Give greed a chance: issuing Social Policy Bonds to achieve social goals

Note: this document is in two sections: this section is a five-page article outlining the main features of the bond concept, applied to female literacy in Pakistan. Section 2 goes into more detail and is 15 pages long.

Give greed a chance: Female Literacy Bonds

There are two ways of trying to solve the worlds' social problems. One way, the way followed by most people and most governments, is to finance institutions, or people, or activities that have as their stated objective the solution of these problems. The other, much simpler but also much rarer, is to reward people for actually solving the problems. That's what they used to do in the Wild West when they wanted to eliminate problems caused by particularly nasty individuals. They put together a cash reward, printed some 'Wanted - Dead or Alive' posters, and let the private sector do the rest. The same principle, more or less, underlies an innovative new financial instrument that can channel the market's incentives and efficiencies into the achievement of social and environmental goals.

Rewarding literacy

Say that instead of wanting to capture someone, or get them bumped off, you had another objective. It could concern anything from the identity of the mayor of your city to the fortunes of your local sports club but let us take something more edifying, like the literacy rate of girls and young women in Pakistan, which you would like to see rise to 95 per cent. You don't know the first thing about Urdu lexicography but you are wealthy, and you believe that raising the female literacy rate in Pakistan would be a good thing in itself, and perhaps even make the world a safer place for yourself and your descendants. So you get together with a few wealthy cronies who are of like mind and put some of your collective millions into an escrow account. This is your donation to the cause of female literacy in Pakistan. You can then call on members of the public to add any spare cash they might have into this account. Then you organise the printing of some fancy pieces of paper, which you call 'Female Literacy Bonds'. These bonds, you promise, shall be redeemed for $10 each once the literacy rate of 14-year old females in Pakistan rises to 95 per cent. You sell them on the open market for whatever price they will fetch. (You keep this sum yourself as a reward for your good behaviour and to cover the costs of printing the bonds and other administrative chores.)

The next step is relatively simple: you just organise some reputable and trusted body to go and test the literacy of Pakistani girls once a year, and … that's it. You don't pay interest on the bonds, you just sit back and when the female literacy rate in Pakistan does reach 95 per cent you give the go-ahead to the escrow account managers and tell them to redeem each bond for $10.

The bonds might at first sell for a fraction of $10 because everybody thinks you're dreaming: female literacy in Pakistan is never going to rise much in their lifetime. Or maybe the bonds will sell for a pittance simply because nobody hears about them, the media being filled with more important stuff like Hillary Clinton's memoirs.

Some people who couldn't care less about female literacy might nevertheless buy the bonds for a couple of dollars each and sit on them, much as they would on a lottery ticket. What happens then? The value of the bonds would fall still further. That's fine. The lower the value of the Female Literacy Bonds falls, the more profit people can make if they buy the bonds *and then do something to raise the literacy rate of girls in Pakistan*. The bonds would be tradable so people can sell them whenever they want. If somebody thought they could do something to raise the literacy level, then they would buy bonds and make a profit on the increase in value as it became more likely that the target will be achieved quickly. They wouldn't have to wait till the objective had been achieved: the market would value their bonds more highly, even before redemption.

Outcomes

There are many funding programmes that distribute cash to favoured activities, organisations or individuals, but Female Literacy Bonds would inextricably link payments to the targeted *outcome*: a female literacy rate of 95 per cent in Pakistan. Unlike programmes run by governments or NGOs the bonds would encourage diverse, responsive and cost-effective projects. Bondholders would gain most by ensuring that literacy is raised quickly. Importantly, you and your public-spirited colleagues would be making no assumptions as to *how* to raise literacy — that would be left to whoever buys the bonds. As issuers you would decide only on the precise definition of the outcome you seek, not on how to achieve it. That would be left up to investors in the bonds, who have every incentive to maximise the increase in literacy for each dollar they spend.

Markets

Why not just give your millions to charity instead? Charities do marvellous work with limited resources, but there are things that they cannot do, however dedicated they are. They cannot use their funds to bribe officials either to do their job properly or to look the other way. They cannot deliberately undermine those in power who are opposed to their cause and determined to resist it. They cannot, in short, play hardball even when doing so would greatly benefit thousands of innocent people. But it is not just a matter of standing up to the thugs, the kleptocrats, the well-meaning idealists, the ill-meaning ideologues, the politicians, the generals, and the men of religion who are doing so much to keep their people ignorant and poor. It is also a matter of bringing financial self-interest - greed, in other words - into the equation. What the bonds do is inject market forces into the achievement of higher female literacy rates.

Markets are the most efficient means yet discovered of allocating society's scarce resources. They have had a bad press because most people associate them with capitalism red in tooth and claw, with extremes of wealth and poverty, or with raping the environment. But there is nothing inherently evil about markets: they can serve public, as well as private, goals, and the bond concept is a way of channelling the market's efficiencies and incentives into the achievement of public goals.

