

Wellbeing, living standards, and their distribution¹

David Gruen, James Kelly and Stephanie Gorecki

Australian Treasury

9 September 2011

¹ Address by the first author to the New Zealand Treasury Academic Guest Lecture series. The authors are grateful to Thomas Abhayaratna, Gerry Antioch, Robert Gardner and Anthony King for helpful comments.

Introduction

It is my great pleasure to be here today in Wellington to talk to you about wellbeing and living standards.

I will start with a brief overview of the Australian Treasury's wellbeing framework and then make some comments on the recently released living standards framework of the New Zealand Treasury, including highlighting some of the differences.

I would then like to focus on a point of commonality — an expressed concern for distributional outcomes — for which such frameworks have the benefit of encouraging us to come out of our comfort zones.

Some observations on the Australian and New Zealand frameworks

The Australian Treasury developed a wellbeing framework about a decade ago to provide some guidance about its mission, which is to improve the wellbeing — rather than living standards — of the Australian people.² It also identified elements that need to be considered in providing thorough analysis, and are particularly relevant to our work.

From an institutional perspective, the process of developing the framework was important in itself, as it required Treasury to think carefully about what people value, and how this relates to policy analysis and advice. Since late last year we've been talking to staff within the Australian Treasury on their experiences with the framework and considering what updating if any is required. Following the consultations, we made some changes to the

² Treasury's mission is to improve the wellbeing of the Australian people by providing sound and timely advice to the Government, based on objective and thorough analysis of options, and by assisting Treasury Ministers in the administration of their responsibilities and the implementation of Government decisions.

framework, to address specific issues raised by staff and reflect more recent developments in the Department's work.

Our framework continues to reflect our view that in addition to income and (material) consumption, a policy relevant assessment of wellbeing also depends on persons' health status, how well-educated they are, the quality of social relationships, and a myriad of other aspects of life that people have reason to value.

Without going into too much detail, the five dimensions given prominence in the Australian Treasury's framework are [slide 2]:³

- The set of opportunities available to people.
- The distribution of opportunities across the Australian people.
- The sustainability of opportunities available over time.
- The overall level and allocation of risk borne by individuals and, in aggregate, the community.
- The complexity of the choices facing people and the community.

Some of these will be familiar to you, as they also appear in your living standards framework in one form or another. One change we have just made, that reinforces this commonality, is to explicitly include sustainability and describe that from a stock perspective. Unsurprisingly, though, there are also some clear differences between the two frameworks.

³ Australia Treasury, 'Strategic Framework 2011-12', p2. For further detail see Banerjee, S and Ewing, R, (2004) 'Risk, Wellbeing and Public Policy', *Economic Roundup*, Winter, pp21-44. For further detail on the framework prior to changes see Australian Treasury (2004), 'Policy Advice and Treasury's Wellbeing Framework', *Economic Roundup*, Winter, pp1-20.

One obvious difference is the greater explicit emphasis the wellbeing framework gives to risk and complexity. One reason why we include these criteria is they directly relate to important aspects of our responsibilities. I am personally concerned with macroeconomic risks. And those parts of the Australian Treasury dealing with taxation and market regulation need to take account of the complexity they are directly responsible for, and they also play a role in our management of risk or otherwise influence choices about risk.

But as importantly, we are concerned about individuals' capacity to make good choices from the set of opportunities available to them — and for this both well-functioning markets and well-developed personal capacities are required.

Perhaps a more intriguing difference between the two frameworks concerns the approach to measurement of wellbeing. Conceptualising and measuring wellbeing are distinct endeavours, and, in this regard, the Australian Treasury's goal in its framework has been modest: merely to bring out our understanding of wellbeing and to identify things that are important in the formulation of public policy advice.

This is because since its development, the framework's intended role has been to provide a broad context and high level direction for policy advice, and not to provide a checklist to be applied in every circumstance, which might result if a list of measures is prescribed. As part of our review of our wellbeing framework, we found however that the intended purpose and usage of the framework has not been clear to many staff, and that some wanted a tool that would deliver concrete answers to the policy questions they encountered.

