Social Impact Bonds: New Product or New Package?*

December 18, 2014

Mark Pauly
The University of Pennsylvania, Wharton School and NBER

Ashley Swanson
The University of Pennsylvania, Wharton School and NBER

Abstract

This note considers a relatively new form of financing for social services, the "social impact bond." Proponents of social impact bonds (SIBs) argue that they present a solution to several problems in funding social services, including performance measurement and the distribution of risk. Using a simple model, we demonstrate that, despite their apparent novelty, social impact bonds in concept need not produce any difference in outcome from standard financing arrangements that allow for private sector social responsibility. We argue that SIBs will lead to greater program success if investors' effort can positively influence outcomes, either directly (effort exerted in production) or indirectly (effort devoted to screening), but are unlikely to do so otherwise. We conclude that, as in the more general theoretical literature, the value of this particular application in terms of funding innovation will be strongly context-dependent.

1 Introduction

The economic analysis of public goods (Samuelson, 1954) properly focuses on the non-excludable and non-rival nature of those goods, characteristics which make their efficient supply in a competitive economy unlikely. In normative applications this has meant that much of the emphasis has been on the demand side—the values citizens place on these goods and the feasibility of public or private methods for organizing demand in some

* We are grateful to Guy David and Joshua Graff Zivin for helpful comments.

way that will offset the inefficiency of competitive equilibrium. Governments use compulsory taxes to avoid the free rider problem, and private nonprofit firms attempt to mobilize community demand in a way that generally reflects more than self-interest.[1]

In this paper, we consider a relatively new form of public good provision that has both public and private dimensions – the social impact bond (SIB). SIBs combine private capital and production with government guarantees – the government and investors write a performance-contingent contract with a nonprofit service provider such that the investors provide the capital for the project and the government pays for success. The first object of this paper is to analyze SIBs in the context of the large literature on methods of engaging private resources to provide public goods. This literature traditionally begins with a neutrality result, and we show that SIBs are no exception; they may make no difference in terms of the likelihood of projects receiving funding or achieving success. Our second objective therefore is to extend this analysis to allow for SIBs which enhance the engagement of private investors in affecting the production process. We show that this extension crucially alters the neutrality result.

In an important paper in the corporate social responsibility (CSR) literature, Graff Zivin and Small (2005) consider a model with for profit firms, nonprofit firms, and altruistic agents. They show that, if direct donation to an altruistic or public goods cause is an option, that donation will always be at least as good if not better (in terms of aggregate

[1] The efficiency of public vs. private methods for organizing demand depends on political economic considerations and agents' utility functions. For example, as discussed in Weisbrod's model of nonprofit firms, there may be large differences in preferences among citizens for the public good: when some citizens attach low values and so constrain political choice, those with higher than average values for the public good may use private methods as a way of altering arrangements to reflect their stronger preferences (Weisbrod, 1991).

supply of capital) than agents trying to work through activities of firms (for profit or nonprofit) producing private goods to affect the outcome. Similarly, we use a simple model to demonstrate that social impact bonds exhibit no advantage over traditional financing arrangements in attracting capital to social programs when they are formed as arm's-length transactions. That is, although the structure of social impact bonds is indeed novel, they contain several ingredients—performance-based contracts and private philanthropy—that are both familiar concepts in social programs and also arguably the source of any improvement in value creation. In the basic arm's-length model, SIBs are no more likely to encourage investment in social programs than a traditional financing model in which service providers contract with government agencies on a performance-contingent basis and rely on traditional capital markets and private philanthropy for up-front capital.

There is much less emphasis in the public goods literature on the supply or production processes for those goods. In Samuelson's classic article, it is simply assumed that labor and capital are to be combined according to some efficient production function to produce any good, public or private. That model is thus silent on whether public good production should occur in private for profit firms, nonprofit firms, or in the public sector. There is no implication that public goods need to be supplied or produced by governmental organizations or government employees, or that they cannot all be supplied by for profit firms, given proper organization of the demand side.

Only the theory that emphasizes attenuation of property rights in nonprofit firms gives a

major role to potential differences in supply side efficiency between sectors.[2] And yet, there is obvious and convincing empirical evidence that the production of public goods—even when demand has been organized and mobilized and is at the ready—need not be efficient. What has been elusive—in theory and in real life—is an effective way to construct a bridge between donor-demanders and efficient production. More precisely, more needs to be added if we are to make normative judgments about when and why the private sector should play a role in producing public or social goods.

In the CSR literature, Graff Zivin and Small's neutrality finding implies that there has to be something added to what transpires in a private firm to give it a role in improving efficiency in giving to social goods. In a series of papers, David Baron (2007; 2009a, 2009b) develops several alternative models with different specifications of that "something different." In some of his papers it is the motivation of entrepreneurs, in others that of firm managers, and in still others private pressure groups attempting to influence firms.

For example, Baron (2007) develops a model in which roles are cast not only for donors willing to devote resources to what they believe to be a good cause, but also for altruistically motivated entrepreneurs who form firms that (among other things) supply these public goods in a way that can be superior to simple donation of money to worthy causes. Later work by Baron (2009a; 2009b) extends similar insights to models in which donors must motivate managers to do their desired good work, or in which social pressure groups alter behaviors of selfish suppliers of private goods to add public

[2] A modest literature on attenuation of capital supply to nonprofits because of the unavailability of equity capital instruments adds some insights as well.

dimensions like redistribution or control of externalities.