Bondholders would have incentives to carry out a wider range of literacy-raising initiatives than either governments or NGOs, and to do so more cost-effectively. Apart from buying off or bypassing the people in authority who block progress toward higher literacy rates, bondholders could lobby the Pakistani government to give a higher priority to literacy in schools, or they could develop literacy-teaching projects of their own. They might finance literacy programmes for TV, or set up village schools, or give prizes to the most literate families in villages. It would be up to bondholders to decide on those programmes that will give them the highest increase in female literacy per unit outlay. The market prices of the bonds, and their changes over time, will generate helpful information as to how fast the objective is being achieved. These prices would be publicly quoted, just like those of ordinary bonds or shares. The literacy of random samples of 14-year old Pakistani girls could be measured annually by pre-selected reading tests. Once the 95 per cent target had been reached and sustained for, say, three years, the bonds would be redeemed.

Some in the Pakistani Government, religious institutions or militant organisations might resent the targeting of such objectives by external agencies in this way. But, while under the current system they can oppose literacy teaching in ways that attract support, under a Female Literacy Bond regime, they would have openly to declare their opposition to female literacy itself. It is precisely this focus on the *outcome*—rather than activities or institutions—that would help strengthen the coalition working to achieve it.

Too many small bondholders could probably do little to build peace. The bonds would most probably end up in the hands of a few large holders, who would have incentives to co-operate with each other, and to finance those projects that they believed would be most effective in raising the level of female literacy.

Female Literacy Bonds for all

- If you are cash-rich but time-poor and know what you want, then you could get together with some cronies and do as suggested above: set up an escrow account and issue your own bonds. If you are less wealthy you could swell the redemption funds by depositing your spare cash into escrow account set up by others.

- If you have more energy than money you could buy some Female Literacy Bonds, and then work to raise female literacy in Pakistan. Your bonds would appreciate in value if literacy levels rose quickly. You could even borrow on the strength of the expected increase in capital value of your bonds, in order to finance literacy-raising projects. You could co-operate with other bondholders and finance those activities that you think will be most efficient in raising literacy.

- If you are already involved in trying to raise literacy in Pakistan you could contact holders of Female Literacy Bonds and if they believe your activities are efficient in reducing conflict they will find it worthwhile to help finance your existing projects.

Top of document

Section 2

The text below elaborates on the article above. It is about 15 pages long, including a technical annex.

How would Social Policy Bonds work?

They would create an interest group — bondholders — who have a strong financial interest in achieving the targeted social objective efficiently, or in paying others to do so. Assume that you and some like-minded friends agree to stump up say £10 million to increase female literacy in Pakistan to 95 per cent. You issue one million Female Literacy Bonds that each become worth £10 when this objective has been achieved and sustained for, say, three years. Because the market would see this objective as unlikely to be achieved in the near future, it would put a low value on the bonds when they are floated. Assume successful bidders pay as little as £1 for the bonds. (This sum would be held by the issuing authority partially to offset the cost of future redemption of the bonds.) Now, they hold an asset that could appreciate in value by 900 per cent if the targeted increase in female literacy in Pakistan were achieved. This provides the motivation for bondholders to do whatever they can to increase the female literacy rate in Pakistan as cost-effectively as possible. .

What objectives?

Most likely you would not want your bonds simply to become a bet on circumstances beyond purchasers' control. So, unless your fortune is very large, and you are very generous, you might not want to issue bonds aimed at achieving and sustaining a stable world climate. If you did issue 'Climate Stability Bonds', any actions bondholders might take would probably be swamped by forces beyond their control, such as governments' compliance with their obligations under the Kyoto Protocol, or natural processes.

The best sort of objective would be one that is relatively self-contained, in the sense that its targeting would not transfer problems to other regions, or create different problems in the same region. The idea is that the objective that you target should not be achieved at the expense of other societal goals.

If your objective concerns a developing country then you are likely to achieve more 'bang per buck', as your pounds will have more purchasing power there. You could make a real difference targeting something like the female literacy rate in Pakistan, or infant mortality in an African country (see immediately below: *Objectives for developing countries*).

On the other hand, you could still make a difference in a developed country, if your objective is focussed. You could, perhaps, target the crime rate in a particular town or region, or the water quality of a particular lake or river, or air quality in a particular town.

Issuing Social Policy Bonds means that you need to be clear about what you actually want to achieve. In principle ends, rather than means to ends, make better targets for Social Policy Bonds. So it will generally be preferable to target the negative effects of problems than the causes of these effects. So, if your concern is infant mortality in Malawi, and malaria is currently responsible for a high proportion of child deaths in Malawi, you might be tempted to issue bonds redeemable only when the number of children with malaria in Malawi has fallen to a low level. But you would do better to issue target infant mortality in general, rather than malaria in particular. If you are concerned about the number of people who are homeless in your area, you should target this number, rather than housing starts, and leave it for bondholders to decide on how best to achieve the desired goal. Or if you would like to improve the water quality of a river flowing through your land, you should not make any single pollutant, such as nitrates (resulting from fertiliser application), the sole target of a bond issue if it were likely that farmers would respond by increasing the use of other pollutants. Instead, water pollution, defined as a combination of pollutants, could be targeted. But even better would be to target some meaningful indicator of the effects of pollution. You might want, for instance, to target the range of fish species found in the river, and issue bonds redeemable only when certain fish had been found in sufficient number, over a sustained period.

