The Welfare State as Piggy Bank

Information, Risk, Uncertainty, and the Role of the State

NICHOLAS BARR

OXFORD
UNIVERSITY PRESS

Chapter 2

The Market and Information

This chapter sets out the economic theory that underpins the rest of the book. It starts (Section 1) with a simple model—the Fisher model—of rational consumer choice over the life cycle. Subsequent sections relax the underlying assumptions, discussing imperfect information in the goods market (Section 2) and insurance markets (Section 3). Section 4 summarizes the implications for policy.

1. THE WELFARE STATE WITH PERFECT INFORMATION

THE SIMPLE FISHER MODEL shows the options available to an individual over time. The horizontal axis in Figure 2.1 shows a person's potential consumption in period 1 (her younger years), the vertical axis that in period 2 (her older years). Suppose she has an initial endowment shown by point a: she can consume C_1 units in period 1 and C_2 units in period 2. However, she can increase her options by trading with other individuals—that is, by saving or borrowing. For example, she could save C_1–C'_1 units of consumption in period 1 in exchange for C'_2–C_2 units in period 2, thus moving to point e.[1]

Thus a person with an initial endowment of C_1 in period 1 and C_2 in period 2, faces a lifetime budget constraint b–b. The consumption pattern that maximizes lifetime utility is shown by point e, which the person attains by saving C_1–C_1' when younger, making possible consumption of C_2' when older. Similarly, with an initial endowment of d on the same budget constraint, the person could move to e by borrowing in period 1.

This simple model is based on a series of assumptions, including a well-behaved utility function, rational behaviour, certainty, and competitive markets. The assumption of certainty is critical. First, it implies perfect information:

[1] If the interest rate were zero, she could save (say) 1 unit in period 1 and consume an extra unit in period 2. If her initial endowment shown by a comprises 7 units in period 1 and 3 units in period 2, she could, by borrowing 3 units, consume 10 units in period 1 and 0 in period 2 or, by saving 7 units, could consume 0 in period 1 and 10 units in period 2, or anywhere in between. Thus her consumption opportunities are shown by a budget constraint with a slope of −1. If the interest rate is 10%, the budget constraint becomes steeper. For each unit she saves in period 1, she receives 1.1 units in period 2. By saving she can therefore move from a to e. Thus the budget constraint b–b goes through the initial endowment point, a, with slope determined by the interest rate. For a simple introduction, see Auerbach and Kotlikoff (1998: ch. 2) and, for fuller discussion, Varian (1999: ch. 10).

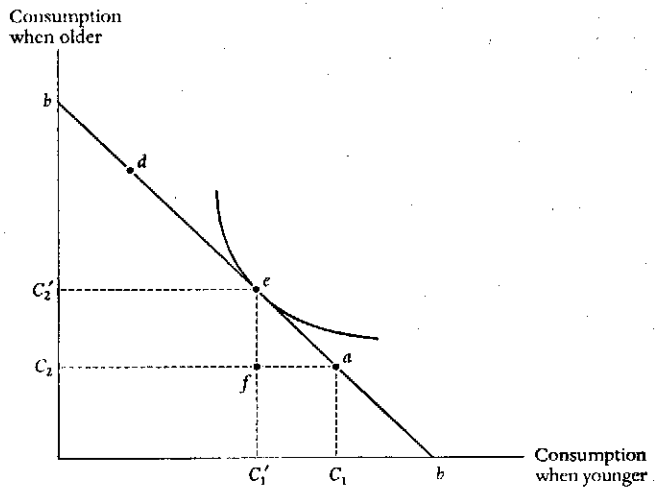

Figure 2.1. *Rational choice in the Fisher model*

consumers are well informed about the quality of the goods and services they buy; for example, people are assumed to be able to make well-informed choices between different pension schemes, different medical treatments, and different educational activities. Secondly, the assumption rules out stochastic outcomes such as risk (where the probability distribution of outcomes is known) and uncertainty (where it is not). The absence of risk means that there is no need for insurance; it also means, for example, that lenders do not have to form a view about the riskiness of applicants for student loans. The absence of uncertainty rules out common shocks such as inflation.

