The Ethical Status of Social Impact Bonds.

Julia Morley
London School of Economics
j.e.morley@lse.ac.uk

April 10, 2017

Abstract

In this paper, I evaluate the ethical status of social impact bonds (SIBs), drawing on work by moral philosophers, in particular that of Debra Satz. In contrast to much of the academic and practitioner literature that focuses on the purported benefits of SIBs, I highlight potential conflicts between commercial and civil society values. Taking a consequentialist approach, I analyse the characteristics and outcomes of such marketised social delivery schemes to establish whether SIBs might represent a move towards unethical or 'noxious' markets. I argue that there is a real concern that SIBs may give rise to unethical markets, and hence suggest that they should be viewed with caution.

Keywords: Ethics, social impact bonds, civil society.

1. Introduction

Are social impact bonds (SIBs) a good way to deliver social care? In addressing this question, many commentators have argued that they offer an opportunity for positive reform of public sector care delivery and provide opportunities for innovations in the public and charitable sectors (Mulgan et al., 2011; Social Finance, 2016; Stoesz, 2014). Some scholars have raised practical concerns regarding the ability of SIBs to improve efficiency and the potential for the creation of the wrong kinds of incentives for social care providers, potentially leading to service providers targeting the easiest results (at the expense of addressing hard cases) of misleading reports of success or 'gaming of results' (Warner, 2013). Others have questioned the 'hype' concerning the the purported benefits of SIBs and the extent to which the specific SIB structures are necessary for achieving these results (Floyd and Gregory, 2015).

Beyond these practical issues about effectiveness, critical scholars have raised more fundamental and value-based concerns about the financialisation of the public sector (Dowling, 2016; Dowling and Harvie, 2014; Sinclair et al., 2014), the colonisation of public sector values by the corporate sector and the fundamental misalignment of incentives between civil society organisations and business interests (McHugh et al., 2013; Nicholls and Murdock, 2012).[1] These concerns about values and incentives also lead to serious ethical questions which are the focus of this paper.

In particular, I will argue that SIBs are, in some cases, a justifiable means for bringing about social change, provided that some critical ethical issues are addressed. These ethical issues largely, but not entirely, derive from the fact that SIBs attempt to generate markets in places where markets previously did not exist. In addition to the question of whether this "marketisation" of certain public (or welfare) services is permissible in ethical terms, there are a number of questions raised regarding the alignment of incentives between the private sector provides and those individuals who are the recipients of the services. Furthermore, there are worries about whether the marketisation of public services satisfies the basic conditions required of markets: informational transparency, capability of transactors making informed judgements, and ability to opt-out of exchange. In addressing these questions, I draw on the theoretical framework proposed

[1] For an excellent review of the literature on SIBs, see Fraser et al. (2016).

by Satz (2010), specifically her concept of a "noxious market". However, I do so with some modifications, for although Satz's framework provides a useful guide for identifying some of the causes of concern regarding SIBs, it does not apply perfectly to such structures for delivering social interventions. A theoretical contribution of this paper, then, is an extension of Satz's analytical framework, which introduces new categories of incentives that risk reducing social welfare through their effects on social exclusion or the misalignment of long-term incentives.

The paper is structured as follows. In section 2, I contextualise the emergence of SIBs, the potential benefits they deliver and the practical difficulties they face. In section 3, I briefly introduce ethical approaches that can be used to evaluate SIBs, focusing on consequentialism. In section 4, I provide an evaluation of SIBs using the framework developed by Satz (2010), and consider the extent to which SIB structures may be associated with poorly informed participants, significant inequalities in power relations and outcomes that harm particular groups within society or lead to potential harm to society overall. In section 5, I provide some concluding remarks.

2. Social impact bonds

Following the rise of new public management in the 1990s, new forms of marketised social interventions have emerged. Impact investing and social enterprises, for example, combine a social mission with the generation of financial returns. SIBs bring together socially motivated private finance and social organisations to provide public services, contracted using outcomes-based compensation. The first SIB, which focused on reducing recidivism in a prison population in Peterborough, was developed in the UK in March 2010. Since the launch of that SIB, 38 further SIBs were developed by March 2015 (Gustafsson-Wright et al., 2015).[2] The SIB structure is not limited to the UK and has been developed in the US as "pay-for-success" (PFS) and in Australia, as "social benefit bonds" (SBB). Additionally, "development impact bonds" (DIBs) and more recently, "humanitarian impact bonds" (HIBs) share common features with SIBs.[3]

[2] Of these SIBs, 4 addressed crime, 18 addressed social welfare (including rough sleeping) and 13 addressed employment.

[3] For more on the DIBs see Center for Global Development and Social Finance (2013) and `http://www.socialfinance.org.uk/impact/international-development/`. For more on HIBs, see `http://www.dalberg.com/conflict-and-humanitarian-aid/`.

Although the use of private organisations to deliver public service contracts is not in itself new, SIBs differ in two key ways from earlier models of public-private provision. First, the upfront funds for the social intervention are provided by the private investors rather than the public sector commissioner. Second, payments to the private investor are dependent on the social outcomes achieved. In the case of the first SIB, aimed at reducing recidivism at Peterborough prison in the UK, the external investors contracted with the Ministry of Justice to receive payments commensurate with the reduction in the re-offending rate for a specified prison population, within a particular period of time.

Several benefits of the SIB approach to social care provision are generally claimed by its advocates, particularly social investment practitioners and the government. Advocates argue that the payment-by-results structure re-aligns incentives for those managing interventions, resulting in the optimising of the allocation of public sector resources. A social investor will be motivated to direct resources to the social organisations whose interventions are most likely to yield the highest social outcomes, thereby resulting in an investment strategy that yields maximal allocative efficiency. The involvement of private sector organisations is also claimed to yield innovative solutions to social problems.

Additionally, the private investor provides the funding upfront, thereby relieving the public sector commissioner of the requirement to allocate funding to the project. This is a significant benefit for resource-poor public sector departments which may not be able to allocate resources for investment in future social benefit given the constraints on current year budgets. As a result of the transfer of the upfront investment to the private sector longer term investment can be effectively contracted out to the private sector. A report commissioned by a social finance intermediary, Social Finance, highlights the role that SIBs can play in enabling long-term investment by the UK National Health Service (NHS):

> . . . there are a range of highly important services that should be provided for more NHS patients that will only become mainstream if we develop a form of funding which moves beyond that provided by the Treasury on a year by year basis. Such services may well be uneconomic if based upon a single year's expenditure and return, but if they were allowed to be judged on a longer period than the one year,

> they would become a significant addition to the health of the nation. (Corrigan, 2011: 3)

Thus, a potential benefit of the SIB structure is its ability to shift the upfront cash requirements from the public sector to private investors. Furthermore, financial risk is transferred from the public sector commissioner to the private sector, as the public sector commissioner pays only for social outcomes achieved thereby reducing the in-year demands on the commissioning department's budget (Mulgan et al., 2011).

