

Chapter 3

Social Policy Bonds

Social Policy Bonds are a new financial instrument designed to inject the market's incentives and efficiencies into the achievement of social and environmental goals. The bonds could be issued and backed by a wide range of bodies, including local or national governments, non-governmental organizations, global bodies such as the United Nations, philanthropists, or ordinary individuals. If issued by government, the bonds' ultimate backers would, of course, be taxpayers. A fixed number of Social Policy Bonds ('bonds') would be issued. They would initially be auctioned to the highest bidders. The bonds' backers would undertake to redeem these bonds for a fixed sum *only when a specified social objective has been achieved*. It is this promise, and the market's view of how easily the objective can be achieved, that would underpin the bonds' market value. The bonds would not bear interest. They would be freely tradeable after issue, and their market value would rise and fall. Social Policy Bonds would therefore differ from conventional bonds in that they would have an uncertain redemption date which, in combination with a fixed redemption value, implies an uncertain yield: holders would raise their bonds' yield by achieving the targeted objective quickly. Once the targeted outcome had been achieved, whoever backed the bonds would redeem them. The rest of this chapter outlines the essential elements of a bond regime. Subsequent chapters look in more detail at their operational aspects.

Social Policy Bonds work by creating an interest group — bondholders — who have a strong interest in achieving the targeted social objective efficiently and quickly, or in paying others to do so. Assume, for example, that an urban authority is prepared to spend a maximum of say \$10 million to reduce the crime rate within its borders by 50 percent. It issues one million bonds that become worth \$10 when the crime rate falls below 50 percent of current levels for a sustained period — say one year. 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 percent if a sustained halving of the crime rate were achieved. This provides the motivation for bondholders to do what they can to reduce the crime rate. Because the bonds are tradable at any time, the identity of bondholders in pursuit of the objective could be constantly changing. This group, essentially a coalition of interests in favour of achieving the targeted outcome, will at any time be made up of those who believe they can achieve the outcome most cost-effectively.

Social Policy Bonds could in principle, be used to solve any social or environmental problem that can be reliably defined and quantified. Key criteria for policy areas within which Social Policy Bonds would show the most marked improvement over current programmes are:

1. Existing policies have objectives that are unstated, obscure, uncoded or conflicting;
2. Problems are large-scale and have multiple causes, many of which cannot be readily identified;
3. Current approaches are either ineffectual or inefficient; and

4. Financial rewards to those involved in achieving objectives are currently uncorrelated to their effectiveness or efficiency in doing so.

Sadly, there are many such policy areas at all levels. At the regional and national levels they include crime, employment, health, education, and air pollution. At the global level they include the possibility of catastrophic climate change and nuclear conflict.

Markets minimise costs

Issuers of Social Policy Bonds would need to give some thought to how much achieving their objective is worth. They could estimate the maximum value of the targeted outcome. One consideration would be the financial impact of solving a social problem. Achieving certain social goals would actually bring about financial savings. A national government, for instance, could make a net financial gain by bringing down the number of unemployed claimants: the state saves unemployment benefit and gains an increase in income tax for each person who leaves the unemployment register and goes into gainful employment. Achieving this particular social goal could therefore generate a net fiscal benefit, even in the short term. For other targeted objectives, such as a lower crime rate, there could also be positive, but less easily quantifiable, net financial benefits, and these may take longer to arise. Other social goals, such as reduced rates of homelessness, or increases in literacy, might increase monetary returns to the government in the long term, but would generate very little net revenue in the short run. And there would be many social or environmental goals whose achievement would impose net financial costs on society in the foreseeable future.

But people and their governments want things other than for financial reasons. A society in which everybody can read, in which people feel safer from crime and breathe cleaner air is surely desirable in its own right. Bond issuers, whether government or private sector, have to decide on how far they will pursue these objectives, and how valuable they are. They would have to take into account the financial and nonfinancial benefits in deciding on the maximum value of each social goal, in advance of issuing the bonds. A bond regime would make this a simpler and more transparent task than the current array of social policies, because people would be asked to value outcomes, rather than activities intended to achieve these outcomes. If government were to use Social Policy Bonds in conjunction with other policy instruments to achieve the same goal, government would also have to decide on the proportion of total expenditure that would be spent on the bond component.