THE CASE FOR THE MARKET. In a world of certainty, there is therefore little need for a welfare state.

- Insurance is unnecessary, since there is no risk.
- People provide for their old age through voluntary saving, and finance their education by borrowing in perfect capital markets. Thus consumption smoothing takes place through voluntary action using private institutions. With perfect information and no external shocks, such behaviour is efficient.
- Transient (i.e. temporary) poverty is also dealt with by borrowing or saving. If the poverty line in period 2 is shown by C_2' in Figure 2.1, someone with an initial endowment of a is poor in period 2, but can deal with the problem by saving in period 1, thus moving from a to e. Dealing with poverty for someone who is not lifetime poor is more akin to consumption smoothing than to traditional poverty relief.
- The only reason for a welfare state in such a world is to provide poverty relief for a person who is lifetime poor—for example, someone with an initial endowment of f in Figure 2.1, whose income is not enough to keep him or her above the poverty line in both periods.

Introducing risk into the model does not change things, since individuals can buy actuarial insurance. For example, a person who does not know how long he will live can convert his pension savings into an annuity. Where individuals and insurers are well informed, such outcomes are efficient.[2]

RELAXING THE ASSUMPTIONS. The rest of the book is concerned with relaxing the certainty assumption. Discussion includes:

- imperfect information in product markets, insurance markets, and capital markets;
- risk and uncertainty, particularly in a context of imperfectly informed insurers;
- external shocks—e.g. demographic change.

Risk and uncertainty imply the need for insurance (Part 2 of the book); and all three factors have major implications for consumption smoothing from middle years to later years to finance pensions (Part 3) and to younger years to finance investment in human capital (Part 4).

2. IMPERFECT INFORMATION IN THE GOODS MARKET

This section summarizes the implications for market allocation of imperfect information and a range of other technical problems.[3]

2.1. *Types of intervention*

Before discussing why the state might intervene, it is useful to outline the methods by which it might do so—regulation, finance, public production, and/or income transfers.

Regulation in some instances, may have more to do with social values than economics (for example, regulation of shop opening hours). But much regulation is

[2] These results could be modelled with an overlapping generations model. Brooks's analysis (2000) of pensions uses four types of people for pensions. A full analysis of the welfare state, including borrowing to finance education, would need at least five types.

[3] Many of the key articles in the literature on imperfect information and other market failures are collected in Barr (forthcoming: vol. i).

directly relevant to the efficient and/or equitable operation of markets, especially where consumer information is imperfect. Regulation of *quality*, mainly on the supply side, includes hygiene laws relating to the production and sale of food and pharmaceutical drugs, and consumer protection legislation more generally. Regulation of *quantity* more often affects individual demand—for example, the requirement to attend school, compulsory social-insurance contributions, and mandatory automobile insurance. *Price* regulation includes minimum wages. There can also be regulation of *total expenditure*—for example, global budget caps for medical spending.

Finance involves subsidies (or taxes) applied to the *prices* of specific commodities or affecting the *incomes* of individuals. Price subsidies can be partial (for example, tuition fees at public universities) or total (for example, free primary education). Similarly, prices can be raised by a variety of taxes, particularly in areas that the state wishes to discourage (for example, a tax on tobacco). Income subsidies raise different issues which are discussed shortly.

Though regulation and finance modify market outcomes, they leave the underlying mechanism intact. More drastically, the state can take over the supply side by producing goods and services itself—for example, school education and (in some countries) most health care. Alternatively, the state can commission and direct private-sector activity—for example, weapons procurement. The distinction between finance and production emerges throughout the book.

The previous interventions all involve direct interference with the market mechanism. Income transfers do not interfere directly, but enable recipients to buy goods of their choice at market prices—for example, elderly people receive a pension with which they buy food.

2.2. *Information and other problems*

As a precursor to information problems, it is useful to clear the intellectual undergrowth by briefly summarizing an earlier literature on market failures.