Another purported benefit is improved efficiency of operations and strategy within the social enterprises delivering the social intervention. This, it is claimed, results from the requirement to demonstrate social impact to investors using social impact reporting (New Philanthropy Capital, 2016). The use of this impact data enables improvements to resource allocation, focusing organisational resources on social-impact generating activities. Such advantages have been discussed in the academic and practitioner literature. For example, Stoesz (2014) has argued that social agencies operating without sufficient evidence, may not only do no good, but may actually do harm and states that:

> As evidence-based policy replaces unconditional entitlements, lawmakers will prefer program activities that are substantiated through field experiments, especially randomized controlled trials, which show lower cost while producing superior outcomes, over predecessors. Accordingly, social service employment opportunities will shift toward more effective programs and away from programs that have not demonstrated comparable efficacy. (Stoesz, 2014: 181)

Finally, the involvement of quasi-commercial investors in social care provision can lead to knowledge spillover from the private to the social sector with associated improvements to efficiency of service delivery. The possibility of knowledge transfers and technical spillover effects at the private-social intersection may thus yield innovation. In particular, the monitoring and advisory role of the private investor may be argued to benefit the social sector organisation and serve to improve strategic and operational practices. As Sara Llewellin, CEO of the Barrow Cadbury Trust and SIB investor is quoted as saying on the Social Finance website:

> "We are delighted with the progress made in the first cohort of the Peterborough Social Impact Bond. As investors, we wanted to prove that by doing something differently, and by being more flexible, we could indeed create a different outcome. An outcome which is a 'WIN, WIN'; a win for the taxpayer as the volume of repeat crime falls and a win for prisoners and their families when they take charge of restabilising their lives." [4]

If, as this statement suggests, the generation of positive results is attributable to the use of a SIB structure, the improvements it delivers to the social welfare of this group of prisoners can potentially be weighed up against the morally questionable profiting by social investors from this group of disadvantaged young men. Such justification of the means of delivery in terms of the ends achieved are associated with a consequentialist approach to ethics.

Consequentialist approaches measure the benefits to individuals and society in terms of utility. Contemporary utilitarians see individuals as 'utility maximisers', who act under constraints and the success of market outcomes is typically assessed in terms of welfare or utility maximisation. Consistent with their reliance on consequentialist approaches to evaluating economic outcomes, many economists argue that market mechanisms will generally improve social welfare through the allocative efficiency that results from the creation of appropriate incentives. The profit motive serves to allocate resources appropriately, meet consumer needs and compensate entrepreneurs for risk (Arnold, 1987). It is this kind of reasoning that provides certain justifications of the use of SIBs, which are intended to improve operational and allocative efficiency. Advocates of SIBs argue that operational efficiency gains result from both the influence of private investors and social investment intermediaries and from the requirement for organisations delivering a social service to measure and report their impact.

Many utilitarian arguments in support of the marketisation are based on the potential improvements that markets can deliver in terms of efficiency of allocation and of resource usage. According to this view markets generate economic growth, which is associated with a better welfare outcome and a higher level of utility for society. Yet some scholars have questioned the extent to which markets can be seen to be competitive and that the satisfaction of consumer

[4](http://www.socialfinance.org.uk/peterborough-social-impact-bond-reduces-reoffending-by-8-4-investors-on-course-for-payment-in-2016)

demand can be taken to be a proxy happiness (Galbraith, 1958; Hirsh).

An alternative view is that the notion of a competitive, entrepreneurial market is merely aspirational as modern corporate structures tend to reward owners of capital who engage in very limited entrepreneurial activity (Brown, 1992). Political economy approaches highlight the importance of market structures and power relations as determinants of wealth distributions, rather than effort or risk-taking by market participants [5] In addition to issues of market structure and power, informational asymmetries may exist, leading to sub-optimal resource allocation. Furthermore, negative externalities may exist which are hard to recognise, let alone measure, as they may not be evident for some time (as in the case of global climate change or the health effects of food products or tobacco). Furthermore, marketisation may lead to the persistence of social norms that inhibit social change, for instance as a consequence of the activities of companies that advocate sexual or racial stereotypes.

2.1. Practical concerns with the marketisation of social care

Taking a market approach to solving market failures raises a number of pragmatic concerns. One such concern is the cost and complexity of organising a social impact bond. In evidence given to a House of Lords Committee, the chief executive of Social Enterprise referred to the "hyperbole around social impact bonds" and claimed that "[t]here are many other forms of finance which can be more useful."[6].[7] A report by the OECD argues that the costs of structuring and managing SIBs have been high:

> They have entailed significant transaction costs that stakeholders should consider before embarking on them. Policy makers should evaluate carefully what is the value added for implementing a SIB for a policy intervention compared to a more traditional approach. (OECD, 2016b)

Given the extremely high level of transaction costs, it is generally argued that the SIB structure will not endure, once political support is removed. Furthermore, the fact that government

[5] For a review of this literature, seeAlesina and Perotti (1994).

[6] `https://www.civilsociety.co.uk/news/social-impact-bonds-incomprehensible-and-irrelevant-lords-committee-told.html#sthash.zI72y6vU.dpuf`

[7] Other forms of financing may include debt or quasi-equity.

departments, think tanks and other agencies expend significant resources in the structuring of SIBs suggests that the true costs of a SIB may be far higher than

Another issue is the difficulty in defining the success of a SIB. It is extremely difficult to establish and demonstrate an organisation's effectiveness in applying its 'theory of change', which specifies how its interventions are expected to result in positive social impact. How, for example, can causal relationships be shown to exist between the actions of the organisation and the claimed social impact observed? Social interventions may yield benefits over many years, which is a problem because it may be unattractive to many social investors who are constrained by short term investment horizons. For example, the social returns associated with the counselling of young children may not be visible for many years, and may require long term studies to establish results. This is problematic if it causes social organisations to focus primarily on activities that yield the highest results, even if this means losing focus on their core mission, ignoring the 'hard cases' or providing exaggerated reports of their success (Joy and Shields, 2013; Lowe and Wilson, 2015; McHugh et al., 2013; Warner, 2013).

Furthermore, it may be difficult to estabish which social intervention deserves credit for a particular outcome, because some outcomes may be causally overdetermined. An outcome is causally overdetermined when several different causal factors result in a particular outcome, each of which is sufficient for the effect and all occur simultaneously. For example, if three organisations are working with one young person to help her find a job, it may be impossible to disentangle the relative causal effectiveness of each organisation's efforts. If all of the organisations report the social impact as their own (and yet each organisation on its own would have been achieved the necessary result), the aggregate social impact reported will be overstated. Social impact reporting would be misleading in such cases. It is also associated with the unnecessary duplication of effort, as three organisations carry out the work that one organisation could achieve on its own, and will thereby lead to an unnecessarily high level of resource usage. It is also possible that outcomes are *underdetermined* if insufficient evidence supports any particular causal influence.

The purported benefits and pragmatic concerns of SIBs discussed so far may appear to be purely descriptive (i.e. 'What discount rate should we use?' and 'Whose interests should we prioritise?'), but such concerns implicitly require us to address ethical issues about SIBs, which are considered below.

3. The ethics of making markets in the social sector

Much has been written about the marketisation of public services (See for example, Humphrey et al., 1993; Nicholls, 2009), in particular the effects it can have on incentives and social welfare. Social impact bonds are somewhat different from public-private partnerships because the service delivery is provided by mission-driven organisations, whether charities or social enterprises. In what follows, I discuss certain ethical considerations that affect even these privately funded pay-by-results contracts.

