

SHELTER FROM THE STORM: REFLECTIONS ON “THE RISK SOCIETY”

“As an economist, it is my duty to talk about costs..” J. B. Brigden ♦

Risk is inherent in life because life is inherently uncertain. As Peter Bernstein has put it, uncertainty simply means that more things *can* happen than *will* happen – and while some of those things are very good, others are bad and some are very bad indeed. How we deal with those bad outcomes is a crucial part of our lives as individuals and as societies.

That governments have a role in helping to deal with bad outcomes is undoubted. But seen in the sweep of history, governments have likely caused at least as much risk as they have alleviated. In the 20th century, governments killed over a hundred million people, often their own citizens, and displaced millions more, destroying their homes, their livelihoods, their hopes and aspirations. It therefore seems paradoxical that we should turn to government for shelter from the storms – yet we do, and perhaps nowhere more so than in Australia.

This has long been so. In his great work *Australia*, published over seventy years ago, W. K. Hancock defined the ideology of Australian democracy as being “the sentiment of justice, the claim of right, the conception of equality and the appeal to Government as the instrument of self-realisation”. All of these, he said, combined in a tendency among Australians to “look upon the State as a vast public utility, whose duty it is to provide the greatest happiness to the greatest number.” But entrenched as those habits are, the “school of hard shocks” of the 1970s and 1980s brought a new sense of balance, and a retreat, however cautious and incomplete, from over-extended government.

At the heart of that change was acceptance that a fair society could only be achieved and maintained in an efficient economy. The logic of efficiency proved a powerful force, at least at a rhetorical level, and though its impact on policy has often been greatly over-stated, it set the terms for policy analysis.

It is therefore unsurprising that the challenge to the balance which emerged from the 1970s and 1980s should be couched in the language of efficiency. More specifically, the prevailing consensus, which is cautious about the benefits of government intervention, is now being challenged by the claim that achieving efficiency requires a government that does more, rather than less, to insure its citizens against economic and social risk.

Advancing that challenge is the main purpose of John Quiggin’s recent and widely-publicised paper on “The Risk Society”, published by the newly-formed Centre for Policy Development.

♦ Quoted in Coleman, W., S. Cornish and A. Hagger (2006) *Giblin’s Platoon: The Trials and Triumph of the Economist in Australian Public Life*, The ANU Press, Canberra, ACT, at p. 59.

There is much to criticise in Quiggin's paper. It is long on assertion and short on evidence. Many of the assertions, when examined carefully, do not hold water; others are, at best, matters of opinion. Perhaps most annoying, the principles at issue and the underlying reasoning are poorly set out. Little, other than a virulent dislike of John Howard and George Bush, is expressed unambiguously, making it difficult to know precisely what it is that Quiggin is arguing for or against. As a result, the reader has to infer the logic and try to make out the key propositions.

Three such propositions can be discerned.

The first is that governments should help individuals bear a wide range of risks. Some of those risks go to accidents of birth – whether we are born rich or poor, clever or less so. Others go to the uncertainties of life – whether we are healthy or sick, do well materially or do poorly, age comfortably or not. Yet others go to the varying fortunes that might face Australia as a whole – such as the risk of cyclical downturns, or of renewed foreign threats. And yet others involve the world as a whole, most notably the risk of climate change.

The second is that the help governments provide – the insurance they give us against economic and social risks – should be substantially greater than it currently is. In other words, it is not only a question of providing *some* help, but of providing significantly *more* help than government now does.

Third and last, Quiggin argues that the help should not only be funded but also supplied by government. The public sector should not only *pay* for services, but also *deliver* them.

It is difficult to disagree with the first of these propositions and no mainstream political movement in Australia does. The hard edge in Quiggin does not come from there. Rather, it comes from the willingness to argue for the second and third of his key contentions – and to argue them on grounds of efficiency. It is this that makes Quiggin somewhat distinctive, and somewhat more interesting, than the bulk of his fellow-travellers: for while they would say that a more far-reaching role for government in tiding people over hardships and difficulties is essential for a just society, Quiggin supports such a role because he thinks it would be efficient – it would make us better off.

It would make us better off, Quiggin contends, because government can insure against those risks at lower cost than we can ourselves. This is not to say that government is the only potential source of that insurance – as the availability of private health insurance makes clear, such a claim would be manifestly untrue. But, Quiggin suggests, those other sources are necessarily either more limited or more costly than the insurance governments could provide. As a result, the best way, from an efficiency standpoint, of protecting from the risks of illness, poverty and ignorance is for governments to compensate those whom those risks afflict. Moreover, because government provision of that insurance would lower its costs, it would allow much more of that insurance to be provided, with further efficiency gains obtainable from government not merely funding the insurance but also directly providing the relevant services.

This is a sharper version of the argument and the contentions than one finds in Quiggin's paper. Whether it does it justice to that paper or not is difficult to tell, as the paper itself is so unclear. But at least it focuses the debate, forcing attention on what are indeed important issues: about who should provide the means by which we manage risk in society, to what extent and through what instruments.

A useful discussion of these issues, and an evaluation of the efficiency claims Quiggin makes, cannot be carried out without a high-level understanding of how economists view risk and insurance. That conceptual framework lies in the background of Quiggin's paper. It is indispensable both for assessing the argument that government is the most efficient insurer of many economic and social risks, and for examining the question of "how much is enough?" – that is, to the extent such government insurance is to be provided, how much of it is desirable. While Quiggin does not make this framework at all explicit, I start by summarising its major features before turning to consider each of the claims Quiggin makes.