Older arguments

Markets are efficient only if a number of conditions hold.[4]

PERFECT COMPETITION arises where there are many buyers and sellers, free entry into and exit from the industry, and a homogenous product. The assumption can fail in various ways. A monopolist maximizes profit by restricting output below the efficient level. Solutions include regulation—for example, the imposition of a maximum price, or a per-unit subsidy (to encourage larger output) clawed back by a lump-sum tax.

[4] For fuller discussion, see Stiglitz (2000: ch. 4).

NO EXTERNAL EFFECTS. This assumption is violated when an act of person A imposes costs or confers benefits on person B, for which no compensation from A to B or payment from B to A takes place. If I am inoculated against a communicable disease, this reduces my chances of catching it, and also benefits others because it prevents them catching the disease from me. Social benefits therefore exceed my private benefits, and the market output will typically fall below the efficient output. The same effect operates in reverse: where an activity imposes an external cost on others (pollution, for example), markets generally lead to output above its efficient level.

A range of solutions exist. Coase (1960) shows that, where the law assigns unambiguous and enforceable property rights, the market itself can resolve the problem through negotiation between the parties concerned. However, this may not be possible. Property rights may not be enforceable (water pollution), or the numbers involved may rule out negotiation (traffic congestion). A second way of dealing with externalities (Meade 1952) is through merger of the affected activities. Where neither solution is possible, intervention may be justified through regulation (for example, mandatory water standards) or an appropriate tax on the activity generating the external cost (for example, a tax on gasoline).

NO INCREASING RETURNS TO SCALE. Increasing returns to scale arise when doubling all inputs more than doubles output. An implication is that competitive pricing generates long-run losses. Two forms of intervention are possible: paying firms a lump-sum subsidy equal to the loss associated with competitive pricing, or taking the industry into public ownership and paying an identical subsidy. The appropriate intervention is, therefore, subsidy or public production, or both.

NO PUBLIC GOODS. Public goods in their pure form exhibit three technical characteristics: non-rivalness in consumption, non-excludability, and non-rejectability. Private goods are rival in consumption in the sense that one person's consumption is at the expense of another's—for example, if I buy an apple, there will be one apple less for everyone else. Excludability means that I can be prevented from eating the apple until I have paid for it. Rejectability implies that I do not have to buy the apple. Not all goods display these characteristics, the classic example being national defence: the arrival of a person from another country does not reduce the amount of defence available to everyone else (non-rivalness in consumption). Nor is it possible to exclude the new arrival by saying that the bombs will be allowed to fall on him until he has paid his taxes (non-excludability). Nor can he reject the defence on the grounds of pacifist beliefs (non-rejectability). Similar considerations apply wholly or in part to roads, broadcast signals, and public parks. Public health, too, has important public-goods attributes: if the water supply is purified, or clean-air legislation enforced, nobody can be

excluded from the benefits; and the structure of laws and, more generally, the rule of law have important public-goods characteristics.

It is a standard result (the classic article is Samuelson 1954) that public goods create one of two problems. The market may be inefficient—for example, if the good is priced at average cost. Alternatively, the market may fail to produce the good at all. In the latter case, if the good is to be provided, it will generally have to be publicly organized. This will involve public funding, but not necessarily public production. For instance, the state may mandate and pay for a water purification programme, but the work might be carried out by a private firm.

The main conclusion is how thin is the 'old-style' economic justification for large-scale, publicly organized welfare-state services.[5] Imperfect competition, externalities, and increasing returns to scale may justify regulation or particular types of subsidy. Only public goods offer a strong efficiency case for public production. The only other *efficiency* argument for extensive public provision is if it is believed that an externality is so strong that it justifies compulsory and/or subsidized consumption of a good by the entire population: examples include publicly organized sewerage in response to nineteenth-century British cholera epidemics and, it can be argued, might also include broader aspects of public health, and basic education. In these cases, the externality is so large that the public-goods analysis applies.

The question, then, is how to explain pervasive public involvement in all countries in insurance and consumption smoothing. One route, discussed in Section 4, is through public choice; the other is through information failures.

