, to increase government accountability by having governments pay only for success, to promote innovation and experimentation in solving complex social problems, and to enable the expansion of programs that work. If SIBs are to fulfill this promise, the government, investor, philanthropic, and service-delivery sectors will need to devise strategies that both manage the risks and exploit the opportunities inherent in social impact bonds' unique characteristics. This is no simple task. The structural tensions embodied in the social impact bond instrument are not easily resolved.

In the end, success will depend on committing to a workable balance between risk and reward; to obtaining reliable, rigorous evidence of effectiveness; to moving beyond metrics focused strictly on government savings; to sustaining effective programs beyond the lives of their SIBs; and to accelerating the trend toward the use of evidence in government procurement decisions.

References

- Barnett, W. Steven. 1995. "Long-Term Effects of Early Childhood Programs on Cognitive and School Outcomes." *The Future of Children* 5, 3: 25-50.
- Bridgeland, John, and Peter Orszag. 2013. "A Moneyball Approach to Government." *Politico* (October 18).
- Butler, David, Dan Bloom, and Timothy Rudd. 2013. "Using Social Impact Bonds to Spur Innovation, Knowledge-Building, and Accountability." *Community Development Investment Review* 9, 1: 53-58.
- County of Santa Clara, California. 2015. "Agreement/Amendment 77051: Consider Recommendations Relating to Implementation of the Chronic Homelessness Pay for Success Project – 'Project Welcome Home.'" Available online at: <http://sccgov.iqm2.com/Citizens/FileOpen.aspx?Type=30&ID=68998>.
- County of Santa Clara, California, and Abode Services. 2015. "Project Welcome Home Fact Sheet." Available online at: www.thirdsectorcap.org/wp-content/uploads/2015/08/150811_SCC-CH-PFS_Fact-Sheet.pdf.
- Garet, Michael S., Stephanie Cronen, Marian Eaton, Anja Kurki, Meredith Ludwig, Wehmah Jones, Kazuaki Uekawa, Audrey Falk, Howard Bloom, Fred Doolittle, Pei Zhu, Laura Szejnberg, and Marsha Silverberg. 2008. *The Impact of Two Professional Development Interventions on Early Reading Instruction and Achievement*. NCEE 2008-4030. Washington, DC: National Center for Education Evaluation and Regional Assistance, Institute of Education Sciences, U.S. Department of Education.
- Glen, Alicia, J.B. Pritzker, and Deborah S. Bayle. 2013. "A New Way to Finance Early Childhood Education." *The Huffington Post*. Website: www.huffingtonpost.com/alicia-glen/post_4978_b_3442237.html.
- Godeke, Steven, and Lyel Resner. 2012. *Building a Healthy and Sustainable Social Impact Bond Market: The Investor Landscape*. New York: Godeke Consulting.
- Goldberg, Steven H. 2012. "SIB Legislation in the US." *The Social Impact Bond Tribune* 1, 1: 8-16. Available online at: www.payforsuccess.org/resources/first-issue-social-impact-bond-tribune.
- Goldman Sachs. 2015. "Fact Sheet: Initial Results Released for First Social Impact Bond for Early Childhood Education Show Success." New York: Goldman Sachs.
- Greenberg, David H., Victoria Deitch, and Gayle Hamilton. 2010. "A Synthesis of Random Assignment Benefit-Cost Studies of Welfare-to-Work Programs." *Journal of Benefit-Cost Analysis* 1, 1: 1-30.

- Gustafsson-Wright, Emily, Sophie Gardiner, Vidya Putcha. 2015. *The Potential and Limitations of Impact Bonds: Lessons from the First Five Years of Experience Worldwide*. Washington DC: Brookings Institution.
- Ioannidis, John P.A. 2005. "Contradicted and Initially Stronger Effects in Highly Cited Clinical Research." *Journal of the American Medical Association* 294, 2: 218-228.
- Liebman, Jeffrey B. 2011. *Social Impact Bonds: A Promising New Financing Model to Accelerate Social Innovation and Improve Government Performance*. Washington, DC: Center for American Progress.
- Lipsey, Mark W., Kerry G. Hofer, Nianbo Dong, Dale C. Farran, and Carol Bilbrey. 2013. *Evaluation of the Tennessee Voluntary Prekindergarten Program: Kindergarten and First Grade Follow-Up Results from the Randomized Control Design*. Nashville, TN: Peabody Research Institute, Vanderbilt University.
- Lipsey, Mark W., Nana A. Landenberger, and Sandra J. Wilson. 2007. "Effects of Cognitive-Behavioral Programs for Criminal Offenders." *Campbell Systematic Reviews* 2007, 6: 1-27.
- McCoy, Dana Charles, Kathleen M. Ziol-Guest, Rui Yang, Hirokazu Yoshikawa, Holly Schindler, and Greg J. Duncan. 2015. "Long-Term Impacts of Preschool Programs on Graduation Rates, Special Education Placement, and Grade Retention." Paper presented at the 37th Annual Fall Research Conference of the Association for Public Policy Analysis and Management, Miami, FL, November 12-14.
- Miller, Cynthia, Johannes M. Bos, Kristin E. Porter, Fannie M. Tseng, and Yasuyo Abe. 2005. *The Challenge of Repeating Success in a Changing World: Final Report on the Center for Employment Training Replication Sites*. New York: MDRC.
- Moazed, Tyler. 2015. *DOL Workforce Innovation Fund Pay for Success Pilot Process Study*. Presentation at the Annual Conference of the National Association for Welfare Research and Statistics, Atlanta, GA, August 25.
- Muschkin, Clara G., Helen F. Ladd, and Kenneth A. Dodge. 2015. "Impact of North Carolina's Early Childhood Initiatives on Special Education Placements in Third Grade." Working Paper 121. Washington, DC: National Center for Analysis of Longitudinal Data in Education Research.
- Nonprofit Finance Fund. 2015. "Pay for Success Learning Hub." Website: www.payforsuccess.org.
- Overholser, George. 2007. "Envisioning a \$1 Billion Social Investment Fund." Speech delivered at America Forward's "Gathering of Leaders," February 12. Available online at: www.thirdsectorcap.org/wp-content/uploads/2015/01/Overholser_Billion-Dollar-Social-Investment-Fund.pdf.
- Overholser, George. 2015. "Response to 'The Payoff of Pay-for-Success.'" *Stanford Social Innovation Review*. Website: http://ssir.org/up_for_debate/the_payoff_of_pay_for_success/george_overholser.