RISK, INSURANCE AND ECONOMICS

Economists view risk much as they view the many other tasks that need to be carried out in society. Bearing risk is an activity, and that activity should be undertaken by those who are best placed to do so. The summary way of expressing this is that risk should be allocated to those who can manage it at least cost – the "least cost risk bearer". Insurance is a way of shifting risk from the person who bears its direct impact onto such a "least cost risk bearer". This yields efficiency gains for two reasons. The first is because it replaces expensive risk-bearing by those who bear the direct impact by less expensive risk-bearing by the insurer. The second comes from the fact that the lower cost of the insurance service thus provided allows us to obtain more insurance, which is something we value.

We value insurance because – despite enjoying a flutter – we dislike uncertainty, at least when the possible loss is sufficiently great to threaten our standard of living. Insurance lets us deal with this by allowing us to transfer income from situations where we value it less because things are going well, to situations where – because an undesirable contingency has arisen – we value it more.

When I take out health insurance, for example, I give up a bit of income each month in the form of the premiums I pay. Since I am healthy and productive when I make those payments, the sacrifice they involve is likely small. Should I fall ill, on the other hand, I get a payment back, in the form of the insurance benefit; given the expenses my illness could cause, and the possible loss of income from working, that extra income is money I value highly.

What the health insurance has done, in other words, is to allow me transfer income from a "state of the world" where the value of a *bit less* income is low to a "state of the world" where the value of *some more* income is high. As a result, it has allowed me to smooth my standard of living over time, which is an outcome most of us prefer to one where that standard of living might fluctuate greatly depending on whether undesirable contingencies did or did not occur.

The concept of smoothing suggests equalisation over time – and that is indeed what *private* or *voluntary* insurance does. Put simply, insurance is not a way of increasing the income I can expect over my life-time, but rather of rearranging it as between different situations, situations that can be viewed as different "states of my luck." From an actuarial perspective, each individual in a private insurance scheme is a tub that stands on its own bottom, where that bottom is defined by the life-time income the individual can expect to earn.

This is clearest in self-insurance, which is what happens when we act as an insurer on our own behalf. In practice, all of us self-insure against many risks. We do that by studying and acquiring skills that allow us to adapt to changing circumstances, by saving some income for a rainy day, or conversely, when our income needs experience a temporary spike, by borrowing against our future earnings. Such self-insurance is likely to be by far the main form of insurance in society. However, there are many limits on how far self-insurance can go.

For example, most people would find it difficult to self-insure against very major risks, such as the risk of needing complex and prolonged hospitalisation. The sums that one would need to set aside would be simply too great. Moreover, even if one could set aside those sums, it would likely be inefficient to do so, as they might lie idle for long periods of time or even never be used. The income sacrifice associated with self-insuring against such risks – that involve high potential costs, but have low likelihoods of occurring – would be so great that absent viable alternatives, most people would simply go without cover. Since the option of getting a bit more income when “bad things” happen is highly valued, an outcome in which many people were left without any cover against those things would be undesirable.

The efficient way of dealing with these “high stakes, low probability” risks is to pool them. All pooling means is bringing together many individual, essentially unrelated, risks. Because these risks are unrelated, it is likely that only a few will eventuate in any period, while many others will not. As a result, the many risks that do not eventuate can help cover the costs of the few that do. Self-insurance is itself a form of pooling, since when we self-insure, we pool our own risks across different periods of time. But there can be very large gains if we extend the insurance pool further.

One way of doing so that occurs in all human societies is through the family. Like self-insurance, insurance within the family is pervasive, and probably dwarfs insurance that is sold through markets. Some of the insurance functions that families provide are trivial, such as making cups of tea when one is ill. Others, such as support during childhood and old age, are crucial. But like self-insurance, insurance within the family has its limits.

One important aspect of these limits is that the income of the family and that of its key member or members are so highly correlated. If I am the primary bread-winner, there is only so much support that my family can promise to provide. In effect, should I lose my job, it is both my income and that of the family unit that has gone – and with them, presumably, much of the help the family can give its members, including me.

This highlights the great benefit of pooling risks with strangers – their bad luck is not likely to coincide with mine. They can continue to pay their premiums when I am a net beneficiary of the insurance, rather than a contributor.

It is for this reason that the development of insurance over the years has involved the construction of ever greater risk pools, which aggregate risks that are only very slightly associated with each other, if at all. Putting together those risk pools is what specialist insurance companies primarily do. The attempt to make those pools even wider, involving risks that are completely uncorrelated, has driven financial innovation not only in the insurance industry but more generally in capital markets for many years.

However, pooling risks with strangers is not easy, and is never costless. This is because while the luck of strangers may be uncorrelated with my own, I also know much less about them. Additionally, I have much less control over them than I have over myself or my family members. Put in economic terms, this is a situation of “asymmetric information” – that is, a situation where market participants differ in the information they have available to them. This information asymmetry creates a number of problems that affect all forms of insurance. Two such problems are of the greatest importance.

ADVERSE SELECTION

The first arises from the fact that we each know more about our own exposure to key risks than anyone else does – for example, about whether our diet and lifestyle makes us likely to fall ill, whether our attitude to work and our relation to our employer makes us likely to lose our job, or whether when we borrow to fund our education, we ever intend to get a well-paying job. When we self-insure, we are forced to take that information into account, however painful that may be. But when we seek insurance from a specialist insurer, we have little incentive to fully disclose what we know if it will result in a higher premium. Insurers therefore offer insurance to people whose risk characteristics they can assess only imperfectly.