Arguments based on imperfect information

Consumers and firms need to be well-informed in at least three ways: about quality, about price, and about the future. A later literature explores these assumptions and the effects where one or more fail.[6]

PERFECT INFORMATION ABOUT QUALITY. Consumers are well informed about many products. However, they might be badly informed—for example, about the quality of a school, or about the appropriate type of medical treatment. Producers might be poorly-informed about the quality of a worker applying for a job, or about the riskiness of an applicant for a loan or for insurance. New (1999, 2000) distinguishes two separate problems: an information problem can be resolved by supplying the relevant information; with an information-*processing* problem,

[5] Other than public goods, it is not possible, for example, to justify large-scale interventions from Bator's (1958) anatomy of market failure.

[6] The quality literature has its roots in classic articles by Arrow (1963) and Akerlof (1970).

in contrast, the necessary information is so complex that, even if it is supplied, economic agents are not necessarily able to make rational choices. New distinguishes three situations in which, even if they are given the necessary information, people might not be able to choose rationally:

- where the benefits of a good or service occur only a long time in the future, where the problem might be regarded as in some sense a failure of imagination;
- where a good or service has a very small potential for harmful consequences for an individual, where the problem is an inability to process very small probabilities;
- where a good or service requires complex information to be digested and processed, where the problem is lack of technical ability.

In the face of such problems, several solutions are possible. The market may develop institutions to supply information: it is possible to buy consumer magazines; equally, one can pay to have a professional survey of a house. When a person buys such products or services, what she is buying is information. In other cases the state may respond with regulation—for instance, hygiene laws for food. This is generally appropriate where the information is sufficiently non-technical for the consumer to understand it. Where information problems are serious and where, as discussed above, there is a long-time horizon, where the relevant probability is very small, or where the necessary information is too technical to be readily understood by the average consumer, the market will generally be inefficient, and public production may be a better answer. We return to the issue in Section 4.

PERFECT INFORMATION ABOUT PRICES. A second strand of the literature analyses the effects of imperfect information about prices and wages. Again, the market may supply the necessary information—for example, car magazines containing price guidelines, or web sites that make it possible to search for the lowest price. Equally, it is possible to have a house or a piece of jewellery professionally valued. If such solutions do not suffice, the state may intervene with regulation requiring sellers to publish prices.

PERFECT INFORMATION ABOUT THE FUTURE. As well as information about quality and price, individuals also need accurate expectations about the future in order to make rational choices over time, this being the decision analysed by the Fisher model. The assumption is broadly true of food (since people know that they will need to eat tomorrow, next week, next month); it is not true of medical care, because they do not know whether or when they will suffer health problems. In principle, the market can cope with risk through insurance, but insurance may be inefficient or impossible.

3. IMPERFECT INFORMATION IN INSURANCE MARKETS

3.1. *Actuarial insurance with perfect information*

The term 'insurance' is used by different people to mean different things: as a device that offers individuals *protection against risk*, or as an *actuarial mechanism*. The first defines insurance in terms of its objective, the second in terms of a mechanism by which that objective might be achieved. Even where institutions are not insurance in the second sense, they might still be regarded as insurance in that they offer protection against risk.

A risk-averse person is someone who prefers a lower income with certainty to an income that is higher on average but with greater variation. Many students, for example, would prefer a definite scholarship of £5,000 rather than face a lottery with a 50 per cent chance of £12,000 but a 50 per cent chance of nothing.

Risk-averse people buy insurance voluntarily because it increases their welfare. If I have to rely on savings to finance my old age, I need to save enough to cover my *maximum* life span. In contrast, if I can insure by buying an annuity, I need to save enough to cover only *average* life expectancy. The difference between the two figures is large. If I need £20,000 per year (roughly UK national average income in the early 2000s), to live comfortably and retire at 65 allowing a maximum age of 110 the necessary pot of savings (ignoring interest) is £900,000 (i.e. 45 × £20,000). On the other hand, if life expectancy at 65 is fifteen years, the necessary pot if I can buy insurance is only about one-third as large. When a person buys insurance, what he or she is buying is certainty.

