

Modelling as Agit-prop: The Treasury's Role in Australia's Carbon Tax Debate

Henry Ergas and Alex Robson¹

Abstract

This paper examines the modelling undertaken by the Commonwealth Treasury of the costs of an Australian emissions trading scheme, published in 'Strong Growth, Low Pollution'. Despite its considerable technical sophistication, we argue that this modelling is primarily an exercise in propaganda: 'the systematic dissemination of selected information to promote a particular doctrine' (Oxford English Dictionary). That propaganda role determined the limited range of questions asked in the modelling, the myriad unrealistic assumptions made in answering those questions, and the limitations imposed on third-party access to the model and data.

Introduction

Modelling has played a central role in the debate over the Rudd government's proposed emissions trading scheme (ETS), as well as the carbon/tax emissions trading scheme actually introduced by the Gillard government. In both cases, large-scale modelling was undertaken by Treasury, with a report on the Rudd government's proposed scheme ('Australia's Low Pollution Future' or ALPF) being released in 2008² and a subsequent analysis, 'Strong Growth, Low Pollution' (henceforth: SGLP), focused on the Gillard government's scheme, being released in 2011.

It should be noted at the outset that the models Treasury has developed have not been released to the public, on grounds that are largely specious.³ Nor have the underlying data sets that inform those models been released. The — presumably intended — result is to make it impossible for those models to be thoroughly tested, or used to examine scenarios other than those Treasury that

¹ University of Wollongong, ergas.henry@gmail.com; Griffith Business School, a.robson@griffith.edu.au

² For an economic analysis of the ALPF, see Robson (2009a) and Robson (2009b). Robson (2007) examines the earlier Shergold Report, and also contains a brief summary of the basic science of global warming.

³ They are scrutinised in more detail below.

has chosen to model. Nonetheless, drawing on the published reports, this paper focuses on the modelling presented in SGLP and assesses the assumptions on which it is based and the consequent relevance of its results.

Our conclusion is that despite its obvious and very considerable technical complexity, the SGLP modelling was primarily an exercise in propaganda — which the *Oxford English Dictionary* defines as ‘the systematic dissemination of selected information to promote a particular doctrine’. It was that propaganda role that defined the questions asked in the modelling, the assumptions made in answering those questions, the interpretation placed on the results, the limitations imposed on third-party access to the model itself and the underlying data sets and the uses to which the results were ultimately put. In subsequent sections we examine each of these aspects in turn.

What questions did ‘Strong Growth, Low Pollution’ ask?

The focus of SGLP was on the costs to the Australian economy of imposing a carbon tax, where those costs were assessed over the period to 2050. SGLP does not examine the benefits of imposing such a tax; indeed, the analysis specifically notes that it does not include the economic impacts of rising global emissions on Australia (with or without Australia’s mitigation actions). As a result, the most that can be said for its results is that they would potentially allow one to gauge the cost side of the ledger, leaving the question of whether there are commensurate benefits to the political process.

With regard to how costs are assessed, while the scheme proposed by the Gillard government is ultimately an emissions trading scheme, the modelling is undertaken as if a carbon tax was in effect, with the extent of that tax depending on the emissions reduction goal being pursued. Those emissions reduction goals are defined in terms of targets for greenhouse-gas concentration levels, with two such targets — at either 550 or 450 parts per million — being modelled. Given those targets, the modelling broadly determines the lowest global carbon price needed to achieve the required emissions reductions and then calculates the change in Gross National Income (GNI) relative to a base case. That base case is defined in terms of a ‘business as usual’ scenario in which the rest of the world undertakes mitigation according to pledges made at the United Nations Climate Conference in Cancún. This scenario differs significantly from the Rudd government’s ALPF modelling, which assumed a ‘reference scenario’ in which

countries introduced no new policies to reduce greenhouse-gas emissions, with atmospheric concentration of greenhouse gases therefore rising to over 1500 ppm CO₂-e by 2100.⁴

The formal modelling of the GNI impacts relative to that base case was undertaken in SGLP using several models, of which two are particularly important. These are the GTEM model, developed by the Australian Bureau of Agricultural and Resource Economics (ABARE), which was used to examine the impact of abatement scenarios on the world economy; and the MMRF model, developed by the Centre of Policy Studies at Monash University, which is used to examine the impacts on Australia.