All these factors would determine how many bonds would be issued. The maximum cost to backers of a bond issue would (ignoring administration costs) equal: the total number of bonds issued multiplied by their redemption value, minus any revenues gained on floating the bonds.

Though they would have to decide on the *maximum* amount they want to spend on achieving their objective, those who issue Social Policy Bonds would not have to work out how much the *actual* cost would be with any accuracy. That would be done by bidders for the bonds in the open market. Assume again that bonds were to be used exclusively in pursuit of a 50 percent reduction in the crime rate, and that the urban authority issues one million bonds, of redemption value \$10.00. If the market valued these bonds on flotation at \$1.00 each, the net cost to the issuers of achieving the targeted objective (ignoring administration costs) would be \$9 million. In other words, the market at the time of issue believes that the cost of achieving

the objective, including its profit margin and after taking into account risk, would be \$9 million.

Now suppose the bond issuers are completely in the dark about how much it will cost to achieve the targeted objective and instead of issuing one million bonds they issue ten million with the same redemption value, \$10.00. They would then be liable for a maximum cost of \$100 million. However, the market would still reckon that it could achieve the objective for around \$9 million. So instead of valuing the bonds at \$1.00 competition between potential investors 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!) *The issuers therefore would not have to estimate with any accuracy how much a targeted objective might cost to achieve, and they would put a cap on their total liability by limiting the number of bonds issued.*

So the Social Policy Bond mechanism ensures that the market, which means people *other than the bond issuers*, would decide roughly how much it costs 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. Under the current system, in fact, many of the people charged with achieving social goals (or, more likely, with supplying certain inputs) have every incentive to inflate the projected cost of their doing so. Under a bond regime, however, the formidable information-processing power of the market would be channelled into minimising the costs of achieving these goals.

Note that the issuing body could add to the number of bonds in circulation after floating at any time, if it wanted to boost the efforts going into achieving a particular social goal. If it wanted, for whatever reason, to *reduce* such efforts, the situation would be a little more complicated. It could buy bonds back from holders, but doing so would reduce the total funds to be spent on achieving the targeted objective, and so would lower the value of all bonds in circulation. People might therefore be unwilling to buy bonds in the first place if they thought there were a high probability of the issuing body's buying some of them back in this way. They would demand some sort of premium for taking that risk. Alternatively, the issuing body could undertake either that it would never buy Social Policy Bonds back or that, if it did, it would pay the market price ruling before it announced its purchase intentions.

What would bondholders do?

Social Policy Bonds targeting the crime rate would rely on the people or institutions that hold bonds initiating or facilitating crime-reduction programmes. Bondholders could invest their own capital, or borrow on the strength of the redemption value of their bonds. They would have an incentive to cooperate with each other to help reduce crime, and to do so as cost-effectively as possible. These people's motivation would come from the expected capital gain they would enjoy as the bond price rises in line with the enhanced probability that the objective will be achieved early.

Consider some of the measures that active investors in bonds targeting the crime rate could put into operation:

- encouraging neighbourhood watch schemes;
- encouraging parents to monitor their children's activities more closely;
- subsidising recruitment of unemployed workers; or
- complementing police patrols with private security patrols.

In many countries, some arms of government already undertake one or more of these activities. And some longer-term projects, like research into the causes of crime, are done by private bodies or universities, independently of government or with only a small contribution from government funds. The crucial difference a Social Policy Bond regime would make is that people would have incentives to seek out and develop those ways of reducing crime that are most cost-effective. A police force, a bureaucracy, or an environmental health department, however well-intentioned, is not currently rewarded in ways that correlate with its success in achieving society's objectives — even if these are explicitly targeted. But under a Social Policy Bond regime the self-interest of bondholders would act so as to encourage those ways of reducing crime that would give whoever funds the bond issue the best return for their outlay. These ways may have been tried before, or tried in different cities, or *they may be new and untried*. Bondholders would be motivated to seek out, invent and use the most efficient methods for the city or country whose crime rate is targeted.

Bondholders need not participate directly in any crime reduction projects. Their role could be one of financing such projects, on the strength of the increase in of their bonds, whether they plan to sell them quickly, or to hold on to them until redemption. They would be motivated by the anticipated supernormal profit arising from early redemption of the bonds.