Following Baron's work in spirit, we expand our analysis of social impact bonds to allow for "something else" – an ingredient via which the SIB structure improves efficiency. Specifically, we relax the assumption that SIB contracts are arm's-length transactions between private investors and service providers. We allow that SIBs may bring "market discipline" to social service provision in a literal sense by engaging private investors with financial and social preferences in the program implementation directly. In the model where investors' effort can impact program performance, the combined self-interest and social preferences of private investors increases social program success and increases the capital available to more risky programs, thereby driving an efficiency wedge between SIBs and traditional financing arrangements.

Ultimately, the interest surrounding innovative ways to provide crucial social services is encouraging. However, several questions remain about whether SIBs have greater promise for efficient investment and efficient performance than with more traditional funding models for public and social services, broadly defined. Social impact investing products are unlikely to facilitate welfare enhancements simply by bundling together features of existing contract structures; however, they may improve program and funding success if they marshal private investors' effort in addition to their financial philanthropy.

The paper proceeds as follows. We first describe social impact bonds and review SIB applications to date. We next present a simple economic setting with self-interested and altruistic investors, traditional capital markets, and social programs with stochastic success and demonstrate again the neutrality hypothesis: with only motivation to provide

public goods or offset social harms in the model, SIBs provide no advantage over alternative organizational structures. They do not add anything. Next, we briefly sketch what might make a difference, very much in the spirit of Baron's approach. In particular, we argue if investors in SIBs are not only socially motivated but also have special skills to contribute to the production or supply side of the public good or externality mitigation in question, then (and only then) might SIBs play a useful role. We conclude by speculating as to the likelihood that the model assumptions are true in the settings currently targeted by SIB proponents.

2 Social Impact Bonds – Background and Applications

Supporters of social funding innovation note that government funding of social programs is rife with weaknesses, including 1) insufficient attention to measurement of outcomes of interest, resulting in persistence of ineffective programs; 2) slow and/or limited adoption of innovative approaches; and 3) service providers' and governments' inability to bear the upfront investment costs and financial risk of performance-based contracts (Liebman, 2011). Social impact bonds have been put forward as a solution to these problems. In an SIB arrangement, the up-front capital is provided by private investors, arguably insulating service providers and government sponsors from financial risk. Further, each SIB specifies a rigorous evaluation procedure at the outset, and the government pays back investors if pre-specified performance targets are met according to a specific timeline; proponents argue that such a focus on performance will bring needed market discipline to the social sector and increase the overall value created by social programs.

The social impact bond has the following structure: private investors provide up-front working capital to a bond-issuing organization, which then gives the capital to one or more service providers to implement the program. Under the SIB contract, the government has agreed to underwrite the project, but will pay only in the event that pre-defined performance targets are met. If the program is deemed successful, it pays the investors, generally a below-market return (Liebman, 2011).

Several SIB projects are currently underway in the United States and elsewhere. In the United Kingdom, the London-based nonprofit Social Finance set up the first SIB, in which private investors (including the United Kingdom lottery and the Macarthur Foundation) have supplied $7.9 million to finance a program to reduce Peterborough Prison inmate recidivism rates. If the program, administered by the nonprofit St. Giles trust, successfully reduces recidivism rates in its treatment group by 7.5% relative to a control group (male inmates are 60% likely to reoffend within one year without intervention), the United Kingdom Justice Ministry will pay back the initial investment plus a small return, which increases if rates are further reduced beyond the target threshold (Liebman, 2011). A similar program is in early stages in New York City—Goldman Sachs will provide $9.6 million to implement a recidivism-reduction program targeted to Rikers Island adolescent inmates provided by MDRC, a nonprofit education and social policy research organization; in this case, 75% of the investment is guaranteed by Bloomberg Philanthropies and the government payer is New York City's Department of Corrections (Chen, 2012).

The SIB projects currently in progress generally have as their goal reduced direct government spending on social services; for example, the Obama administration's 2012

budget included $100 million for pilot programs in seven program areas, including workforce development, education, juvenile justice and care of children with disabilities (whitehouse.gov, 2013). In the United States and elsewhere, SIBs have gotten substantial attention from the health policy world. Paul Corrigan of Britain's National Health Service London Strategic Health Authority has argued that SIBs would be an effective way of financing programs to manage diabetes and asthma for NHS patients (Corrigan, 2011). In 2012, the New Jersey Assembly passed the NJ Social Impact Bond Act, which seeks to attract private funding for public programs aimed at reducing the cost of providing healthcare to low-income and uninsured residents while maintaining quality and access (NJ Fifth Legislative District, 2012). In California, Social Finance is in early stages of a program to reduce emergency room visits and hospitalizations among Fresno children with asthma (Jensen, 2013).

SIB advocates argue that the built-in evaluation procedure in SIBs implies that they can be used to pilot and evaluate innovative programs (Loder, et al., 2010). Further, they argue that, since government payment is performance-contingent, government agencies are insulated from the risk of program failure, which will help with political viability, and that private provision of working capital prevents service providers from bearing financial risk (Liebman, 2011).

It is important to note that performance-based contracting between government agencies and social service providers is not a novel concept. In a 2010 survey, Urban Institute researchers found that 89 percent of surveyed human service nonprofits had government contracts or grants that required them to report to funding agencies the results, outcomes, and impact of programs and services, and 17 percent had performance-based payments

(based on outputs, not inputs) specified in the contracts (Boris, et al., 2010). Such performance contingencies have the benefit of improving incentives by aligning principal and agent objectives. However, they can be associated with problems of agent (service provider) solvency as service providers' financial well-being can be exposed to risk based on client heterogeneity, economic fluctuations, etc., which may limit the applicability of performance payments. The SIB structure is intended to solve this problem by providing the upfront working capital to service providers from private investors; the investors take on the service providers' risk. On the other hand, as we argue below, there are other ways to marshal private philanthropy and insulate service providers from performance-based contract risk.