Objectives for developing countries

Social Policy Bonds would in many ways be at least as well suited to application in developing countries. There are several reasons for this:

- Governments and non-governmental organisations (NGOs) would find it more difficult to go through unconventional channels, or to engage in the hardball tactics that might be necessary. Public sector employees in developing countries are generally not well paid and more likely to be corrupt than in most developed countries. Aid donations, whether from overseas governments or from NGOs or charitable organisations, can fail to reach their intended beneficiaries, or come with strings attached. Social Policy Bonds could bypass, co-opt or undermine such obstacles. Bondholders would have incentives to get value for money, and they will be prepared to bribe or otherwise play hardball with those in authority if it is the most efficient way of achieving your targeted objective.

- The very severe social problems and the enormous social changes that are occurring in developing countries mean that there are plenty of outlets for well-intended, outcome-focused philanthropy. Developing countries are urbanising rapidly, with all the social dislocation this entails. Infant mortality rates are high, and there is a great deal of poverty, unemployment and squalor. Many children are outside the educational system altogether and standards in state systems, while variable, are generally low. Environmental problems are especially severe in developing countries, and can be a threat not only to inhabitants of these countries, but to parts of the planet's entire ecosystem.

Social Policy Bonds could therefore target broad health, educational and environmental objectives, where improvements could come quite rapidly, and whose achievement could bring very large net benefits to large numbers of developing country citizens.

There is clearly something wrong with the current aid programmes run by western governments. In the long run, if privately-backed Social Policy Bonds were more successful in actually helping solve the problems of people in developing countries, governments could possibly issue Social Policy Bonds themselves. They could replace their - often misdirected - aid to developing countries, and instead back bonds. Instead of aid being given on a government-to-government basis, funds would thereby bypass corrupt politicians and officials, and the institutions they control. Funds aimed at solving global environmental problems, such as climate change, could similarly reward those who undertook worthwhile projects, rather than corrupt governments. Or corrupt governments could even choose to buy Social Policy Bonds targeting problems in their own countries themselves. Their financial self-interest would encourage them to modify their behaviour in favour of targeted objectives. The example of Female Literacy Bonds, described in the first section of this page, illustrates how this might work.

Indicators

Your choice of objective will be influenced in part by how readily it can be *targeted by quantifiable indicators*, whose progress accurately corresponds with progress toward the desired social outcome. Obviously it would be unsatisfactory if, for instance, your female literacy objective could be achieved if all literacy-raising resources were concentrated on few dozen Pakistani schoolgirls. Your Female Literacy Bonds should therefore include some provision for random testing of literacy from a sample drawn from all Pakistani girls: in other words, the identity of the girls whose literacy is to be tested for bond redemption purposes should not be known in advance.

Depending on your chosen objective, there would be incentives for people to do their own information gathering. For instance, if you targeted female literacy in Pakistan, depending on the sum riding on the achievement of this goal, potential investors in the bonds would carry out their own literacy tests, or at least take educated guesses as to the likely results of the literacy tests that you have stipulated will decide whether the bonds are to be redeemed or not.

Time period

Social Policy Bonds are versatile: you could incorporate bonus provisions if the goal you target is achieved by a specified date. Or you could stipulate that bonds would not be redeemed unless the targeted objective were achieved by a certain date, or that they would be redeemed for a sum that would diminish over the time it took for the objective to be achieved. The market would factor all such penalties or bonuses into the bond price.

Social Policy Bonds are efficient

Cascading incentives

Governments and NGOs already work towards social and environmental objectives. But the crucial advantage of a Social Policy Bond regime is that it would give people incentives to seek out and develop those ways of achieving the targeted goal that are most cost-effective. Under a bond regime the self-interest of bondholders would act so as to encourage those ways of reducing crime that would give you, as the backer of the bonds, the best return for your outlay. These ways may have been tried before, or tried in different cities, or they may be new and untried.

Bondholders need not participate directly in any goal-achieving projects. Their role could be one of financing such projects on the strength of the redemption value of their bonds, or on the strength of any increase in the value of their bonds. Their motivation would arise from the anticipated supernormal profit arising from early redemption of the bonds.

Innovation, diversity and quick to respond

Social Policy Bonds would encourage investigation of new activities. They would encourage investigation of local circumstances, on the basis that doing so could lead to more efficient ways of achieving targeted outcomes than a uniform approach. This is where they might show advantages over the ways in which governments typically do things. Most governments opt for a uniform approach. But the solutions to many social and environmental problems are *not* always known in advance, and the optimal choice is seldom a one-size fits all, top-down, government-dictated policy. More often, it is a matter for investigation and experimentation, and a wide variety of approaches is essential. Bondholders might find, after a bit of experimenting with different approaches, that certain activities work better than others under certain conditions. They would have incentives to take the best of these approaches, and apply them where their return would be greatest, and they would recognise that, for certain objectives, a mosaic of diverse activities would be most efficient.

