

Review of proposed New Zealand Emissions Trading System

Dr Suzi Kerr
Motu Economic and Public Policy Research

November 4, 2007

Summary

The Government's proposed Emissions Trading System (ETS) is overall an impressive piece of policy development. The design embodies much of the best international thinking: it responds to lessons learned from economic theory and from the evaluation of existing emissions trading systems and other environmental markets in unprecedented ways. The proposal takes significant steps toward achieving the Government's political goals without severely compromising the long-run effectiveness, efficiency or true equity of the system.

In particular the choice of points of obligation in the liquid fuels, stationary energy and industrial processes sectors that emphasise comprehensive coverage and low compliance costs have advantages for the environment and for efficiency, as well as making compliance relatively transparent. This is in contrast to most previous emissions trading systems (notably the European system) where the point of obligation was influenced by sectoral lobbying and the choice made led to serious implications for coverage, for treatment of new entrants and exiters - which created perverse incentives, and for (mis-)allocation of free units.

The comprehensiveness of the system is also an international first. It creates an architecture that makes any future NZ commitments credible. It avoids many of the problems of National Allocation Plans in the European Union system because no inter-sectoral emissions allocation decisions need be made – all emissions are covered by one system. It is the first system to even seriously contemplate including non-CO₂ emissions from agriculture or sequestration/removals in forests.

These characteristics, if effectively sustained and carried through to demonstrate an environmentally effective, economically efficient, politically feasible domestic policy could be one of New Zealand's most significant contributions to international climate change mitigation efforts.

That said, several very important aspects of the proposal require further development: the environmental contribution – short and long run; the need for clear thinking on interrelated leakage and allocation issues; how to achieve a smooth, low-risk transition; and the use of science in regulation (simplicity and modelling). These issues are exacerbated by the comprehensive nature of the proposed system and are in some cases issues where international literature and experience is less helpful. These are the issues on which my review will focus.

This review draws heavily on discussions with the Climate Change Policy Dialogue group and the short papers prepared for that group. These can be found at www.ecoclimate.org.nz/ETS.htm. Any opinions expressed are however mine alone and I am responsible for any errors and omissions.

1. Overall Goals - Environmental Impact

With adequate access to international markets and with strong efforts to ensure compliance, the currently proposed ETS architecture will lead to New Zealand's compliance with our Kyoto obligations.

The scheme will not lead to any particular level of domestic emissions, however the price signals on gross emissions and for forest sequestration are likely to lead to a lower level of net emissions than business as usual. The level of reduction will depend on the international price, access to the international market, and the short run responsiveness of each sector. Domestic reductions are likely to be small at \$15 per tonne CO₂-e with a liquid market, while responses at \$50 are likely to be significantly higher even before 2013. While no comprehensive modelling can provide strong signals on likely New Zealand abatement efforts, past experience with emissions trading systems (such as the US Acid Rain programme) suggests that responsiveness is usually underestimated in advance of a credible price signal. In any case, New Zealand's domestic reductions are not themselves that important on a global scale.

In an environment with no effective international enforcement, the development of international cooperation is what really matters for long-term climate stability. Successful cooperation requires good technical support, open communication, fair sharing of costs globally, credible monitorable commitments, leadership, and that we steadily build and maintain trust that all countries are all contributing fairly. While there are many examples of successful local cooperation without formal enforcement, addressing climate change is a challenge on a scale humans have never faced before. We need to be careful, smart and generous as we tackle it.

The ETS will make a contribution to the global effort in several ways that go beyond our own domestic emission reductions. New Zealand will have a relatively robust architecture that will enable us to make credible commitments in post 2012 negotiations. If we say that we can take a, for example, 20% cut in our assigned amount units, we will be able to achieve that, even if we most do it by enabling emission reductions in other countries. If our example helps other countries develop credible and flexible domestic policies, they too will be able to make firm future commitments. This will facilitate the difficult international cooperation-building process. Currently many countries may be hesitant to commit to stringent cuts if they do not believe that others will follow through with their simultaneous commitments. The lack of credible domestic policies has stymied implementation of agreed Kyoto targets.