Using those models, SGLP examines two quantitative scenarios for each abatement target: a first is one in which the rest of the world pursues one or the other of the emissions reductions goals, while Australia remains in a 'business as usual' policy — that is, does not implement a carbon price of its own; and a second in which both Australia and the rest of the world implement a carbon price, that price being determined internationally. In other words, SGLP does not quantify the cost to Australia of the scenario in which Australia implements a carbon tax or emissions trading scheme, but large parts of the world do not. It does, however, present a brief, informal discussion of that scenario, which we comment on below.

The modelling neglects a number of costs that are likely to be associated with the Government's actual policy. For example, in seeking to compensate households for the increase in prices due to the carbon tax, the Government has reduced some average personal income-tax rates, but in doing so has increased marginal tax rates for many taxpayers.⁵ The modelling ignores the welfare costs associated with these higher marginal tax rates, as it treats compensation to households as lump-sum transfers.

In addition, the modelling ignores the economic losses likely to be associated with the Clean Energy Future Fund and other outlays funded by the carbon tax and so reflects only the costs due to the tax, rather than those arising from any wasteful expenditures the tax revenues permit.

In short, SGLP is not a cost–benefit appraisal of the government's scheme; it is only an assessment of its costs. It does not compare those costs to alternatives, most importantly that of adapting to climate change rather than avoiding it. And, in undertaking its assessment, it only quantifies scenarios in which the

4 On this point, see the Centre for International Economics (2011).

5 On this point, see Williams (2011).

rest of the world is credibly committed to global emissions reduction, and then looks at the impacts on Australia of participating or not participating in a global carbon market.

The assumptions

The absence of a scenario in which Australia has a carbon pricing scheme but many other countries do not reflects the core assumptions in SGLP about the global environment. Those assumptions are: that all countries will implement the abatement commitments made at the 2010 United Nations Climate Change Conference held in Cancún, Mexico, although those commitments are not legally binding; that in doing so, at least the industrialised countries will implement some form of carbon trading mechanism, and those mechanisms will be sufficiently linked that a global carbon price will emerge as of 2015–16; that the resulting unified price will reflect (and equalise) participating countries' marginal cost of abatement; and that over time other countries, notably China, will join this market.

A number of other crucial assumptions are related to these core assumptions. In particular, it is assumed that there is some very low transactions-cost mechanism that provides credible 'offsets'; that is, permits that involve payment for emissions avoided, typically in developing countries. At the same time, market participants can not only 'bank' permits (that is, purchase permits now but use them at a later date) but also 'borrow' them. Moreover, both banking and borrowing can, on net, occur between periods, with the important consequence that some scheduled abatement can be deferred to future periods.

Before proceeding, it is useful to consider the import of these assumptions about the international environment. Put simply, there is very little prospect of their corresponding to reality. In particular, there is no reason to believe a global framework for emissions abatement, with credible, binding targets, will emerge between now and the end of this decade. Even were such an agreement to eventually be reached, there is no prospect of a global, unified, carbon price emerging between now and 2015–16, as (given implementation lags) that would require an agreement to be in place now. As for major developing-country emitters, notably China and India, making binding commitments to significant emissions reductions (relative to business-as-usual, which still involves large increases in emissions relative to 2001), this seems unlikely, all the more so if measures are taken (as is currently proposed) to curtail the use of difficult-to-police instruments such as the Clean Development Mechanism (in

which developing countries receive payments not to emit). Finally, there is no precedent for 'borrowing', which was specifically prohibited under the Kyoto protocol.

Treasury's justification for those assumptions is that it is simply taking countries at their word. This is not accurate, however, as Treasury's assumptions go far beyond any specific commitments countries have given. For instance, while China has at times suggested that it will consider a credible and binding commitment to multilateral abatement at some stage in the future, it has never gone beyond noting the possibility. India, for its part, has committed even less.

But even if some countries had made non-binding commitments to stabilise emissions levels at some lower level, it is unclear why it should be assumed that those commitments would be realised in full. Rather, just as in examining options for our defence planning and our trade policy, a risk-adjusted approach should be used that explicitly accounts for the fact that outcomes are uncertain, including the extent to which countries abide by their Cancún commitments. Yet Treasury's modelling assumptions not only rule these risks out, but rule in scenarios — such as the emergence of a global carbon price in 2015–16; an eventuality that was entirely implausible even when the modelling was carried out.

