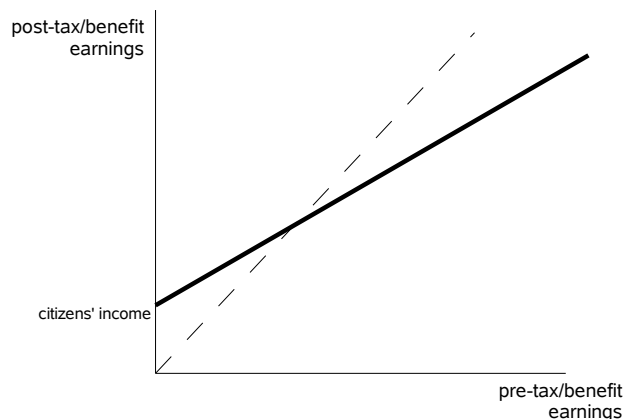



Should the redistributive and insurance functions of the social security system be separated? Is the legitimate function of the social security system to redistribute income over the lifecycle, between individuals, or both?

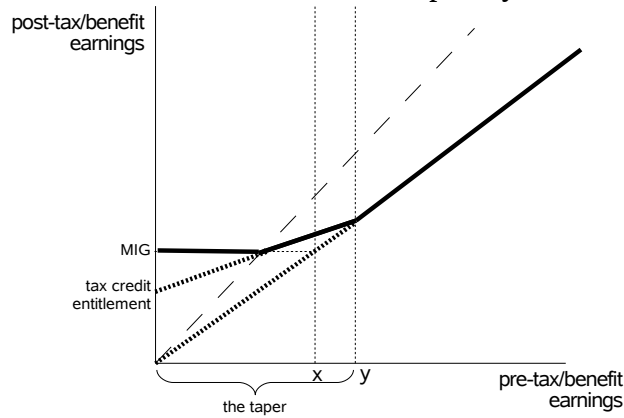
- We should perhaps begin by considering the main institutional arrangements that can be considered to fall within the scope of the question. Although government involvement in the provision of healthcare, education or transport does contribute to "social security" in the literal sense, the Social Security System is usually considered to consist primarily of the various cash benefits or transfer payments. Included in this list would be income support, unemployment benefit, state retirement pensions, disability benefit, incapacity benefit, child benefits of various kinds (although we will deal with these in the second presentation below). We will follow this narrower definition here.
- Before considering how the Social Security System should organize the relationship between its intra-generational and inter-generational redistributive, and insurance functions, it is helpful to summarize the justifications in economic theory for state intervention in each of these 3 areas.
- Intra-generational redistribution - Assuming as a given that pre-tax/benefit earnings from the market are unequally distributed, and that there is broad agreement that a more equal distribution of income is, other things equal, more desirable (this "other things equal" is of course the key bone of contention between different ideological positions), there is a *potential* gain in social welfare (and if all individuals agree on the value of equality, possibly even a Pareto-improvement), from state intervention in the market economy to create a more equal distribution of post-tax/benefit earnings. This immediately raises the question, however, of why *voluntary* redistribution from charitable giving does not result in the optimal level of redistribution without state intervention. A number of reasons can be given:
 - (1) Individuals do not agree about the distribution of income that is fair. (For example, some people might say that people who have completed more education deserve a higher income to compensate them for years of drudgery in education whilst others took the easy way out and dropped out. Others might take the opposite view and see going to university as a consumption choice that was subsidized by the state and therefore should be paid for in a lower income later on. Which view is taken depends upon how education is subjectively valued; does it give a positive or negative utility during consumption?) People who disagree with the consensus position may not voluntarily redistribute enough and so need to be forced.
 - (2) People do not fully take into account the social welfare function in the manner that they should. For example, the s.w.f. may be $W=U(Y_1)+U(Y_2)$ but individual 1 may, for instance, be maximizing $V_1=U(Y_1)+\theta U(Y_2)$ where $\theta < 1$, so that they do not place sufficient valuation upon the utility of the other person.
 - (3) If people care about the *average* distribution of income, then a more equal income distribution can have the properties of a public good. Richer people do not gain utility by giving away their own income to one or more poorer people, because this has a negligible effect on the overall income distribution. On the other hand, if all the rich people including them are simultaneously forced to redistribute, then this will raise their utility because the gain in utility from the more equal income distribution outweighs the loss in utility from the loss in income. This richer people are in a prisoners' dilemma situation exactly analogous to, for example, an environmental externality (e.g. individuals are not willing to cut back their own emissions because the effect on the air quality is negligible but their utility can be increased if they are forced, along with everyone else, to cut back their emissions).
 - (4) It is, of course, likely that all the the above factors contribute to the need for state-enforced redistribution of income.