One area in which the current policy falls short is that it does not create a clear mechanism for implementing a more stringent target. For example, allocations are being given to sectors out to 2025 in the form of tonnes. These may consume a higher percentage of New Zealand's total allocation post 2012 than is currently anticipated. An alternative would be to designate post 2012 allocations in the form of 'shares' of the post 2012 target. A more stringent target for New Zealand may also come in the form of the change in the 'rules' of the agreement – for example how greenhouse gas emissions are measured. These could reflect new science. Some of these could have significant impacts on the effective stringency of New Zealand's target. We need to create mechanisms that will appropriately translate any international changes through to domestic policy to avoid unnecessary uncertainty and renegotiation of the ETS architecture.

On the positive side, the inclusion of forestry and agriculture in our ETS will provide a strong incentive to develop and adopt new practices and technologies that reduce net emissions in these two important sectors. These will be available to developing countries that may, as a result, be more willing to commit to control their emissions. We might want to consider complementary policies to enhance this effect. The proposed new research funding will address this in part but not entirely.

Three other design features affect the environmental impact of the ETS: the acceptance of Clean Development Mechanism units; the acceptance of Assigned Amount Units (including so called 'hot air'); and the treatment of leakage. Each of these has counter-intuitive environmental effects

The system proposes to accept virtually unlimited use of Clean Development Mechanism (CDM) credits. The advantage of this is that it keeps good faith with the international agreement process. Developing countries ratified Kyoto on the basis that there would be benefits to them from participation – the main mechanism for realising those benefits was the CDM. It was considered critical to involve them in some way, even if their commitments were extremely weak, and to provide a mechanism that starts developing their domestic capacity to mitigate greenhouse gas emissions. If we want them to take on and abide by significant commitments in future (and this is critical for long term climate stability) we must build their trust by abiding by our own past commitments.

CDM credits do, however, carry potentially large environmental risks, largely because of interaction between the difficulty in estimating 'what would have happened otherwise' baselines and widespread corruption. It is impossible to objectively predict emissions from a firm or whether a more GHG efficient technology would have been used in a given time period; commercial decisions are too complex. At the same time, economic gains to project organisers from predicting both high emissions and low take-up of efficient technology and hence creating a high emission baseline are potentially large. This invites bias and corruption and suggests that a large percentage of CDM credits may not represent real reductions. Rather than addressing this by banning the purchase of CDM credits however, it is best approached through providing an alternative more robust pathway for developing country participation in the 'post 2012' negotiations.

Excess units from Russia, Ukraine and others pose a similar problem in some ways. We (and the Europeans who are now refusing to buy AAUs) agreed to give them AAU allocations that exceeded business as usual. If we now act in bad faith and refuse to honour those units in the market, that sends a negative signal both to those countries and also to developing countries that will also seek relatively generous allocations as a condition of their entry into binding targets. The environmental implications of these generous allowances can be offset in the post-2012 agreements by all countries taking on relatively more stringent targets. If any country wanted to prepare for a potential reduction of allowances post-2012, they could do so by voluntarily retiring some of their own AAUs. What counts for the environment is cumulative emissions. Those who are willing to take on the costs of a higher level of mitigation will do so, rather than trying to force this cost on one specific group of countries. If the countries with 'hot air' know that future commitment periods will be more stringent, they will not lower the price before 2012 by selling large numbers of units – they will bank them.

One alternative that offers an apparent middle ground is 'green AAUs' where payments for AAUs are linked to spending on the environment in the host country. While this sounds attractive, we are unlikely to be able to confirm that the money is additional to what would have been spent and that it is well spent; this alternative suffers from many of the problems of the CDM. We would be better buy unlimited AAUs and, if we feel concerned about the environmental implications, separately make additional environmental investments in ways that we judge are valuable – for example, New Zealand may have a comparative advantage in development and diffusion of low-GHG-intensity agricultural practices.

