WORKING PAPER

WORLD RESOURCES INSTITUTE

Biofuels and the Time Value of Carbon: Recommendations for GHG Accounting Protocols

LIZ MARSHALL

Marshall, Liz. 2009. “Biofuels and the Time Value of Carbon: Recommendations for GHG Accounting Protocols.” WRI Working Paper. World Resources Institute, Washington DC. 13 pp. Available online at http://www.wri.org/publications.

World Resources Institute
10 G Street, NE
Washington, DC 20002
Tel: 202-729-7600
Fax: 202-729-7610

www.wri.org

April 2009

SUMMARY

The quantification of the carbon dioxide emissions impact associated with land-use change for biofuels production is complicated by the fact that the carbon costs from land-use change and the avoided emissions from substituting biofuels for fossil fuel in transport occur over an extended period of time. Estimating the net carbon impact therefore requires a method for aggregating the increased and avoided emissions that play out over time into a single figure. The choice of accounting method can have a significant impact on the resulting net emissions measure for specific land-use options such as biofuels production. This in turn will influence the relative desirability of different land management scenarios for a given piece of land. Traditional cost-benefit analysis regularly uses discounting to compare and aggregate monetary units over time. However, extrapolation of this approach to assess physical units of carbon dioxide emissions released or avoided in the future is not straightforward. Selection of an appropriate discount rate for physical carbon units requires a consideration of multiple additional variables. These include rates of carbon accumulation and decay in the atmosphere and estimates of the marginal damages arising or avoided from changes in atmospheric carbon stocks.

Accounting recommendations for quantifying the emissions impact of land-use change for biofeedstock production:

1. Ideally, a GHG accounting method for land use change associated with biofeedstock production should explicitly analyze the expected damages associated with those flows over time. The corresponding monetary units associated with this damage can then be discounted to determine how the impacts of future flows compare to those of the present.

2. There is little theoretical justification for discounting physical carbon flows. Discount rates used for physical carbon units are not analogous to monetary discount rates such as interest rates or the social rate of time preference. They therefore should not be selected based solely on an extrapolation of how those financial discount rates are usually applied.
3. The "project horizon" should be considered independently of the longer atmospheric "impact horizon" when selecting appropriate discounting horizons. In the context of biofuels production, the "project horizon" refers to the period of time over which feedstock cultivation will occur (and benefits from displaced transport fossil fuel realized). The "impact horizon" refers to the period of time over which impacts of increased or decreased emissions are felt in the atmosphere.
4. The impact horizon should be applied as a rolling target that is measured relative to the year of emissions, which can occur at any point over the project horizon, rather than as a fixed target that is measured relative to year 0 of the project. Atmospheric impacts are therefore fully accounted for, whether the emissions or emissions savings occur at the end of the project or at the beginning.
5. When it is necessary to bypass the full-cost accounting suggested in #1, selection of a next-best discount procedure for carbon units may need to consider: a range of possible discount rate values beyond those normally used for financial discounting (including zero or negative numbers); different discount figures for the two distinct time horizons; and non-constant numbers such as declining discount rates for the longer impact horizon.
6. Salvaged carbon from acreage reversion or revegetation should not be considered as part of the GHG accounting protocol for land-use conversion for feedstock production. Carbon benefits associated with revegetation are not guaranteed when acreage is initially converted to biofuels production, and should more appropriately be considered a benefit associated with a future form of land-use change should such conversion occur.

INTRODUCTION

In the United States, biofuels have been promoted as an alternative to petroleum-based fuels. The anticipated benefits include boosting rural economies, promoting energy independence, and reducing the greenhouse gas emissions associated with the transport sector. In recent years, however, experts have raised concerns about the magnitude of greenhouse gas (GHG) emissions arising from land-use change associated with the production of biofuel feedstocks. These concerns raise questions about the claims of GHG benefits associated with biofuel production and use (Searchinger, 2008).

There have been two major legislative efforts to support the use of biofuels in the U.S. as a transport fuel alternative. In 2007, California governor Arnold Schwarzenegger signed into law an executive order calling for a "Low Carbon Fuel Standard" requiring that California's transport fuel supply contain 10 percent less carbon dioxide than an equivalent amount of gasoline by 2020. Also in 2007, the Energy Independence and Security Act of 2007 (EISA 2007) expanded the existing 2005 Renewable Fuels Standard (passed as part of the Energy Policy Act of 2005) to require that additional amounts of renewable fuels be blended into the nation's fuel on an increasing schedule through 2022. In response to concerns about the GHG content of the biofuel, however, EISA 2007 departed from the original 2005 Renewable Fuels Standard (RFS) and added a GHG content threshold above which biofuels would not be considered as qualifying for the RFS.

The regulatory processes supporting the California executive order and EISA legislation therefore require quantification of the GHG content of the biofuel in order to ensure that they satisfy the GHG requirements of the law. Both the California Air Resources Board and the U.S. EPA are in the process of designing the necessary quantification methodologies for this purpose. In both cases, GHG "content" is broadly interpreted to mean a life-cycle-based measure of all GHGs emitted throughout the production and transport of the fuel and its major inputs. The carbon dioxide released when biofuels are combusted as fuel is not included in such accounting because it is "short-cycle" carbon that was absorbed from the atmosphere as the biofuel feedstock grew. However, other GHG emissions arise at several stages in the production of biofuels that must be accounted for, including the GHG emissions associated with clearing or converting land, growing and fertilizing the feedstock, transporting the feedstock, and converting the feedstock into fuel.