Assuming that we have an inequality averse policy-maker, the issue is how best to achieve this intra-generational redistribution. We will consider the advantages and disadvantages of a minimum income guarantee versus a citizens income, or negative income tax:



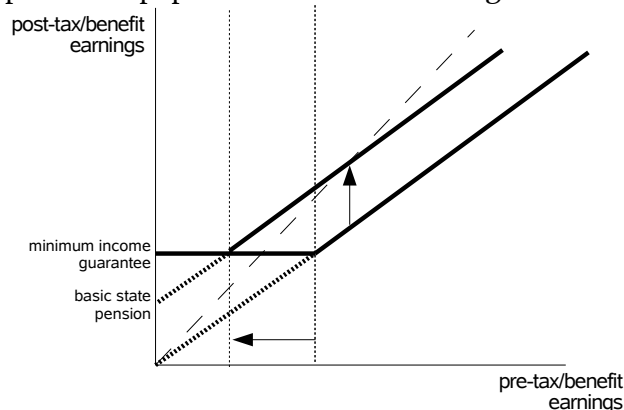
- The main strength of the **minimum income guarantee (MIG)** is that it is extremely well-targeted; assuming, for sake of simplicity, that the state can eliminate fraud, all of the money goes to the neediest (i.e. those below the MIG line). The weakness of such a system is its effect upon incentives. Anyone who is currently earning less than x , the amount of income required to reach the MIG, after tax has been deducted, has no incentive to earn any more at the margin, because anything they earn is taken away in reduced benefits. Using a minimum income guarantee to try to hit a relative poverty target (say 50% of average income) will mean that anyone who would be earning below this level will give up work altogether. This would be very inefficient (because labour resources are being wasted) and probably very expensive for the state. This justifies why income support in the UK is not uprated in line with average earnings. The main aim of current government policy is, by increasing other benefits, to render income support increasingly a safety net of last resort for the very needy.
- A **citizen's income system (CI)** avoids the incentive problem by paying a certain amount in negative income tax to every citizen, and then taxing all earned income at the same rate. This means that even people who earn a very small amount are able to gain post-tax income from extra hours they work. Thus there is no benefit-trap for the poor. The weakness is that a citizen's income is not well-targeted; everyone, including millionaires, would receive it. For a given level of funding, therefore, the achievable citizen's income would be far lower than the achievable MIG, because a great deal of revenue would be "wasted" in payments to the more well-off.
- **Tax credit schemes**, like the child tax credit and pensioner tax credit in the UK, are essentially a hybrid between MIG and CI schemes. Whereas the MIG is withdrawn at 100% (i.e. a £1 increase in post-tax earning results in a £1 decrease in income support) and the CI is withdrawn at 0%, tax credits are withdrawn at somewhere in between these two figures. (For example, the current primary pre-tax withdrawal rate for the UK child tax credit is 37%, meaning that when a family on child tax credit earns £1, 37p is taken in reduced benefit, 22p is taken as income tax and 11p as national insurance contributions, leaving a 30p marginal gain in post-tax income.) By

withdrawing at less than 100%, the taper becomes longer than in the case of the MIG, but the distortion in incentives is less extreme. On the other hand, tax credits are better targeted than the citizens' income (because above pre-tax income y at the end of the taper, no payments are made), and so the tax credit entitlement can be higher, for a given level of funding, than the CI. As the recent UK experience has borne out, their downside is administrative complexity.