The New Zealand framework similarly states that it is intended to be used as an input to the policy process, rather than an analytical tool. But from the living standards paper, it is

evident that your framework is more ambitious by more directly addressing measurement related issues. For example, in its direct recognition of subjective wellbeing measures, and the extended coverage of measures. Related to this, the detailed discussion of the different stocks can give a reader a strong sense of policy directions.

So it will be interesting for me to see how you find staff use, or want to use, your living standards framework. Based on our experience, being clear about how you expect it to be used and getting a sense at some point on how staff are travelling may be worthwhile. It would then be interesting to compare notes.

A shared concern for distributional issues

If we look at the wellbeing and living standards frameworks together, it is striking that many of the dimensions identified concern the distribution of material and non-material things, to use the language of the living standards paper.⁴

In particular, we can be concerned with distribution:

- within a generation (and over that generation's lifetime);
- between generations, which we both label 'sustainability'; and
- between uncertain or contingent states, which in the Australian framework alone is separately identified as 'risk'.⁵

⁴ New Zealand Treasury (2011), 'Working Towards Higher Living Standards for New Zealanders', New Zealand Treasury Paper 11/02.

⁵ Identifying 'risk' as concerning distribution may seem odd to some. We can think of risk in a number of ways, but consider, for example, that a person's decision as to whether to take out health insurance is a decision as to whether to distribute resources (premiums) from their healthy state to their unhealthy state (payouts). Alternatively, consider the Rawlsian original position, in which a veil of ignorance generates uncertainty as to status and resources that will be available to people.

For the rest of my presentation I'd like to focus my attention on the more conventional, intra-generational, distributional issues. My reason for doing so is that despite the formal emphasis we place on distribution in thinking about wellbeing, or living standards, it is all too easy to be reticent about explicitly discussing or advising on distributional issues.

We have found that one benefit of having a wellbeing framework has been to emphasise concern over distributional issues, and encourage internal debate about them.⁶ So I was struck by page 27 of the living standards paper, which states: [slide 3]

'Where normative approaches ask what the distribution of living standards should be, positive approaches ask what the distribution is. ...

Treasury takes a positive approach to distribution as opposed to a normative, value-based one. ...'

Here, it seems to me, is reticence proclaimed a virtue. As the living standards paper notes, there are many different theories of distributive equity. So there is good reason for being cautious about being lost in the normative jungle.

But there is no avoiding that jungle. Among other things, some normative theories are built into much of the economic literature on which we rely as policy advisers. The optimal tax literature is one obvious example, built on welfarist, predominantly utilitarian foundations. I will give an example of this later on.

⁶ For a public example of this debate, see Henry, K (2007), 'Addressing extreme disadvantage through investment in capability development', at <http://www.treasury.gov.au/contentitem.asp?NavId=&ContentID=1327>, and Henry, K (2009), 'How much inequity should we allow?' at <http://taxreview.treasury.gov.au/content/Content.aspx?doc=html/speeches/05.htm>.

And, as with other choices that communities, governments and parliaments ultimately make, policy advisers can seek to improve the quality of the public debate and decision making.

Making sense of distributional measures

It also does not seem possible to me to simply communicate ‘what the distribution is’ without engaging with normative issues.

Implicit in the distributional measures that we see are normative judgements about *among whom* it is that we’re concerned, such as households within a country. There are also judgements about *what* it is that is of distributional concern. Typically, the focus is on cash income or wealth. Intimately related to the latter are judgements as to the *appropriate distributive rule or concern*. This can be a concern over inequality or over absolute or relative poverty. [slide 4]

Many of the distributional measures have been around some time, and are reported on frequently. For example, charts showing income decile comparisons, gini coefficients, and relative poverty are in the Annex to the living standards paper. Taken as a whole though, those that focus on equality or inequality seem to have had little direct impact on policy advice and government decisions — at least in Australia. Measures of adequacy seem to have had greater impact — and as should become clear, I think there are good reasons for this.

But without pretending to know why in practice equality focussed distributional measures have had less impact, it could in part be a function of three related things: [slide 5]

- the ‘what’ that is of concern is measured too narrowly;

- there can be a disconnect between these measures and the reasons why we may be concerned about distribution; and
- there can be a disconnect between these measures and some common ideas of fairness.