The notion that there is something morally wrong with profiting from the provision of social care arises most naturally if one is committed to a deontological system of ethics[8] Deontological approaches reject those actions that are considered to treat people without due respect. One should not profit from the misery of others. This is an important question to ask when evaluating SIBs. Slavery is an obvious case of a practice that fails to show due respect to individuals who are deprived of their freedom and exploited, and this would be rejected on grounds of this lack of respect for personhood. Deontologists might also reject other cases that are less clear cut, in which an action that would be judged immoral nevertheless ultimately deliver benefits to society. The deontological approach would thus also reject the argument that the end justifies the means, for example the sale of a kidney by someone who is desperate for cash to someone who will otherwise die. Other examples of morally questionable ways of obtaining economic gain, highlighted by moral philosophers such as Sandel (2012), include the case of companies which have an insurable interest in their employees, known as 'janitor's insurance'. Sandel (2012) discusses the case of Walmart, which collected an insurance payment that was paid out following the death of an assistant manager at one of its stores on which it had taken out life insurance — with his family receiving no part of this sum (Sandel, 2012: pg. 131). A deontologist would ask, does this action fail to show adequate respect for the employee, in that it treats his death as a

[8] This is because such actions may conflict with Kant's second Categorical Imperative (known as the Formula of Humanity) which states: "Act in such a way that you treat humanity, whether in your own person or in the person of any other, never merely as a means to an end, but always at the same time as an end."

means to profit generation? In summary, deontologists reject certain kinds of actions as being morally impermissible, even if they are undertaken for good (or at least defensible) reasons and bring about an overall increase in well-being, because they do not treat persons with respect.

Applying a deontological approach to the case of SIBs the main concern is that social investors, particularly those investing in SIBs, profit directly from human suffering, even though they may not be primarily treating these individuals as a means to an end. Consider the case of the first UK social impact bond, intended to reduced the recidivism rate among a cohort of young prisoners in Peterborough. The level of profits earned by investors for this SIB was driven by the percentage reduction achieved in the re-offending of these young men. On the face of it, the ethics of such profit generation is importantly different from the case of janitor's insurance, because the Peterborough Bond was intended to *benefit* the young prisoners as well as providing financial returns to the investors.

Nevertheless, the very possibility of a company making a profit from the Peterborough SIB relies on the existence of prisoners whose misfortune is essential to the ability of social investors to make a financial return. As with the case of janitor's insurance, profits are generated as a result of the involuntary participation by the individuals whose outcomes crystallised particular commercial pay-outs, whether for the employer, in the case of janitor's tax, or for the social investor, in the case of SIBs. A Kantian might object to the lack of choice by some participatants, because it undermines their autonomy. An application of moral reasoning based on deontological ethics might then lead to a view that SIBs cannot be deemed to be moral activities because of the fact that they are predicated on the profiting by investors on the misery of others, in a way that may fail to reject the autonomy of a participant, whether it is taking a gamble on likelihood of a prisoner reoffending or on a drug offender staying clean for a certain period of time.

The ethical status of SIBs is quite different, however, if we consider alternatives to deontological ethics. The main alternative ethical approach, for our purposes, is consequentialism, which establishes the moral legitimacy of actions and practices purely by analysing their consequences.[9] Historically, the most common type of consequentialist theories, and hence that

[9] I set aside those approaches falling into the tradition known as 'virtue ethics' simply because its primary focus is a different question from the present concern. Whereas both consequentialist and deontological theories focus on the question of "What kind of actions should we do?" (that is, the ethics of conduct), virtue theories focus on the question of "What kind of people should we be?" (that is, the ethics of character). For a discussion of the

most frequently incorporated into traditional economic theory, is some variant of utilitarianism. Utilitarianism focuses solely on the maximisation of aggregate utility (originally understood by Bentham as the quantity of pleasure or happiness, but which, in contemporary economic theory, is now understood as the satisfaction of individual preferences).[10] Actions performed for reasons that would commonly be viewed as morally questionable may nevertheless be considered ethical if they increase utility. Thus, the view that profiting from human misery is morally unacceptable whatever the consequences might well be disputed by utilitarians who, adding up the benefits generated and the costs incurred by an action or practice, find that the overall balance is a net positive. Since the neo-liberal turn in contemporary politics in the 1980s, there has been increased interest in the scope of markets in public policy (Hood, 1991; Humphrey et al., 1993; Kurunmäki and Miller, 2011). This has resulted in an increased ethical focus on the marketisation of public services and the application of markets to out-of-market issues. Such approaches are justified by consequentialist reasoning, as proponents of the making of markets for social care do not distinguish between different *types* of product. They may therefore argue that markets in body parts, such as kidneys, are not fundamentally different from markets in any other good (Hippen, 2005). Consequentialists favour transactions of this nature if they increase social welfare, for example if an individual could benefit financially from selling one of her kidneys while someone with poor health would benefit from buying a kidney.

Some moral philosophers have criticised consequentialist theories, arguing that for certain activities, such as the privatisation of public services, the justification that 'the end justifies the means' should in fact be inverted because the end cannot be evaluated without reference to the means by which this was achieved. We might consider that the means are, to some degree, constitutive of the ends and thus different ways of carrying out an action can determine the nature of the outcome. In particular, they highlight the importance of the identity of the person or agency carrying out a public service of this kind. Using this line of argument, Dorfman and

implications of marketisation for virtue ethics, see Graafland (2009). Although there are, of course, multiple points of intersection between 'what we should' do and 'the kind of person we should strive to be', for the purposes of this paper I will restrict my attention to questions about the ethics of conduct.

[10]Those who think that mere aggregation fails to capture every aspect of moral significance often favour some version of *prioritarianism*, the view that the goodness of an outcome depends upon the well-being of all individuals, but where extra weight is given to those who are worse off (Parfit, 1984). This departure from egalitarianism means that prioritarianism and utilitarianism, although both aggregative consequentialist theories, can disagree in their evaluation of actions. For our purposes, we shall only concentrate on utilitarian theories.

Harel (2013) use the example of the justice system and the military stating that:

> "Political practices such as punishment and war promote a variety of desirable ends (such as deterrence and security), but the realization of these ends often depends on the agent who performs the task ... The trouble with instrumentalist theories of punishment and war is that they fail to identify some of the important ends of these practices, in particular those ends that can only be realized by public agents." (Dorfman and Harel, 2013: 102)

We can see how the means might at least partially constitute the ends by contrasting the case in which the state engages in a military intervention using its armed forces and one in which it uses mercenaries. By analogy, in the social sector, we might consider it unethical for private sector agents to provide certain types of social interventions, for example those related to criminal justice. The very *idea* of justice requires it be performed by the state, rather than by a private company. If we take this argument seriously, it throws doubt on the justification for delegating social care of certain kinds to the private sector, and certainly to giving the private sector policy-making decisions.

A common criticism of utilitarianism is the difficulty of responding to its distributional effects. In her book, *Why some things should not be for sale: The moral limits of markets*, Satz (2010) provides a guide to the characteristics of a moral market. Satz identifies the necessary characteristics of an ethical market.