In summary, I believe that New Zealand needs to consider now how the ETS will respond to changing science and international conditions; we should consider how we can enhance the internationally beneficial effects of our agricultural innovations; and I agree with the government's decision to allow purchase of a wide range of international units. In contrast, I have concerns about the environmental effects of leakage; I discuss these in section 3.1.

2. Good design characteristics that should be protected and enhanced

2.1 Point of Obligation

For liquid fuels, stationary energy and industrial processes sectors I believe the upstream points of obligation chosen are appropriate because they maximise comprehensiveness while minimising compliance/monitoring costs.

In the energy sector an upstream point of obligation is potentially under threat. Two arguments are used in favour of going 'downstream'. The first is that some firms may want to market themselves as offsetting their emissions with particular types of NZUs – for example those associated with indigenous forest regeneration. The counter argument to this is that they could achieve this by visibly estimating their own 'footprint' and providing NZUs from their chosen sources to their energy suppliers as part of their purchase contracts.

More concerning is the second argument for going downstream, that people respond more to a requirement to be a point of obligation than they do to a price signal. This is a frequently used and often successful argument and has led to complex and non-comprehensive systems overseas. The counter argument is first that if firms and individuals are not responding appropriately to the price signal they face the full costs of this at least in the short run – the environmental target will still be reached. Second, if we want to address this problem, we should understand the true causes of the lack of response. It seems implausible that a government regulation designed to achieve one goal – to ensure compliance with an emissions trading system – is the best way to help people make good decisions in a new world where carbon is priced. If the problem is education and information, there must be better ways to provide this. If the problem is management structures, the firms themselves should find ways to improve decision-making rather than rely on a government regulation to determine their internal decision-making processes.

In all previous systems, the point of obligation and the point of free allocation have been identical which has obscured lobbying regarding points of obligation. Given that this link is being broken in New Zealand – the decisions about the best point of obligation may be able to be kept in the technical realm where they belong.

For forestry the only feasible point of obligation, the landowner, has been chosen for deforestation but potentially the trade-off between comprehensive coverage and compliance costs has not been managed as well as it might have been for post-1989 forests. The level of accuracy and verification required for carbon monitoring in forests creates high compliance costs. The effects of this have been addressed by making involvement voluntary for post-1989 forests so the system is not comprehensive. It could alternatively have been dealt with through a simpler monitoring option for small properties to increase participation and lower the monitoring costs for those involved. Large areas of New Zealand could potentially be reforested with plantations or indigenous forest but in individually small blocks that may not be able to benefit from this high compliance cost system.

The question of point of obligation is still an open one in agriculture, which poses some especially difficult challenges. The trade-off is essentially between a relatively easily managed system where processors are points of obligation but where emissions are a function of output only, versus a complex, difficult system where points of obligation are on farms and a full range of options to reduce emissions per unit of output are available. The government's expressed preference for the processor level is administratively understandable but would only exacerbate the leakage problems I discuss in section 3.1 and means that most of the important opportunities for true emissions reduction will be lost.

2.2 Comprehensive coverage

The scheme's comprehensive coverage has economic benefits because it increases the range of actions through which we can respond to greenhouse gas implications. As free allocations are reduced, comprehensive coverage spreads the economic burden of New Zealand's mitigation effort widely and on the basis of where the costs are induced. It also provides coherence to our overall climate change policy. Comprehensive coverage minimises the number of discretionary decisions the government needs to make about the appropriate level of emissions from certain activities. This is a distinct advantage in a small country where we have limited analytical capability. We should resist all pressure for exclusions and regularly review any exclusions made (such as the de minimus and voluntary participation rules).

2.3 Limited free allocation

The decision to make no free allocations in the liquid fuel sector where costs are passed on is logical and unprecedented. The planned phase-out of all other free allocation and identification of allocation with the transition period only are also extremely useful components of the scheme. Awareness of the phase-out may avoid renegotiation of free allocations once the rules are initially set. The one-off nature of these allocations is consistent with the idea that free allocations are used to compensate people whose capital assets fall in value because the introduction of a price on carbon –i.e. backward looking compensation. This is also consistent with the decision to make no free allocation to new entrants (who by definition are not existing owners of capital).