One activity that bondholders might indulge in is lobbying government. So, in the crime example, they might press for longer prison sentences, thinking that these would deter potential criminals or keep convicted criminals out of circulation. Such lobbying, of course, already goes on because government is always making decisions that create winners and losers. Under a Social Policy Bond regime the source of this sort of pressure, and the motivation for it, would be more transparent than under the current system and it need not pose any different problems. (The next chapter looks at the subject of lobbying in more detail.)

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 from 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. Some of these could be apparent at the time of issue: for instance, one of the determinants of crime is demography. Specifically, the greater the number of young male adults, the larger the number of crimes tends to be, all other factors being equal. Demographic variables like this, and others that can be anticipated, would help determine the market price of the bonds at the time of issue. But other influences cannot be anticipated. So, for example, the market price of bonds targeting property crime could fall if, say, there were a string of power failures that led to looting. Or it could rise on the capture of a ringleader of a particularly successful gang of burglars or car thieves. The price of bonds targeting air pollution could rise or fall with climatic conditions, volcanic eruptions, or the price of oil or coal. The value of bonds targeting unemployment could rise or fall with financial data, such as the exchange rate (making the country a more or less attractive venue for overseas investment), or interest rates (making firms more or less likely to lay off employees).

As with other investments, risk and uncertainty would be important determinants of the bonds' market price. Bonds targeting more remote objectives (cutting crime by 80 percent say) would be riskier than those whose outcomes were closer to current levels (cutting crime by 20 percent). And there would also be uncertainty attached to the Social Policy Bond mechanism itself, especially in the early years of a bond regime, as it would be untried and unproven.

As with shares and other tradable 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. All this information would generate extremely useful insights into the relationships between circumstances, events, social problems and desired outcomes.

Giving investors in the bonds the chance to benefit from price movements is 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.

Cascading incentives

As the bonds are traded, they will tend to flow towards those who are most able to help solve the targeted social problem. In fact, though, trading of 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 could hold the bonds from issue to redemption. The

key point is that the bond mechanism would ensure that the people who allocate the finance have an incentive to do so efficiently and to reward successful outcomes, rather than merely to pay people for undertaking activities. 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 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' (see 'The Free Rider Question' in the next chapter). So there would be a powerful incentive for all bondholders, tacitly or overtly, 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 could share information, trade bonds with each other or collaborate on objective-achieving projects. They could also set up payment systems to ensure that people, bondholders or not, were mobilised to help achieve targeted objectives. This might mean that bondholders pay people not according to how much they actually help contribute to an outcome – which can be difficult to determine – but according to how much bondholders *estimate* they are contributing to the outcome, or to more measurable variables as spending or outputs. But while there might on those occasions be no direct link between payment and efficiency in achieving the overall outcome, bondholders would have strong incentives to strengthen that linkage where it is worthwhile to do so. If they failed to do that they would not maximise their own rewards from holding Social Policy Bonds.

To summarise then: 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 owns the bonds, aggregation of holdings and the cooperation of large bondholders would ensure that people who help achieve social goals are rewarded in ways that maximise their efficiency.

Objectives and indicators

For a Social Policy Bond regime to be effective, clarity and transparency of objectives are essential. The targeted objective must be *carefully defined* so that targeted changes either actually are, or are strongly and inextricably correlated with, what society wants to achieve. For instance, numbers of reported crimes could be targeted if the objective were to achieve a

safer urban environment. But this indicator may be unsatisfactory if, for instance, the crime rate became so high that people did not bother to report minor assaults or burglaries to the police. A more appropriate indicator might be derived from responses to victim surveys. Because the bonds target outcomes they demand clear thinking and transparency as to exactly what it is that society is aiming to achieve. Is reduced crime itself the real objective? Or reduced fear of crime? Or some combination of the two? Or, looking at employment: is lower unemployment our real objective? Or lower expenditure on unemployment benefit? Or higher employment? Is it worthwhile aiming to reduce, in particular, unemployment amongst 16–24 year old males? Or ethnic minorities? Or the unemployment rate in particular regions? Note that it would clearly be unsatisfactory to redeem the bonds the moment a targeted fall in unemployment has been achieved. The objective should be a *sustained* lower level of unemployment, and that is how it would have to be defined when the bonds are issued. Social Policy Bonds would give such clarity a higher priority than it enjoys under the current policymaking regime.