Until recently, the possibility for significant carbon dioxide emissions associated with land-use conversion for feedstock production was largely neglected. Although early life-cycle accounting of the GHG emissions of biofuels emphasized the importance of considering the underlying land-use change that enables feedstock production, few had the tools necessary to quantify those impacts (Farrell et al., 2006; Hill et al., 2006; Zah et al., 2007). Recent research reports, however, have suggested that the potential magnitude of these emissions is significant; Searchinger (2008) estimates that in some circumstances, the carbon dioxide release associated with indirect land-use impacts may be comparable to the carbon dioxide release from the entire rest of the fuel's life-cycle, thereby nearly doubling estimates of the carbon content of the final fuel.[1] To address such concerns, EISA 2007 specifically instructs that the quantification methodology developed to support its GHG performance goals include "significant indirect emissions such as significant emissions from land use changes." Both regulatory efforts are therefore driving the development of tools to quantify the GHG emissions from land-use change for feedstock production, and along the way unearthing and grappling with significant methodological issues associated with carbon quantification.

One such quantification complication arises because carbon emissions from land-use change, and the avoided emissions from substituting biofuels for fossil fuel in transport, are ongoing over time. Efforts to quantify the net emissions associated with land-use change and attribute those emissions to current biofuel production or biofuels policy, therefore, must design an accounting methodology that allows them to compare and aggregate such emissions into a single figure that can be compared across fuels and across other policy options for reducing GHG emissions.[2] This paper will address the methodological accounting challenges that such ongoing emissions, and emissions reductions, pose for quantification protocols as well as the complexity of the strategies required to overcome them.

QUANTIFYING THE CARBON IMPACTS OF LAND USE CHANGE OVER TIME

When land is converted from forest to agriculture, there is a large initial carbon loss in the form of above-ground biomass.[3] This carbon may remain somewhat inert if it is redistributed as wood products but is immediately released if the land is cleared through burning. Soil carbon is then released over time as the land remains in cultivation, until the soil reaches a lower soil carbon equilibrium under the new production regime. The "carbon cost" of conversion is an aggregation of these factors, which play out over time, corrected for whatever above-ground carbon sequestration capacity the new land use is able to offer. Conversely, the carbon benefits of the new land use are fairly constant over time, and they comprise the displaced petroleum emissions gained from using the biofuel grown on that land.

Calculating the net carbon costs of land-use conversion requires that regulators somehow aggregate these costs and benefits over time in order to be able to compare the two paths in similar units. For such aggregation, it is necessary to select two important analytical parameters: 1) a time frame for analysis that sets temporal boundaries within which emissions or benefits are counted, and 2) some sort of weighting scheme that allows the analyst to compare one unit of emission (or displaced emission) that occurs today with a similar unit that occurs at variable points in the future. These parameters are familiar in economic analyses. In cost-benefit analysis they are referred to as the "amortization period" and the "discount rate", and they represent critical policy decisions about the relevant time frame for analysis and how much future emissions (or savings) will "matter" relative to those today.

Unfortunately, there are few practical or theoretical guidelines that can be used to help select an appropriate set of parameters for aggregating the ongoing GHG emissions costs and benefits associated with biofeedstock production. Consider these parameters in the case of land conversion for biofuels production. In theory, converted land could be used for hundreds of years for the production of ethanol to displace fossil fuel. Should we therefore consider the total benefits of land conversion to be the yearly benefits of petroleum displacement aggregated over hundreds of years? Given changes in transport and fuel technology as well as expected changes in the availability of oil, it is unlikely that the plot of land will remain in continuous biofuel production, with a continuing impact on petroleum displacement, for that many years. But how does one choose the "appropriate" time frame for analysis? Do we expect 20 years of carbon benefit from the production of biofuels on that land? Or 30 years? Or 100 years?

Although the carbon costs of land conversion are often concentrated in the first few decades after conversion, as above-ground biomass degrades and soil carbon is released, clearly the estimated *net* benefits of land conversion will be highly sensitive to the amortization period.[4] In an attempt to avoid this issue, some experts have presented their benefit/cost comparisons in terms of "payback time." Payback defines the time period that land would have to be used for feedstock production in order for net GHG impacts to become positive (i.e., the point at which GHG benefits from displaced petroleum exceed the GHG costs associated with fuel and feedstock production). While this is a very useful figure for intuitively illustrating the magnitude of the problem, it does not preclude the need to make some sort of judgment about what time frame is relevant for assessing and comparing the costs and benefits to determine net benefits from a land conversion decision.