Growing Unequal?

The OECD's *Growing Unequal?* report of 2008 illustrates these points, notwithstanding that it is an impressive piece of work.⁷

The report looked at income inequality and poverty in OECD countries. It focussed primarily on measures of inequality of disposable cash income and relative poverty for households. It also looked at the effects of taking account of consumption taxes and in-kind benefits provided by government and some wealth measures.

Cash income as a measure of 'what' it is that is of distributional concern is quite narrow. For example, it does not account for either home production or the value of leisure. It focuses on annual income in preference to consumption (which deviates less over lifetimes) or a lifetime measure. Nor does it reflect the consequences for those unemployed of not having a job beyond the loss of disposable income.

The distribution of health outcomes is also not factored in: good health and long life would likely merit large weights in any assessment of living standards and quality of life. Further, the distribution of rights or liberties is not reflected. These have probably become more equal in recent decades, and, as the living standards paper suggests, are also likely to be important of themselves for wellbeing.

⁷ OECD (2008), *Growing Unequal?: Income Distribution and Poverty in OECD Countries*.

The distribution of self-reported subjective wellbeing is also not considered, whether mood happiness, life satisfaction, or other subjective states. Some would advocate doing so because they see positive subjective states as an appropriate objective of policy. Whether you accept that or not, it would be worth examining the distribution of subjective wellbeing for the reason suggested in the living standards paper — as a useful cross check against the narrow measures of resources used in the distributional statistics presented.

Why are we concerned with inequality?

Whether disposable cash income — either extended to include in-kind benefits or not — is too narrow a measure depends of course on what your objective or concern is. So why are we concerned about inequality?

The OECD gives a few reasons separate from just fairness. These included concerns over the sustainability of the polity, potential benefits from reducing resistance to reforms that are fuelled by concerns over inequality, and that inequality can directly generate unhappiness.⁸

Maximising a social welfare function

You probably won't be surprised to hear that a desire to maximise a social welfare function doesn't rate a mention in *Growing Unequal?*. Such an objective doesn't resonate with most people. But if you are not surprised, I would suggest you should be a little disturbed that it's not even mentioned in passing. It underlies much of economics, including that concerning

⁸ A separate argument is that income inequality has a negative effect on health, most recently popularised by Wilkinson, R and Pickett, K (2009). But it remains unclear whether this claim is supported by the evidence, for example see Deaton, A (2003) 'Health, Inequality and Economic Development', *Journal of Economic Literature* Vol. XLI, March, pp.113-158.

optimal tax-transfer systems for achieving distributional outcomes. And its central tenets have some persuasive proponents.⁹

Social welfare functions add up, or count in some way or other, a measure of the welfare of each individual in a designated group. Of course, many types of social welfare functions have been put forward, and there are widely differing views as to whether they are appropriate, what they mean, how they should be interpreted, and even whether they are meaningful.

If your welfare function is Rawlsian in form, so that you are concerned with maximising the utility or the resources of those with the least, equality per se is not relevant. Only the conditions of the people at the bottom are relevant, and the rest of the distribution ignored. Inequality or relative poverty measures will tell you little of what you need to know.

On the other hand, if you are a Benthamite utilitarian, then you want to maximise aggregate utility,¹⁰ the sum of individual utilities, and you consider utility to be ‘happiness’ simply understood.¹¹ If we assume that there is diminishing marginal utility of consumption, then a more equal distribution of income (including leisure) will help equalise the marginal utility of income and hence — subject to efficiency costs — maximise aggregate utility.

Note that the utilitarian’s concern is to equalise marginal utility, not levels of utility between persons. But if people are alike and have equal capacity to convert resources or income into

⁹ For example, see Kaplow, L (2008) *The Theory of Taxation and Public Economics*, Princeton University Press.

¹⁰ More typically today, utility is seen as reflecting the fulfilment of desires or preferences. But if so interpreted it is unclear how individual utilities can be added up or what average utility is — see Sen, A (1999) *Development as Freedom*, Anchor Books.