This creates particular problems for markets where taking on insurance is voluntary – that is, where individuals can choose whether or not to insure and hence form part of the risk pool. If “bad risks” can masquerade as “good risks” and pay only the “good risk” premium, then if the pool is to break even, the premiums charged to the genuinely “good risks” must be higher than the actuarial cost of insuring them (as they have to subsidise some of the high risk cases). However, this reduces the incentives the genuinely “good risks” have to buy insurance (as against self-insuring), unleashing a process in which the quality of the insured pool deteriorates as its coverage of the “good risks” shrinks.

This is the phenomenon known as “adverse selection”: the fact that the relevant risk characteristics of people who voluntarily purchase insurance differ from the risk characteristics of those who do not in ways that are not fully observable by insurers. In the extreme, adverse selection can eliminate a voluntary insurance market completely.

Voluntary health insurance markets, for example, rely on risk pools in which those insureds who do not experience an adverse event “subsidise” medical services for those who do. Health funds that attract only high risk members incur higher costs than they would if an unbiased sample of the population insured with them, and must raise their premiums to stay in business. If these higher premiums must be charged to high and low risks alike, then those premiums will discourage healthy, low risk individuals from insuring, which in turn leads to a further deterioration of the risk pool. Left unchecked, this process could lead to a downward spiral in which voluntary plans were not financially sustainable.

Usually, adverse selection does not lead to the disappearance of voluntary insurance. But it can lead to what seem like inefficiencies in its provision, at least compared to the insurance that would be provided in a world where insurers had complete access to information about insureds.

For example, insurers spend considerable resources screening applicants for insurance – resources that would be unnecessary if they simply covered the pool as a whole. Additionally, they use exclusions from coverage so as to reduce their exposure to conditions that applicants are likely to know more about than they are. Seen from the standpoint of a perfect world, those exclusions are inefficient, because the benefits of being able to insure those conditions would outweigh the costs. And lastly, even with the best designed voluntary insurance scheme, the need to cover the costs of undetected bad risks will cause some of the good risks to self-insure in whole or in part, despite the fact that they would be better off pooling their risks in a scheme that was actuarially fair.

It is worth stressing that these are inefficiencies relative to what would happen in a perfect world. Moreover, even given that (obviously impossible) benchmark, their extent should not be overstated. There are many insurance markets – going from life insurance to property and liability insurance – which function relatively well despite the challenges posed by adverse selection. This occurs for two reasons.

First, insurance providers are relatively adept at identifying, classifying and pricing risks, as well as at putting together very wide risk pools, both for insuring and for reinsuring risks. The search to maximise profits has proven to be as powerful a driver of innovation and efficiency in this market as in so many others.

Second, while there is adverse selection in insurance markets, there is also considerable “beneficial” or “advantageous” selection. This is the phenomenon in which the best risks, being those people who are most cautious (and hence drive more carefully, take better care of their health, and so on) place the greatest value on insurance. They are therefore most likely to remain insured, even if doing so requires paying some contribution to the costs of risks that are not as good as their own.

As a result, voluntary insurance markets in the “real world” may be far closer to tolerable levels of efficient pooling than simple textbook models – which emphasise the problems of adverse selection – make out.

MORAL HAZARD

The second consequence of asymmetric information is no less troublesome than adverse selection – indeed, there are important respects in which its costs can be even higher.

People who are insured alter their behaviour. The insurance itself means they have less incentive to avoid the contingencies that entitle them to the benefit; and once one of those contingencies has occurred, they have less incentive to mitigate its effects. This is the phenomenon referred to in insurance as “moral hazard”.

Like adverse selection, moral hazard arises from the difficulty of observing and controlling strangers. When I self-insure, I have no incentive to increase the risk to which I am exposed, either by making bad contingencies more likely to occur or by failing to reduce their consequences should they occur. Equally, when I am insured by my family, the tightness of the links that characterise family life is likely to make it costly for me to act in ways that increase our risk exposure. But when I am in an insurance pool with millions of others, the costs to me of acting in ways that increase risk for the pool as a whole may be very slight, or even negative. In those circumstances, I may well act in ways that impose substantial net costs.

Unemployment insurance, for example, may reduce my commitment to staying in a job, and lower the diligence with which I attend to its requirements, making it more likely that I will lose my job and take up the insurance – a form of moral hazard “before the event” (often referred to as *ex ante* moral hazard). Equally, once I am unemployed, the fact of unemployment insurance may reduce the time I invest in seeking another job (as against lying in bed or going to the beach), or make me more choosy than I should be in considering the jobs on offer, prolonging the period for which I receive the insurance payments (a form of moral hazard after the event, or *ex post*). Both *ex ante* and *ex post* moral hazard impose costs on the insurer (in the form of higher payments, which then require higher premiums), on other insureds (whose premiums must rise to cover the greater payments made to me) and on society as a whole (which suffers from the loss of output while I am unemployed, as well as from any distortions associated with raising the revenue to cover the transfers I receive). Those costs are likely to greatly exceed whatever pleasure I gain from spending more time lying in bed or going to the beach.

Moral hazard need not only alter my behaviour – it can also alter the behaviour of those who provide me with services. This is especially likely when the insurance payments are relatively open-ended.