Satz (2010) describes the failure of moral markets, what she labels "noxious markets". She highlights, as an example, the development of a market for toxic waste advocated by the Harvard economist, Lawrence Summers. Summers describes this market as a means of improving welfare for less developed countries. Satz points out that increasing *aggregate* utility in such a way may nevertheless lead to a morally problematic *distribution* of welfare, in which a reduction in utility for some individuals may occur. Such variations may result from asymmetries in information by decision-makers who may not have full knowledge of the likely consequences of the transaction, or because of an imbalance in the power balance between market participants. In addition, incentives may play an important role in generating such unequal distributions of utility if decision-makers are not subject to the negative consequences of their decision. Furthermore, issues concerning individual agency may arise because the government in the less developed country may make decisions on behalf of particular individuals, yet not take into account the impact of their decision on the welfare of particular individuals. Satz (2010) notes that Summers does not consider the effect of such schemes on the utility of those individuals affected by the toxicity of the waste due to their geographical proximity. She states:

> A market exchange based in desperation, humiliation or begging or whose terms of remediation involve bondage or servitude is not an exchange between equals." (Satz, 2010: 93)

Satz's argument is that the power of individuals participating in a market is the key issue in determining whether a market is a "noxious market". She does not reject a consequentialist approach, but requires certain characteristics and outcomes to exist if a market is to be considered to be moral. If these are not satisfied, she claims that the market is "noxious". She sets out the criteria by which a market should be judged as either moral or noxious using two characteristics and two outcomes, which are summarised in Table 1 below.

Market characteristics:	**Weak agency** Inadequate information about consequences of market transactions	**Vulnerability** Powerlessness, poverty and destitution of market participants
Market outcomes:	**Extreme harm for the individual** Damages welfare and/or agency interests of the individual	**Extreme harm for society** Undermines democracy or civil society

Table 1: Characteristics and outcomes of a noxious market taken from Satz (2010): 98

As the table shows, the foundations upon which a market is built reveal whether or not it will likely be viewed as noxious. As a first requirement of a non-noxious market, participants must be well-informed, so that they can make decisions that maximise their utility. Whereas the standard theory of expected utility requires decision-makers to be hyper-rational, we can realistically relax this assumption in practice to say that agents must have well-informed expectations of the likely outcomes of their actions. This is consistent with work by moral philosophers and social welfare theorists who have argued that certain kinds of activities, such as healthcare, do not lend themselves to marketisation without the risk of moral problems arising. In this regard,Arrow (1963) argues that a private healthcare market is not efficient from a welfare perspective. Although individuals would like to insure themselves against ill health they cannot do so because they are not sufficiently well-informed about their future health and may consequently make imprudent decisions that are detrimental to their long-term utility. This lack of information need not result from the intentional actions of others (for example, the covering up of scientific research into the effects of smoking cigarettes by tobacco companies) but may result from the cognitive limits of market participants. Markets characterised by "very weak or highly asymmetric knowledge and agency" of participants are likely to be noxious. Asymmetries of this kind will often be more pronounced if there is a significant time lag between initiation and completion of the contract as knowledge of the utility associated with the completion of the contract may be incorrect at the time the contract is initially agreed. An example of such information asymmetry is the case of a woman selling her right to a child she will bear in the future, as has been with human surrogacy, as the woman may not be cognisant of her future feelings for the child at the time she entered into the contract.

Second, Satz argues that no individual should have significantly more power than any other. According to neo-classical economics, optimal resource allocations occur in markets in which all agents have the same market power. In perfectly competitive markets, all agents are price takers (and are not able to shift market prices) and the level of returns made by individuals reflects their investment and their willingness to take on risk. By contrast, in a market where individuals or companies form alliances to control price or output, we would expect to see supernormal profits being made by firms and a suboptimal level of economic efficiency. Another characteristic likely to lead to poor moral outcomes relates to the equality of power of the market participants. A vulnerable participant may enter into a transaction which is exploitative, and may thus be considered morally wrong, even if it improves his or her welfare at that point in time. If a destitute person who desperately needs money to feed her children sells all her assets at below market price due to her desperate circumstances, she is acting in a way that maximises her welfare, but her lack of power as a market participant results in what may be considered to be a noxious market. The problem stems from the fact that she has no real choice about whether or not to enter into the transaction and her powerlessness renders the transaction - and hence the market that facilitates it - to be noxious.

If we consider these first two criteria (relating to information and power) with respect to some empirical data from the trade in human organs, it becomes evident that this market might violate the requirements Satz sets out. Using evidence of outcomes for vendors of kidneys in India, Greasley (2013) argues that markets for organs do not improve outcomes for vendors. Naqvi et al. (2007) analyses a set of interview data with 239 individuals who sold kidneys in Punjab, Pakistan, finding that 93% of vendors were motivated to sell a kidney to repay debts. The authors state:

> Many of the vendors, being illiterate, did not know the amount of money they had to pay back. . . . The vulnerability and exploitation of the vendors are reflected by the fact that none received the promised amount of money. Expenses were hidden from them and all were short charged when deductions were made from the agreed sum to pay for travel, food and hospital stay.

They found that the average sale price agreed was $1737, but vendors received only $1377 after charges for hospital and travel expenses. After the transaction, 88% reported no overall improvement in their financial position and 98% reported worse health.

Satz does not only require that these two market characteristics be in place, but also stipulates particular market outcomes, relating to individuals or society as a whole, that must not arise if a market is to avoid being considered noxious. Harmful outcomes for individuals might include transactions that have deleterious effects for them subsequently, for example a transaction that provides ammunition for a genocidal civil war or leaves people destitute or starving. Harmful outcomes for society occur if the existence of a particular market shapes individuals' preferences or capacities in a way that is detrimental to social welfare in the future. We may think of the effect of an unregulated commercial market for media which exerts negative effects on the identities of women or minorities. Similarly, as Sandel (2012) observes, a market in the trade of votes is immoral because it would undermine democracy.

4. An ethical evaluation of SIBs

When viewed through the lens of consequentialism - specifically utilitarian ethics - SIBs may be acceptable if they operate within, or construct, moral markets. Given this, a consequentialist evaluation of SIBs would focus on the benefits that they deliver to society and the costs that they incur. These benefits and costs and may be proximate, that is arising directly and visibly from the SIB or they may be ultimate, in that they arise more indirectly and possibly over a longer time horizon.

The proximate benefits, discussed above, include skills and knowledge transfers at the private-public interface and the transfer of some of the operational risk and upfront funding requirements from the public sector to the private sector. However, contracting costs for SIBs are very high, and are often implicitly borne by the public sector. A cost-benefit analysis might lead us to ask whether other, potentially better, ways of structuring social interventions exist that deliver similar improvements in welfare, without the need for profits flowing to private investors? Essentially this question asks whether the spillover of knowledge from the private sector to the social sector, or the transfer of risk from the public sector to the private sector, can be achieved in some other way? For example, this might be achieved through the placement of trustees and advisors from the private sector on the boards of charities and social enterprises and the increased use of interns from social investment organisations in these social organisations. Even if a proximate cost-benefit analysis appears to suggest that SIBs are acceptable from a utilitarian perspective, the extension to a longer term view of costs and benefits and the questioning of the characteristics demanded by Satz for a moral market, may raise ethical problems for SIBs.

4.1. Applying Satz's taxonomy to SIBs

The extent to which SIBs present ethical problems is likely to depend not only on the proximate costs and benefits they bring about, but also on whether they satisfy Satz's criteria for noxious markets. In what follows, I analyse the SIBs market with respect to Satz's taxonomy.

4.1.1. Weak agency — informational weakness of participants

The example provided by Satz (2010), and also Arrow (1963), concerning the ethical problems resulting from informational weaknesses in the market for healthcare is a good starting point for evaluating the ethics of SIBs. Indeed, both suppliers of social care and its recipients are unable to predict with any accuracy the causal effectiveness of an intervention, largely because evidence of prior interventions is tends to be noisy (due to difficulties in establishing neat categorisations of concepts, such as degrees of mental illness or homelessness) and potentially skewed by the existence of confounding factors, (such as changes in clinical definitions or in economic conditions). Even if past performance of interventions can be measured accurately, its effectiveness can be seen to be context dependent. By this, I mean that the application of the same intervention subsequently, in a different context, may have significantly different effects than those demonstrated for a past application. An intervention to reduce homelessness in for one particular group, in geographical location and at one particular time may yield different results from an intervention with different subjects in a different location and at a different time period. The fact that results vary with the context of application leads us to question the extent to which any participant can claim to be fully informed.