But what arguments can one use to establish a "relevant" time frame? Some analyses have defaulted to a 30-year analytical time-frame (Righelato, 2008). Use of that time frame is sometimes based on the argument that ethanol plants are estimated to have a 30-year life-span, so feedstocks are likely to be produced for at least that long once land is converted. This number is, however, a shot in the dark—plants could be repurposed or decommissioned, and land use can be changed in response to changing market conditions and policy and the increasing scarcity and value of environmental goods and other ecosystem services—but at least it provides a concrete boundary for the GHG analysis. A report commissioned and released by the U.K.'s Renewable Fuels Agency, on the other hand, implicitly lobbies for a much shorter analytical time frame, arguing that net carbon benefits should become positive when aggregated over a much more conservative10-year pay-back period (Renewable Fuels Agency, 2008).

While the amortization period determines the time range over which costs and benefits will be considered, the discount rate selected determines how a unit of emissions at the beginning of that time frame will be weighted relative to a unit of emissions at the end of the time frame in aggregating carbon costs and benefits. A discount rate of zero weights all carbon emissions across time equally, while a high discount rate suggests that carbon emissions (or avoided emissions) in the future should be considered less harmful (or less beneficial) than carbon emissions or avoided emissions today. The 30-year studies above chose a discount rate of zero, but the EPA is considering use of a non-zero discount rate in its GHG calculation methodology under RFS II. Unfortunately, selection of an "appropriate" discount rate for such purposes is as problematic as selection of an amortization period.[5]

Discount rates in cost/benefit analysis are generally used to capture observed decision-making behavior in capital markets. It is argued that in a world with scarce resources for investment, we should compare growth rates of other capital investments in deciding our optimal investment paths over time. The discount rate therefore captures some measure of the opportunity cost of *not* investing in other capital-improvement activities and instead investing in the project under consideration. That opportunity cost should also reflect a risk premium arising from the uncertainty associated with future outcomes of that investment decision (Howarth, 2005).

Because discount rates are generally used in the context of investment decision-making to reflect the "time value of money", they are usually applied to monetary units, such as costs or benefits, rather than to physical units such as tons, million metric tons of carbon equivalent (MMTCE), or lbs per acre. Although the practice of using discounting to estimate the "time value of carbon" in assessing carbon mitigation options is becoming more common (Stavins and Richards, 2005), a great deal of disagreement exists about the validity of applying discounting principles to carbon units. In an early analysis of carbon discounting, Richards (1997) concludes: "(T)he choice of whether and how to treat the time value of carbon emissions reductions depends very much upon the policy context for which the analysis is designed."

To understand the practical implications of incorporating a discount rate into GHG accounting methodologies, consider the question of temporary carbon storage. Put simply, is there any reason to invest in mitigation projects that will capture carbon today and then release an equivalent amount of carbon in 50 years? Ideally, this study would be conducted as a cost/benefit analysis, with explicit inclusion and comparison of emission cost and benefit functions over time. It would be *de rigeur* to include a discount rate in such an analysis, though interested parties may never agree on what that discount rate should be.

In practice, however, explicit cost and benefit functions for carbon emissions are often not available to analysts, nor are the resources to develop them.[6] GHG accounting methodologies therefore instead address whether a "net carbon benefit" exists by focusing on the physical carbon unit itself. In the temporary storage case described above, a discount rate of zero would yield a net carbon benefit of zero, suggesting that such a project would be neither beneficial nor harmful from a greenhouse gas perspective. A positive, non-zero discount rate, on the other hand, would yield a positive carbon benefit. The decision about whether the estimated carbon benefit would be "worth" the cost of the mitigation project then would depend on additional analyses about project cost and comparison to other mitigation options.

When transferring the discounting practice over to physical units, it is important to recognize that, despite a failure to include explicit benefit and cost curves in the analysis, the estimated time value of carbon is nevertheless a function of underlying curves that are assumed to drive changing "carbon values" over time. In such studies, the discount rate must therefore capture more than just the "time value of money" dynamic generally associated with discounting practices. An appropriate physical carbon discount function form and rate must *also* reflect very complicated relationships among variables such as the rate of change of the damages produced by atmospheric GHG stocks (which reflects changing assumptions about available mitigation technologies), the persistence rate of GHGs in the atmosphere, initial GHG stock levels, etc. (Richards, 1997). Simple extrapolations from default monetary or market discount rates, or even the lower "social rates of time preference" often used in intergenerational analyses, are not appropriate except under very restrictive assumptions about the shape of the marginal damage curve from carbon emissions and its relationship to atmospheric stocks.

Dissecting the physical discounting process

The purpose of comparing physical carbon emissions in the future to physical carbon emissions in the present through some sort of discounting procedure is essentially to evaluate how the value of the damage caused by a unit of emissions in the future will compare to the value of the damage caused by a unit of emissions today. The process of applying a discount rate to carbon tonnage is a "short cut" to information about how the value of damages changes over time that skips a series of important steps related to translating physical impacts into economic impacts. To clarify what we are doing when we discount carbon tons, it helps to first describe the step that we are skipping. It is only by understanding the steps that we are leaving out that we can understand whether, and how, physical carbon discount rates can be used as an appropriate short cut to achieve the accounting objectives.