Moreover, SIB critics are skeptical that performance contingency improves political viability of innovative programs, noting that many governments involved in current SIB projects pre-fund the programs for budgetary reasons (McKay, 2013) and that SIBs thus far have been cautious and have used only generally available proven best-practices methods to meet performance targets (McKay, 2013; Pratt, 2014).[3] In this paper, we consider explicitly whether SIBs can foster financing for riskier projects that might not otherwise be undertaken; we leave the details of government budgeting, as well as the question of whether SIBs will shift funding toward programs with more easily-measured outcomes and away from programs with complex outputs, for future work.

3 Analytical Framework

[3] In the context of a proposed SIB for prisoner reentry in Maryland, the State's Office of Policy Analysis argued that, in fact, given the difficulty of linking the evaluation of a social program to a highly complex contract centered on an outcome payment, the government may actually increase its operational risks in undertaking a SIB (McKay, 2013).

Consider a model in which the government (citizens organized collectively) values a successful social program (e.g., reduction in recidivism rates) at b, but for political reasons cannot make funds available for the program prior to proven success.

Suppose that a service provider, nonprofit or for-profit, can write a performance-contingent contract with the government to implement the program and will achieve success (and be paid) with probability p. The service provider has no reserves of its own to fund program implementation, and thus must seek financing for working capital from outside sources. The up-front cost of the program, which takes one period to implement, is X. For now, we assume that p is fixed, exogenous, and commonly known to all agents, ignoring for the moment the fact that this model obviates the usual need for a performance-based contract (moral hazard in principal-agent settings).

Finally, suppose that there are a large number of private investors, some of whom are altruistic; both kinds of investors hold diversified portfolios and so are risk-neutral with respect to any single project. The government as well can be assumed to be risk neutral. The non-altruistic investors require the market return r per dollar invested and do not care which type of company they invest in, while the altruistic investors require the return $a < r$ for social projects, but will only invest at below-market rates or donate to nonprofits.[4] This implies that the altruistic investor values social program success at v such that

[4] *See* Pauly (1973) and Glaeser and Shleifer (2001) for a discussion of the role of nonprofits and for-profits in enterprise.

$$a = r - {}^{(p*v)}/_{X}.$$[5]

In this baseline model, the expected benefit of the program is $p * b - X$. Therefore, from a social perspective, the program should be implemented if and only if $p \geq X/b$. Assume this is so. We then can distinguish several cases.

If $p * b \geq (1 + r) * X$, then the program can be implemented by a competitive for-profit service provider issuing equity to altruistic or non-altruistic investors.

Assume instead that $(1 + a) * X = p * b < (1 + r) * X$, so that the altruistic investors would be just willing to provide complete funding for the project, while non-altruistic investors would not. In this case, the nonprofit organization can provide the service and the altruistic investors are willing to invest.

However, the nonprofit cannot issue conventional equity, so there may be a barrier to financing. Here, the social impact bond model provides a solution; the nonprofit and altruistic investors can contract jointly with the government, specifying the payment rule and financing (with at least the required sum X paid to the nonprofit up front). In an SIB, the investors would receive a nominal return per dollar invested of c, and would be willing to invest so long as $p * (1 + c) * X \geq (1 + a) * X$. For example, this is satisfied with $(1 + c) = b/X$, implying $p * (1 + c) * X = p * b = (1 + a) * X$. Thus, the

[5] Here, we note that a (and, in turn, v) are surely program-specific. This term captures the valuation of program success by any altruistic investor *over and above* the financial return to investing in the project. It may be based on pure altruism; it may also be a function of image motivation (*see* Ariely, et al. (2009)). For our purposes, this distinction does not impact results, and we remain agnostic about what drives the discrepancy between a and r (and between investors' preferences and social preferences – we simply allow them to differ with $0 \leq a \leq r$).

conditions for implementability are equivalent to those that would pertain if nonprofit providers were able to issue equity – the social impact bond provides an opportunity for efficient investment of private dollars in a desirable social program.[6]

What this analysis ignores, however, is that nonprofit organizations are not so constrained in their financing options as the model supposed. In fact, nonprofit organizations quite often offer bonds in traditional capital markets. While these bonds promise to return interest and principal, default is possible if the nonprofit firm does not earn enough to make those payments. In such cases, the bonds are rated as lower quality, and sometimes as "junk." In this stylized setting, the nonprofit can offer several products.

First, it can issue a bond with the promised return of b in case of program success and 0 (default) in case of failure. That is, this bond promises a nominal return per dollar of r', so that the expected return is $p * (1 + r') * X$. Given our above assumptions, this bond would be acceptable to altruistic investors as long as $p * (1 + r') \geq (1 + a)$. For example, $r' = \frac{b}{X} - 1$ would satisfy this condition and the program would be implemented (essentially mimicking the SIB but without involving investors in the government contract directly). Thus, efficient investment is available in the more traditional form of socially responsible investing, in which investors accept lower expected returns when investing in companies or funds with social goals.