Yet, this informational weakness is likely to be as challenging for social policy makers, as profit-seeking SIB structures. As such, it might seem that the Satz informational requirement is not relevant for evaluating SIBs. I view them as relevant because of the structure of agency in the delivery of social interventions using SIBs. Satz views the market in fairly simplistic terms, with consumers and suppliers coming together as market-makers, but I argue that the reality of the SIB market is rather more complex. On the supply side, the decision-makers include not only the organisations that deliver social interventions (service providers) but additionally the commissioner, the investors, intermediaries, evaluation consultants and legal advisers. The *process* of decision-making is thus more subtle than described by Satz, with a variety of inputs from different agents.

As a result of the fact that agency is distributed across a number of players, it is hard to know exactly how to evaluate the requirement that decision-makers are fully informed. The connection between the investor and front line staff delivering social interventions will be punctuated and the decisions made will be between investor and intermediary, intermediary and management of the social organisation and finally, management at the social organisation and the staff delivering interventions. Even if there is a sense in which the group decision is fully informed, there is no single individual who is in possession of all the facts.[11] Consequently, it seems unlikely that Satz's informational requirements will be satisfied.

Furthermore, incentives in decision-making will vary between the different agents participating in a SIB. Thus, the 'decision-maker' described by Satz becomes little more than a metaphor, much as the market can be viewed as a metaphor for a whole series of processes motivated and constrained by a multitude of different factors. The primary objective of a social organisation entering into a SIB is to deliver on the contracted outcome requirements (such as reducing recidivism) in order that they satisfy their contractual obligations. If we consider what exactly is involved in this overall objective, however, we can see that many other incentives exist and may in fact compete with this primary goal of delivering outcomes. Reporting results (rather than just achieving them) is important. As a result, the social organisation will be motivated to provide evidence of achieving the required outcome to evaluators.[12] This may result in organisations focusing on the easiest targets or "shape their provision around the terms of the contract rather than the needs of the client" (McHugh et al., 2013: 250).

In addition, the organisation may be motivated to juggle the demands of different external stakeholders such as other funders (maybe donors) and regulators. Social organisations will also face demands from internal stakeholders, such as front-line staff, the finance team and trustees or board level advisors. Once all these different, and sometimes competing, demands are taken into consideration, it becomes evident that the use of a simple model that abstracts from the complexity of multiple agency may be inadequate for evaluating the ethical status of SIBs.

Even if we were to abstract from the problem of sequential decision-making and that of the multiple incentives facing agents, further issues with respect to the rationality of agents present difficulties with the use of Satz's taxonomy to establish the ethical status of a market. Satz's model implicitly assumes some level of individual or organisational rationality: agents

[11]The issue of judgement aggregation has been addressed by social choice theorists and philosophers (List and Pettit, 2011) but pragmatic problems in group decision-making have been highlighted (see for example, Alexander and Morley, 2017).

[12]Power (1997) has highlighted the increasing demands on organisations to provide auditable evidence.

are presumed to make rational decisions. Yet an extensive literature, spanning psychology economics, sociology and management theory has highlighted the fact that most decision-makers are in fact boundedly rational (Simon, 1955; Kahneman and Tversky, 1974, 1979; Tversky and Kahneman, 1986; Simon, 1988; Gigerentzer et al., 1999). Boundedly rational agents make choices that are inconsistent with rational choice theory; they may be subject to various biases, framing effects, inconsistent time preferences and may rely on heuristics.

Others have argued that any kind of consequentialist action (that is forward looking, rational or boundedly rational decision-making) may be inappropriate for explaining much of human behaviour. Humans tend to be swayed by proximate influences, such as others in their environment and may be driven by fads and fashion (Bikchandani et al., 1992; Shiller, 1995).[13] Others have argued that the very notion of decision-making is unclear, as rule-following, automated cognitive processes and habit may drive behaviour (Scott; Baldwin, 1988; Aarts et al., 1998). Where do these non-rational decision theories leave the Satz taxonomy of noxious markets? The requirements she makes for the informational characteristics of an ethical market seem to be inapplicable to many real markets, and certainly appear far too simplistic for SIB structures.

4.1.2. Vulnerability — the lack of power of participants

Those who are subject to social interventions via SIBs are by definition vulnerable. The objects of SIB-based interventions are those who suffer from the effects of poverty, mental health problems, homelessness and drug abuse, and as such lack market power. This inherent vulnerability undermines the moral legitimacy of the market according to Satz's model. Although some SIBs do give subjects the choice of opting in or out of programmes (Nicholls and Tomkinson, 2013: 19), not all do so, and it is not clear that the subjects are fully informed and capable of making that decision. This failure to satisfy Satz's requirement that market participants should not be vulnerable signals an ethical weakness in marketising this area. As with the informational issues discussed above, however, the vulnerable nature of some participants does not itself single SIBs out as unethical, for most social interventions will be made on subjects who have little or no choice in the matter.

[13] For an excellent review of the literature on 'herding' behaviour see Raafat et al. (2009).

However, an important difference between SIBs and other public sector interventions, is that SIBs are intended to yield a profit for investors.[14] A crucial piece of information that may not be made clear to the individuals who are intervened upon by the SIB is that they are effectively a profit-centre for private investors. In 2011, the Economist ran an article about the Peterborough SIB. It carried a photograph of a prisoner in a cell with the caption "Fancy a flutter?", thereby capturing the ethical problems associated with gambling on the outcome of a human experiment, based on a subject who is imprisoned and potentially subject to a number of social and psychological problems. Betting on the invidious position of the prisoner may be considered to be showing insufficient respect for this person (which would be treated as unethical by many subscribing to a deontological approach) while also highlighting the relative imbalance of information and power between gambler and the prisoner.

The fact that corporations such as Goldman Sachs have invested in similar SIBs demonstrates the likelihood of a power imbalance between the agencies involved in social care and the recipients of that care. While the social enterprises and charities that deliver front-line interventions as part of a SIB may be mission-driven and motivated by moral concerns, the same cannot be said to be true of corporations motivated by profit-generation and serving the demands of its shareholders.[15] The participation of profit-driven corporations may be problematic from an ethical perspective. Where opportunities for profit-making conflict with moral objectives, we would expect the directors of a corporation to act in the best interests of the company, which is usually interpreted as generating shareholder wealth by prioritising activities that yield financial returns (Friedman, 1953).

The prioritisation of profit over moral activities by investors in SIBs matters because the profit motive of investors may drive the method and timing of delivery, thereby subjecting those upon whom the social interventions are made to potential short-termism and a focus on subjects who are most likely to demonstrate results, rather than those who are most deserving (known as 'creaming' or 'cream skimming'). For instance, in the case of a SIB addressing recidivism, such as those in Peterborough, UK and Riker's Island, US, we should be aware that those driving the

[14]As Baliga (2013) has noted, private prisons and SIBs have common features in their profit-seeking approach to providing social care.

[15]Goldman Sachs invested in a SIB developed by New York city focused on recidivism of prisoners at Rikers Island prison.