Beyond those core assumptions, myriad technical assumptions are, of course, built into the modelling. Particularly contentious are those related to the rate at which new, low-emissions, technologies become available. In some instances, that rate appears to be determined exogenously — that is, technologies simply appear and become available for use. Particularly significant in this respect are the assumptions relating to the timing of the commercial viability and deployment of carbon-capture-and-storage (CCS) technology. The modelling assumes that CCS is deployed by as early as 2024 in some scenarios. This has a substantial impact on Australia's income, given the size of our coal endowments. However, there is little evidence to suggest that this assumption of commercial deployment in less than 12 years is realistic.

As well as such exogenous technical change, the modelling also provides for induced innovation; that is, for the development of low-emissions technologies as a response to the carbon price. That linkage is effected through 'marginal abatement cost' (MAC) functions, which relate the marginal cost of abatement in each industry to the level of the carbon price. These functions are relatively 'aggressive', implying a strong response of technological innovation to the carbon price. However, this raises a host of issues about the credibility of the carbon price as a signal to innovation (discussed in Ergas 2012).

Another important modelling issue is what cost this induced innovation involves — in other words, how the cost of updating the capital stock is modelled, assuming innovation is largely embodied in new capital goods. As best one can tell, those costs are determined differently in the GTEM and MMRF models. That points to a further set of assumptions which are required to link the international and domestic scenarios and to close the respective models (that is, fix the constraint that then determines their solution). While there is no simple way of doing this, and the choice of closure and linkage can significantly affect the results, SGLP provides no information as to precisely how that linkage is made.

Finally, it is worth noting that both models lack any form of forward-looking behaviour. For example, once the carbon price is determined, it is assumed to rise over time in a manner consistent with the Hotelling valuation principle; that is, in line with an assumed rate of interest. However, the fact of those increases does not ‘bring forward’ any form of adaptation: agents do not, in other words, anticipate those increases and adjust to them in a forward-looking way. Rather, they are entirely myopic, which seems entirely inconsistent with the assumption that the Hotelling valuation principle holds and is also at odds with modern modelling practice.

The results

If one accepts the modelling assumptions, the results are not implausible. However, the most striking feature of the results is that they highlight the implausibility of the assumptions themselves. In other words, while there appears to be internal validity (as best one can tell, the model ‘works’), the model lacks external validity — that is, a meaningful correspondence to the world as it is and is likely to be.

The model suggests, for example, that China will bear a very large share of the costs of achieving global abatement: by 2050, China’s per-capita GNI is modelled to be 5–10 per cent lower than it would otherwise be, while US per-capita GNI is essentially unchanged. The situation is even worse for the former Soviet Union, whose per-capita GNI is modelled to decline by 9–12 per cent. If these results are to be believed, it is unclear why defection would not be a strongly dominant strategy for these countries — that is, why they would not be significantly better off reneging on any commitments to abate. The assumption of stable, binding agreement therefore seems implausible.

Moreover, even those results depend on countries deferring planned abatement to periods beyond the modelling scenario. In other words, countries emit but offset those emissions against promised cuts in un-modelled future years, thus

reducing the cost of achieving abatement targets in the modelling period. Nor is that on a small scale: rather, by 2050, there is a global debt to the future of about 8 per cent of the total modelled abatement task. But that begs the question of whether such 'borrowing' would be consistent with the global policy framework (as noted above, borrowing from the future was specifically prohibited under the Kyoto convention); and even more so, of why the promises to abate in future would be credible. Yet if they are not credible, it is unclear why innovators would develop very-low-emissions technologies whose viability depends on the extremely high carbon prices that would be needed to clear that backlog of deferred abatement.

Equally, the results for Australia highlight the outputs' sensitivity to the assumptions. Though no direct comparison is provided in the documentation, the marginal cost of abatement seems significantly higher in SGLP than it was in the 2008 modelling (see Centre for International Economics 2011). This likely partly reflects the Gillard government's decision to exclude agriculture (which was included in the Rudd proposal) from its scheme, as well as the somewhat higher share of mining in the Australian economy in 2010 than in 2008. Be that as it may, the result is that a much higher share of Australian abatement in SGLP occurs by buying permits from overseas than was the case in the 2008 modelling. However, this obviously makes it crucial whether low-cost permits are indeed available in global markets and, if so, whether they are credible (that is, whether the abatement they promise in some other country actually occurs). The modelling suggests the bulk of these permits will be purchased from the undefined 'rest of the world', which is largely composed of relatively poor developing countries. At the moment, however, the market for such permits is extremely problematic, and there is no reason to believe the many obstacles that lie in the path of its development will be removed in the near future.