Figure 1 and the following sections will attempt to clarify those steps and to introduce a more precise vocabulary for the discussion of carbon discounting in the context of GHG accounting for biofuel-related land-use change. Creation of a common understanding of the complex concepts that underlie carbon discounting and how they impact the value of the discount rate and other time-related parameters that should be used in carbon discounting is a critical step toward using discounting appropriately in GHG quantification methodologies.

There are two distinct time horizons illustrated in Figure 1: the "project horizon" and the "impact horizon." In the context of

Figure 1. Distinct time horizons exist when evaluating GHG emissions impacts over time.

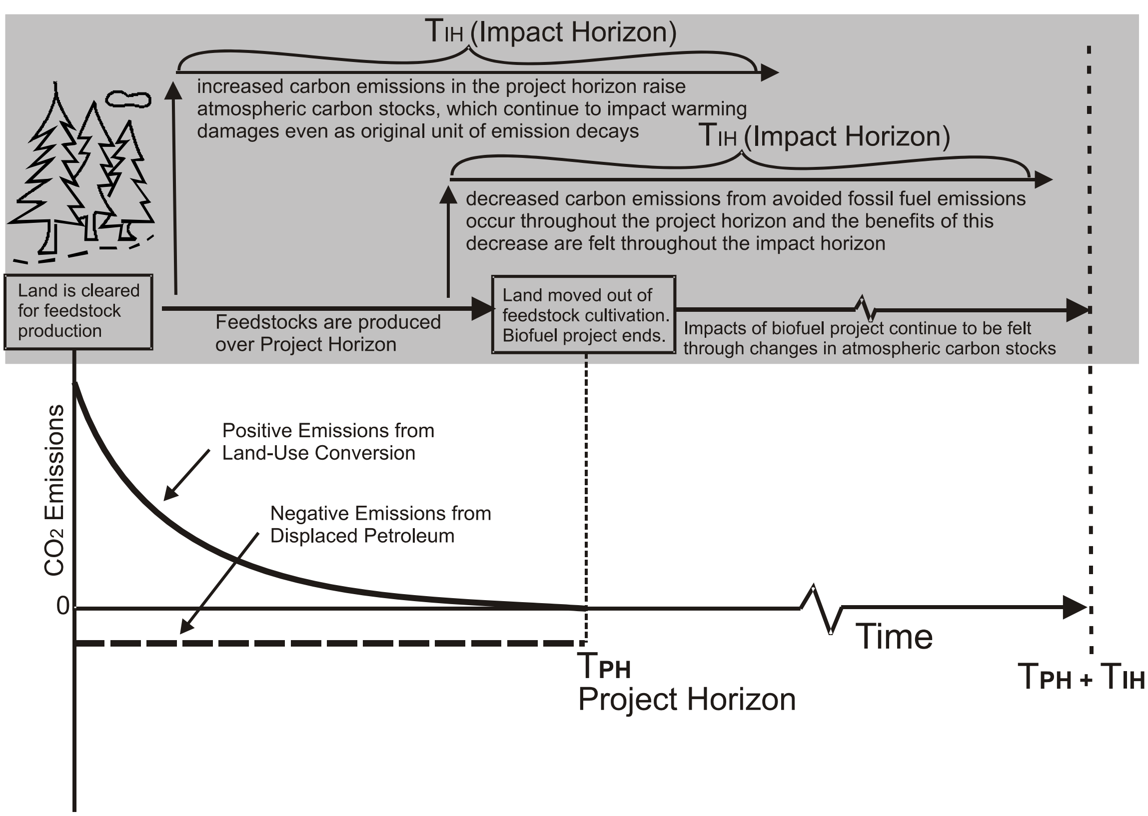

land-use conversion for biofeedstock production, the project horizon refers to the period of time over which biofeedstock production on that land will result in avoided petroleum fuel use. This is, in a sense, the "lifetime" of the biofuel project that is driving the initial land-use conversion to biofeedstock production, or the length of time that biofeedstocks will be produced on that land before the land moves into some other use.

The "project horizon" is a planning construct. It represents a prediction about how long converted land is likely to remain in feedstock production. That prediction captures the period of time over which benefits from reduced emissions due to biofuel production on that land will continue to be generated through avoided petroleum use. There are several factors that could shorten the expected cultivation time, including: the advent of alternative transport fuel technologies such as electricity, the commercialization of waste-sourced biofuels to replace crop-based biofuels, and policy changes such as reduction or elimination of subsidies to biofuels or biofeedstocks.

The "impact horizon" on the other hand, is largely a physical construct that reflects how long a unit of emissions, once it enters the atmospheric carbon stock, continues to significantly contribute to warming and the damages caused by that warming. Because greenhouse gases persist in the atmosphere and produce warming over time, the damage created by a unit of emissions in any time period includes a stream of warming potential into the future. The "impact horizon" is likely to be much longer than the "project horizon" because, although the emissions reductions associated with biofuel production will cease as soon as the land is moved out of feedstock production, the atmospheric benefits of those reductions continue. Similarly, the atmospheric impacts of the carbon dioxide emissions from the initial conversion will continue to be felt long after the land has moved into other uses. The distinction between these two time periods reflects the momentum of land-use decisions made within the project horizon by acknowledging the persistence of emissions in the atmosphere and the cascading impacts of those emissions over time on the damages expected from global warming.[7]

Appropriate GHG accounting for biofuels-related land-use change must recognize the distinction between these time horizons. Designing a quantification scheme around a single time horizon that equates the impact horizon with the project horizon creates tension in the establishment of an appropriate length for that single horizon; extending the single horizon allows one to capture the implications of persistent carbon in the atmosphere, while shortening it makes it more reasonably reflective of how long land is likely to stay in cultivation. In fact, the time scales of the two horizons are completely different and should be treated as such.