Government has real difficulties in investigating new approaches in its social and environmental programmes. This is partly because government is generally more interested in preventing failure than in rewarding success. In many areas of social and environmental policy it believes it should carry out only those activities that it can plausibly justify on the basis of a past record. These need not be very efficient, or even partly efficient. As far as many government bodies are concerned they need only to have been tried in the past and not to have been publicly identified as disastrous. This is not a strategy designed to optimise performance; rather it is a strategy that minimises the perceived risk of failure. It leads to the continuing of inefficient, unimaginative activities, whose main recommendation is that they have been done before. As the persistence of social problems attests, these activities are not always very successful.

Neither can government readily try different ideas in different regions, partly because then it would have to face criticism from people who had experienced the less successful ideas. So government generally adopts a uniform approach. In some policy areas, such as education or the environment, it is too easy for central government to override the wishes of local authorities, while local authorities themselves are tempted to override the policies of, say, individual schools when it comes to educational matters. But smaller policymaking bodies, be they local authorities or individual schools often want to employ diverse approaches, and these approaches might well be optimally efficient *in the local circumstances* at achieving desired outcomes.

In one area, female illiteracy might be a very obvious and direct result of poverty. In another part of the same country it might be a result of the particular views of a powerful cleric. Different causes of the same problem demand different approaches, but a uniform regime, of the sort often imposed by central planners in government, often fails to take these differences into account. In a Social Policy Bond regime, bondholders would put maximisation of their return per unit outlay, which would encourage diverse solutions as long as doing so will maximise such returns.

Where objectives are stable, but the manner of achieving them is likely to change over time, a bond regime might be preferable to conventional approaches. Female Literacy Bonds, for instance, could encourage special TV broadcasts, rather than rely on the traditional solutions usually favoured by central planners.

How much should you spend?

Unlike spread betting, your maximum cost outgoings would be capped. The most you would have to spend would be the total number of bonds issued multiplied by their redemption value, minus any revenues gained on floating the bonds. You would also have to pay administration costs, including the costs of verifying whether or not your chosen target has been achieved.

After issuing Social Policy Bonds you can look forward to a quick return in the form of the revenue gained from selling the bonds. If your bonds sold for virtually nothing, then your maximum outlay would be the total cost of redeeming the bonds. This could happen if the market believes the objective you have targeted to be so remote that there is no realistic chance of their being redeemed at any time in the future.

Competitive bidding for the bonds would mean that you can be sure of getting the very best return for your outlay. The way markets work means that you would have to decide only on the *maximum* amount you are prepared to spend on achieving your chosen objective. You would not have to estimate the *likely* cost with any accuracy.

To see this, first assume, for simplicity's sake, that the funds you allocate for bond redemption are the sole funds going into the solution of a particular social problem, and that the actual cost of solving that problem are less than the sum you are prepared to donate.

Say you issue bonds in pursuit of a 50 per cent reduction in the road accident rate in your area, and that you issue ten thousand bonds of redemption value £10 each. If the market decided that the issue value of these bonds were £1.00, the net cost to you of achieving your chosen objective (ignoring administration costs) would be £90 000. In other words, the market at the time of issue believes that the cost, including its profit margin, of achieving the objective would be £90 000.

But suppose you are feeling extravagant - or careless - and you decide to issue not 10 000 but 100 000 bonds, each with the same redemption value, £10. You would then be liable for a maximum cost of £1 million. However, the market's assessment of how much it would cost to halve the road accident rate would not change. The market would still reckon that it could achieve your objective for around £90 000. So instead of valuing each bond at £1.00 it would bid up the issue price of the bonds to around £9.10. (Social Policy Bonds would be an unusual financial instrument, in that the more that were issued, the higher would be their value!) *You, the issuer, therefore would not have to estimate with any accuracy how much a targeted objective might cost to achieve*. If you issue a large number of bonds, you will simply receive a higher price for each bond that you issue. And you could put a cap on your total liability by limiting the number of bonds issued.

But what if the sum you donate to the redemption fund is so small that it the market thinks it will not help to achieve a very remote objective? In that case, the price you receive for the bonds on issue will be close to zero and, unless you either issue more bonds, or external events make the likelihood of the goal being achieved more likely, the bond price will stay near zero. If you, or the public, contribute more to the bonds' redemption fund, then the potential profitability of achieving the goal will rise, and so will the bonds' value - but you, as the issuer, will not gain any financial reward from that. But the chances of the goal being achieved will rise.

In essence, the Social Policy Bond mechanism ensures that potential investors in the bonds the market will decide roughly how much it would cost to reach a specified social outcome. They would do this when they bid for the bonds at issue and at all times afterwards. (This fact, and the would-be bondholders' incentive to minimise their costs, contrast with the current system in which the costs of achieving particular outcomes, if they are calculated at all, are not widely known, nor subject to competitive bidding. That is why issuing bonds could be a better idea than offering a contract to achieve a specified outcome by tender. Under the current system, in fact, many of the people involved in achieving social goals have every incentive to inflate the projected cost of their doing so.)