SIB might be motivated to focus on those prisoners who are most likely to stay out of prison for the stipulated time, and thereby ignoring the 'hard cases' who are least likely to respond (a practice known as 'parking'). A recent OECD report into SIBs referred to this problem of incentives leading to gaming. It stated:

> There is a fear that this strong focus on results can change the public service ethos or lead to a narrow mechanical determinism in service delivery. It may be possible to game the results by selecting clients that are easiest to reach ("cream skimming") while leaving those that would be most expensive without service ("parking"). (OECD, 2016a: 16)

Although the report found no evidence of such gaming activity in the two SIBs it analysed (one of which was the Peterborough SIB), it acknowledged that placing incentives on social organisations to demonstrate the achievement of specific targets might have unwanted consequences:

> Focusing too much on narrow metrics may lead to unintended consequences and gaming in fields where cooperation and holistic approaches may be required. (OECD, 2016a: 20)

The existence of incentives for gaming activities renders the SIB ethically questionable, particularly because the subjects are not in a position to opt out of the intervention and select an alternative that is not profit-based.

4.2. Outcomes associated with SIB markets

Satz specifies certain inputs, or structural characteristics of the market, that she views as necessary for an ethical market, but also two key outcomes. These are that no harm should be done to individuals, and that no harm should be done to society overall. Although these are perhaps less obviously relevant for SIBs, I consider them below in turn, drawing on work by Michael Edwards, an academic with extensive experience at the Ford Foundation, World Bank, Oxfam and Save the Children.

4.2.1. Harm for the individual — Damages the welfare or agency of the individual

Satz argues that the profit motive may lead to the preservation of a suboptimal societal status quo or even the indoctrination of particular groups in ways that undermine their power and life chances. Advertisers may use racial or gender stereotyping as a means to develop brand value and thereby generate future financial returns. Such activities may undermine the long term agency and utility of particular groups, without their knowledge. In the same vein, Edwards (2008) highlights the fact that many social problems can only be improved in the very short term by activities that are consistent with the incentives of profit-seeking organisations, and that a true resolution of such problems often requires fundamental societal change that is inconsistent with corporations' agendas.

Efficiency in the delivery of social care may be improved by knowledge spillovers and particular incentives through SIBs, but this form of profit-led social policy may overlook the underlying cause of social problem. For example, job creation SIBs may address certain problems related to structural unemployment, but may fail to address the fundamental problems of economic growth, class issues, poverty and poor educational attainment among lower socio-economic groups that are the root cause of youth unemployment.

One stakeholder group that is often neglected in discussions concerning the ethics of social investment and SIBs, is that of staff at social enterprises and charities. How might their welfare be affected if the organisation for which they work is engaged in a SIB? Just as prisoners receiving social care from a SIB may feel uncomfortable at the prospect of becoming a 'profit centre', staff at the social enterprise that is delivering the service might also feel uncomfortable at the prospect of working to deliver profits to external financial stakeholders through their interactions with clients. Such profit-based activity might be viewed as inconsistent with a worker's social identity that is based on the performance of charitable actions motivated by moral sentiments.

Even if staff view the profit-motive as acceptable, another component of a SIB, the measurement of impact and outcomes, might cause motivational issues. SIBs require impact and outcome measurement for contracting purposes, as well as the use of impact reporting to enable investors and intermediaries to manage the ongoing performance of the investee. It has been argued that the excessive use of impact measurement and impact language sometimes exerts a negative effect on staff motivation because descriptions of their activities that use the language of impact measurement can appear inconsistent with their social identity and undermine the moral worth they assign to their role (Morley, 2017). The classic comparative study of blood donation schemes, one of which used volunteers and one which paid donors for their blood, highlighted the importance of intrinsic (compared with extrinsic, or financial) motivation (Titmuss, 1971) and was followed by a number of studies addressing intrinsic motivation.[16] Scholars have identified a significant positive correlation between working in a non-profit and job satisfaction, finding that, "non-profit work is institutionally different than for-profit work, providing employees with the possibility to enjoy intrinsic benefits from work" (Benz, 2005: 168). Others have identified a model of trust that is important to the motivation of staff in social organisations, and argued that,

> The introduction of incentives of almost any kind is considered by the trust models' proponents to be actually damaging to professional motivation. For those who are motivated to perform effectively by internal factors, such as their sense of professional duty and their altruistic concern for the welfare of the people they were serving, the use of external rewards or penalties at best leads to discouragement and demoralisation, and at worst actively promotes damaging, self-interested behaviour. Extrinsic motivation drives out intrinsic motivation. (Le Grand, 2010: 59).

If staff motivation is damaged by the taint of profit-making or by the use of excessive impact measurement, the welfare of staff may be reduced, and this in turn may, in the worst case, compromise relationships with clients and ultimately damage the the effectiveness of the organisation.

4.2.2. Harm for society — Undermines civil society

Satz argues that the marketisation of certain social activities can damage democracy and civil society. She gives the example of the buying and selling of votes in a general election which would ultimately undermine the democratic process. This criterion for ethical markets may not initially appear to be as relevant for SIBs as the earlier issues relating to information, vulnerability of participants and the causing of harm to individuals or specific sub-groups in society. Yet it

[16] The study concluded that paying for blood reduced not only the quality but the supply of blood because some donors' motivation was potentially damaged by marketisation.

is evident that as a free market ideology takes hold, non-market based solutions may become increasingly unfashionable. Charitable structures may be considered anachronistic and those with commercial or financial experience may be viewed as having more relevant expertise in many cases than those with issue-specific experience.

The privatisation of care delivery that existed prior to SIBs left commissioners with discretion about the exact composition of social care that was to be provided by private subcontractors. By contrast, the SIB structure enables investors and intermediaries from the private sector to dictate elements of care delivery, in particular, performance management. Investors who stand to lose their entire investment if targets are not achieved will inevitably demand a say in the strategy employed by the front line organisations. They will be motivated to change the approach used if they are not seeing results, and potentially change the organisations that deliver the service. The UK government acknowledges the need for investors to control the delivery of the SIB, because of the risk they are bearing and suggests that this can be achieved in various ways, such as taking a seat on the board of the service delivery provider or "replacing providers if outcomes aren't being achieved".[17]

On one hand, this interface between the different groups involved in a SIB is necessary for the transfer of knowledge and skills from the private investors to the social care providers, and for generating the conditions for innovation and improved efficiency. On the other hand, the delegating of strategic and operational decision-making about social care to profit-seeking organisations represents an effective move towards a privatisation of policy decisions. While this privatisation of social policy may work for the limited number of SIBs that have been developed to date, the worry is that this transfer of power over social care to private interests, if extended, may undermine certain elements of civil society. Given that many of the issues addressed by SIBs are what we might label 'out of market' problems, it is questionable whether the corporate sector is capable of delivering consistently better, and fairer results.