If my house burns down, or my car is wrecked in an accident, my insurance will pay me a fixed amount. Should I fall ill, on the other hand, insurance will typically cover some part of the costs of my treatment and rehabilitation. Those costs, however, depend on a wide range of factors, and are not generally capable of being determined in advance of my episode of illness occurring and being dealt with. The payment I receive is therefore at least somewhat open-ended.

This open-ended nature of the payment then creates scope for moral hazard *both* by me and by those providing me with services. In effect, so long as compensation to those service providers exceeds their additional cost of providing the service (that is, the price a supplier receives exceeds its incremental cost of supply), the providers have an incentive to “over-service”. Even when suppliers do not have a financial interest in over-servicing, they may be induced to do so (for example, in the form of prescribing unnecessarily expensive procedures) if they are not exposed to the financial consequences that entails.

The fact that insurance removes or reduces the incentive for me to shop around can accentuate the extent and impact of this “supply-side” moral hazard. For example, why would I travel a bit further to access a more efficient health care provider, if all of the benefit of my so doing will accrue to the insurer? Knowing that, service providers have incentives to charge more, give less or both.

Moral hazard tends to be more problematic than adverse selection for two reasons.

The first is that adverse selection is to some degree offset by beneficial selection – the tendency of especially cautious people, who are the best risks, to especially want to be insured. However, there is no offsetting tendency for people, once they become insured, to want to act more cautiously as a result. The fact that insurance reduces the cost of being a bit less careful to all insureds affects all insureds, regardless of how good or bad their risks would otherwise be – just as a reduction in the price of bread increases the demand for bread from all consumers, regardless of how much or how little bread they consumed before the price fall.

This leads to a second point which is that moral hazard alters behaviour in the entire pool, while adverse selection tends to impose costs at the margin. As a result, the efficiency costs of moral hazard may be high relative to those imposed by adverse selection.

THE ROLE OF GOVERNMENT

In short, insurance is never perfect: it is either too narrow, as in self-insurance, imposing costs in the form of unduly concentrated risks (a kind of adverse selection), or too wide, as in large pools grouping millions of anonymous participants, creating costs in terms of moral hazard and more generally slackened effort. These inescapable imperfections form the backdrop against which Quiggin's key claims need to be assessed.

The problem in assessing these claims is that Quiggin does not state them with any precision or specificity. He argues as if each proposition was unbounded – government is *the* least-cost risk bearer for *all* risks; in covering those risks, it should cover them *entirely*; and it should be the one that not only funds but also provides *all* the services meeting these risks involves. To demolish, indeed ridicule, these propositions would be easy: they are caricatures of arguments, rather than arguments as such. However, such a criticism would not engage with the difficult task, which is to specify the limiting principles that Quiggin's analysis fails to recognise, much less set out.

In discussing those limiting principles, I proceed as follows. I start by considering Quiggin's claim that governments are generally likely to be the least-cost risk bearer; I then turn to his claim that in taking on the relevant risks, governments should make the coverage of those risks far more generous and comprehensive than it now is; and finally, to the claim that governments should not merely fund but also provide the insurance services.

GOVERNMENT AS THE LEAST COST RISK BEARER

In considering whether governments are generally the least-cost risk bearer, it is useful to start by noting that for private insurance to work at all, there are a number of conditions that must be met. Three of these conditions, which are commonly referred to as conditions of insurability, are especially important, as they each create circumstances which have been commonly used to justify government involvement.

The first is that there must be a verifiable insurable event. All this means is that it must be possible to define and distinguish the situations in which I pay a premium and do not receive payment from the insurer, from those situations in which I do receive such payments. The "good" and "bad" states must, in other words, be capable of being determined, and of being determined at reasonable cost.

This last clause is important because many premiums are small while individual payments by the insurer may or may not be. If high outlays – in terms of work by assessors and expenses incurred in the settlement of disputes – were needed to verify whether “good” or “bad” states of the world prevailed, then the resulting costs would likely make the price of the insurance exceed its value. Extreme cases in which this would occur are readily imagined – for example, an insurance against unhappiness would not be feasible, given the difficulties that determining whether someone was or was not unhappy would inevitably involve. But the same holds for some more relevant contingencies, such as pains and disabilities associated with muscular-skeletal conditions.

A second requirement is that the insurable event be uncertain. If it is certain that I will experience an adverse event, then the insurance premium required is simply the value today of the payment that will need to be made when that event occurs. Insurance here is merely a form of saving, perhaps assisted by some subsidies or tax concessions which make saving in this form more attractive than putting money in a bank. Moreover, even if events are uncertain, but the likelihoods of their occurring are high, then insurance tends to be unattractive, because of the high administrative costs it involves relative to those involved in other forms of saving. As a result, insurance through the pooling of many individual risks works best when each of the individuals in the pool faces a low probability of the “bad” event occurring.

This can create significant challenges. These typically arise for individuals for whom it is highly likely that they will experience an adverse event against which they would ideally like insurance. For example, there are medical conditions that, thanks to developments in genetics, are capable of being diagnosed well ahead of any capacity for treatment, making those most at risk effectively uninsurable. Additionally, there are some chronic conditions, such as mental illness (which can develop early and last for life), for which insurance cover is difficult to obtain once they have first occurred, as there is little uncertainty about the likelihood of future costs being incurred.

Finally, a third requirement is that it be possible to find a large number of individuals whose risks are uncorrelated, in the sense that it is highly unlikely that they will all eventuate at once.