If your escrow account is open, in the sense that your initial contributions could be swelled by public donations, then the sum for which each bond is redeemed could continue to be a fixed sum or it could be a share of the total fund once the time for redemption has been reached. If the former, a bigger escrow account would mean issuing more bonds at auction. If the latter, the same number of bonds in circulation would each be worth the same share of a larger redemption fund, so the bonds would rise in market value.

Practical aspects of a bond regime

Trading the bonds

Social Policy Bonds, once issued and sold, must be readily tradeable at any time until redemption. The operation of such a 'secondary market' would be critical to the way Social Policy Bonds work. Many bond purchasers would want or need to sell their bonds before redemption — which might be a long time in the future. With a secondary market, these holders would be able to realise any capital appreciation experienced by their holdings of Social Policy Bonds whenever they chose to do so. This would make the bonds a more attractive investment in the first place.

Such capital appreciation would arise from upward movements in the market price of the bonds. Of course, these prices could move in either direction. Major determinants of the bond price would be:

- how remote the market believes the targeted objective is to being achieved;
- market perceptions of risk and uncertainty; and
- the relative attractiveness of other investments.

These and other determinants would vary with time. Note that the market's valuation of the bonds would be influenced not only by efforts that bondholders make toward achieving the targeted goal, but by external factors.

Like shares and other financial instruments, the prices of Social Policy Bonds would be in constant flux. New information affecting the prices would become available day by day. As well as external influences on the bond prices, people would carry out research aimed at determining the value of the bonds as an investment. The effects of all these data on the bonds' market value would give useful insights into the relationships between circumstances, events, social problems and desired outcomes (see *Annex*).

Giving bondholders the chance to benefit from these price movements would be essential. Apart from making the bonds more attractive at issue, a healthy secondary market would be important for another crucial reason: some investors may be able to speed up only one, or a few, of the processes necessary for the targeted objective to be achieved. Once these investors had made their contribution and seen the capital value of their bonds increase in line with the increased probability of the bonds' early redemption, they might have no wish to speculate on the speed at which the remaining processes would be carried out. Other groups of active investors, who could have greater expertise in performing these later processes, must be given an incentive to use their expertise to accelerate attainment of the targeted objective. The possible capital appreciation of bonds bought from previous owners and sold at a still higher price (or redeemed) would provide this incentive. The new owners would, if they were successful in these later stages, realise this capital appreciation.

As the bonds were traded, they would tend to flow towards those who were most able to help solve the targeted social problem. In fact, though, trading in the bonds would not always have to occur. Large bondholders might simply decide to subcontract out the required work to many different agents, while they themselves would hold the bonds from issue to redemption. The important point is that the bond mechanism would ensure that the people who allocate the finance had an incentive to do so efficiently and to reward successful outcomes, rather than merely to pay people for undertaking an activity. At the limit just one single investor could buy all the bonds. If this buyer were determined to hold on to the bonds until redemption, then the bonds would function as a sort of performance-related contract, with the issuers paying only when the objective had been achieved. The buyer could contract out most, or all, of the work required to achieve the objective, with the incentives generated by the bonds for speedy accomplishment cascading down from the bondholder to those subcontracted to do the work. If this bondholder, for whatever reason, were to become inefficient in pursuit of that objective, or were simply to lose interest in it, then he or she could simply sell the bonds to more efficient and more highly motivated investors.

Too many small bondholders?

Too large a number of small bondholders would probably do little to help solve targeted social problems by themselves. If there were many small holders, it is likely that the value of their bonds would fall until there were aggregation of holdings by people or institutions large enough to initiate effective problem-solving projects. In much the same way as share privatisation issues the world over have turned out, the bonds might end up mainly in the hands of large holders, be they individuals or institutions. Between them, these large holders could account for the majority of bond holding. Even these bodies might not be big enough, on their own, to achieve much without the cooperation of other bondholders. They might also resist initiating projects until they could be sure that other holders would not be 'free riders'. So there would be a powerful incentive for all bondholders to *cooperate with each other* to help solve the targeted problem. They would share the same interest in seeing targeted objectives achieved quickly. So they would share information, trade bonds with each other and collaborate on objective-achieving projects. They would also set up payment systems to ensure that people, bondholders or not, were mobilised to help achieve targeted objectives. Bondholders would either trade bonds, or make incentive payments to ensure that any proceeds from higher bond prices, or from redemption, would be channelled in ways most likely to stimulate speedy achievement of the targeted objective. Large bondholders, in cooperation with each other, would be able to set up such systems cost-effectively.

Regardless of who actually owned the bonds, aggregation of holdings and the cooperation of large bondholders would ensure that people who help achieve social goals were rewarded in ways that maximise their efficiency.

What about free riders?