¹¹ One conclusion you can draw from the subjective wellbeing literature is that happiness cannot be simply understood — for example, is it mood happiness as it is experienced, or as you look back, or is it life satisfaction?: Fleurbaey, M (2009) ‘Beyond GDP: The Quest for a Measure of Social Welfare’, *Journal of Economic Literature*, vol. 47:4, pp1029-1075. And what of other subjective states such as a sense of purpose that also appear to matter to people in their own right: Benjamin, D, Heffetz, O, Kimball, M, & Rees-Jones, A (2010) ‘Do People Seek to Maximise Happiness? Evidence from New Surveys’, NBER Working Paper 16489.

things or functions that they value, and then to convert things that they value into happiness, then equalising income after tax and transfers will help equalise individual utilities as well as marginal utility. Again, this is subject to any trade-off with efficiency.

One-well known critique of this line of reasoning has been provided by Amartya Sen.

[slide 6]

People don't have equal capacity to convert income into things or functions that they value.

A blind person for example may require more income to do or achieve certain things.

Others may not have the education or skills to make the best use of the income they have. If you're concerned with maximising aggregate utility you may want to give such a person *less* income than others, as they may be an inefficient utility generator. But if you are concerned with equalising individual utilities you may want to give them more than an equal share of income or resources.

So Sen suggested that we look at people's capabilities: the outcomes or functionings they can achieve —such as good health and their ability to participate in society as well as consumption of goods and services. If you conceptualise income in a very broad way as both Treasuries do, then the distinction between income and resources on the one hand, and capabilities on the other, begins to disappear. Differences in people's ability to convert resources or income into something they value would be reflected in your measure of resources or income.

Sen is less keen on going the next step to focus on happiness or subjective well-being. But there can also be blurring on that side as well. If what people value ultimately is happiness or some other subjective state, then some may be less efficient than others in converting

resources, or capabilities, into happiness. It would seem strange, for example, to care about a person's mental illness if it negatively affects the tangible outcomes they can achieve, but not care if it leaves their capabilities unchanged but affects their mood happiness or sense of self.

As a general rule then, and as researchers in the field have long understood, broader measures of resources and income are called for than disposable cash income. While noting that it is a formidable task, the 2009 Stiglitz-Sen-Fitoussi report on the Measurement of Economic Performance and Social Progress argued that measures of the distribution of what it called 'full income' should be part of the research agenda.¹²

In essence, such a measure could be expected to reflect the underlying broad productive capacity of individuals or households, akin to the stock concept emphasised in the living standards paper. It is also akin to the optimal tax literature's concern with taxing an individual's productive capacity (and so avoiding disincentives to paid work), through finding proxies for that otherwise unobservable capacity.

Should we care about inequality because it is unfair?

Reasons for caring about equality can also be derived from overarching theories of justice. Rawls is well known for constructing a thought experiment to argue that fairness, or justice, requires the equality of rights and liberties and that we choose those institutions that maximise the resources available to those with the least.

So perhaps we should care about inequality of incomes because it is simply unfair or unjust.

¹² The report suggests full income would include leisure, household production and imputed rent as well as financial income and in-kind benefits: Stiglitz, J, Sen, A, Fitoussi, J (2009), *Report by the Commission on the Measurement of Economic Performance and Social Progress*, at <http://www.stiglitz-sen-fitoussi.fr/en/index.htm>.

This is a natural line of thought, and in explaining why we should care about inequality of cash income, *Growing Unequal?* places emphasis on social concerns (page 283). But other common views around ‘fairness’ are largely ignored in *Growing Unequal?*, or even viewed as adding to the negative consequences of inequality (page 131).

In 2003, the economist James Konow provided a survey of the empirical evidence — essentially surveys, vignettes and social laboratory experiments — on fairness.¹³ One of his goals was a positive analysis of normative theories of justice. That is, how the normative theories of say Rawls or of utilitarians, and also Pareto principles, stack up with the views of people in general.

To illustrate, consider the following stylised question that Konow records as being asked of 211, presumably American, respondents: [slide 7]

2. Jane has baked 6 pies to give to her two friends, Ann and Betty, who do not know each other. Betty enjoys pie twice as much as Ann. In distributing the pies, what is fairer:

A. 2 pies to Ann and 4 to Betty, or

B. 4 pies to Ann and 2 To Betty, or

C. 3 pies each?

[Slide 8]

¹³ Konow, J (2003) ‘Which is the Fairest one of All? A Positive Analysis of Justice Theories’, *Journal of Economic Literature* Vol. XLI, December, pp.1188-1239.