When my bad luck and that of strangers is highly correlated, pooling does not share risks but adds them up, making the total amount the pool would have to cover, should the undesirable contingency arise, very large indeed. For the pool to be able to do so, it would need to hold very large amounts of capital in reserve, forcing premiums to levels that may be too high to attract the revenues required. This is much the same problem as occurs with insurance within the family, but here it is the wider group that cannot afford to collectively self-insure.

There are types of risks for which likelihoods of occurrence are correlated between individuals. Mass epidemics, large scale fluctuations in the macro-economy and global events such as climate change are all cases in point. If I come down with a highly infectious disease, it is likely that many other people will too; if there is a severe downturn in the world economy, it is likely to affect us all; and were significant climate change to occur, it would likely affect the entire planet. These are not risks where “bad things” might happen to me without also happening to you. The fact that risks of this type are so difficult to estimate in a statistical sense makes the difficulty of covering them all the greater (as it makes the reserves that need to be held against the possibility of many insured risks simultaneously going bad larger and hence more expensive).

Relative to each of these requirements for insurability, governments may have advantages which could allow them to provide insurance where it would otherwise be unavailable, or to provide that insurance at lower cost.

To begin with, governments have powers of compulsion that can help deal with problems of verifiability. For example, it would be difficult for a private insurer to be as invasive as governments routinely are in determining whether people qualify for welfare assistance. Additionally, and more innocently, there are some cases where governments may be especially well-placed to determine insurable events. Often, this arises because that determination is intrinsically bound up with some other activity of government.

For example, the administration of the tax system means it is likely significantly easier for government to identify whether former students have reached particular income thresholds, and can hence pay back some of the costs of their education, than it would be for private lenders or insurers. Equally, the administration of the criminal justice system can make it easier for governments to identify and help compensate victims of crime. In these cases, combining an insurance function with those wider administrative processes helps achieve “economies of scope” (the gains that come from undertaking activities jointly) and hence can be efficient.

Second, governments may be more readily able to insure “high probability” risks that would otherwise be effectively uninsurable.

Consider, for example, those medical conditions that are essentially chronic, in the sense that once they present, they are very likely to recur. Making these conditions insurable almost always involves an element of redistribution, at least once the condition is discovered, and hence moves from the sphere of private and voluntary insurance to that of social insurance. (In private insurance, each individual’s expected income is left unchanged; in contrast, social insurance transfers expected life-time income between individuals, in the sense that it alters the income different people can expect to earn over their full span of activity). As such, there is inevitably a degree of government involvement, as there must be some cross-subsidy from those who pay for the insurance to those receive it. That cross-subsidy would come unstuck in a purely voluntary system and must be effected through regulation or through the tax system.

Third and last, governments may be able to pool risks that private insurers cannot. That ability arises from the power governments have to tax, which allows governments to effect two forms of mandatory pooling.

In the first, the power to tax allows governments to coerce people into insurance pools that cover the population as a whole.

For example, I may not believe I stand much risk of being unemployed, and hence would not take out voluntary unemployment insurance, especially if the premium does not reflect my own low risk. However, the government can (and does) compel me to do so. Additionally, using that power of compulsion, governments can make an insurance scheme work even when that scheme could not survive were membership voluntary. What this means is that the insurance can be redistributive, in a way that would unravel in a voluntary arrangement. In particular, low risks can be forced to pay premiums (in the form of taxes) that are higher than the cost they impose on the insurance pool, while high risks pay less than the costs they impose. Adverse selection would unwind the resulting cross-subsidy in any voluntary scheme, as the low risks left the scheme – but exit is obviously more difficult when the cross-subsidy is being enforced through taxation.

Second, the power to tax can allow governments to deal with situations where risks are highly correlated at one point in time, but are not so correlated over time.

Consider, for example, a catastrophic contingency, such as an outbreak of a highly contagious disease. While such an event has the potential to severely reduce incomes in the current generation, it is likely that incomes would eventually recover. By engaging outlays that would ultimately be recouped from future taxes, government could effectively pool this generation's risk with the risk of generations in the future.

In short, governments can have advantages in dealing with each of the conditions of insurability. However, while these advantages can be significant, they are by no means unbounded.

In seeing this, it is useful to begin by noting that private insurers have strong incentives to manage risks efficiently and profitably. One result of this is that they excel in developing and managing risk pools that are extraordinarily wide and deep.

Historically, of course, insurance pools were not so comprehensive, especially across national borders. Rather, it was governments that had the most diversified portfolios, as they derived their revenues from a national economy as a whole. This meant that they could spread risk more widely than could nationally-based insurers, including not merely over the economy at a particular point in time, but also from previous generations of activities (by accumulating savings from those activities' revenues) and from future generations of activities (by borrowing against the revenues those activities will generate).

Today, in contrast, private insurers and reinsurers operate in highly integrated global capital markets. Within those markets, they spread risks over many national economies, using instruments that tap a very wide range of funding sources. Moreover, they spread those risks to parties who choose to bear them, and hence can organise their affairs to minimise the costs holding those risks entails.

As a result, it seems likely that the real advantage of governments today is less in the breadth of the risk-spreading than in the ability that governments and governments alone have to tax – i.e. to compel involuntary membership of insurance pools, and use that involuntary membership to set premiums that are redistributive.

This is important because the power to tax comes at a high cost.