Testing the models

SGLP gives very little indication of sensitivities, but it is clear that even were one to accept the characterisation SGLP gives of the future global policy environment, the results are highly sensitive to the modelling assumptions. For example, as shown in Centre for International Economics (2011), the global price depends crucially on the marginal abatement costs of the lowest-abatement-cost country, so that small changes in the range of countries that participate in the scheme and in their cost structures could change the results very materially. Equally, the estimated global costs of abatement are very sensitive to the level and structure of the MAC functions — that is, to the implied elasticities of induced innovation with respect to the carbon price — which do not seem to be based on any form of empirical testing. Finally, the precise way the linkage of the global and Australian models is carried out is likely to have a significant impact on the results.

On top of this, any sensible analysis would test how the costs to Australia change as the assumptions made on the global policy environment are changed. In particular, given that it is unrealistic to treat global agreement as an inevitability, a proper policy analysis would examine the costs to Australia of adopting a carbon tax/ETS when the rest of the world (or at least, major parts of it) do not.

This is discussed in SGLP, but entirely qualitatively and almost parenthetically. In a short, informal discussion, the report claims that ‘going it alone’ would impose lower, not higher, costs than would be borne under coordinated global action.⁶ The reason given for this is mainly that world demand for our resources would be stronger than it would be under the coordinated-action scenario. On that basis, the report implies Australia would be well-advised to abate in any event, as it would not bear especially high costs from doing so.

But this approach seems wrong-headed. Thus, each cell in Table 1 below lists the basic structure of payoffs to Australia in four scenarios, and illustrates the comparisons that are relevant for policy considerations. The benefits to Australia when the rest of the world mitigates are B. As Australia accounts for 2 per cent of global emissions, it is reasonable to assume B does not depend on whether or not Australia engaged in mitigation. Treasury’s modelling concludes that C, which is the cost to Australia when Australia doesn’t mitigate and the rest of the world does, is less than K, the cost to Australia when both we and the rest of the world mitigate. In other words, conditional on the rest of the world mitigating, mitigating imposes a cost on Australia. Treasury does not estimate Z, which is the cost of ‘going it alone’ (that is, in the absence of global mitigation, if Australia mitigates, we will incur costs of Z), but in defending unilateral mitigation, claims $Z < K$.

Table 1: Structure of payoffs to Australia in various scenarios

		Rest of World	
		Mitigate	Don’t
Australia	Mitigate	B-K	-Z
	Don’t	B-C	0

Source: SGLP 2011.

Note: B is the assumed benefit to Australia of global mitigation which is assumed to be the same regardless of whether Australia does or does not mitigate. K is the cost to Australia when we mitigate along with the rest of the world, while C is the cost when we don’t mitigate but the rest of the world does. Z is the cost of unilateral mitigation by Australia.

6 Page 84 of SGLP states that ‘Some commentators suggest Australia is ‘going it alone’ and, as a result, Australia’s mitigation costs will be far greater than reported because impacts on domestic competitiveness will be far greater. This argument is misconceived. If the extent of global action is less than assumed, then Australian mitigation costs will be lower, not higher, than reported for two main reasons. First, less stringent world action would strengthen export demand and output for our energy exports. Second, if global action is less than assumed, world carbon prices will be lower, making it less expensive to source abatement overseas.’

However, from the perspective of rational decision-making, this is simply the wrong comparison. Rather, the relevant comparison for deciding whether Australia should mitigate when the rest of the world does not is between the costs in that scenario (in other words, the payoffs in the right-hand column of Table 1) — not that between the costs of Australia mitigating with and without coordinated global action (the top row of Table 1). For not mitigating to be ruled out as the dominant strategy for Australia, we would need $Z < 0$, which is farfetched.⁷

Even putting that aside, the report's suggestion that unilateral action is less costly to Australia than coordinated action also seems implausible — especially if the key difference is that in the coordinated-action scenario our resource competitors abate while in the unilateral scenario they don't. It seems likely that the report's claim — which is not backed by any modelling results — reflects unrealistically low estimates of substitution elasticities between Australian and foreign natural resources.