Selection of an appropriate discount rate or rates once relevant horizons have been identified

As illustrated above, there are two distinct time horizons that must be considered in such analyses. Each of the distinct time horizons has its own associated stream of impacts and its own challenges for aggregating those impacts over time. Each separate aggregation procedure requires careful consideration of an appropriate discount rate for that aggregation (Figure 2).

Consider first the "impact horizon", which encompasses the path of warming impacts that result when a unit of carbon is emitted, regardless of when that emission occurs. The objective of aggregating over that time horizon is to associate a unit of carbon emissions in a given period with a single measure of

Figure 2. Discounting emitted carbon tons is a short cut for estimating two distinct rounds of "implicit" discounting, or carbon weighting.

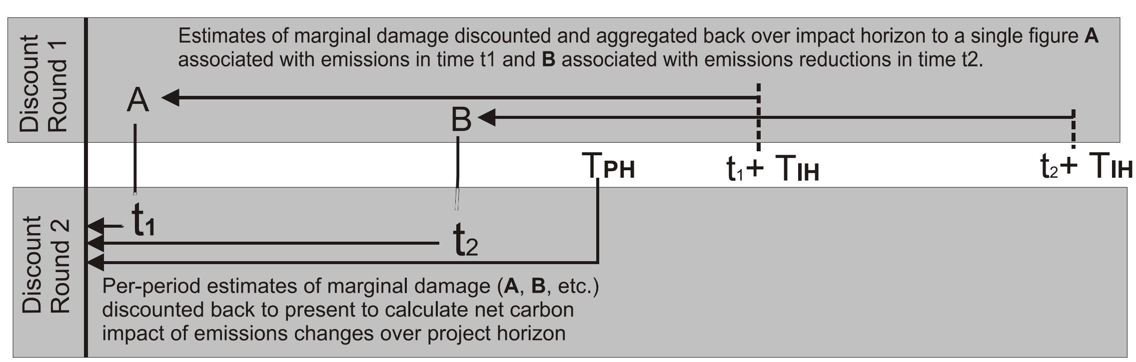

damage that reflects the "cost" of that emission over time, or, conversely, the "benefit" of preventing that emission in that time period. There are several variables that affect the path of damage over time that is expected from a unit of emissions. One of these is the rate at which atmospheric carbon decays over time as carbon is re-absorbed into biotic sinks such as forests and oceans. The way in which this decay is represented varies, with some authors using a fixed decay rate applied to atmospheric stocks (Richards, 1997) and others using an exponential decay function that reflects a declining rate of carbon decay over time (Fearnside et al., 2000). In both cases, this variable reflects the purely physical dynamic of the persistence of carbon in the atmosphere over the impact horizon and translates a unit of emissions into an atmospheric carbon stock impact over time.

The second relationship defining the path of damage expected from a unit of emissions is the relationship between carbon stock and the damage expected from that stock. This relationship translates the physical stock dynamic described by the decay function into a measure of the cost implications of that stock response and moves the "impact horizon" into the realm of economics. Although there are many simplifying assumptions used in different analyses of carbon stock damage over time, such as the assumption that marginal damages are not stock-dependent at all or that they are linearly related to stock, the reality of this relationship is likely more complicated than such assumptions suggest. Although such simplifications improve the analytical tractability of the problem, they are difficult to justify for any other reason.

So in any time period, a unit of emissions is associated with a path of expected damages over time that reflects both the impact of that unit on atmospheric carbon stocks over time and the impact of those carbon stocks on damages from global warming over time. Integrating that damage path over the impact horizon produces a single value for the expected costs associated with a unit of emissions in a given time period. Because these impact figures are monetary, one might also include an economic discounting term in that aggregation procedure in order to reflect the "time value" of the cost and benefit numbers. (Failure to use a discount rate can be considered simply a special case of discounting where the discount rate chosen is equal to zero.)

Once a path of emission damages has been condensed into a single cost number associated with a unit of emissions (or a single benefit number associated with an avoided unit of emissions) in each time period, the second round of aggregating over time occurs. In the second round, the objective of the aggregation is to calculate a single total present value of all the carbon emission costs and avoided emission benefits that occur over the project horizon. Unlike the first round of aggregation, this is a fairly straightforward process of discounting cost and benefit figures over a finite time horizon using economic discounting.

It is quite likely that appropriate discount rates will differ between the project horizon and the impact horizon. Selection of an appropriate discount rate for the impact horizon should consider the relevant biophysical variables described above, and the emerging literature on declining discount rates and the role of uncertainty in discounting over long periods (Guo et al., 2006). The discount rate used over the shorter project horizon, on the other hand, may reflect the higher interest rates used to capture market opportunity costs over shorter investment horizons. The result of such an analysis could be very different discounting structures applied to the two distinct time horizons.