If altruistic investors are for some reason not interested in bonds of this kind, the nonprofit could also sell bonds to non-altruistic investors and make up the shortfall in the required return with donations from the altruistic investors equivalent to the donation

[6] Here, our use of "efficient" indicates that private dollars are invested whenever altruistic investors value the project (in expectation) more than the opportunity cost of investing.

embodied in the social impact bond. That net donation in the case of the SIB is $(r - a) * X$. Call this amount D and imagine it is made as a lump sum donation to the nonprofit firm that will receive the government contract. The nonprofit then borrows X from non-altruistic donors, promising to pay them $(1 + r') * X + D$ if the program succeeds, with $r' = {}^{1 + a}/_{p}$, and D if the program fails. Since $p * (1 + r') * X = (1 + a) * X$, it follows that the expected return to the non-altruistic investors is $(1 + a) * X + (r - a) * X$, or $(X + r * X)$, their required return. The value of r' is the same as the value of c in the SIB. However, the "donation" in this case is paid to the non-altruistic investors whether or not the project succeeds. (Investors could agree to receive nothing if the project fails, but then the nonprofit would have to pay ${}^{D}/_{p}$ if success occurs; it would have to "insure" the donation.)

Another equivalent option would be a bond paying b in case of success and g in case of failure, the g being supplied by an altruistic investor acting as guarantor as in the Bloomberg arrangement described in Section 2. The altruistic investor is willing to guarantee the bond if $(1 - p) * g \leq (r - a) * X$ and the non-altruistic investors are willing to purchase the bond if $p * b + (1 - p) * g \geq (1 + r) * X$. The conditions on the model parameters for an implementable bond to exist are identical for implementability of a SIB.

We thus have established a neutrality result: with fixed, exogenous, and commonly-known *p,* the condition for implementability of the social program is simply that $p * b \geq (1 + a) * X$, the same condition under which altruistic investors would purchase equity in the service provider. Nonprofit organizations cannot issue equity, but the

program can be implemented either under a variety of conventional debt arrangements or under an SIB; in the former case, private philanthropy can make up for the shortfall in market returns in a number of implicitly arm's-length transactions, including a simple up-front donation. This is a close analog to Graff Zivin and Small's (2005) finding of corporate social responsibility neutrality when corporate giving to social programs is a perfect substitute for personal giving. If the SIB is even $\varepsilon > 0$ more costly to set up (than using any other combination of debt and philanthropy described above), due, for example, to complicated legal arrangements, it is strictly dominated by more standard financing. Hence, some modification to the model is needed to justify preference for an SIB over standard methods.

4 Investor Effort

What is missing from the above model is a characterization of the "market discipline" SIBs are purported to bring to social services provision when engaging private investment. Here, we formalize this notion by suggesting one mechanism by which private investors involved in an SIB may have a non-neutral impact on program efficiency.

Suppose that p is not fixed, but rather an increasing function of effort by an agent outside the service providing organization (e.g., through monitoring of the service provider's efficiency) $p(e) = \overline{p} + m*e$ with $m > 0$.[7] Suppose also that only some altruistic investors

[7] As noted by Conning and Morduch (2011), large-scale social investors may provide both cash grants and specialized technical assistance. In the case of microfinance, social investors "take seats on the microlenders' boards, help shape governance, and transfer banking knowledge. It is the combination of capital and active engagement that has done

are willing and able to provide effort of this sort; managers of nonprofit firms are not, and non-altruistic investors are not.

Consider the case where $X \leq \overline{p} * b < (1 + a) * X < (1 + r) * X$. Now, we find that the range of traditional debt-plus-philanthropy options described in the previous Section may fail to implement the program, which even with zero effort is efficient from a social perspective.

First, a bond with the promised payment of b (with default in the event of program failure) would not be accepted by altruistic investors because an arm's length debt arrangement does not allow for creditor involvement in issuer activities and the altruistic investor's expected return absent effort is by assumption lower than her required return. Obviously, since $r > a$, the service also cannot be provided by a for-profit with non-altruistic investors.

There are multiple ways of engaging altruistic investors' effort in the production process. For example, we can assume that an altruistic agent's effort can impact program success if she is an entrepreneur in a SIB, or if she participates in program implementation as a board member. We demonstrate in a simple example that, while debt-plus-donation arrangements and SIBs may both improve program success and increase the set of implementable programs, the former may underperform relative to the latter. Specifically, a bond with the promised return of $b + D$ in case of program success and D in case of

---

the most to encourage additional private commercial funding and effectively expand the reach of markets."

failure (with D supplied by the donor-board member) may not facilitate program implementation as often as the SIB.

Suppose that the altruistic agent, either as board member or SIB investor, has disutility of effort $f(e)$. Her expected payoff as donor-board member, if the program is financed, would be $\min\{(\overline{p} + m * e), 1\} * v - D - f(e)$. She will exert effort $\hat{e}$ such that the marginal return to effort $v * m$ equals the marginal cost $f'(\hat{e})$. If, for example, $f(e) = e^2$, then effort supplied will equal ${}^{v * m}/_{2}$. She will be willing to donate D to this arrangement if $(\overline{p} + m * \hat{e}) * v \geq D$ – in our specific example, this requires that $\overline{p} * v + {}^{(m * v)^2}/_{2} \geq D$.

The wedge between this debt-plus-donation contract and the pure debt to altruistic investors contract is thus:

$$ {}^{(m * v)^2}/_{2} $$

from the altruistic agents' perspective. The other requirement is that $b * (\overline{p} + m * \hat{e}) + D \geq (1 + r) * X$, so that the additional effort relaxes the constraint on obtaining the remaining non-donated capital from outside investors as well.[8] Put more succinctly, the above simply says that all productive effort should be engaged in production if possible.

[8] This assumes that the altruistic investor/donor values program *success* in addition to her own return on investment; other models could assume that the altruistic investor/donor is willing to accept a lower return on investment regardless of probability of success. For example, an image motivation model in which altruistic investors value being observed to donate to particular causes could be applicable in some settings. In such a case, the debt-plus-donation contract would not induce additional effort relative to the pure debt contract.