Targeted objectives should also be capable of being *targeted by quantifiable indicators*, whose progress accurately corresponds with progress toward the desired social outcome. As well, objectives should, in general, be as *broad* as possible, so that one particular objective cannot be achieved at the expense of other societal goals.

The last point needs elaboration. Consider the application of a bond regime to environmental problems. Assume the concentration of atmospheric lead were to be targeted by a bond issue. It might be that targeting lead in this way would cause people to increase their use of other polluting substitutes — and these could be at least as dangerous as the original levels of lead. One way of dealing with this problem could be to aim initially at an unambitious reduction in the lead level. Depending on the effects of such a reduction on the use of offending substitutes, other bonds could then be issued, either targeting the level of lead, or targeting the level of offending substitutes. But a better approach would be to target, more comprehensively, atmospheric pollution. This could be expressed perhaps as an index of atmospheric pollutants, weighted according to their lethality and other factors (see box, *What to target?*).

What to target?

Breadth of objective Social Policy Bonds lend themselves to targeting combinations of objectives. It would probably be unsatisfactory to target, say, atmospheric lead as the sole target of a bond issue targeting air pollution if it were likely that polluters would respond increasing emissions of other toxic chemicals. Instead, all atmospheric pollutants could be made the target of a single bond issue: the bonds could target an index encompassing all pollutants, weighted according to their lethality. Targeted objectives should, in principle, be as large-scale and broad as possible

Ends or means? In principle ends, rather than means to ends, would make better targets for Social Policy Bonds. Thus, it would be preferable for bond issuers to target, for example, homelessness, rather than housing starts, and leave it for bondholders to decide on how best to achieve that targeted goal. Similarly, it might be preferable to target not air pollution, but such indicators of environmental status as human, animal and plant health, perhaps in conjunction with more subjective

indicators like the views people have about the quality of their environment, as measured by questionnaire responses. Bonds could be issued whose redemption value is on a sliding scale, reflecting the perceived adverse environmental impacts of the targeted range of pollutants.

Spatial distribution Bonds aimed at improving national averages of such indicators as pollution would be adequate sole policy instruments only if society were unconcerned about the distribution of pollutants. Otherwise bonds targeting pollution could be made redeemable only on the condition that pollutant thresholds shall not be breached in any part of the country concerned.

Redemption structure: time Bond issues could provide bonus payments for achievement of the targeted goal by a specified date. Or issuers 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 potential penalties or bonuses into the bond price.

Redemption structure: partial fulfilment Rather than simply pay a lump sum upon achievement of a fixed goal, bonds could be issued that pay rewards for partial achievement of a long-term goal.

Similar concerns, perhaps less clear-cut, could arise when targeting regional problems. If bonds were issued targeting the number of unemployed people of working age in northeast England, say, then bondholders might attempt to solve the problem by paying the unemployed of that region to move somewhere else. This might, of course, be seen as a social benefit. But if not, provisos could be written into the bond issue, such that they would not be redeemed if the population in the north-east fell below a certain level, or if the unemployed population in other regions rose above a certain level. In general, objectives that are complementary and that, if not pursued jointly, could conflict, should be targeted by a single bond issue.

Successful targeting: how not to do it

Even under the current system, policymaking by national and supranational governments is inevitably, though often only implicitly, aimed at improving some highly aggregated targets and indicators. Policymakers can identify problems in a society of more than a few thousand mainly by quantifying them: they have to use numbers to monitor progress toward solving these problems. There are perils in doing so, not all of them obvious. You would think, for instance, that reduced infant mortality correlates strongly with an improvement of society's well-being. But the bald figures may not be all they seem:

In a tragic sort of way, inferior prenatal care could actually boost average life expectancy while lowering health care costs. Adequate prenatal care may reduce the incidence of miscarriage, especially in the second half of pregnancy. Had my wife's perinatologist not detected her dilating cervix in the 22nd week of pregnancy, we would probably have lost our daughter. And she would have been a miscarriage statistic, not an infant mortality statistic.¹

Or, to take another example, it was found that British five-year cancer survival rates are lower than in other countries. Can one conclude that the UK's National Health Service is performing poorly and that physical well-being would be improved if only survival rates could be improved (as one journal did²)? Not necessarily: as one commenter pointed out, an alternative explanation is that Britain does not waste money on extensive testing procedures for diseases it cannot cure.³ For such diseases, and many cancers are among them, earlier diagnosis merely serves to raise the number of years between the identification of the disease and death—it does not affect expected mortality.