- Overholser, George, and Caroline Whistler. 2013. "The Real Revolution of Pay for Success: Ending 40 Years of Stagnant Results for Communities." *Community Development Investment Review* 9, 1: 5-11.
- Palandjian, Tracy, and Jeff Shumway. 2015. "Response to 'The Payoff of Pay-for-Success.'" *Stanford Social Innovation Review*. Website: http://ssir.org/up_for_debate/the_payoff_of_pay_for_success/palandjian_shumway.
- Popper, Nathaniel. 2015. "Success Metrics Questioned in School Program Funded by Goldman." *New York Times* (November 4), p. B1.
- Rangan, V. Kasturi, and Lisa A. Chase. 2015. "The Payoff of Pay-for-Success." *Stanford Social Innovation Review*. Website: http://ssir.org/up_for_debate/article/the_payoff_of_pay_for_success.
- Redcross, Cindy, Megan Millenky, Timothy Rudd, and Valerie Levshin. 2012. *More Than a Job: Final Results from the Evaluation of the Center for Employment Opportunities (CEO) Transitional Jobs Program*. New York: MDRC.
- Rudd, Timothy, Elisa Nicoletti, Kristin Misner, and Janae Bonsu. 2013. *Financing Promising Evidence-Based Programs: Early Lessons from the New York City Social Impact Bond*. New York: MDRC.
- Schaeffer, Sam, Jeff Shumway, and Caitlin Reimers Brumme. 2015. "After Pay for Success: Doubling Down on What Works." *Stanford Social Innovation Review*. Website: http://ssir.org/articles/entry/after_pay_for_success_doubling_down_on_what_works.
- Social Finance, Inc. 2012. *A New Tool for Scaling Impact: How Social Impact Bonds Can Mobilize Private Capital to Advance Social Good*. Boston: Social Finance, Inc.
- Social Finance, Inc. 2014. *Investing in What Works: "Pay for Success" in New York State, Increasing Employment and Improving Public Safety*. Boston: Social Finance, Inc.
- Tai, Vicki, Andrew Grey, and Mark J. Bolland. 2014. "Results of Observational Studies: Analysis of Findings from the Nurses' Health Study." *PLoS ONE* 9, 10. Available online at: <http://journals.plos.org/plosone/article?id=10.1371/journal.pone.0110403>.
- Taubes, Gary. 2007. "Do We Really Know What Makes Us Healthy?" *New York Times Magazine* (September 16).
- Third Sector Capital Partners. 2013. *Case Study: Preparing for a Pay for Success Opportunity*. Boston: Third Sector Capital Partners.
- Third Sector Capital Partners. 2014. *Pay for Success and Social Innovation Finance*. Boston: Third Sector Capital Partners.
- Third Sector Capital Partners. 2015. "Massachusetts Juvenile Justice PFS Initiative." Website: www.thirdsectorcap.org/portfolio/massachusetts-juvenile-justice-pfs-initiative.

Vera Institute of Justice. 2015. *Impact Evaluation of the Adolescent Behavioral Learning Experience (ABLE) Program at Rikers Island: Summary of Findings*. New York: Vera Institute of Justice.

Wood, Benjamin. 2015. "Preschool Paying Off for Goldman Sachs and Utah Kids, According to United Way." *The Salt Lake Tribune* (October 7).

About MDRC

MDRC is a nonprofit, nonpartisan social and education policy research organization dedicated to learning what works to improve the well-being of low-income people. Through its research and the active communication of its findings, MDRC seeks to enhance the effectiveness of social and education policies and programs.

Founded in 1974 and located in New York City and Oakland, California, MDRC is best known for mounting rigorous, large-scale, real-world tests of new and existing policies and programs. Its projects are a mix of demonstrations (field tests of promising new program approaches) and evaluations of ongoing government and community initiatives. MDRC's staff bring an unusual combination of research and organizational experience to their work, providing expertise on the latest in qualitative and quantitative methods and on program design, development, implementation, and management. MDRC seeks to learn not just whether a program is effective but also how and why the program's effects occur. In addition, it tries to place each project's findings in the broader context of related research — in order to build knowledge about what works across the social and education policy fields. MDRC's findings, lessons, and best practices are proactively shared with a broad audience in the policy and practitioner community as well as with the general public and the media.

Over the years, MDRC has brought its unique approach to an ever-growing range of policy areas and target populations. Once known primarily for evaluations of state welfare-to-work programs, today MDRC is also studying public school reforms, employment programs for ex-offenders and people with disabilities, and programs to help low-income students succeed in college. MDRC's projects are organized into five areas:

- Promoting Family Well-Being and Children's Development
- Improving Public Education
- Raising Academic Achievement and Persistence in College
- Supporting Low-Wage Workers and Communities
- Overcoming Barriers to Employment

Working in almost every state, all of the nation's largest cities, and Canada and the United Kingdom, MDRC conducts its projects in partnership with national, state, and local governments, public school systems, community organizations, and numerous private philanthropies.