A, with more given to the person whose enjoyment will be greatest, is consistent with a utilitarian perspective: choice A will generate more aggregate enjoyment, or average enjoyment, than B and C. It received support from 40 per cent of those asked.

B, with more pies given to the person who enjoys them the least, corresponds with equalising enjoyment (or utility) between individuals. It reflects what is often called Sen's weak equity axiom. It received a positive response from only 4 per cent.¹⁴

So it was 'C', a simple two way split of the pies, that received majority support of 56 per cent.

Based on considering a wide range of other such examples structured to cast light on different perspectives on fairness, Konow set out what he saw as the three principles underlying what people consider a fair distribution. These were: [slide 9]

- *Need*, concerning the satisfaction of basic needs.
- *Equity* (or just deserts), concerning proportionality (for example, of income to work effort or risks deliberately taken) and individual responsibility.
- *Efficiency*, concerning maximising aggregate income, consumption or even subjective outcomes like enjoyment.

Reflecting the behavioural economics literature, he noted that context also matters, and sets the stage for the playing out of these competing principles. He also noted that people appear to trade off the three principles rather than giving precedence to one over the

¹⁴ Konow notes that other evidence supports Sen's weak equity axiom. In particular, he gives example where respondents support the unequal distribution of food to achieve equal health outcomes.

others. That is, people are generally pluralists, not monists or absolutists like some utilitarians I happen to know.

Strikingly, Konow argues that equality or egalitarianism is itself not an underlying principle. Rather he describes it as the default rule that people apply when they lack information on need, equity or efficiency. Perhaps it would be better thought of as a default principle, that is applied in the common situation in which information is often either incomplete, ambiguous or plainly not available this default rule becomes prominent.

Of the principles that Konow identified, thinking of equity in the sense of desert or proportionality, and as reflected in Robert Nozick's concerns for procedural fairness, is probably what economists — and the authors of *Growing Unequal?* — find most uncomfortable. We do of course see rewarding effort and risk as important from an efficiency perspective, but for that reason only.

Desert though features heavily in public discourse. Gregory Mankiw sees differences in the weights that people place on equality and just deserts as explaining major political fault lines in the United States.¹⁵ It also underlies arguments for equality of opportunity rather than equality of outcomes.

But as for equality where we always need to ask about 'equality of what', so we should ask 'fairness of what'. One way to interpret Konow's findings is that proportionality reflects a concern for inequality in respect of a broad measure of income, by taking account of leisure and risks taken — rather than just cash received or goods and services consumed.

¹⁵ Mankiw, G (2010) 'Spreading the Wealth Around: Reflections Inspired by Joe the Plumber', *Eastern Economic Journal*, 36, pp285–298.

So the point for me is much the same as before. Whether we're concerned about equality or community ideas of fairness, we have to continue to try to move beyond narrow measures of income such as cash income.

The Australian Treasury's perspective

So what is the Australian Treasury perspective on the distribution concern identified in its wellbeing framework? Related to that, and the issue I raised previously of how should you use such a framework, how do we go about operationalising it in our advice?

Needless to say, in a large organisation that has evolved over time there is no single answer to either of these questions. But there are distinctive features to how we have approached these issues that I suspect are shared by New Zealand Treasury.

The Australian Treasury's underlying concern

Our revised framework points to the Australian Treasury's underlying concern as relating to need in the sense of all in the community being able to lead a fulfilling life and participate fully in the community. Such a perspective would I expect have general community acceptance, and while it does not advocate more egalitarian or redistributionist policies that a government may be minded to implement, it does not preclude them either. [Slide 10]

Such a perspective or objective is of course concerned with more than just the distribution of cash income over a defined period.¹⁶ And it does not neatly fit in either the absolute or relative poverty camps: it treats both senses of poverty as relevant.