The inadequacies of such narrow targets bring to mind Goodhart's law, named after a former chief economist of the Bank of England, which says that whatever is adopted as a target ceases to be a relevant target once it has been adopted. It is essential to target indicators that are broad enough to be inextricably linked to well-being. In the current, institution-led target-picking environment, it is too easy to be misled by figures taken out of context. The examples given above, and targets such as A&E throughput rates (see chapter 1: 'Gesture politics') are irrelevant to the health of the population. Of course, defining broader health indicators and measuring them is more difficult than measuring the length of stays in A&E departments, but with sufficient ingenuity - of the sort that is currently applied to gaming the system - it can be done. To put it bluntly: if a government is concerned about the health of the population, then it should target the health of the population.

What we could do

Perhaps more compelling as targets than indicators averaged over a large population are indicators of more basic levels of social welfare. It does seem to be mainly at lower standards of living that the correlation between a quantifiable indicator and social welfare is strong and therefore valid as a guide to policymakers. Indeed it is perhaps government's historical successes at the basic level that have led it to apply the same policy approaches into the middle and upper range of well-being, where they will not be so effective. At higher levels of income, wealth or standard of living, it does appear that the correlation between well-being and objective criteria breaks down, or at least that extra resources generate diminishing returns.^{4 5} But it is not just from the point of view of efficiency (improvement in well-being per dollar) that the case for government intervention to help the poor should be made. The stronger case is that it is the poor or otherwise disadvantaged who are most in need of government intervention. It would be difficult to argue against the targeting of, say, improved basic levels of education and health in the industrialised countries, or broad indicators of poverty in the developing countries.

The Social Policy Bond concept could also be applied to the supply of public goods and services, such as reduced crime rates and a cleaner environment. Moreover, the bonds are versatile in that they can reward the absence, within a specified time period, of negative outcomes such as nuclear conflict, other manmade disasters, natural catastrophes or the damage they cause. For instance, a global body under the auspices of the United Nations could issue Social Policy Bonds ('Nuclear Peace Bonds') that would be redeemed only if there were fewer than 1000 deaths resulting from explosions of nuclear devices before 2050.

To simplify discussion this book looks at Social Policy Bonds with a straightforward payment structure; specifically an all-or-nothing objective, upon achievement of which a fixed sum is paid. But in practice, bonds could be issued that rewarded partial achievement of a long-term

goal, or gave extra rewards for early achievement of goal. A perfectly competitive market would factor the risk-adjusted possibility of such payments into the flotation price making them conceptually equivalent to the all-or-nothing bonds discussed here, but there may be practical advantages in issuing these or other variants of the basic Social Policy Bond concept.

¹ *Prenatal and Neonatal Care*, Andrew Marx, 'Sensible Knave', an online blog, 2 October 2005; <http://sensibleknave.blogspot.com/2005/10/prenatal-and-neonatal-care.html>, sighted 17 July 2008.

² *The NHS's Cancer Failure*, 'Prospect', February 2006; http://www.prospect-magazine.co.uk/article_details.php?id=7307, sighted 17 July 2008.

³ *NHS Survival Rates*, Peter Sugarman (written 14 February 2006), letter published in 'Prospect' April 2006; http://www.prospect-magazine.co.uk/article_details.php?id=7381, sighted 17 July 2008.

⁴ See *Happiness is a warm vote*, 'The Economist', 17 April 1999.

⁵ "The point where more wealth ceases to imply more happiness is around \$10,000 per capita annually", says Robert Wright, summarising work in this area in *Will globalization make you happy?*, 'Foreign Policy', September/October 2000. Available online at http://www.foreignpolicy.com/story/cms.php?story_id=1704.