The removal of free allowances from those who exit is inconsistent with a compensation motivation – it seems to be a crude way to address leakage issues. Not only is this unlikely to be the best approach to address leakage, but it also creates perverse incentives to keep old firms alive at the expense of new ones even if the new ones could be more efficient.

Also less positively, there is no clear link between the relatively high level of the allocations and any analysis of the share of the future burden that will fall on the owners of this capital (primarily owners of firms including shareholders and the pension funds etc that invest in shares) relative to workers who will in some cases bear significant adjustment costs, or to consumers in general who will bear all the long term costs as well as facing capital losses on their energy inefficient houses, cars and appliances.

All these issues, and those below under leakage, need to be discussed, analysed more, and, where possible, informed by credible empirical evidence during the discussions with stakeholders to be held during 2008.

3. Issues that need significantly more attention

3.1 Environmental losses through leakage?

'Leakage' is an increase in GHG emissions in a country with weak climate policies as a result of domestic climate policies in another country that lead to decreases in that other country's production. If significant leakage occurs, some firms and workers in New Zealand will incur economic costs while the global environment suffers.

For New Zealand as a whole, allowing leakage allows us to meet our Kyoto obligations at lower cost than we would face if the global agreement were complete (i.e. if all our competitors were also regulated): we simply 'export' some of our highest intensity production. At the same time avoiding leakage is likely to be expensive. This argues for caution in addressing leakage.

Some argue that New Zealand should focus its efforts on the international negotiations when it wants to improve environmental outcomes and then simply comply with the resulting agreement in the lowest cost way possible. I differ.

New Zealand can contribute to controlling climate change in several ways. First, we can take actions that lead to real reductions in emissions. While our total reductions are small because we are small, our mitigation, tonne by tonne, is just as valuable as anyone else's. What is different for us, however, is that our influence through policy innovation, could be disproportionately large. If we can design a domestic policy that leads to genuine reductions in global emissions and does not allow significant leakage, this could provide a precedent that others can follow.

The risk of leakage is a considerable political barrier to participation in most countries. While the ultimate solution to the leakage problem is full global participation, this is unlikely to eventuate in the near future. Providing a model for how countries can reduce the impacts of leakage may help more complete global participation. A policy with environmental integrity will also give us better credibility than a policy that complies with Kyoto but only in the letter of the law.

It is difficult to define products and processes that are vulnerable to significant leakage so any policy to address this will inevitably be crude and imperfect. That does not mean that it should not be addressed in at least some cases. In New Zealand several industrial processes (not entire firms) and agriculture are potential targets for leakage policy.

Leakage arises because the marginal cost (i.e. the cost of producing one more) of producing specific products and processes rises because of their embodied greenhouse gases. It is not addressed by policies that freely allocate units based on any historical measure.

Leakage can be addressed in three ways: border adjustments to equalise treatment of traded goods; output-based allocation of units instead of fixed allocations; and progressive obligations.

Border adjustments work by exempting exports of specific products from an obligation to cover their embodied emissions and simultaneously imposing obligations to cover the embodied emissions of those products when they are imported. Domestic producers are put on a level playing field with their international competitors. Border adjustments are the cleanest approach and involve no subsidy to import-substituting goods so impose lower costs on taxpayers. They can discriminate among products based on the climate policy in the originating or receiving country.

Border adjustments can also be used to protect exports by rebating emission units for the average emissions of exported products to the farmer or processor. This clearly has fiscal costs but they are no greater than any other policy that addresses leakage in the export sector. It has the advantage that it does not subsidise output that is domestically consumed or consumed in countries that face climate regulation.

On the down side, there are likely to be WTO implications relating to the use of border adjustments as they will 'discriminate' against products from countries with weak climate policies even if they have ratified Kyoto. New Zealand may not want to be the country that faces the inevitable challenges. We should monitor and support efforts by other countries to develop this option and ensure that our legislation does not preclude use of border adjustments at a later time.