The supply of insurance has an easy intuition. Suppose that there are 100 of us; that we decide to fly to Paris to see a football match, that each of us has a suitcase whose contents are worth £1,000, and that we know from long experience that on average 2 per cent of suitcases get lost. Thus each of us faces a potential loss, L, of £1,000, which occurs with a probability, p, of 2 per cent. In those circumstances, it would be possible to collect 2% × £1,000 = £20 from each of the 100 people, i.e. £2,000 in total; when we arrived in Paris, we would find which two people had lost their suitcase, and pay each £1,000 in compensation.

More formally, the actuarial premium for the ith individual, π_i, is defined as:

$$\pi_i = (1+\alpha)p_iL \tag{2.1}$$

where p_iL is the individual's expected loss, and α is the loading the insurance company charges to cover administrative costs (for example, sending an expert to assess the damage) and competitive profit. π is the price at which insurance will be supplied in a competitive market.[7]

[7] For fuller discussion, see, in ascending order of difficulty, Burchardt and Hills (1997: ch. 1), Stiglitz (2000: ch. 12), Barr (1998a: ch. 5), Culyer (1993), and Rees (1989).

The price of insurance thus depends on (*a*) the degree of risk and (*b*) the size of the potential loss. Car insurance premiums are high for a driver who is young, or who lives in a high-crime area (both factors leading to a higher probability of loss); and they are high for someone who drives a Rolls Royce or Mercedes (because the potential loss is large). A middle-aged person with a good driving record driving a small, elderly Ford pays a much lower premium.

This, broadly, is the way in which private insurance operates. Note that thus far there is no need for state intervention. As discussed earlier, a rational risk-averse person facing a known risk (for example, about how long he will live) will buy actuarial insurance, which the market can and will supply. Thus the simple Fisher model extends easily to this case.

3.2. *Information problems*

Insurance along the lines of Equation (2.1) is efficient only if a number of conditions hold. Where they fail, private insurance may be inefficient or impossible.

INDIVIDUAL RISK, NOT COMMON SHOCK. Insurance requires a predictable number of winners and losers; in other words, the probabilities in Equation (2.1) are independent. This condition holds for individual risk such as age at death or the likelihood that one's car will be stolen. With a common shock, in contrast, if one person suffers a loss, so does everyone else. As discussed in Chapter 7, an inflationary shock can adversely affect all pensioners. Actuarial insurance is not able to cope with this problem: common shocks are uninsurable.

RISK, NOT CERTAINTY. Insurance addresses risk. Thus p_i in Equation (2.1) must be less than one. If $p_i = 1$, it is certain that the insured person's car will be stolen; hence, there is no possibility of spreading risks. As a result, Equation (2.1) simplifies to:

$$\pi_i = (1+\alpha)L \tag{2.2}$$

and the insurance premium exceeds the insured loss. In those circumstances there are no welfare gains from joining a risk pool. Thus medical insurance for the elderly is problematic, because the probability of requiring medical care is very high; for the same reason, actuarial insurance cannot cover medical problems that the individual already has at the time he or she applies for insurance, for whom the probability of ill health equals one. Pre-existing conditions, in short, are generally uninsurable.[8]

[8] The problem has both efficiency and equity aspects. It creates inefficiency in the sense of a missing market—risk-averse individuals who would like to buy insurance but are unable to do so. For the same reason, there is also inequity.

The two conditions just discussed relate to the fundamental nature of insurance—what it is and what it is not. The remaining conditions show how insurance can fail for reasons explicitly related to imperfect information.

RISK NOT UNCERTAINTY. If p_i in Equation (2.1) is unknown, the insurer cannot calculate a premium, making actuarial insurance impossible. In other words, insurance can cope with risk (where the probability is known) but not with uncertainty (where it is not). The problem can arise, first, where the insured event is rare (for example, early satellite launches); with few observations, any estimate of the probability will have a large variance. The probability might be unknown, secondly, because of complexity. Thus actuarial insurance against future inflation is difficult or impossible, because the probability of different levels of future price increases cannot be predicted.[9] Equally, there is no way today of estimating the extent of risk from exposure to BSE ('mad cow disease'). The problem can arise, thirdly, where the insured event has a long time horizon. An example discussed in Chapter 5 is the probability that someone aged 25 today will require long-term care in extreme old age.