Unfortunately, it is not possible to test any of these issues as Treasury has not released the models and data sets on which it relied. Initially, Treasury claimed this was because of objections from ABARE; however, ABARE had previously signalled its intention to release the model. Subsequently, in replying to questions from Senator Mathias Cormann in the Senate's Estimates Committee (on 17 October 2011), the Executive Director of ABARE made it clear that the failure to disclose the models and data sets was a government decision, not one made by ABARE. Since then, Treasury has refused Freedom of Information requests to disclose the models and data on the basis that doing so would involve an undue compliance burden, as those models and data are not kept in a form that would allow their disclosure.

This is a deeply troubling justification. Assuming it is true, it implies Treasury does not adhere to good modelling practice, which requires managing information resources in a way that allows those resources to be audited, permits results to be replicated and avoids unnecessary costs when models and data sets need to be updated.

7 Rodriguez-Neto (2008) argues that the theory of repeated games may 'save the day' in relation to climate change policy and that folk theorems, which (broadly speaking) state that under certain circumstances cooperation may emerge as an equilibrium in an infinitely repeated game, may be important. But why should the theory of infinitely repeated games be more applicable than dynamic game theory, in which the players' strategies and payoffs change over time? In such games, it is possible that cooperation may actually be hindered if players place a sufficiently high value on the future (Skaperdas and Syropoulos 1996). Moreover, folk theorems do not state that cooperation *must* be an equilibrium of an infinitely repeated game. Indeed, defection in every round of an infinitely repeated game can still be an equilibrium. Therefore, whilst folk theorems are interesting, by themselves they provide little concrete guidance in the climate-change policy debate.

To make matters worse, the SGLP documentation does not fully provide the information needed even to indirectly test the results. Relatively few model parameters are actually disclosed and even when they are, they are frequently incomplete. For instance, while the parameters of the MAC functions are listed, the operation of these functions is smoothed in the early years — but quite how that is done is not explained. And, at times, the documentation simply refuses to disclose even estimated parameters: the constant partial equilibrium expenditure and own-price elasticities of electricity demand, to take a striking case, ‘are deliberately not presented so as to avoid any misinterpretation’ (SGLP: 160). One rather wonders how likely it is that anyone who knows what a constant partial equilibrium is needs to be protected from misinterpreting it.

How the results were used

Its obvious technical sophistication notwithstanding, SGLP is first and foremost an exercise in propaganda, which the *Oxford English Dictionary* defines as ‘the systematic dissemination of selected information to promote a particular doctrine’ — the doctrine being that of the utility of the Gillard government’s carbon tax. The selective aspect of the information being disseminated is evident in the failure to release scenarios that could undermine the case for that policy: most obviously, the scenario specifying the costs to Australia of a unilateral carbon tax. It is also highlighted by the refusal to give access to the models and data to possible critics of the policy.

At the same time, the fact that the modelling was undertaken by Treasury (which was clearly constrained in its ability to comment publicly on Ministers’ claims about its modelling) meant there was little constraint on misrepresentation by government of its results. For example, the foreword to SGLP by Treasurer Wayne Swan and Climate Change Minister Greg Combet asserts that ‘The modelling finds jobs continue to grow under a carbon price’; however, the model assumes the labour market clears (and given population growth, it would be extraordinary if the labour force fell), so the fact that employment increases is hardly a result of the modelling — rather, it is a result of the assumption that the labour market clears in a scenario where the labour force continues to rise.

Similarly dubious claims have been made in relation to the effective carbon price paid as a result of the issuing of free permits to selected, trade-exposed, industries. Averaging the zero price of gifted permits with the price of the permits producers have to buy, Minister Combet and Trade Minister Craig Emerson have argued that the effective carbon tax will be ‘as low as \$1.30 a tonne’, so that the harm to production will be minimal. But these claims are fallacious. Production decisions obviously do not depend on the average price,

but on how much firms can save by reducing production. If avoiding one tonne of emissions saves \$23 in carbon tax payments, then it is that \$23 — the price at the margin — that matters.