- Inter-generational redistribution - The argument for income redistribution does not, by itself, justify an explicitly inter-generational aspect to redistribution in the form of state pensions. The state pension is effectively a citizens' income scheme only available to people over pensionable age. What justifies this? Why not just have a universal MIG, CI or TC scheme available to all? (We will leave aside the issue of support of children for the moment, although it is clear again that the current UK tax credit schemes are significantly more generous to working age poor people with children than those without):

- (1) Economic growth - Later generations will likely be richer than earlier generations. In order to keep pensioners out of relative poverty, it is therefore necessary to redistribute some wealth from current workers to pensioners. This argument could justify having a minimum income guarantee, citizens income or tax credit entitlement which rises in line with average earnings. However, it does not justify why pensioners should have a different kind of scheme (i.e. a citizens income style scheme whilst the working age population only have tax credits and income support).
- (2) Incentive issues - Whereas most pre-tax/benefit income for working age people comes from labour income, most for pensioners comes from savings. The savings/investment system is of vital macroeconomic importance. Having a state pension, even if it is less than the MIG from income support, reduces the disincentive of working age people to save, because less of their savings is withdrawn as reduced benefits when they reach retirement age. In other words, the combination of a basic state pension along with a minimum guarantee reduces the length of the taper along which there is no incentive to save, thus reducing the number of people on the population for whom saving "will not pay":



- (3) Targetting - If pensioners are more likely to be poor than the working age population, then the state pension scheme is better targetted than a universal CI scheme. (The rationale for disability benefits can be similarly justified if disabled people are more likely to be poor.)
- (4) Myopia - If people do not save enough of their own volition due to failure of full rationality, then it may be welfare improving for the state to force them to save (in a sense, to redistribute to themselves later in life). However, this may not justify state pensions, but rather forced enrolment in private pension schemes.
- (5) Insurance against inflation - Unanticipated inflation will eat away at the value of non-inflation indexed assets. State pensions index-linked to inflation or average earnings can overcome this problem. However, the same could be achieved by making inflation-indexed government bonds available to private pension funds.
- (6) Dynamic efficiency issues in overlapping generations (OLG) models - See other handout for details on this.

- Insurance - In order to theoretically justify the need for state provided **social insurance**, the reason for the failure of the private insurance market in a particular area must be specified. Insurance is a form of trade that improves welfare by pooling risk. This can be shown mathematically by considering N individuals each with a random income Y_N each with an identical probability distribution with mean μ and variance σ^2 (σ is the standard deviation). Suppose that individuals are risk averse so that $E[U(Y_N)]$ (the expected utility of Y_N) is increasing in μ and decreasing in σ^2 . Suppose now that the N individuals pool their income so that every individual now receives an income of $\bar{Y} = \frac{(Y_1 + Y_2 + \dots + Y_N)}{N}$. The expected

value of \bar{Y} is $E[\bar{Y}] = \frac{(E[Y_1] + E[Y_2] + \dots + E[Y_N])}{N} = \frac{N\mu}{N} = \mu$. The variance of \bar{Y} ,

meanwhile, is $Var[\bar{Y}] = E[\bar{Y}^2] - E[\bar{Y}]^2 = E\left[\left(\frac{(Y_1 + Y_2 + \dots + Y_N)}{N}\right)^2\right] - \mu^2$. If incomes are

uncorrelated, then $E[Y_N^2] = Var[Y_N] + E[Y_N]^2 = \sigma^2 + \mu^2$ and, for distinct i and j ,

$E[Y_i Y_j] = \mu^2$. This gives us $Var[\bar{Y}] = \frac{N\sigma^2 + N^2\mu^2}{N^2} - \mu^2 = \frac{\sigma^2}{N}$. By pooling risk, the

riskiness (as measured by the variance) of the pooled income is reduced. This mathematical result is key to understanding why insurance works, and is profitable for a private provider. (See Barr chapter 5 for more.)