Many people might purchase Social Policy Bonds with the idea of doing nothing but holding on to them until they could sell them at a profit. Such passive investors would have no intention of doing anything to help achieve the social goal targeted by their bonds. Some of them could be casual purchasers who would buy the bonds with the same intent as they would a lottery ticket. They would hope to hold bonds until their redemption, or until their market value had risen sufficiently high for them to enjoy a worthwhile capital gain. Other passive investors might be speculators who thought that the likelihood of the targeted objective being achieved quickly were greater than the rest of the market believed it to be — in other words, that the bonds were underpriced.

Casual purchasers and speculators would want to become 'free riders', hoping to benefit from any increase in the bond price without actually participating in any objective-achieving projects. They would not do much to help achieve targeted goals. However, markets for the bonds would work to limit the benefits from these people's passive investing. To see this, assume that most of a particular issue of bonds were held by would-be free riders. Then very little, if anything, would be done to help achieve the targeted objective. As the objective became more remote, the value of all the bonds would fall. And as the bonds lost value, they would make a more attractive purchase for people who *were* prepared actively to help achieve the targeted objective. So free riders would be tempted to sell, even at a loss, rather than see the value of their bonds continue to fall. Some history of falling bond prices would tend to make free riding on Social Policy Bonds less appealing with future issues. Free riding then would become a self-cancelling activity. There are other reasons why bondholding would be unattractive to potential free riders:

- Individual free riders would have no incentive to collude with other free riders, because the more they did so, the more remote the targeted objective would become, and the further would the value of their bonds fall. This would act so as to limit any free riding activity to small players.

- As with other financial instruments, small players would have to pay higher transaction costs than the bigger institutions — the ones that would be most likely to initiate objective-achieving projects.

- Small players also would not have access to the research that would enable big players to value the bonds accurately. Therefore they would be at a disadvantage in the market.

Note also that even if free riders were to gain from holding Social Policy Bonds, they would be doing so only because their bonds had risen in value as a result of a targeted objective becoming closer to being achievement. As well, attempted free riding would have positive effects: it would add liquidity to the bond market.

In short, there are grounds to believe that free riding would not seriously undermine the operation of a Social Policy Bond regime, mainly because it is unlikely much free riding would occur, and partly because even if it did occur, it would not impede the operation of the bond mechanism.

Defining the redemption terms

You should take some care in defining the terms under which the bonds you issue shall be redeemed. The bonds would work by generating financial incentives for people to achieve particular goals. Unfortunately some people might try to fulfil the objective by complying with the letter of your redemption terms, rather than the spirit. Say you issue Female Literacy Bonds that, as in our example above, target the literacy rate of girls and young women in Pakistan. Relying on the Pakistani Government's literacy tests might be unwise: that Government or its agents could come under great pressure from unscrupulous bondholders to falsify the results of any reading tests that would determine whether a literacy objective had been reached or not. The solution? You would stipulate, as a condition for redeeming the bonds, that any reading tests would be undertaken by an impartial body. You might be able to identify a trustworthy agency already carrying out robust literacy tests that with a little customising could fit your criteria. You should to be convinced that its test results will be an accurate indicator of female literacy in Pakistan.

The bonds might also induce people to modify behaviour in ways that, while not illegal, would undermine what they were trying to achieve. To take the female literacy example again: bondholders might decide that the most effective way of raising female literacy would be to persuade all schoolteachers in Pakistan to drop the teaching of, say, arithmetic, to females and spend all their time on reading programmes instead. You might think this would be a worthwhile trade-off, but what if bondholders instead convinced charities to stop distributing food aid, or family planning supplies, in order to teach literacy? You could anticipate this by incorporating provisos in the bonds' redemption conditions. For example, you could stipulate that Female Literacy Bonds shall not be redeemed if real spending on other teaching projects fell below 100 per cent of the levels prevailing when the bonds are issued.

If higher levels of literacy were targeted, bondholders may be tempted to lobby in favour of easier reading tests. Again by judicious specification of the targeted objective could forestall the problem: the bonds could stipulate the exact reading test to be used, or that the test would have to be certified as appropriate by a specified panel of impartial literacy experts.

Social Policy Bonds and government

Government has the power to pass laws that would affect bond prices, or its actions could influence bond prices in other ways. They could come under great pressure from holders of Female Literacy Bonds, for instance, to increase spending on literacy programmes. Bondholders would lobby for such changes and they would benefit in obvious pecuniary ways if they were successful. But this would be no bad thing: the source of the pressure, and the motivation for it, would be easy to identify, and anyway lobbying is a legitimate activity. There is no reason why bondholders, in common with other pressure groups, should not lobby politicians. They might of course be doing so mainly out of financial self-interest. But existing pressure groups are also self-interested and, provided your objective is responsibly chosen, bondholders' self-interest will be channelled into valuable social benefit.

Such lobbying, of course, already goes on because governments are always making decisions that create winners and losers. People become wealthy by exerting influence on politicians under the current system, but they and their effects on behaviour are not always identifiable. As now, under a Social Policy Bond regime it would be up to politicians to weigh the evidence for and against any course of action promoted by lobbyists, with due regard to the lobbyists' motivation. The sources of this sort of pressure, and the motivation for it, would be more transparent than under the current system so bondholders' lobbying need not pose any significant problems.