Taxes alter, and usually distort, the incentives to work, save and invest. They lead us, for example, to spend less time in paid employment and more time in activities that are untaxed, such as leisure or home improvements. They reduce our willingness to save, as saving is still taxed more heavily than consumption. And they make us less willing to invest in risky ventures, as the higher incomes we would obtain should our ventures come good are taxed far more heavily than any relief we obtain when such ventures do poorly. All of these mean that our society is less efficient, and hence poorer, than it could be.

Measures of the extent of these costs of taxation – the “excess burden”, as economists call it – vary, but most public finance economists regard as reasonable a figure of twenty to thirty cents in the dollar at low to moderate levels of taxation. Each dollar raised by taxation, in other words, causes distortions that cost society twenty to thirty cents. Moreover, the costs per dollar of taxation would be even higher were tax rates to rise significantly, as the unit costs of taxation rise more than proportionately with the tax rate.

Such an increase in tax rates would be needed were the government’s role as insurer to be significantly expanded. Government would, in particular, need to extract higher taxes so as to fund the increased benefits that were being paid to net beneficiaries. Those higher rates of contribution would induce taxpayers to take evasive action. They might, for example, work less, or retire early, or move themselves and their businesses overseas, thus entirely opting out of the risk pool.

All of these responses are forms of adverse selection: the good risks leave insurance pools in which they have to subsidise bad risks. To do so, they impose costs (for example in the form of reduced productive effort) that offset some – perhaps all – of the efficiency gains the increased insurance was aimed at providing. The fact that in most government schemes one can stop paying premiums (say, by retiring early or otherwise reducing one’s taxable income) without foregoing benefits makes the extent of this response all the greater.

As a result, it is simply incorrect to believe – as Quiggin seems to – that governments can avoid or preclude adverse selection: it simply bites them in a different form. Moreover, unlike private insurance – where adverse selection is to a degree offset by beneficial selection – there are no taxpayers who choose their activities more efficiently (for example, in terms of the balance between work and leisure) when faced with revenue-raising taxes. The entire reaction to the required tax rates is therefore an efficiency loss.

To these efficiency losses must be added those from moral hazard.

Governments do not have a magic bullet that allows them to avoid moral hazard, any more than they have one that allows them to avoid adverse selection. Rather, the best that can be said for governments as managers of moral hazard is that they can bring draconian measures to bear when and where they choose to do so.

That governments can do so is obvious from the use of police powers to prevent alleged “roting” following adverse media publicity. Voluntary insurers could not (and in practice do not) so readily seize the financial records, monitor the homes, and tape the conversations of alleged cheats. This greater ability and demonstrated willingness of governments to use coercive means may make it easier for them to curb fraud, albeit at a potentially high cost to civil liberties.

But this “advantage” governments have (if it can be called that) in dealing with moral hazard is more than offset by the fact that governments do not have the same binding constraints on their expenditure as do private insurers. Rather, government budget constraints are often “soft”, as spending levels can be adjusted to accommodate political pressures. Moreover, taxes are fungible, so that the structure of spending can also be shifted as political pressures dictate, in ways that a voluntary insurer (who after all has sold a particular service and hence must deliver that service) could never do. These looser budget constraints, along with the absence of a profit motive, weaken the need and the incentive to control moral hazard.

Experience suggests that the balance that emerges from the scope to use draconian powers on the one hand, and soft budget constraints on the other, can be very perverse indeed. In that balance, what seems like moral hazard by the intended beneficiaries of the insurance is treated exceptionally harshly, while moral hazard by politically powerful suppliers is accommodated, if not encouraged. Rather than the aggregate costs of moral hazard being reduced, their distribution is shifted from the weak to the strong. The weak are the intended beneficiaries of the insurance – the ill, the poor, the victims of crime. The strong are those who are supposed to supply them with services: the hospital staff; the teachers; the police – and the bureaucracies responsible for each of the insurance services provided. The coercive powers of the state are used to frighten and punish the first, while soft budget constraints are used to try to placate the second.

This is a peculiar but recurring version of the “nanny state”: one that flogs the children, while treating the nannies to oysters and champagne. But however popular this version of the “nanny state” is with the nannies, it is inefficient on at least two counts.

First and most obviously, accommodating supplier moral hazard unnecessarily increases costs, which is inefficient in itself and causes added inefficiencies as it feeds through into higher (and hence more distorting) taxes.

Second, it is questionable whether the State’s coercive powers actually do much to reduce moral hazard – and hence avoid its wider economic costs – among the direct beneficiaries of the public insurance.

In effect, the bulk of moral hazard does not involve fraud. It is far more difficult than fraud to define and detect, as it involves subtle shifts in behaviour that are not easily shown to be culpable conduct: such as being less diligent at work because of access to unemployment insurance, or being less vigilant about diet and exercise given subsidised access to medical care. Short of a police state, these forms of behaviour are not readily controlled.

In brief, whatever advantages governments have in dealing with the preconditions for insurability are accompanied by substantial and inescapable costs – the costs of the taxes needed to fund the outlays; and the costs of moral hazard by suppliers. These costs are all the more substantial when the insurance provided by government displaces reasonably (though obviously not perfectly) efficient private suppliers of insurance services.

This does not mean governments are *never* the least-cost supplier of insurance; but it does mean that they are only likely to be that least cost supplier in exceptional conditions.