Now consider the SIB model. In this case, the altruistic investor receives payout $(b + v) * (\overline{p} + m * e) - f(e)$ after investing, and thus will choose effort $\tilde{e}$ such that $f'(\tilde{e}) = m * (b + v)$. For example, in the case above where:

$$f(e) = e^2, \qquad \tilde{e} = {}^{m * (b+v)}/_{2} > {}^{m * v}/_{2} = \hat{e}.$$

Hence, greater effort will be provided under the SIB than under the debt-plus-donation setup *or* the pure debt contract.[9]

This result can be seen in Figure 1, in which we plot the expected payoff to the altruistic donor/board member (donating the minimum lump sum D required to make the non-altruistic investor indifferent between investing and not investing) and to the SIB investor contributing all of the required program capital X. At lower effort levels, the SIB investor is worse off than the donor because the probability of program success is low and she is supplying more capital to the project up-front. However, the return to effort is substantially higher for the SIB investor because she values both her intrinsic payoff to the social program v *and* the potential financial payout b, so that the optimal effort supplied by the SIB investor is higher than that of the donor; her payoff at the optimum is higher as well even though supplying the additional effort is costly.

Further, the SIB expands the set of implementable programs relative to those that can be implemented by the debt-plus-donation setup by decreasing the threshold for $\overline{p}$ such that altruistic investors are willing to invest/donate as well as exert productive effort. Some comparative statics can be seen in Figures 2 and 3.

[9] This result is due to our assumption of convex effort cost – if the disutility for effort were linear, the comparison of models would be discontinuous depending on the parameter values.

In Figure 2, we show the relationship between, on the y-axis, the minimum $\overline{p}$ such that the social program is implementable, and, on the x-axis, the donor valuation of the social program v. As we see in the Figure, by incentivizing higher effort than the donor model, the SIB expands the set of potentially successful programs. At $v = 0$, the altruistic donor/board member is not willing to donate to any social program, so only programs with $\overline{p} * b \geq (1 + r) * X$ will be implemented. However, the SIB investor is willing to invest in some less-promising projects because her productive effort pushes them over the threshold into implementability. As v increases, the set of implementable programs expands – the presence of altruistic investors/donors increases the likelihood of investment in social programs – and expands in such a way that the SIB model always dominates the donor model. Indeed, if $v > 0.6$, the baseline probability of success under the SIB model can equal $\overline{p} = 0$ and the program can still be implemented – the effort expended by motivated SIB investors is sufficient to increase the probability of program success in equilibrium such that they, anticipating the productive effort they will expend on the project, are willing to invest. At the assumed parameter values, the altruistic *donor's* valuation v would need to be 4-5 times as large to implement a $\overline{p} = 0$ program.

In Figure 3, we show the relationship between the minimum $\overline{p}$ such that the social program is implementable, and, on the x-axis, the parameter m capturing the productivity of effort. We see that the SIB model and debt-plus-donation models are equivalent if $m = 0$. If effort is not productive, then we are back in the setting described in Section 3, where we obtained our neutrality result – the presence of altruistic donors/investors expands the set of implementable programs beyond the setting shown in Figure 2 with $v = 0$, but the less-promising $\overline{p} = 0.2$ program described in Figure 1 cannot be

implemented. As m increases, we expand the set of implementable programs further under both models, but moreso under the SIB model than under the debt-plus-donation model – the additional value captured by the SIB investor when the productivity of effort increases leads to greater effort supplied in equilibrium and pushes more projects over the implementability threshold.

In sum, if altruistic actors can contribute productively to program success, then both a debt-plus-donation model and an SIB model increase the average program's success and expand the set of programs that are implemented in equilibrium, but the dual incentive provided to SIB investors (monetary and altruistic) drives a wedge between the two models and the SIB model dominates from an efficiency standpoint. This result has a certain intuitive appeal—if an agent is interested in both investment potential and social value, there may be a multiplier effect such that she will exert more effort than if she were interested in the social value alone.

This suggests that the SIB's inherent improvement over alternative models, if any, lies in involving altruistic investors in the process of organizing the provision of the service without violating the non-distribution constraint imposed on nonprofits. A similar alternative mechanism via which SIBs may offer an improvement over the more standard debt-plus-donation model could involve effort or investment on the part of the altruistic agents *prior* to program implementation, in the identification of more promising projects. For example, if altruistic actors have access to screening technology that allows them, at a cost, to observe the underlying success probability of a given program with greater precision, then the ex post success probability among *implemented* programs would be

higher under the SIB model than under the debt-plus-donation model, as investors in the former model would have a greater incentive to invest in the screening technology than in the latter. This result hinges crucially on this technology not being available to altruistic investors in an arms-length bond transaction – SIB investors must be in some sense "insiders" to obtain non-neutrality.

Put another way, in the case of the typical nonprofit firm, it may be the case that the firm's management have neither the desire nor the ability to make effort to increase p beyond $\overline{p}$ and that donors' incentive to exert effort flows only from their private value of the social good v, whereas the SIB has investors with both the power to affect p and a greater incentive to do so. Indeed, it is worth noting here that the usual benefit of performance-contingency is as a solution to an agency problem when effort is non-verifiable. If the probability of program success is fixed and exogenous as in Section 3 (or the equilibrium outcome of managerial effort on the part of the nonprofit's staff), then the SIB's performance-contingency does not solve any agency problem. This Section provides a rationale for performance-contingency to be productive in the SIB framework while still insulating nonprofit service providers from downside risk.