A further condition is that all agents—the person buying insurance and the person selling it—must be equally well-informed. Where the condition fails—the problem of asymmetric information—actuarial insurance is inefficient or non-existent. There are two classes of problem: adverse selection and moral hazard.

ADVERSE SELECTION. Efficiency requires that high-risk individuals pay a higher premium than low-risk individuals. Thus the efficient premiums for people with a low and high risk of loss, p_L and p_H, respectively, are:

$$\pi_L = (1+\alpha)p_L L \tag{2.3}$$
$$\pi_H = (1+\alpha)p_H L. \tag{2.4}$$

With automobile insurance, for example, someone who is twice as risky pays roughly twice as high an insurance premium.

Adverse selection arises where the purchaser can conceal from the insurer the fact that he is a bad risk, and is thus an insurance-market manifestation of 'lemons' (Akerlof 1970). The individual knows he is a 'lemon' (that is, a bad risk), but can conceal the fact from the insurer, hence the description of adverse selection as 'hidden knowledge'.[10] Private insurers in the USA, for example, believe (rightly or wrongly) that elderly applicants for medical insurance are disproportionately people who have something to hide. The problem can also arise if health care is an important part of employer benefits: firms with the best health-care packages will tend to attract workers with health problems, thus reducing the firm's competitiveness.

[9] As discussed in Chapter 7, the government can issue indexed bonds to deal with inflation; that, however, is not private insurance, but tax-funded state intervention to assist private insurance.

[10] Akerlof's competitive analysis was extended by Rothschild and Stiglitz (1976) to cover strategic behaviour by firms. For further discussion, see, in ascending order of difficulty, Atkinson (1989: ch. 7) reprinted in Barr and Whynes (1993: ch. 2), Culyer (1993), and Rees (1989).

The problem is not that people differ in their riskiness, but that insurers are less well informed than the buyer about the applicant's risk status. Insurers can respond in two ways, though, as discussed in more detail in Chapter 4, Section 2, in the context of medical insurance, neither approach is a complete solution. One option is to seek a 'pooling equilibrium' in which the insurer charges a common premium, $\bar{\pi}$, based on average risk, $\bar{p}$. Two potential problems result: the outcome might be inefficient because low risks, who have to pay a higher-than-actuarial premium, buy an inefficiently small amount of insurance and, conversely, high risks an inefficiently large amount; or the market may fail entirely as low risks opt out. A second way in which insurers attempt to get round adverse selection is to try to separate high and low risks through self-selection (a 'separating equilibrium') by offering policies that incorporate incentive structures whereby a customer's choice reveals his or her true probability—for example, by offering only limited cover for the first year of the policy.[11] Two problems can result: such a separating equilibrium might not exist; or, where it does, it is inefficient.

A third possible approach is to enforce a pooling equilibrium by making insurance compulsory. Thus low risks are not able to reduce the amount of insurance they buy nor to opt out; and high risks are not able to buy an inefficiently large amount of insurance. If preferences do not differ greatly across individuals, the welfare loss from such compulsion may be small.

MORAL HAZARD. A second class of asymmetric information, moral hazard, arises where the insured person can influence the insurer's expected loss, $p_i L$, without the insurer's knowledge (hence the characterization of moral hazard as 'hidden action').

Pauly (1974) analyses individual expenditure on a preventive activity, z, which reduces the probability of the insured event. The efficient level of z is where its marginal cost is equal to the marginal reduction in insured losses. But, if the individual bears none of the cost of inaction and the insurer cannot monitor a person's preventive activity (hidden action), the individual's incentive is to spend little or nothing on z. Inefficiency arises because, if monitoring is not possible, people behave differently if they are insured. Four cases should be distinguished.[12]

1. Endogenous p_i, but at substantial psychic cost. An example is suicide: the probability is endogenous, but only at high cost to the individual. Because

[11] See Ravallion and Datt (1995) for analysis of such self-selection in different contexts.

[12] For fuller discussion of moral hazard, see Stiglitz (1983). Rees (1989) or Culyer (1993).