Output-based allocation uses free allocations linked to products' levels of output so that the impact of the policy on the marginal cost of production is reduced. Each extra unit of production is associated with the free allocation of some extra NZUs so while marginal costs may rise with the policy, at least some of the cost rise will be offset.

This approach does not need to perfectly offset the cost rise, and the level of NZUs per unit of output does not need to be based on any measure of 'best practice', which is inherently hard to determine. It should be explicitly recognised as a simple crude way to reduce the impact of leakage rather than an attempt to get a perfect policy. It could be a percentage of historical emissions per unit of output in New Zealand as a whole. Those who have relatively low emissions intensity already will benefit relative to those with high emissions intensity. If the percentage is set low enough, no firm, even if their emissions intensity is really low, will be receiving an output subsidy so there should be no trade implications.

This policy could operate within an absolute cap on emissions. The output of the chosen products by individual firms could be used to share a fixed pool of free units rather than having a fixed amount of units per unit of output. This would hold fixed the total transfer of compensation to the sector thus avoiding any influence of sectoral level rent seeking on the decision-making process. One option for the level of the fixed pool would be the current agreed free allocations – this would hold the fiscal cost at the current level. If we will incur the fiscal cost anyway, we may as well get leakage reduction benefits. Any changes in the capped pool of free units (particularly as we approach 2025 and the free allocations are phasing out) would tradeoff fiscal cost against the effects on leakage.

Under this approach firms would face a muted incentive to reduce output (which is the point of the policy and is environmentally favourable), but an efficient incentive to reduce emissions per unit of output. In the agriculture sector, this could simultaneously address the question of how to freely allocate units that intend to compensate for capital losses (loss of land value).

Progressive obligations are a less efficient version of an output-based allocation approach. They require firms to cover only a share of their emissions. They do mute the incentive to reduce output (which is good where leakage is an issue): a firm that produces one more unit of output only needs to cover a share of the additional emissions so marginal costs don't rise as much. Progressive obligations do not, however, provide efficient incentives to reduce emissions per unit of output. A firm that reduces its emissions intensity receives only a share of the benefit of that reduction because they only pay for a share of their emissions. Also, total emissions are not capped; this imposes fiscal risk on government.

Any policy used to address leakage should be simple and closely targeted. It should also be designed to phase out as other countries regulate their emissions.

3.2 Managing economic risk in the short term

The Government frequently assumes an international carbon price of around \$15 per tonne (e.g. in its calculation of the Government's liability) or \$25 per tonne (in the general equilibrium modelling) and does not predict significant carbon market access problems. The Government uses general equilibrium modelling (Infometrics 2007) to predict an extremely small overall impact on the economy. What if these assumptions are wrong?

They are almost certainly wrong for several reasons. Different price outcomes would not have linear or symmetric economic effects so we cannot focus only on the expected price even if that were unbiased. Very high prices would be very costly, while very low prices would not produce an equivalent gain. To address this we can both do supplementary analysis to try to reduce uncertainty and build policies that are more resilient to extreme outcomes.

Why are they likely to be wrong? All models (including those used to predict international carbon prices) are wrong and models trying to predict the effects of new policies in a highly uncertain international environment are even more likely to be wrong than usual. This is not a criticism but simply the reality of modelling. What modelling is good at is providing indications of likely values, insight into the drivers of those values, and some idea of the possible range of outcomes under reasonable assumptions.

Price predictions and predictions of access to international credit markets need to be presented with clearer uncertainty bounds. The price is driven not only by mitigation costs curves in all countries involved (about which information is pretty limited) but, more importantly, by enormous uncertainty about which countries will participate on what terms. The price could change dramatically if progress is made on a more stringent post 2012 agreement. Our access to AAUs is heavily controlled by political decisions by a few governments over whom we have no control so the domestic price may be driven by domestic supply and demand rather than international forces.