More recently, the Government has relied on the modelling to justify its assumption in the 2012–13 Budget that the carbon price in 2015–16 will be \$29. However, the carbon price path in SGLP is simply the result of applying the Hotelling valuation principle to the price needed to achieve the ultimate abatement target;⁸ there is no empirical basis for assuming prices will in fact follow the price path the Hotelling principle determines, all the more as it is well known not to hold in markets for exhaustible resources. But Treasury has not demurred from the statements made by Minister Combet and Treasurer Swan that claim justification in Treasury's modelling. The lack of robustness of this estimate is emphasised by the fact that the Government continues to state that this 2015–16 permit price will remain at \$29 even though the Government announced that the Australian scheme will now be linked to the European ETS.⁹

Conclusion

While technically sophisticated, the modelling done for the Rudd and Gillard carbon price schemes was undermined by its partisan ends. Even its technical quality has been flawed by reliance on assumptions that serve those ends; moreover, the inability of third parties to replicate the results, or test the modelling sensitivities, limits the weight that can be placed on it.

Particularly egregious is the decision to assume speedy global agreement on cutting carbon emissions. By the time of the SGLP modelling it was apparent that this was highly unlikely; given that fact, the least one could expect is rigorous testing of whether Australia should nonetheless implement a carbon price, and a discussion of what such a unilateral approach would cost. That is all the more important as unilateral abatement would clearly yield, at most, trivial environmental benefits.

Some inkling of the magnitude of these costs can be derived by looking at the estimated reduction in Australian national income that occurs even with coordinated global action — for each dollar in revenue the carbon tax raises,

8 For the abatement target to be reached, price must rise to a threshold level. Given that level, the Hotelling valuation principle pins down the rate of increase in price.

9 Paradoxically, the assumed equality of the autarky and free-trade price also suggests that there are no economic gains to Australia linking to Europe, which begs the question as to why such links are being pursued in the first place.

national income declines by around two dollars.¹⁰ In other words, the carbon tax has an average excess burden four times greater than that of the most distorting tax identified by the Henry report.¹¹ These results imply that absent any environmental benefits, the carbon tax is extremely distorting in terms of its average excess burden — and that average excess burden would presumably be even higher if unilateral action causes an even greater decline in Australian national income.

That the AEB is high is unsurprising, as the carbon tax is a tax on an intermediate input and falls especially heavily on industries in which Australia has a comparative advantage. Given the emphasis Treasury has placed on average excess burdens in criticising mining royalties, one might have expected the risk of the economy bearing so great an AEB for little or no benefit to figure in Treasury's discussion; instead, it is completely lacking.

In short, this is modelling for a specific political purpose. Its technical quality is unknowable, as the models and data have not been released. Even were it high, however, that would not make it more useful to the public interest.

References

- Centre for International Economics 2011, 'Notes on "Strong Growth, Low Pollution": Modelling and Related issues', Paper prepared for the Australian Industry Greenhouse Network, September.
- Ergas, H. 2012, 'Policy Forum: Designing a Carbon Price Policy: Using Market-Based Mechanisms for Emission Abatement: Are the Assumptions Plausible?' *Australian Economic Review* 45: 86–95.
- Robson, A. 2007, 'The Prime Ministerial Task Group's Report on Emissions Trading A Solution in Search of a Problem', *Agenda* 14(3): 19–25.
- 2009, 'ETS Forum: Why No Cost–Benefit Analysis?' *Quadrant Online*, 8 August, at: <http://www.quadrant.org.au/blogs/doomed-planet/2009/08/alex-robson>
- 2010, 'What Does the Government's CPRS Modelling Tell Us?' in Moran, A. (ed), *Climate Change: The Facts*, Institute for Public Affairs, Melbourne.

10 Using, for example, a 1 per cent discount rate applied to data from Treasury's modelling output to 2050, the present value of carbon tax revenue is \$570 billion, whilst the present value of GDP losses is about \$1 trillion.

11 Modelling undertaken for the Government by KPMG for the AFTS Review estimates the average excess burden of a gambling tax at 0.54.

- Rodrigues-Neto, J. 2008, 'Climate Change Policy: A Theorist's Plea to Take Heed of Game Theory and Ambiguity Aversion', *Agenda* 15(2): 85–90.
- Skaperdas, S. and Syropoulos, C. 1996, 'Can the Shadow of the Future Harm Cooperation?' *Journal of Economic Behavior and Organization* 29(3): 355–72.
- Williams, M. 2011, 'Why Increasing the Tax-Free Threshold is a Bad Idea', *IPA Review* (September): 27–8.