In *Small Change*, Michael Edwards argues that business cannot be relied upon to resolve social issues. This is because profit-seeking organisations will inevitably focus on improving the existing system — a system that is consistent with their own long term success — rather than

[17] https://www.gov.uk/guidance/social-impact-bonds.

making fundamental changes that may conflict with their desire to maintain the status quo. He argues that in order to resolve hard social problems we should "want to transform economic power structures and ways of living together, rather than just using markets as instruments to deliver social goods" (Edwards, 2008: 28). We would not expect corporations, whose shareholders demand financial returns in the short term, to engage in activities that might damage their markets or inflate their costs. The problem with delegating policy to business is that fundamental conflicts of interest exist between maintaining brand value and predictable future cash flows and making significant changes to power relations and the capitalist system. Businesses are not likely to be incentivised to effect the kinds of fundamental change to structures within society which may be necessary if deep-seated social problems are to be solved. Instead, business involvement in social care can legitimately be seen to provide innovation at the margins, to improve some aspects of operational efficiency and to extend access to individuals who would potentially have been excluded. While this result contributes to society in a positive way, it is limited and may restrict future options for social improvements. The conflict of interest between commercial interests and those who want fundamental social change lie at the heart of the problem with extending SIBs. While the private sector can do good at the margins, if allowed to dominate social care policy, it may undermine the advancement of civil society values and improvements to the welfare of certain groups within society.

4.3. Additional concerns

Thus far, I have considered the extent to which SIBs raise issues within the taxonomy of the market characteristics and outcomes proposed by Satz (2010), and have also extended her requirements, for example considering the potential complexity of decision-making given the SIB structure and considering the practical likelihood that decision-making is not rational, but possibly boundedly rational, rules-based or simply following fads or fashions. Furthermore, I considered the staff within the organisations as market participants, whose welfare, motivation and relationships with clients, may be affected by the moral status of the organisation for which they work, or even by the way the organisation describes their work using social impact reporting. Other issues, which do not fall quite so neatly within Satz's taxonomy may also be relevant. I propose two further issues relating to the possible secondary effects of SIBs, that may lead to undesirable ethical outcomes for particular groups or for society.

Issues of exclusion: The first issue relates to the exclusion and ultimate damage to particular groups. If SIBs (and social investment more generally) became more dominant as a means of financing social interventions, might some deserving recipients be excluded and thereby disadvantaged? Two types of recipients come to mind. The first, are those for whom the effect of social interventions is hard to measure objectively, or that is revealed only after an extended period of time. For these groups, SIBs are not appropriate, but if SIB structures become prevalent, it may lead some social enterprises to move away from serving these clients. A second concern might arise because commercial SIB investors have reputations at stake and may be motivated to preserve the value of their brand. As a result, certain groups may be rejected because of their negative associations, for example groups working with paedophiles or violent offenders. It might not be easy for organisations working with such groups to gain access to resources, using these kinds of mechanism because social investors might balk at the prospect of risking their brand reputation through such an association.

Secondary incentives that worsen social problems: Due to the misalignment of long-term incentives between profit-seeking commercial interests and charity, there are also risks that the market systems which replace public values will result in actions that worsen social problems. Marketisation of healthcare and prisons may be problematic if commercial investors are motivated to take actions that are unethical in order to increase the net worth of their shareholders. An example is private prisons in the US. Some have argued that providers have lobbied for for increased rates of incarceration and the lengthening of custodial sentences (Sarabi and Bender, 2000) while others are more agnostic about the direction of causality, but note the "incarceration binge" in the US that accompanied the rise in private prisons. It would certainly be unethical if private interests had caused an increase in incarceration, but even if they have simply profited from it, their moral position would be questionable.

5. Conclusion

SIBs offer a tantalising opportunity for innovation and for improvements in allocational and operational efficiency of social care provision. Yet, they also raise signficant ethical concerns which I addressed in this paper. Using a consequentialist approach, and drawing on the framework developed by Satz (2010) to establish if a market is noxious, I evaluated the ethics of marketising social care using SIBs.

First, I analysed the ethical status of SIBs in terms of two market characteristics: the extent to which participants are informed and the extent to which they lack any significant power imbalance. I argued that information concerning the causal effects of social interventions, is often weak, but that it is not necessarily any different for SIBs than for other provision of social interventions, whether through the public sector or civil society organisations. However, subjects of SIB interventions may not be fully informed about the differences in treatments associated with opting in or out of a SIB cohort, and may not understand that investors are profiting from them. Furthermore, given the power imbalance between the external investor and the participant, and the vulnerable nature of many of the participants, I argued that these individuals may not be capable of giving informed consent regarding participation in SIB interventions. Additionally, I extended Satz's taxonomy by arguing that the view of decision-making presented in her framework fails to reflect the complex structure and varied incentives within SIBs.

I then considered Satz's requirements concerning market outcomes for groups of individuals and for society overall. I argued that SIBs may fail to satisfy the requirement not to cause harm to groups of individuals because the objectives of profit-seeking organisations may be inconsistent with actions that improve the welfare of certain groups. I argued that *internal* stakeholders, in particular staff, may be negatively affected by the activities of SIBs. This is because the intrinsic motivation of staff may be undermined if the social organisation engages with profit-seeking organisations and begins to use the business-style of language associated with reporting social impact. With respect to Satz's requirement that ethical markets should not exert any negative impact on the welfare of society overall, I suggested that the partial delegation of policy-making to SIBs may be problematic. Profit-seeking investors may be motivated to achieve a 'quick fix' and hit the required targets for improving certain social problems, but are unlikely to attempt to shift fundamental social structures that may be required as an effective, long term solution. This would be of particular concern if SIBs were to become the dominant means of organising social services.

Finally, I extended Satz's taxonomy by suggesting secondary effects relating to the marketisation of social care that are relevant for an ethical evaluation of SIBs. I suggested that issues of exclusion may arise if SIBs become prevalent for social care delivery. Certain social issues do not lend themselves to measurement, for example issues related to mental health that require long term interventions and where the results of such interventions are hard to demonstrate objectively. In addition, some social issues may be excluded from SIB funding because the issue is viewed as socially unacceptable, or even immoral, by the SIB investors' stakeholders. For example, a SIB to help paedophiles might fail to attract interest from social investors whose other stakeholders might balk at their association with this group. I argued, further, that the extension of SIBs might be associated with secondary incentives that are inconsistent with increased social welfare due to the misalignment of incentives for policy-setting between private interests and overall social welfare.

The main thrust of this paper, then, is that SIBs have much to offer, but that we should be somewhat circumspect, as they demonstrate the potential to generate ethical problems. These problems will be exacerbated if they are extended to cover a significantly greater proportion of social interventions. Future research could usefully focus on methods of regulation to address these ethical issues.

References

Aarts, H., Verplanken, B., and van Knippenberg, A. (1998). Predicting behavior from actions in the past: repeated decision making or a matter of habit? *Journal of Applied Psychology*, 28(15):1355–1374.

Alesina, A. and Perotti, R. (1994). The political economy of growth: A critical survey of the recent literature. *The World Bank Economic Review*, 8(3):351–371.

Alexander, J. M. and Morley, J. (2017). Group agency, collective attitudes and intragroup deliberation. Working paper.

Arnold, N. S. (1987). Why profits are deserved. *Ethics*, 97:387–402.

Arrow, K. J. (1963). Uncertainty and the welfare economics of medical care. *The American Economic Review*, 53(5).

Baldwin, J. D. (1988). Habit, emotion, and self-conscious action. *Sociological Perspectives*, 31(1):35–58.

Baliga, S. (2013). Shaping the success of social impact bonds in the united states: Lessons learned from the privatization of us prisons. *Duke LJ*, 63(2):437–479.

Benz, M. (2005). Not for the profit, but for the satisfaction? – evidence on worker well-being in non-profit firms. *Kyklos*, 58(2):155–176.

Bikchandani, S., Hirshleifer, D., and Welch, I. (1992). A theory of fads, fashion, custom and cultural change as information cascades. *Journal of Political Economy*, 100:992–1026.