Of course, effort is not the only consideration. We implicitly have assumed in this Section that the comparison of efficiency across financing options is holding production processes internal to the nonprofit fixed. However, the management of a nonprofit firm probably has different and more specific managerial skills for the problem in question than general investors (altruistic or not). We have assumed that altruistic investor effort can only raise p, but in reality, investors and managers can have conflicting viewpoints in

operational matters and such conflict may be detrimental to program success. As between SIB and the nonprofit firm model, there may be a tradeoff between effort to raise p and skill in raising p, with the ideal depending on which influence is stronger.

Finally, the above neutrality result may rely on unrealistic assumptions about nonprofits' access to capital markets. If there are administrative costs to defaulting on bonds sold to non-altruists (even if they are offered as risky bonds), but lower such costs with the SIB model, that feature may favor SIBs. Similarly, nonprofits may have more limited access to traditional capital markets than the model assumes. As noted by Fisman and Hubbard (2003), there is wide variation across nonprofits operating in different sectors in their abilities to access debt financing. For example, in a review of nonprofit tax filings covering 1987-1996, 81% and 87% of health and housing organizations, respectively, had accessed loans, but only 18% of crime/legal-related nonprofits and 53% of job-related nonprofits had done so.

5 Some Empirical Hypotheses

This analysis suggests that SIBs are most valuable in specific situations, and one might further conjecture that those situations are ones in which they will be used. There are two main implications: (1) SIBs will emerge when other sources of capital have relatively high cost; and (2) SIBs will emerge when investors in them have special skills that can be deployed to improve program effectiveness.

On the first point, the attraction of SIBs may partly be technical: they may offer an investment vehicle that better taps the broader capital market than what can be done either by governments or nonprofit firms (see, e.g., Fisman and Hubbard, 2003). As we

are aware of no mechanism via which SIBs would better attract altruistic investors' capital than traditional bond markets, we leave this point for future research.

On the second point, it is surely the case that investors in many early stage social impact bonds are more engaged than in the typical atomistic capital market. Not only do they express concern about the social problem being addressed and the current inability of government or private firms to deal with it, they may also offer specific ideas about methods and techniques to solve the design and management problems. We do not have sufficient data on SIBs to analyze this engagement empirically, but research on impact investing more broadly has found strong representation of fund managers with specialized expertise in the impact investment sector. For example, the Intellecap Impact Investing Network is an angel investing network that provides both capital and managerial expertise to social impact investees (Huppe and Silva, 2013). On the other hand, SIB advocates such as Judith Rodin at the Rockefeller Foundation argue that, as the industry matures, SIBs will become more attractive to investors with no specialized expertise, thus diluting the potential advantage of SIBs discussed here (Rodin and Brandenburg, 2014).

6 Conclusion

Viewed purely as a vehicle for financing projects paid under performance contracts, we have shown that social impact bonds need not be different from other more conventional methods of finance used by nonprofit firms, and so need not offer any novel efficiency advantage. The innovation in social impact bonds, if there is one, must come from an hypothesized greater ability to solve the agency problem. In a sense, this is tautological:

if the funders of social impact bonds are altruistic financiers with management skills in the project at hand—say, dealing with prison or diabetic populations—that arrangement will be the best. But the relationship between willingness to make altruistic efforts and skills in doing so is generally unknown, and surely variable. At the extreme, altruistic and financial goals may come into conflict during program implementation. Consider, for example, the microloan investment sector, which, other than being explicitly focused on for-profit enterprise is fairly similar in structure to the SIB form. There were an estimated approximately $40 billion in microfinance loan portfolios in 2008 (Rhyne, 2010), but the large Indian sector was in severe crisis in 2011 due to accusations of aggressive loan collection practices, over-lending to indebted customers, and exploitative interest rates (Conning and Morduch, 2011). This critique has been applied broadly to a variety of settings involving business intervention in social services; in the area of corporate social responsibility, consider Friedman's (2005) comment that "Whole Foods has no special competence in deciding how charity should be distributed." Given the likely variable effect of effort on outcomes, the desirability of SIBs relative to other ways of financing and managing the supply of publicly demanded services is an open empirical question. Expanding the set of investment arrangements in itself is doubtless desirable, but not all innovations necessarily work out well.

7 References

Baron, D. 2007. "Corporate Social Responsibility and Social Entrepreneurship," *Journal of Economics and Management Strategy* 16(3), 683-717.

Baron, D. 2009a. "A Positive Theory of Moral Management," *Journal of Economics and Management Strategy* 18(1), 7-43.

Baron, D. 2009b. "Managerial Contracting and Corporate Social Responsibility," *Journal of Public Economics* 92, 268-298.

Boris, E., et al. 2010. "Human Service Nonprofits and Government Collaboration." *Urban Institute,* October.

Chen, D. W. 2012. "Goldman to Invest in City Jail Program, Profiting if Recidivism Falls Sharply." *The New York Times* A14.

Conning, J. and J. Morduch. 2011. "Microfinance and Social Investment." *Annual Review of Financial Economics* 3(2) 407-434.

Corrigan, P. 2011. "Social Impact Bonds: A New Way to Invest in Better Healthcare." *Social Finance.* September.

Fisman, R. and R. G. Hubbard. 2003. "The Role of Nonprofit Endowments." *in* Edward L. Glaeser (ed.) *The Governance of Not-for-Profit Organizations.* Cambridge, MA: *National Bureau of Economic Research. S!*

Friedman, M. 2005. "Rethinking the Social Responsibilities of Business. A Reason Debate, Featuring Milton Friedman, Whole Foods' John Mackey and Cypress Semiconductor's TJ Rodgers." *Reason*. October.