As $N \rightarrow \infty$, $Var[\bar{Y}] \rightarrow 0$, and so a large number of individuals, by pooling risk, are able to swap a risky prospect for a certain one with equal expected value. The **certainty**

equivalent Y_1^{ce} of the risky prospect Y_1 is defined as $E[U(Y_1^{ce})] = E[U(Y_1)]$. If individuals are risk averse then $Y_1^{ce} < \mu$. A large number of individuals is unlikely, due to

administrative costs, to be able to pool risk without some overhead cost. However, a private insurance provider would be able to offer to buy each individual N's risky income Y_N for an amount X such that $Y_1^{ce} < X < \mu$, since this would be preferred by the risk averse consumer to their original risky prospect Y_N and would potentially allow the insurance provider to cover costs. We can therefore define the insurance overhead α such that

$X = (1 - \alpha)\mu$ (this assumes that the insurance company is able to pool a very large number of individuals' income). The insurance company will have $\alpha\mu = \mu - X$ to cover its costs C . Provided $C < \mu - Y_N^{ce}$, therefore, insurance will be privately provided. This result,

however, depends upon a number of assumptions:

- (1) There is some degree of uncertainty - for example, a bad state of the world with payout $\underline{\mu}$ which occurs with probability 1 cannot be insured against because then $\underline{\mu} = Y_N^{ce}$. Furthermore, as the worst possible outcome becomes more likely, $\underline{\mu} \rightarrow Y_N^{ce}$ and so it becomes less likely that $C < \mu - Y_N^{ce}$.

- (2) The incomes of the different individuals are independently distributed. If when one person has a below average income, so does everyone else, it becomes very difficult for private insurance to operate because private insurers will make big losses in many states of the world. This is why private insurance often does not cover "acts of God". (Note it is always possible just by really bad luck for private insurers to make massive losses, but the law of large numbers makes this extremely unlikely.)
- (3) Insurers can ascertain the size of average payouts - In other words, insurers must know μ with a reasonable degree of accuracy. When adverse selection or moral hazard operates in an insurance market, this condition does not hold because μ differs between individuals and is known only to them (adverse selection) or depends upon their decisions which are unobserved by the private insurance company (moral hazard). This creates the problem that the value of μ becomes endogenous to the type of insurance contract offered or the price charged for it. (See Barr chapter 5.)
- Private insurance markets are prone to fail when the above conditions are not met. Examples of events which are in various ways not privately insurable against include unemployment (because when there is an increase in the unemployment rate due to macroeconomic fluctuations, μ suddenly increases in an unpredictable way; a private unemployment insurance system would therefore probably require periodic bail out with extra government funds to keep it from going bust), sickness and disability (this is privately insurable *before* it happens - e.g. private life insurance being the most extreme case - but once someone is incapacitated, the bad outcome - that their income is low because they can no longer work - is now certain and hence uninsurable).
 - The failure of the private insurance market in a particular area does not necessarily imply that the state should step in as a direct social insurance provider. For example, it is possible to imagine a workable solution to disability insurance where people are compelled to take out private disability insurance, in the same way as they are required to take out private 3rd party car insurance. However, there is a case to be made that often the most efficient solution is a compulsory government scheme. Akerlof's market for lemons¹ and Rothschild and Stiglitz's model of failure in a competitive insurance market² (both of which are on the core micro reading list) both provide good arguments from microeconomic theory why a compulsory centralized insurance scheme might sometimes be necessary for efficiency.
 - Given we accept that the social security system has both redistributive and insurance functions, the question is whether these two aims are best met by an institutional separation between the insurance and redistributive branches. The current UK system can be seen as partially separated in that certain benefits (e.g. state pensions, jobseekers' allowance, sickness and invalidity) are somewhat dependent on previous contributions whilst others are purely based on need (child benefit, income support, disability benefit). The actuarial connection is, however, weak in that national insurance contributions are proportional to earned income whilst benefits are flat rate (but uprated in line with inflation or average earnings). The general trend is for the actuarial connection to be further weakened. For example, the requirements for qualifying for the basic state pension have been relaxed so that more people will qualify for the full pensions. Jobseekers' allowance is now virtually unconnected with previous contributions. The original system introduced in 1946 and in existence until 1976 was much more like a quasi-actuarial system in that contributions and benefits were both flat rate. This meant that the whole system had to be tied to contribution and benefit levels which were affordable by the average worker. The 1976 reforms, by making contributions proportional to income, enabled the system to grow further, but NI contributions have become more like an additional income tax. The three directions of possible further development we will consider will be:
 - (1) Completely drop any remaining actuarial features by unifying the national insurance system and income tax systems together.