Complications in the application of monetary discount rates to physical carbon units arise when "current value" estimates of marginal damages from a unit of carbon emissions are expected to change over time. "Current value" estimates are estimates of marginal damage expressed in terms of the value at the time of emission. In the scenario illustrated in Figure 2, these values correspond to the values A and B. These values have been calculated using a discount structure from the time of emission forward, but that value has not been discounted back to the present.[8] If A=B for all time periods in the project horizon, then regardless of the discount rate structure applied to the impact horizon, the appropriate discount rate to apply to carbon units is whatever discount rate is selected as theoretically appropriate for the project horizon discount procedure illustrated above.[9]

The assumption of constant marginal damages is a very limiting case, however. There are many possible causes of non-constant marginal damages over time. These possible causes include atmospheric carbon degradation rates that vary with atmospheric carbon stock and paths of marginal damage that vary non-linearly with atmospheric carbon stock. The former dynamic would exist, for instance, if greater atmospheric carbon levels result in faster dissipation of carbon from the atmosphere through carbon fertilization impacts, or impacts of increased carbon on absorptive capacity of terrestrial and ocean carbon pools. Non-linear marginal damages exist if the impact of an equivalent change in atmospheric stock is expected to

Figure 3. Discounting rounds with a fixed impact horizon.

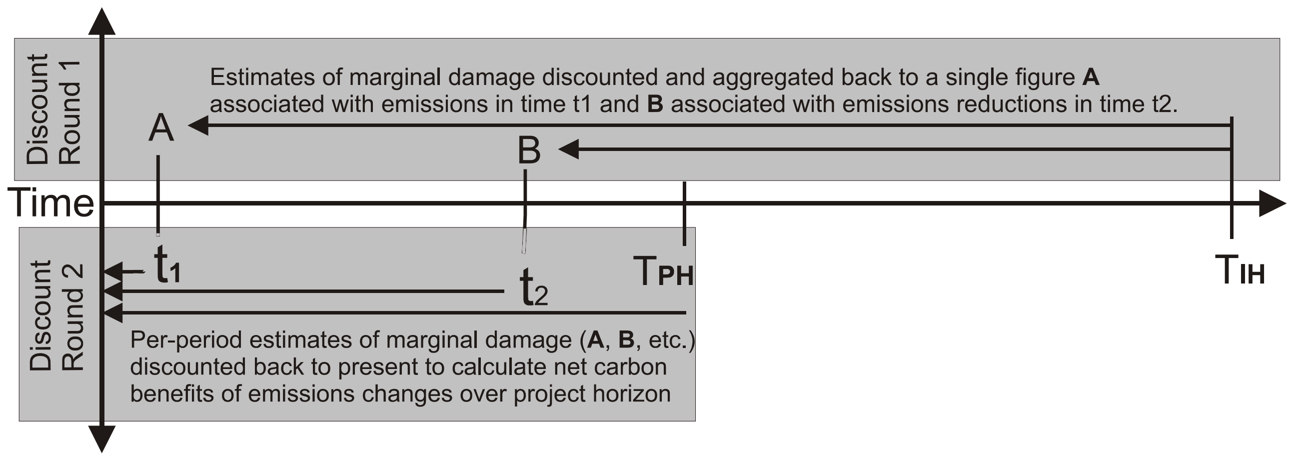

vary depending on the original stock level. Catastrophic atmospheric carbon thresholds are an extreme example of non-linear impacts; damages that are assumed to be a quadratic function of atmospheric stocks are another.

In a theoretical exploration of the concept of discounting physical units, Richards (1997) arrives at the following generalizations (which have been reworded to fit the context described here):

- If the marginal damages from emissions are growing over time (i.e. if B > A in Figure 2), then the discount rate chosen for the project horizon will be higher than appropriate for application to carbon units.
- If marginal damages are growing over time at a rate equal to the discount rate that has been chosen as appropriate for the project horizon, then the appropriate discount rate to apply to physical carbon units is zero.
- If marginal damages are growing very quickly over time, then emissions reductions later in time have higher value than earlier reductions, and the appropriate discount rate to apply to carbon units may even be negative.

The increasing marginal damages over time can be caused by a rapidly increasing atmospheric carbon stock, or by a marginal damage function with rapidly increasing damage as a function of stock. Either of those scenarios will cause marginal damages to increase rapidly over time, which causes the appropriate carbon discount rate to fall below the "project horizon" discount horizon, and possibly even fall below zero. A negative carbon discount rate will bias the analysis toward projects with current emissions over those with later emissions (or with later reductions over those with current reductions).

Note that if marginal damages are increasing at a non-constant rate, it is likely that an appropriate carbon discount rate will also be non-constant. In the scenario where marginal damages from emissions are assumed to be increasing at an increasing rate with atmospheric carbon stock, for instance, an appropriate physical carbon discount rate structure is one with a discount rate that declines over time at a decreasing rate.