When they assess the value of the bonds, potential investors would take into account possible changes in legislation and their potential influence on the speed at which the targeted objective could be achieved.. And it would be up to potential investors in Social Policy Bonds to take into account likely or possible changes in the legislative environment when bidding for the bonds.

There would be, and need be, nothing to prevent government agencies, as competitive suppliers of objective-achieving services, from buying any bonds that you issue, and participating as active investors in the bonds.

You should check the legal status of Social Policy Bonds before you issue (or purchase) them. Some US states, for example, might see the bonds as lottery tickets, and their redemption as a lottery prize. This could affect their viability. See http://www.arentfox.com/post/forum/csmessages/427.html.

Social Policy Bonds and existing institutions

Few bodies charged with achieving social goals are currently paid in ways that encourage better performance. Nevertheless some charities or NGOs are the main sources of expertise for solving social problems and some of them are bound to be cost-effective. How would your issuing, say, Female Literacy Bonds, bear on their operations? The decision would be up to bondholders. They might investigate the activities of these bodies and help to finance those that were most cost-effective. Or they might find it more efficient to set up their own bodies devoted to achieving the objective that you are targeting.

If your bonds represented a large potential source of funds, existing organisations might themselves react by reviewing the results of their programmes and projects. Incentives for efficiency would focus their attention on the cost-effectiveness of their operations. If they could convince bondholders of their efficiency they would stand a greater chance of receiving more funding from them.

What happens once an objective has been achieved?

Once your objective is close to being achieved, you could float a new set of Social Policy Bonds aimed at maintaining the achieved outcome or at further improvements. Sustaining the outcome beyond the period specified in the original bond issue would probably be cheaper than achieving it, while further improvements targeted by a second bond issue would most likely cost less, in terms of benefit per unit outlay, than those achieved by the first issue. There are two main reasons for this:

1. Assume that a bond issue aimed at reducing the level of some indicator from *x* led to its reaching a level of *y*. Most probably it would take more than a withdrawal of this funding for the indicator to revert back to *x*. Why? Say the indicator represents the female literacy rate in Pakistan. Once it has risen to 95 per cent and stayed at that level for the time stipulated in your bonds' redemption terms, it is quite possible that some of the activities funded on the strength of the bonds will come to an end. But expectations will have risen, in the general population, and especially amongst those whose literacy has been raised. As well, some investment will have gone into literacy-raising infrastructure and systems that will persist after the initial bonds are redeemed. Similarly, if your chosen target were air pollution, maintaining the lower levels of pollution achieved by your first bond issue could be cheaper than achieving it in the first place. In both cases, people would have invested some know-how, or physical capital that cost less, per unit benefit, to keep running than they did to set up.

2. Less specifically, it is likely that general improvements in productivity, mainly arising from technology (including information technology), will continue to occur in our economies, and that bondholders would make use of them.

Epilogue: status of Social Policy Bonds

The Social Policy Bond idea has had an unusual fate for an unusual idea. It has been in the public domain for about 14 years, and it has not been so far been adopted anywhere, to my knowledge. But neither has it been dismissed outright. In April 2002, I presented a paper on the bond concept to joint meeting of the Agriculture and Environment Committees at the Organisation for Economic Cooperation and Development (OECD) in Paris. At the meeting, delegations from most of the OECD's member countries made comments on the paper. These were mostly along the lines of "this is very interesting — but unworkable in practice." But one of the delegates perhaps articulated the deeper feelings of those present when he said "if this gets adopted we'll all be out of jobs!"

Most of the delegates at that meeting were government employees, and government has been more wary of trying out the bond concept than the private sector. Certain private individuals are now proposing to issue their own Social Policy Bonds. They are considering floating bonds for projects as diverse as boosting voter registration, raising literacy in developing countries and developing open-source software. Most of the funds we contribute to the solution of social and environmental problems are allocated by government, but it is the private sector that is mainly concerned about cost-effectiveness.

Annex: Pricing efficiency

Social Policy Bonds, in addition to their other advantages, would mean that the cost to you, the issuer, of achieving the targeted outcome would be minimised and capped. And if bondholders fail to perform funds for the cause will remain in escrow until the objective is achieved. If you stipulated that the objective had to be achieved by a certain date, and if the objective had not been achieved, you could reclaim the funds yourself.

A potentially valuable benefit to policymakers, or NGOs, is that the market for the bonds would be elegantly efficient in conveying information about the cost of achieving objectives and how this cost varies with time and circumstances.