Glaeser, E. L. and A. Shleifer. 2001. "Not-For-Profit Entrepreneurs." *The Journal of Public Economics*. 81(1): 99-115.

Graff Zivin, J. and A. Small. 2005. "A Modigliani-Miller Theory of Altruistic Corporate Social Responsibility," *Topics in Economic Analysis and Policy 5,* Article 10.

Jensen, R. 2013. "'Social Impact Bond' Buzz Heads West." *The Bond Buyer*. *Available at* http://www.bondbuyer.com/issues/122_8/fresno-california-pilot-project-social-impact-bonds-1047541-1.html.

Liebman, J. B. 2011. "Social Impact Bonds: A Promising New Financing Model to Accelerate Social Innovation and Improve Government Performance." *Center for American Progress*.

Loder, J. et al., 2011. "Financing Social Value: Implementing Social Impact Bonds." *London: The Young Foundation*. January. *Available at* http://youngfoundation.org/files/images/01_10_Social_Impact_Bonds_FINAL.pdf.

McKay, K. 2013. "Evaluating Social Impact Bonds as a New Reentry Financing Mechanism: A Case Study on Reentry Programming in Maryland." *Department of Legislative Services, Office of Policy Analysis.* Annapolis, Maryland. January.

New Jersey Fifth Legislative District. 2012. "Fuentes Social Impact Bill Advanced by Assembly Panel." *Available at* http://www.politickernj.com/61532/fuentes-social-impact-bill-advanced-assembly-panel.

Pauly, M. V., and M. Redisch. 1973. "The Not-For-Profit Hospital as a Physicians' Cooperative." *The American Economic Review* 63(1) 87-99.

Pratt, J. 2013. "Flaws in the Social Impact Bond/Pay for Success Craze." *Nonprofit Quarterly.* April.

Rhyne, E. 2010. "The State of Global Microfinance: How Public and Private Funds Can Effectively Promote Financial Inclusions for All." *Subcommittee on International Monetary Policy and Trade Committee on Financial Services U.S. House of Representatives*.

Rodin, J. and M. Brandenburg. 2014. *The Power of Impact Investing: Putting Markets to Work for Profit and Global Good.* Philadelphia, PA: Wharton Digital Press.

Samuelson, P. 1954. "The Pure Theory of Public Expenditure." *The Review of Economics and Statistics*, 36(4) 387-389.

The White House, 2012. "Paying for Success." *Office of Management and Budget*. http://www.whitehouse.gov/omb/factsheet/paying-for-success.

Weisbrod, B. 1991. *The Nonprofit Economy.* Cambridge, MA: Harvard University Press.

8 Figures

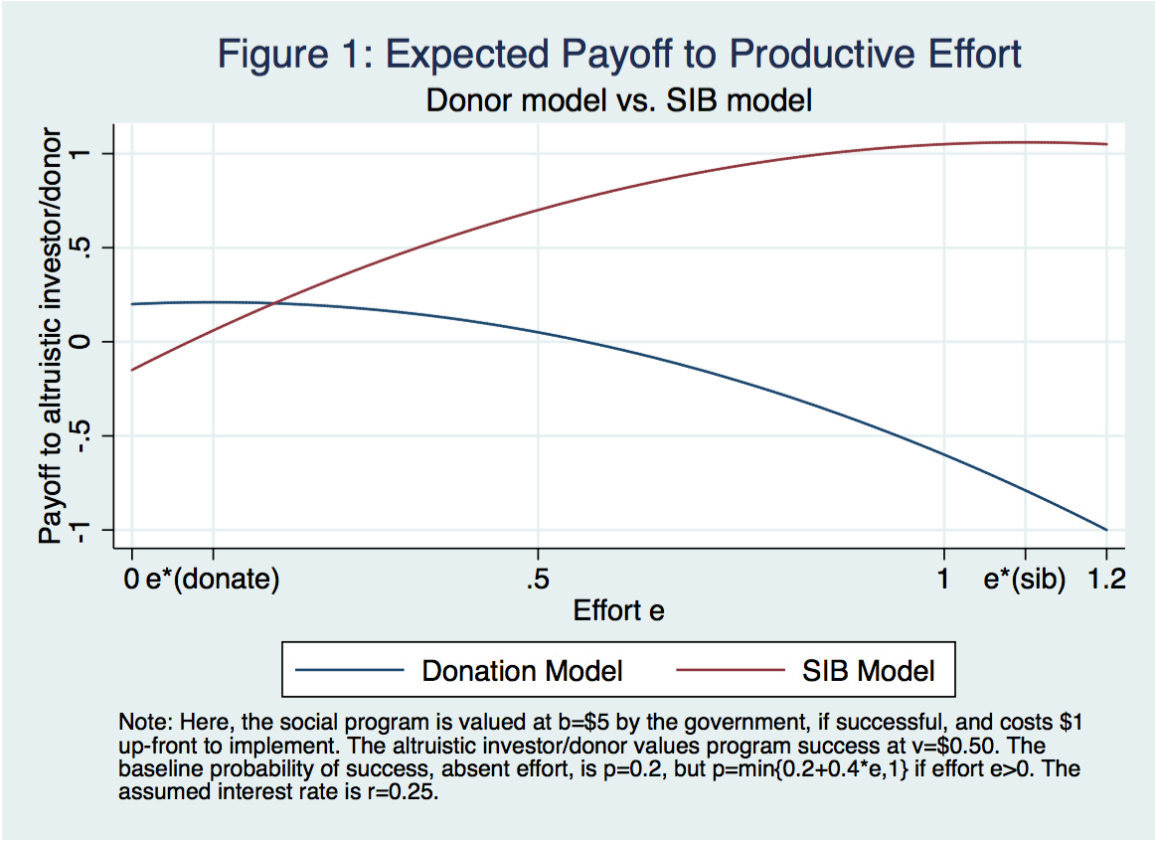

Note: Here, the social program is valued at b=$5 by the government, if successful, and costs $1 up-front to implement. The altruistic investor/donor values program success at v=$0.50. The baseline probability of success, absent effort, is p=0.2, but p=min{0.2+0.4*e,1} if effort e>0. The assumed interest rate is r=0.25.