ADDITIONAL NOTES ON ACCOUNTING STRUCTURE AND VARIABLES

Rolling versus fixed impact horizon

Note in the graphs above that the impact horizon is depicted as a rolling horizon. In other words, the impacts of a unit of emissions are measured over the same number of years, regardless of whether that emission takes place at the beginning of the project horizon or at the end. The alternative scenario would be a "fixed horizon." A fixed impact horizon is measured relative to year 0 in the accounting methodology, rather than relative to the year in which the emission occurs, so that the impacts of emissions in later years are measured over fewer years than the impacts of emission in earlier years. For a fixed impact horizon, Figure 2 would be modified to appear as in Figure 3.

The problem with establishing a fixed impact horizon is that this methodology will automatically favor projects whose emissions are deferred to the end of the project horizon.[10] This bias occurs because the impact of emissions occurring at the end of the project horizon is measured over fewer years than the impact of emissions occurring early in the project horizon; it is an artifact of the measurement truncation that does not reflect a legitimate difference in damage incurred between early and late emissions. In the context of emissions quantification for biofuel production projects, this bias means that the early carbon costs associated with the initial conversion will be weighted relatively more heavily than the later benefits associated with displaced carbon emissions from avoided gasoline use. Although such a result may emerge analytically from use of certain marginal damage functions or from use of a non-zero discount rate, there is no theoretical justification for artificially exacerbating that effect through use of a fixed impact horizon. For that reason, impact horizons should usually be measured on a rolling basis as shown in Figures 1 and 2.[11]

The dissipation rate of GHGs in the atmosphere

The path assumed for carbon decline in the atmosphere can significantly impact the appropriate carbon discount rate through its impacts on the path of marginal damages expected from a unit of emissions. Although for analytical ease it is tempting to characterize carbon decline as a fixed proportion of stock, as Richards (1997) does, in fact the precise path of decay is more complicated than that. The 1996 IPCC revisions, for instance, described an atmospheric carbon decay model with a more rapid decline in early-year atmospheric carbon than prior reports had. Fearnside et al. (2000) found that using the revised stock decline model significantly increased the value of temporary carbon sequestration, suggesting that a higher carbon discount rate would be appropriate with the revised expectations about stock decay.

Although not relevant to the land-use related carbon accounting issue discussed here, which focuses specifically on carbon impacts, the importance of dissipation rates in influencing appropriate discount rates becomes particularly important when comparing mitigation efforts across GHGs with different atmospheric persistence rates. The existence of different "half-lives" in the atmosphere suggests that there might be significant differences in the "appropriate" emissions discount factor applied across different GHGs.[12]

Damage functions, technological change, and the risk of catastrophic "phase shift"

One of the defining characteristics of the damage functions associated with atmospheric carbon stock system is the potential for irreversible change in the form of melting ice caps, changing ocean current patterns, etc. when certain atmospheric carbon stock and warming levels are reached. Although this risk is often ignored as a simplifying assumption in analyzing future costs of climate change, the existence of irreversible tipping points or "phase shifts" implies that GHG emissions from the present cannot be fully mitigated by a comparable level of sequestration once that phase shift has occurred. The potential for irreversible change is one of the significant determinants of the expected damage function for GHG emissions that must be considered in determining how to compare current to future emissions, and is one of the most convincing arguments for the need to make some sort of distinction between current and future, or pre-change and post-change, emissions.

In any scenario with an increasing risk of catastrophic system change, or phase shift, as atmospheric carbon stocks increase, the possibility that current emissions may expedite such a collapse must be considered in determining how current GHG emissions compare to future carbon emissions. The appropriate discount rate will depend on the assumptions made about this risk and about exogenous changes in technology that can help reduce that risk. This argument reflects the "buying time" justification for carbon discounting, which states that current emissions should be considered more important than future emissions because in the future there will be more technological options for mitigating carbon emissions. According to that argument, weighting current carbon emissions more heavily than future emissions therefore "buys time" for mitigation technology, such as carbon capture and storage, to be developed and implemented.

This argument, however, is critically dependent on the premise that technological improvement will increase quickly enough to out-pace increases in marginal damage arising from increasing atmospheric stocks. That premise reflects embedded assumptions about the relationship between stocks and marginal damages and the rate of change in available mitigation technology. As described by Richards (1997), cases can exist where rapid growth in marginal damages over time leads to the conclusion that later carbon emissions are more important to present expectations of damages than current emissions are. In such cases, later reductions are considered more valuable—and later emissions more harmful—than their "current" counterparts, suggesting that a negative carbon discount rate may actually be the most appropriate for capturing the behavior of expected emissions damages over time under certain scenarios.

Reversion of feedstock production acreage

Several researchers have raised the possibility that revegetation of land after feedstock cultivation could lower the net carbon impact of land conversion for biofuel production by re-sequestering some of the carbon originally released (Delucchi, 2008). Some stakeholders argue that it is an error to neglect this possibility in GHG quantification for biofuels, as a failure to account for this "salvaged carbon effect" would result in an overly large carbon cost associated with initial land conversion.