General equilibrium modelling is useful for understanding likely long-term structural shifts but it tells us nothing about adjustment costs or macroeconomic impacts (out of equilibrium). It assumes that overall employment is unaffected by the policy. Also the model used is not a particularly sophisticated one (though it is probably the best in New Zealand). Further, it assumes that free allocations are provided on the basis of output and hence will lower marginal costs; this is not in the current design. This mutes the price impact significantly in sectors with free allocation - up to 90% in extreme cases. The costs and structural shifts could be significantly higher than the Infometrics report implies.

Some of the other analysis in the Framework document may also be misleading. One example is the impacts on the agricultural sector. These are presented as though the marginal cost will have no effect on production and are presented relative to payouts rather than profit which may be a more meaningful long-term measure. A simple calculation¹ suggests that for dairy farmers a \$50 price in 2002 would have led to an 11% fall in revenue but a fall in profits of nearly 40%.

In several places the free allocations are implicitly or explicitly treated as though they will lower marginal costs, e.g. 'moderate this price impact', and hence reduce the size and rapidity of changes in production. In other cases the price impacts are presented in such a way that they are implicitly compared to price variations that we experience without the policy. These costs, however, are not the same as fluctuations in exchange rates or even commodity prices because they can be expected to persist and even grow; thus this comparison may be misleading. It is like comparing climate change to weather variability.

¹ Hendy, Joanna, Suzi Kerr and Troy Baisden (2006) 'Greenhouse gas emissions charges and credits on agricultural land: what can a model tell us?' Motu working paper 06-04.

3.2.1 What supplementary analysis would be useful?

On prices, we should not only consider the expected scenario (even assuming the Government is not being optimistic) but also consider the possibility of very negative scenarios (high prices and limited market access) that could have disproportionately severe costs. It would be useful to more closely consider scenarios where prices are driven by domestic unit availability alone not because they are entirely realistic but they would set a clear bound on the outcome in a worst-case access scenario.

We need to consider the potential size of structural adjustment costs and to what extent they could lead to macroeconomic effects. Models cannot do this formally but comparisons could be made between the effects of this change and previous changes that New Zealand has experienced such as the 1980s reforms. It is possible that these changes are much smaller (though they have similarly broad impact in that they affect all sectors) but that cannot easily be seen from the analysis provided so far.

Two relatively simple comparisons would be the size of structural shifts, and sizes of the regional employment shifts², predicted in a general equilibrium model compared to the shifts actually experienced between 1985 and 1990. These types of comparison would be only indicative but could either reassure people that this policy will have a much smaller impact and/or suggest the type and scale of transition and macroeconomic cost we might anticipate.

Even with fixed expected (average) price, the potential for significant price uncertainty is problematic in an immature credit market. Some firms, particularly points of obligation, will face very large changes in gross financial flows which they will need to learn to manage. If prices were stable this would not be so hard but this is unlikely to be the case. Secondary financial instruments such as options and hedges and futures markets will not be well developed in the early years so firms will be unable to manage their risk through the markets.

For example, the value of carbon flows could easily surpass the value of timber flows in forestry – foresters can protect themselves against this volatility by banking NZ units, but this may not be their best or only option as their business moves towards carbon management rather than forestry. It may however be difficult to implement any other management strategy in a highly uncertain environment. These firms may face significant risks in decision-making during a transition and learning period.

To understand the potential impacts of these risks it would be useful to have explicit analysis of the financial flows involved, both expected and the potential range, relative to existing financial flows in sectors likely to be heavily affected.

The type of information these supplementary analyses would provide could be used in two ways: first to inform discussions about how to manage the transition and second to inform our position in post 2012 negotiations. I will discuss only the former.

² Possible method in Kerr, Suzi and Joanna Hendy (2002) "Regional Employment Impacts of the Kyoto Protocol" Prepared for the Ministry for the Environment.

3.2.2 Implications for transition into ETS

If the adjustment costs might be high, is gradual entry of sectors the right way to deal with transition costs and uncertainty? Gradual entry effectively provides temporary exemptions to some sectors and moves all cost and risk to government/taxpayers. To a large extent it probably just delays the learning period. The argument that it gives these firms 'time to prepare' is probably pretty weak. History suggests that people only seriously consider how to respond to an emissions trading (or tradable permit) market when they are forced to.