Figure 2: Range of Implementable Programs

Donor model vs. SIB model

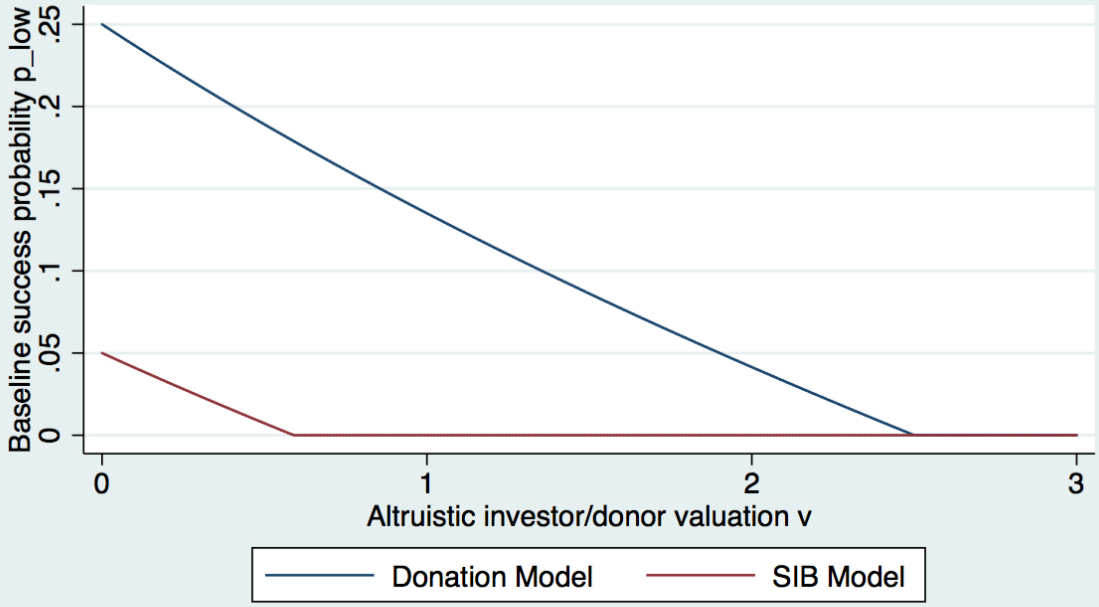

Note: Here, the social program is valued at b=$5 by the government, if successful, and costs $1 up-front to implement. The probability of success, given effort, is p=min{p_low+0.4*e,1} if effort e>0. The assumed interest rate is r=0.25.

Figure 3: Range of Implementable Programs

Donor model vs. SIB model

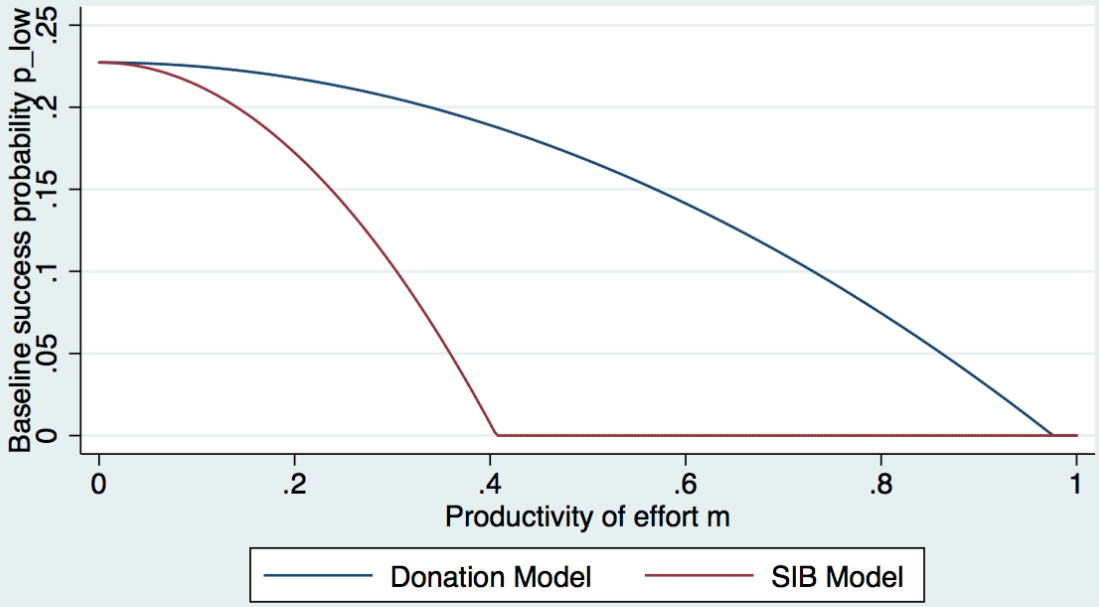

Note: Here, the social program is valued at b=$5 by the government, if successful, and costs $1 up-front to implement. The altruistic investor/donor values program success at v=$0.50. The probability of success, given effort, is p=min{p_low+m*e,1} if effort e>0. The assumed interest rate is r=0.25.