In particular, the benefits of government action are only likely to exceed the costs in three sets of cases: where governments are uniquely well-placed to determine whether an insurable event has occurred because they can combine that task with some other part of their ordinary operation (as in HECS); where the risks involved are otherwise uninsurable because they lack the requisite degree of uncertainty (as in recurring mental illness); or where the risks are of such a scale, and so difficult to predict statistically, that even the widest and deepest capital markets cannot pool them efficiently (as in terrorism insurance).

THE EXTENT OF INSURANCE PROVISION

This leads to the second of Quiggin's claims: that in taking on all the risks for which he believes governments are the least-cost insurer, governments should make coverage against those risks comprehensive. As with his other claims, this one is not stated at all clearly; but what is clear is that he does not set out any form of test that would limit the extent of coverage. So here too, the claim is made as an unbounded truth.

However, stated as such an unbounded truth, Quiggin's claim is plainly incorrect, at least as far as efficiency is concerned. In effect, the existence of moral hazard means that complete insurance coverage is *never* efficient. The insured themselves must always make some contribution to costs, so as to have an incentive to keep risks under control. That is why voluntary insurance schemes involve threshold and deductibles, while most government-provided schemes have an element of co-payment. The greater the likelihood of inefficient moral hazard, the greater those deductibles and co-payments need to be. Removing them, as part of an across-the-board increase in benefits, would not promote efficiency, but rather reduce it.

Moreover, the tax costs of such an across-the-board increase in benefits could be very substantial. The increased taxes required to fund that increase would have two effects.

The first is a once-off reduction in wealth. This is simply the flip-side of the shift of effort away from productive, taxable activities to other activities, such as leisure – even though those other activities are in fact less valued than the activities that are being cut back.

The second is a reduction in the rate of growth of wealth. For example, a smaller economy may be less able to achieve economies of scale in innovation and in other activities that cause productivity to increase over time. Even more importantly, an economy in which strong “insiders” in the labour market feel less exposed to unemployment – because social benefits are generous and open-ended – is likely to experience stronger inflationary pressures during cyclical upswings. This would make growth episodes shorter and recovery slower, as those who are out of work take longer to accept new jobs. The combined effect would be economic growth that was less rapid on average and less sustained.

It is likely that the costs of these reductions in income would fall most heavily on those with relatively low incomes. They are the ones who are most vulnerable to unemployment; and they are the ones who would suffer most from the inevitable contraction in benefits, when that contraction ultimately came.

Of course, even taking that on board, there may be areas where the current extent of coverage might not be quite right. In some of those areas, the benefits of somewhat greater coverage might exceed the costs. However, that conclusion (that greater insurance would be worth the added cost) could not properly be drawn without a much deeper, more empirically informed, analysis of each of the areas Quiggin glances at – health, education, income support – than Quiggin provides.

These are hardly areas that have been unexplored. In health care, for example, government insurance can help overcome adverse selection by making membership of the insurance pool compulsory; but it does so at the cost of increased moral hazard, both by insureds (especially if the insurance covers forms of treatment that are largely elective) and even more so by service providers. Equally, in education, supply-side moral hazard (such as poor teaching) is difficult to avoid in schemes that do not provide subsidies through fully transferable vouchers (and even in some voucher schemes). And as for income support arrangements, they invite both adverse selection and moral hazard, in extreme cases encouraging the intergenerational transfer of “welfare dependency”.

Given these impacts, it is by no means obvious – and certainly is not demonstrated by Quiggin – that we should provide *more* insurance across the board than we now do; there are some areas (such as mental health) where that may well be true, but in many others (such as support to firms in difficulty, assistance to single mothers and “sit-down” payments to Aboriginal communities), there may be a compelling case for less.

As a result, taken as a general proposition, Quiggin’s claim that government-provided insurance against economic and social risk should be comprehensive is simply wrong. Rather, it is important to provide some incentive for insureds and for suppliers to limit the extent of moral hazard. It is also important to keep the costs of taxation to manageable levels. Together, these factors make the presumption against any large-scale expansion in provision all the stronger.

SUPPLY AS WELL AS FUNDING

Finally, we come to the third of Quiggin’s claims – that governments should not only fund protection against economic and social risk, but also directly provide the services, such as education, health care and welfare support, through which that protection would be made available.

Quite why Quiggin believes this is never made at all clear. As best one can tell, he thinks that non-government supply undermines the efficiency of government supply. Put, in economic terms, this amounts to the claim that supply of the services at issue is a “natural monopoly”, that is, a situation where supply by a single provider incurs lower costs than supply by two providers or more.

The more sophisticated form of this argument, though not well put by Quiggin, can be exemplified in the case of schools. Non-government schools, it is said, “cherry pick” the state schools, taking the students who come from the best homes or who for other reasons are more able and ambitious. The result is not only to worsen the pool that is left behind, but also to deprive the remaining children of the better children as classmates and mentors. A vicious spiral can set in, in which what might otherwise be good performance in the state schools becomes ever worse. At the end, those children who have no alternatives are left with a poor education and limited life-chances. That limitation in life-chances is not only unfair but also inefficient – if the children who were left behind could borrow at actuarially rates to finance their education, the gains from that education would outweigh the costs incurred.

It is difficult to dispute the proposition that assortments matter. Who we study with, play with and more generally get to know doubtless affects the quality of education and our life chances. But it is an empirical issue, and one on which Quiggin provides no evidence, whether in practice non-government schools cause that mix to worsen. However, even if they did (and I do not believe they do), the suggested cure seems wrong-headed.