¹⁶ This is not to say such concerns should be irrelevant to public policy. Transfer systems will always have a concern for people who are very cash poor over even short periods of time, such as a fortnight, and face liquidity or borrowing constraints.

This view informed, for example, the recent review of the Australian tax and transfer system by a panel chaired by Dr Ken Henry and supported by a Treasury secretariat. In considering the appropriate level of assistance to demographic groups with little capacity to work or expectation that they should work, it suggested transfers be 'sufficient to provide an adequate standard of living, based on an accepted community standard.'¹⁷ Adequacy or need, rather than reducing inequality per se, was the concern.

In the living standards paper I find some resonance with this view in the statement that (page 28):

[New Zealand] Treasury's advice on distribution has tended to emphasise the inefficiencies that result from having living standards distributed in ways that prevent some people from fully participating in the economy and society.

It also finds some resonance in the *Growing Unequal?* report's indicators of material deprivation, identified as 'one area where consumption data hold special promise' (page 298). To quote the report:

today, most authors define material deprivation as 'exclusion from the minimum acceptable way of life in one's own society because of inadequate resources' (page 179).

By looking at whether people are able to achieve a given set of outcomes or functionings, and also avoid negative outcomes such as financial stress or poor environmental conditions, material deprivation takes a broader range of resources into account than cash income or wealth alone, including some factors affecting the conversion of resources into outcomes.

¹⁷ Commonwealth of Australia (2010), *Australia's Future Tax System: Report to the Treasurer*, Part One, at page 59.

These other factors can be personal, such as a lack of self control, or external, such as the climate.

Looking at material deprivation generates some different results from relative income poverty measures that we often default to. For example, using individual level data to test for a common set of seven deprivation items such as inadequate heating, poor environmental conditions and an inability to make ends meet, the OECD found that many of those who are cash income poor based on a relative poverty measure are not materially deprived.

Related to this, the age profile of material deprivation is different from that associated with relative poverty. Whereas relative income poverty has a U-shaped relationship — declining in middle age and increasing in old age — material deprivation, when measured as the share of an age group experiencing two or more deprivation items, generally declines with age.

As for the other measures of distribution, more work remains to be done. But such measures would seem to provide important information that measures based on income or wealth alone lack.¹⁸

Distributional concerns and policy advice

To be concerned with distributional issues is not to say that we should factor such concerns into our advice on all issues.

¹⁸ The Australian Bureau of Statistics is currently undertaking a Low Consumption Possibilities research project (see <http://www.abs.gov.au/AUSSTATS/abs@.nsf/Lookup/1504.0Main+Features5Jun+2009>), that combines data on the income *and* wealth of households, rather than looking at income *or* wealth alone. While not a measure of material deprivation per se, it is likely to indirectly capture some of the same outcomes, for example, in respect of the position of the aged. In a similar vein, distributional analysis based on consumption or expenditure rather than income also appears to better reflect observed deprivation: Meyer, B and Sullivan, J (2011) 'Viewpoint: Further results on measuring the well-being of the poor using income and consumption', Canadian Journal of Economics, Vol. 44 No. 1, pp52-87..

In economics there is a long tradition of trying to separate consideration of efficiency issues from those of equity or distribution. For example, the Kaldor-Hicks criterion that only requires that the losers from a change could potentially be compensated.

For policy makers, it seems generally sensible to adopt such an approach and focus on efficiency concerns except where affecting distribution is the primary goal. So, for example, the recent review of the Australian tax system argued that:

‘The transfer system, together with progressive personal taxation, is better suited to this task, and should be the primary means through which the government influences the distribution of income in the economy ...’¹⁹

Separating efficiency from equity considerations

This was in the section of the report dealing with the taxation of consumption. But experiences with consumption taxes illustrate the difficulty, both political but also analytically, you can have in trying to separate efficiency and distributional issues. For example, when the introduction of a GST in Australia was being debated in the late 1990s, the question of whether to include basic food in the GST base was one of the major points of contention.

Anthony Atkinson and Joseph Stiglitz, in their seminal 1976 paper, found that where there is an optimal direct (labour) income tax in place, any given distributional goal can be achieved more efficiently using the tax-transfer system alone.²⁰ This finding, and similar results from

¹⁹ Commonwealth of Australia (2010), *Australia’s Future Tax System: Report to the Treasurer, Part Two*, Volume 1, at page 276.