Take, for example, the objective of lowering some index of water pollution in a region from 50 to 40 units. Assume that you and some wealthy, like-minded residents have issued one million bonds targeting water pollution, each redeemable for £10 once this lower level has been attained. The *maximum* cost to you and the rest of the consortium of achieving this objective would then be £10 million. But if the bonds, when issued, fetched £5 each, then the market would be saying that it thought it could achieve this objective for just £5 million. It wouldn't say *when* it thought it could achieve that objective, but that could be inferred from market behaviour and the market value of the bonds compared with other financial indicators. But what if the bonds sold for virtually nothing, and the market value of the bonds failed to move from that floor? That would mean that you had miscalculated: in the market's view there would be no realistic chance of the objective being achieved for an outlay of £10 million in the foreseeable future. You could respond in different ways:

- You could wait for new technology to arrive, or for circumstances to change in other ways, such that the market would see the objective as becoming more easily achievable, and the value of the bonds would consequently rise. Or

- You could issue more bonds, with the same specification, also redeemable for £10. You could do this in stages, gauging the market reaction to each new tranche of bonds, which would tell you the maximum cost of achieving the objective.

Either way, you could be reasonably sure that it would be getting a good deal, expressed as 'reduction in water pollution per unit outlay'. This important benefit is worth spelling out in more detail. Valuing the *benefit* of achieving an environmental outcome is up to you as the issuer of the bonds, and will be largely a subjective process. But minimising the *cost* of your desired outcome is a different matter.

The maximum cost of issuing Social Policy Bonds is easy to determine: it is simply the total number of bonds issued multiplied by the redemption value *plus* administration costs *minus* any revenues gained on floating the bonds. And potential bondholders, as we have seen, would have every incentive to minimise that cost by competing against each other when they bid for the bonds at issue.

But the bond mechanism would not merely minimise the *total* cost to issuers of achieving a specified objective. It would also indicate the *marginal* cost of achieving further improvements, and this could be useful both to issuers (in case you issue further bonds once the initial objective has been achieved) and others with an interest in furthering similar objectives, such as government or NGOs. Say the one million water pollution reduction bonds were to sell for £5 each. This would indicate that the present value of the expected maximum cost, including bondholders' profits, of reducing water pollution from 50 to 40 units would be £5 million. You, or others with an interest in reducing water pollution, might then suppose that you could afford to be more ambitious, and aim for a further fall in pollution to 30 units. You could issue a million additional bonds redeemable when this new lower concentration were reached. These would (probably) have an initial market value of less than £5, reflecting the (probably) diminishing returns involved in lowering water pollution. The point is that, in pricing the bonds, the market would be displaying an informed view of the *marginal* cost of the water pollution targets. So if the bonds targeting the new level of 30 units were to sell for £4 each, then the maximum cost of achieving that objective would be £11 million, being equal to: £5 million (paid out when the level fell from 50 to 40 units) plus £6 million (paid out when the level fell from 40 to 30 units). The marginal cost of a 10-unit drop in water pollution would thus have been revealed to rise from £5 million to £6 million. Would it be worthwhile aiming for a further fall to 20 units? Following such water pollution-targeting bond issues policymakers would have robust information about the cost of doing so.

This is, of course, a simplified example and in fact the bond market would continuously update its pricing information. Say that improvements in technology, of the sort that might be stimulated by a large initial water pollution targeting bond issue, made it much cheaper for farmers to reduce their water pollution emissions. Bondholders may, for example, have financed successful research into new varieties of grasses that exhibit better uptake of nitrogen fertiliser that would otherwise pollute rivers. How would the market react to such a development? Once the new varieties' effectiveness had been revealed, the value of all the bonds would rise. Instead of being priced at £5 and £4, the two water pollution issues of the example might sell for £8 and £7. The total cost to the issuers of redeeming these bonds would not change: it would remain at £11 million (though redemption would most probably occur earlier). But the market would be generating new information as to the likely cost of future improvements in water quality. The market would now be expecting reductions of 10 units of water pollution to cost £2 million (from 50 to 40 units), and £3 million (from 40 to 30 units). The new grass varieties would have reduced the costs from £5 million and £6 million (respectively). So the cost of any further pollution reductions would also fall, and by following market price movements you, or others interested in reducing water pollution, could gauge approximately by how much.

These figures are hypothetical, but they do indicate the role that markets for Social Policy Bonds could play in helping people decide on their spending priorities. The importance of this sort of market information can hardly be exaggerated. The failure in history of central planning can plausibly be attributed to the absence of market-generated information.[i] Market prices reflect all of the information used by all who transact, or choose not to transact, in the market. Central planning fails in comparison with a market economy because it encounters the limits of human beings' calculating capacity: no individual or group of individual planners knows or feasibly can know all the dispersed information that is embodied in prices. Even with a sound incentive system in place — and the centrally planned economies had some fearsome systems — without the information that only markets can generate the computational task of organising an efficient allocation of resources is too great. Prices incorporate and simplify all of the dispersed information implicit in getting a product or service to the marketplace. Markets for Social Policy Bonds would continually generate and reveal this information to policymakers and all those involved in achieving social and environmental outcomes — probably for the first time on a systematic basis. *A Social Policy Bond regime would combine market information with incentives to use it efficiently: the synergies arising could be of enormous benefit to society as a whole.*

Top

Top of section 2

---

[i] See Hayek, F A, 'The Pretence of Knowledge', in his *New Studies in Philosophy, Politics, Economics and the History of Ideas*, University of Chicago Press, Chicago, 1978.