In effect, analytically, what happens when “opting out” leads to a worsening in the quality of the remaining pool is not very different from adverse selection generally. The good risks leave, because in doing so they secure benefits that exceed the costs they bear; but in the process, they impose some cost on those who, for whatever reason, cannot leave the pool. The most efficient way to deal with this is not through a quantitative restriction – a prohibition on exit – but by ensuring the price signals are correct, including by removing the obstacles that currently trap some children in pools which (if they had the means) their parents would not have chosen. While not without implementation problems of its own, a fully portable voucher scheme, where the value of the voucher is properly calibrated to take account of mix effects and socio-demographic factors, offers the best hope of achieving this goal.

The solution, in other words, is not to eliminate competition but to ensure that it is fair and neutral – i.e. allows all forms of supply to compete on their merits.

This is especially the case given the very high costs that restrictions on exit would impose. Most immediately, they would over-ride the educational preferences of, and likely materially reduce the educational quality obtained by, those parents and children who had opted for non-government schools. There is no reason why the preferences of these parents and the welfare of their children should be given a lower weight than those of other parents and children.

Additionally, such a strategy of deterring exit from government supply would only further increase the risks of supplier moral hazard. Given an entrenched monopoly over public funding, the incentives for state schools to focus on providing quality education would be weakened, at least to the extent to which those incentives depended on the threat of losing students to alternatives. At the same time, the market power of school administrators, teachers and the state educational bureaucracies would be enhanced, making it even more likely that costs and quality would prove difficult to control.

Last but by no means least, entrenching government providers would reduce and restrict the scope for diversity and innovation. This is of obvious importance because there is still so much scope for improvement in the way services such as health and education are provided. These improvements are gains in dynamic efficiency – that is, they give us the ability, over time, to be more productive. These gains are likely to be more important than all the other forms of efficiency combined, and yet it is they that would be most directly threatened by giving government providers the exclusive rights to state funding.

As a result, quite in contrast to Quiggin's argument, there is a compelling case for moving to a system in which voucher-type schemes dominate, and governments, when they fund services, allow full and open competition in their provision. Obviously, there are many important design issues involved in developing efficient voucher schemes; but the benefits they can offer are clear. Competition would impose more direct disciplines upon suppliers, as the failure to provide "value for money" led to their being displaced; it would increase scope for innovation, which is crucial to addressing issues such as the treatment of chronic illnesses; it would make the government less vulnerable to "hold up" by monopoly suppliers and their unions; and importantly, it would give greater dignity and choice to the beneficiaries of the services the government is funding.

This also has implications for each of the areas discussed above where governments may have advantages in dealing with the conditions of insurability. More specifically, it may be far wiser for governments to consider themselves primarily as reinsurers, rather than directly stepping into insurance functions.

For example, with risks that are difficult to insure because they are not sufficiently uncertain – such as risks associated with recurrent chronic diseases – it would be more sensible for governments to fund the carriage of those risks by competing private sector suppliers, than for government itself to control the case-management. Equally, with risks that are difficult to insure because of their correlated or catastrophic aspect, it is the reinsurance of the risks that is required, rather than the definition and marketing of policies and the management of claims. As a result, the efficient scope of the "insurance state" would turn out to be far more limited than Quiggin contends.

CONCLUSIONS

The efficient management of risk in society is an area where there are no solutions, only trade-offs. Those trade-offs arise because there are costs, as well as benefits, from expanding any form of insurance, and especially so for insurance that is government provided. The problems that insurance faces – problems that include adverse selection and moral hazard – cannot be wished away, and afflict government insurance in forms that are generally no less virulent, and are in some respects more virulent, than those which beset insurance in voluntary, private markets.

When account is taken of both the costs and the benefits, and of the strong pressures for efficient operation that bear on private providers, there is no reason to conclude that government is generally the least-cost risk bearer – indeed, in many instances, the opposite will be true. Equally, there is no reason to think that a substantial increase in the scope and generosity of public provision would increase efficiency – as a general proposition that is at best unproven, at worst manifestly incorrect. And finally, the case for enforcing government *de jure* or *de facto* exclusivities over service supply is completely unconvincing.

All of this should not be vastly surprising – rather, it would seem to capture the consensus view in Australian economic policy for a decade or more. That too is unsurprising, as that consensus view is founded on some solid economics. What is surprising is that Quiggin should have come to conclusions that depart from it so strongly.

He does so because he completely ignores the costs his proposals would involve. The distortions associated with taxation do not even get a passing reference; as for moral hazard, it is noted and dismissed without explanation or justification. There is, of course, no easier way to make a case: ignore all the factors that might tell against it. But while that is easy, it is hardly what sound policy analysis is about.

Ultimately, for Quiggin, the economics are merely the road to politics – the continuation of political argument by other means. That is the wrong way around. An identification of the costs and benefits, and the mechanisms by which they operate, should come first; the political choices, over means and over ends, can then build on that.

That does not, however, exhaust the problems with Quiggin's approach to the politics. Rather, there too, the difficulty lies in knowing what it is – other than an abhorrence of John Howard and a "Nightmare on Elm Street" view of the United States – that he is actually holding out to the community. What is that he is promising? Is it that in his revitalised social democracy, government will remove the risk individuals face?

If that is the promise, then it is a promise falsely offered. Government can no more deliver on it than it can hold back the tides or move the moons. To claim otherwise is to promise only ultimate disillusionment, with much disappointment along the way.