²⁰ Atkinson, A & Stiglitz, J (1976), ‘The design of tax structure: direct versus indirect taxation’, *Journal of Public Economics*, Vol. 6, pp55–75.

related branches of the optimal tax literature, have been highly influential, and have been adapted to analyse other important policy questions.

As there are many excellent discussions of the Atkinson-Stiglitz and related theorems and their policy relevance in the literature, I will only touch on some aspects particularly relevant to us today.

First, the theorem can take into account the presence of subsistence levels of consumption, such as basic non-discretionary items like parts of household expenditure on food, water and shelter. However, what is seen as a non-discretionary item is likely to change over time and vary across the population. For the theorem to hold, governments must be able to observe and act on these basic needs through lump sum transfers. But what government can do in practice is likely to be more constrained.

This leads to a second (more general) point regarding assumptions about the behaviour of governing authorities. While there is a growing body of work that takes into account the effect of political economy constraints, arguments like Atkinson and Stiglitz typically assume that governments' act consistently to maximise social welfare. In the Australian GST debate, critics questioned whether governments could be trusted to compensate properly and then not unwind compensation over time.

Finally, the theorem also relies on making normative judgements rather than just arguments relating to efficiency. Underlying the theorem is an objective of maximising a social welfare function which is essentially utilitarian. You may think the set of social welfare functions consistent with achieving the Atkinson-Stiglitz result to be acceptable, but this is a

normative position. Much of the debate in Australia involved normative claims inconsistent with this viewpoint.

These three observations also relate more broadly to the subsequent economic analysis of optimal taxes. To undertake economic analysis involving both efficiency and equity objectives, we need to have a sound grasp of both the positive and the normative underpinnings. The optimal tax literature, by beginning to formalise and empirically test these underpinnings, has given policymakers some valuable insights into how these underpinnings trade-off against each other. But that literature cannot provide definitive answers.

Achieving distributional outcomes outside of the tax-transfer system — work and education

Government action to achieve distributional outcomes is of course not limited to using the tax-transfer system, regardless of its optimal configuration.

And it is important that it not be limited to the distribution of cash. Governments and their advisers are more aware these days of the potential negative consequences of transfer payments in entrenching disadvantage, and the benefits for some groups of attaching labour force participation requirements to transfers made to them, supplemented where appropriate by active labour market programs.

This is not only because of the potential improvements in the distribution of disposable income in the medium- to long-run through earlier attachment to the labour force. It is also because, in general, we see work having benefits in its own right beyond the income provided. The living standards paper points to some evidence that supports this view.

If we look at the related issue of labour market deregulation, one view is that while the movement in recent decades to more deregulated labour markets in countries reduced

income inequality by improving employment outcomes, it widened wage inequalities at the same time. As a consequence the net effect on income inequality has been ambiguous.

Yet if work has benefits separate from income, the calculus is much more likely to be net positive.

We are also now more aware of the importance of human capital, of education and good physical and mental health for the outcomes people achieve. Better education is associated with improved participation in the labour market, higher lifetime incomes and longer lives. Assistance provided through better education outcomes therefore avoids some of the incentive constraints of cash transfers.

Also of importance for us, however, are the non-income benefits associated with education — that it gives people greater capacity to convert other resources or income into positive outcomes, and more sensibly choose between them. Better education is associated with improved health, lower rates of incarceration and increased engagement in civic life.

Improving educational outcomes for those otherwise disadvantaged is likely then to improve distributional outcomes, regardless of the distributional rule you favour. It is also likely to sit well with general community notions of fairness.

Conclusion

The Australian Treasury has found having a wellbeing framework a positive experience. Not because it tells us to forget our economic frameworks, but because it reminds us that they are richer and broader than is often assumed.

We have also found that having a framework is not the end of the matter. How you make use of your framework is also important. Critical here will be the example of the senior leadership of the New Zealand Treasury.

Today I have focussed on distributional issues. I expect one advantage you'll find of your living standards framework is that it encourages you to debate these and some other issues more than you have in the past.

I hope I have contributed to that debate.