Free allocation (through lump-sums) can reduce the shock of transition by providing a capital buffer both for those whose capital value has fallen and for those whose gross financial flows are rising. Even if the units are not given to those who are points of obligation, if they are part of the same supply chain the freely allocated units can be included in contractual relationships between the point of obligation and the firm that is allocated units.

A lower, more stable price during the adjustment may be a better option than either late entry or free lump-sum allocation for dealing with adjustment costs. This could be achieved by closing the domestic market to sales (purchases would be allowed) and prohibiting banking of NZUs at least until 2012, then introducing a safety valve where the government would offer unlimited NZUs for sale at a fixed price. If some sectors could come in earlier than currently planned, the cost to government of the price cap would be offset by early devolution of more emission obligations. To protect foresters from price risk, some banking would need to be allowed in forestry; the amount banked by each landowner would match the accumulation of future liabilities but no more.

4. Forestry

The overall structure of the forestry part of the ETS is good. I would suggest close attention to four issues as the details are developed further: simplicity of monitoring; consistency between treatment of plantation and indigenous forestry; more targeting of compensation for deforestation liabilities; and an explicit link between the afforestation grant scheme and the ETS.

It is always tempting to monitor carbon sequestration and storage as accurately as scientifically possible. The Permanent Forest Sink Initiative (PFSI) is an example of a system that involves considerable complexity. We need to recognise that accuracy comes at a cost. It creates high compliance costs, which will reduce participation by small and less sophisticated landowners (including many Maori land blocks). There is significant potential for sequestration (both plantation and indigenous) in small blocks. Complexity creates a lack of transparency and potential for manipulation because data and calculations cannot be easily replicated. These costs need to be weighed carefully against the gains. A two-tier system is a possible solution, with simple but conservative monitoring for small blocks and more accurate monitoring for larger properties.

There needs to be consistency in the treatment of plantation and indigenous forestry. This reduces the potential for perverse incentives that could lead to loss of indigenous forest. It also reduces complexities if plantation forest is converted into indigenous forest, whether this transition occurs gradually by abandonment or at the point of harvest, .

The current approach of providing free allocation to offset deforestation liabilities treats the forestry industry as though it is one entity, which is obviously untrue. Compensation needs to be more closely targeted toward those who will face large losses. The difficulty is identifying who these people are.

One approach to this, which targets those with the highest per hectare losses (but offers no compensation for those with low losses), is a tradable deforestation permit system, where the government issues the same number of units as currently planned but makes them usable only to offset deforestation. The government can sell them through an auction where they offer to sell the fixed number as well as extras that would be added to the auction pool at the international price. The price paid by buyers would be the average cost (those already provided free by government would be priced at zero), so that the government makes no money from the auction (but also faces no fiscal risk). Once auctioned they can be traded and banked.

Those who value the ability to deforest most (i.e. those who lose most by not being able to deforest) will buy the units. While they will probably pay a positive price, it will be significantly less than the full cost if they needed to buy NZUs so they will be partially compensated. If the true desire to deforest is actually below the amount the government has agreed to cover, the price of deforestation permits will be zero and those who deforest will face no liability. The downside of the system is that the incentive to deforest will be inefficient to the extent that the price of deforestation permits is below the market price. This is a 'revelation mechanism', which leads people to voluntarily reveal their true value – this year's Nobel Prize in economics was awarded for this type of mechanism.

The motivation for the afforestation grant scheme is to provide a simpler mechanism to reward carbon sequestration. The need for this could be reduced if the standard ETS mechanism for forestry were made as simple as possible. If it is still desirable to reduce the financial risk and complexity of participation in the ETS, the grants offered should be explicitly linked to the rewards the landowner would have received if they had done the same thing with their land and received benefits through the ETS. The government will simply be acting as a broker who provides an alternative sequestration 'contract' with up-front payments, less monitoring and more restrictive conditions.