1 George A. Akerlof, "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism", The Quarterly Journal of Economics, 1970

2 Rothschild M., Stiglitz J. "Equilibrium in competitive insurance markets", Quarterly Journal of Economics, 1976

(2) Return to a more quasi-actuarial system by making the size of benefits paid out of the national insurance fund proportional to income.

(3) Continue with the current system.

- Option (1) has the attraction that it would simplify the administration of the tax system. Currently people on the basic rate effectively pay 33% income tax at the margin (22% income tax plus 11% national insurance) whilst people on the higher rate pay 41% effective income tax at the margin (40% income tax plus 1% national insurance). (This leaves aside a few complexities such as the fact that NI contributions are currently levied on a weekly rather than annual basis.) This system could be unified into a single income tax schedule. Disability benefit, the basic state pension, unemployment benefit and others would then all simply be flat rate and based on need-orientated characteristics (age, disability, being unemployed) and not upon previous contributions. The possible argument against this is that breaking the (albeit already weak) link between NI contributions and the payment of benefits like the basic state pension will worsen labour market distortions. This is probably not a very strong argument, but then neither is the argument for the benefits simply from unifying the system in this way, since the current set-up works reasonably well anyway.
- Option (2) would take the UK welfare state more in a continental European direction. The size of unemployment benefit and state pensions would then be linked to previous contributions, so that people who had earned and paid in more in the past would get a higher unemployment benefit or a higher pension (like a defined-benefit occupational scheme). Similar conditions would apply to incapacity/sickness benefit. The argument for this position would come from the efficiency reasons why compulsory social insurance is necessary. People are unlikely to be able to perfectly smooth their income throughout the lifecycle and in different possible states of the world using only private market-based insurance due to the reasons already given for the failure of private insurance markets in many circumstances. Labour market distortions from payroll taxes would also be reduced because people would know that on average they get back what they pay in to the system. The argument against is that actuarial connections between contributions and benefits will only ever be imperfect, and so labour market distortions will remain. In order to keep benefits for the poorest at the same level, meanwhile, the proportion of GDP taken by the national insurance system would have to be massively increased. The state would be taking a larger chunk of national income as "payroll tax", and these resources would not be efficiently targeted upon the neediest individuals. The European systems face the problem that they will gobble up increasing proportions of GDP as the population ages, unless they are reformed to make them less generous. This of course forces a policy-maker to choose between redistribution and quasi-actuarial insurance as the guiding logic.
- Option 3 is clearly the messy solution, but if there are no strong arguments to go in direction 1 or 2 then it may be optimal, given that the current system is reasonably functional, to stick with the current partial separation of the insurance and redistributive functions in the form of the separate NI contributions and NI fund.