Brown, G. (1992). Are profits deserved? *Jounral of Business Ethics*, 11:105–114.

Center for Global Development and Social Finance (2013). Investing in social outcomes: Development impact bonds. Technical report.

Corrigan, P. (2011). Social impact bonds: A new way to invest in better healthcare. `http://www.socialfinance.org.uk/wp-content/uploads/2014/07/A-new-way-to-invest-in-better-healthcare.pdf`, Social Finance.

Dorfman, A. and Harel, A. (2013). The case against privatisation. *Philosophy and Public Affairs*, 41(1):67–102.

Dowling, E. (2016). In the wake of austerity: social impact bonds and the financialisation of the welfare state in britain. *New Political Economy*, http://dx.doi.org/10.1080/13563467.2017.1232709.

Dowling, E. and Harvie, D. (2014). Harnessing the social: State, crisis and (big) society. *Sociology*, 48(5):869–886.

Edwards, M. (2008). *Small Change: Why Business won't Save the World*. Demos. Berrett-Koehler Publishers Inc.,.

Floyd, D. and Gregory, D. (2015). The report of the alternative commission on social investment. Technical report, Commissioned by Esmee Fairbairn Foundation.

Fraser, A., Tan, S., Lagarde, M., and Mays, N. (2016). Narratives of promise, narratives of caution: A review of the literature on social impact bonds. *Social Policy & Administration*.

Friedman, M. (1953). *Essays in Positive Economics*. University of Chicago Press.

Galbraith, J. K. (1958). *The Affluent Society*. Houghton Mifflin.

Gigerentzer, G., Todd, P. M., and the ABC Research Group (1999). *Simple Heuristics that make us smart*. Oxford University Press.

Graafland, J. J. (2009). Do markets crowd out virtues? an aristotelian framework. *Journal of Business Ethics*, 91:1–19.

Greasley, K. (2013). A legal market in organs: the problem of exploitation. *BMJ Journal of Medical Ethics*, 40(1–2):(online).

Gustafsson-Wright, E., Gardiner, S., and Putcha, V. (2015). *The potential and limitations of impact bonds: Lessons from the first five years of experience worldwide*.

Hippen, B. E. (2005). In defense of a regulated market in kidneys from living vendors. *Journal of Medicine and Philosophy*, 30(6):(online).

Hirsh, F. *Social Limits to Growth*. Cambridge, Mass,. and London Harvard University Press.

Hood, C. (1991). A public management for all seasons? *Public Administration*, 69(1):3–19.

Humphrey, C., Miller, P., and Scapens, R. W. (1993). Accountability and accountable management in the uk public sector. *Accounting, Auditing and Accountability Journal*, 6(3):7–29.

Joy, M. and Shields, J. (2013). Social impact bonds: The next phase of third sector marketization? *Canadian Journal of Nonprofit and Social Economy Research*, 4(2):39–55.

Kahneman, D. and Tversky, A. (1974). Judgment under uncertainty: heuristics and biases. *Science*, 185:1124–1131.

Kahneman, D. and Tversky, A. (1979). Prospect theory: an analysis of decision under risk. *Econometrica*, (47):263–291.

Kurunmäki, L. and Miller, P. (2011). Regulatory hybrids: Partnerships, budgeting and modernising government. *Management Accounting Research*, 22:220–241.

Le Grand, J. (2010). Knights and knaves return: Public service motivation and the delivery of public services. *International Public Management Journal*, 3(1):56–71.

List, C. and Pettit, P. (2011). *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford University Press.

Lowe, T. and Wilson, R. (2015). Playing the game of outcomes-based performance management. is gamesmanship inevitable? evidence from theory and practice. *Social Policy & Administration*. SPA-2014-0090.R1.

McHugh, N., Sinclair, S., Roy, M., Huckfield, L., and Donaldson, C. (2013). Social impact bonds: A wolf in sheep's clothing? *Journal of Poverty and Social Justice*, 21(3):247–257.

Morley, J. (2017). The impact of "impact": The effect of social impact reporting on staff identity and motivation at social enterprises and charities in the uk. *Unpublished working paper*.

Mulgan, G., Reeder, N., Aylott, M., and Bo'sher, L. (2011). Social impact investment: The challenge and opportunity of social impact bonds. Technical report, Young Foundation, London.

Naqvi, S. A. A., Ali, B., Mazhar, F., mirza Naqi Zafar, and Rizvi, S. A. H. (2007). A socioeconomic survey of kidney vendors in pakistan. *Transplant International*, 20(11):934–939.

New Philanthropy Capital (2016). What makes a good charity? Technical report, http://www.thinknpc.org/publications/good-charity/.

Nicholls, A. (2009). We do good things, don't we?: 'blended value accounting' in social entrepreneurship. *Accounting Organizations and Society*, 34:755–769.

Nicholls, A. and Murdock, A. (2012). The nature of social innovation. In Nicholls, A. and Murdock, A., editors, *Social Innovation: Blurring Boundaries to Reconfigure Markets*. London: Palgrave Macmillan.

Nicholls, A. and Tomkinson, E. (2013). The peterborough pilot social impact bond. Technical report, Said Business School.

OECD (2016a). Social impact bonds: State of play & lessons learnt. Technical report.

OECD (2016b). Understanding social impact bonds - working paper. Technical report, OECD.

Parfit, D. (1984). *Reasons and Persons*. Clarendon Press, Oxford.

Power, M. (1997). *The Audit Society: Rituals of Verification*. Oxford University Press.

Raafat, R. M., Chater, N., and Frith, C. (2009). Herding in humans. *Trends in Cognitive Sciences*, 13(10):420–428.

Sandel, M. J. (2012). *What money can't buy: The moral limits of markets*. Penguin, London.

Sarabi, B. and Bender, E. (2000). Prison payoff: The role of politics and private prisons in the incarceration boom. *National Criminal Justice Reference Service*, NCJ 20795.

Satz, D. (2010). *Why some things should not be for sale: The moral limits of markets*. Oxford University Press.

Scott, J. F. *Internalization of Norms: A Sociological Theory of Moral Commitment*. Englewood Cliffs, NJ: Prentice-Hall.

Shiller, R. J. (1995). Conversation, information, and herd behavior. *American Economic Review*, 85(2):181– 185.

Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69:99–118.

Simon, H. A. (1988). Rationality as a process and as product of thought. In Bell, D. E., Raiffa, H., and Tversky, A., editors, *Decision Making: Descriptive, Normative and Prescriptive Interactions*, pages 58–77. Cambridge University Press.

Sinclair, S., McHugh, N., Huckfield, L., Roy, M., and Donaldson, C. (2014). Social impact bonds: Shifting the boundaries of citizenship. In Farnsworth, K., Irving, Z., and Fenger, M., editors, *Analysis and Debate in Social Policy*. Social Policy Review 26.

Social Finance (2016). Social impact bonds: The early years. Technical report, Social Finance.

Stoesz, D. (2014). Evidence-based policy: Reorganizing social services through accountable care organizations and social impact bonds. *Research on Social Work Practice*, 24(2):181–185.

The Economist (2011). Performance bonds: Who succeeds gets paid.

Titmuss, R. (1971). *The gift relationship: From human blood to social policy*. NY: Random House.

Tversky, A. and Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, 59:S251–78.

Warner, M. (2013). Private finance for public goods: social impact bonds. *Journal of Economic Policy Reform*, 16(4):303–319.