It is certainly true that managed reforestation of retired feedstock acreage could recover a significant amount of lost carbon and that even unmanaged land abandonment might result in a slight recuperation of carbon losses. However, in the absence of post-project polices that guarantee that lands will be revegetated or rehabilitated, there is no assurance that "salvaged carbon" will be reclaimed. It is also possible that land would be converted to food production, grazing, or development, and additional losses could be incurred at that time. Because post-project land-use policies would be difficult, if not impossible, to implement and enforce, it is more appropriate to consider post-project salvaged carbon value as part of a second, independent land-use change that occurs when the biofuel project itself has terminated. We do not believed, therefore, that this "salvaged carbon" should be included in the quantification of the carbon associated with biofuels-related land-use change.

It is worth noting additional concerns about the argument that loss of biomass-based GHG sequestration is reversible and can therefore be "undone" at the end of the project horizon with revegetation of the land area used. Research in the Amazon suggests that land-use activity in the forest increases risk of forest fire, causing additional carbon losses in neighboring forests, and that such fires increase the forest's susceptibility to further burning (Nepstad et al., 2008). Such land-use changes are also associated with irreversible changes such as fragmentation of existing natural habitat, expansion of degraded "edge" habitat, and loss of native species and biodiversity. The potential for irreversible change along other social and environmental dimensions highlights the need for a more comprehensive definition of the sustainability of biofuel production than that captured by the GHG requirements alone.

Assumptions about endogeneity versus exogeneity of atmospheric carbon stocks in selection of an appropriate discount rate

The process of selecting a carbon discount rate that is described here is highly sensitive to the assumptions made about the behavior of atmospheric carbon stocks. In selecting an appropriate carbon discount rate, one can assume that atmospheric carbon stocks are changing either exogenously or endogenously. An assumption of exogeneity means that the analyst assumes some path of atmospheric carbon change that is outside the control of the analyst, and therefore the analytical task is simply to explore the marginal impacts of a small set of projects on that "predetermined" path of atmospheric carbon stock. This is, in a sense, a private, project-based approach to selection of a discount rate. The analyst takes the aggregate impact of all other emissions behavior on atmospheric carbon stocks as "given", and then explores the impact of their own emissions paths within that scenario in order to select an appropriate discount rate to apply to carbon units for their project.

A more "social" approach to selecting a discount rate is to assume that atmospheric carbon stocks are changing endogenously. In that case, it is assumed that the whole path of atmospheric carbon change will respond to the discount rate chosen, so that there is no portion of the stock change that is "given" and independent of the discount rate chosen. Clearly, such an assumption is only appropriate if all projects involving carbon emissions evaluate their emissions over time using the same discount rate and therefore make emissions decisions using the same expectations of future atmospheric carbon stock behavior and the damages that could arise.

The generalizations provided in this paper about the relationship between non-constant marginal damages and an appropriate carbon discount rate are based upon an assumption of exogenously changing atmospheric carbon stocks. Given the global debate about atmospheric warming, and the uncertainty surrounding the various impact and cost relationships described in this paper, it is unlikely that anything approximating a "consensus" carbon discount rate will be arrived at anytime soon. However, any game theorist can tell you that the behavior emerging from a scenario in which all players make decisions taking everyone else's behavior as exogenous is often not the "best case" outcome. It is therefore critical when discussing project-based selection of a discount rate to understand how that process fits into the broader, social context of selecting an optimal emissions path and discounting structure.

In this context, imagine that projections of rapidly increasing atmospheric carbon stocks lead project managers or regulators to select lower carbon discount rates that bias project analysis towards projects with current rather than later emissions. The logic underpinning that bias is that since later atmospheric carbon stocks will be large and will cause later emissions to have large damage impacts (regardless of what the project does), it is better from a project perspective to emit carbon now, when impacts are smaller. But if everyone exercised that logic in planning projects with carbon impacts, the result could be a relatively rapid escalation of atmospheric carbon and a rapid incurrence of increased damages.

If instead all carbon players cooperated in choosing a carbon discount rate, the selected discount rate, and the path of expected atmospheric carbon and damages, could be completely different. In such a case, the appropriate carbon discount rate, if any, would depend upon the relative values of the marginal costs of control versus marginal damage and how they change over time (Richards, 1997). The choice of an appropriate social discount rate emerges from the broader choice of an optimal path of emissions that balances costs and benefits; the two decisions are made simultaneously. The narrower choice of a project-based discount rate does not address the issue of an optimal aggregate emissions path over time or how the discount rate chosen will affect that path. The two approaches are not completely independent, however. If a project manager is reasonably certain that all other carbon emitters will use a consensus discount rate, that assumption impacts the assumed exogenous carbon path, which impacts the private discount rate chosen. As regulators tackle the question of carbon discounting over a broader array of projects and sectors, it will become important to understand how this decision-making interaction occurs and its implications for aggregate carbon emissions over time.

CONCLUSIONS AND POLICY IMPLICATIONS

The extension of discounting to physical carbon units lacks the theoretical foundation that justifies its use with monetary units. Use of a simple discount factor on physical units is, in fact, a poor substitute for an explicit representation of the complex warming and damage dynamics associated with those units. It is therefore always preferable to avoid the use of physical discounting and instead include a full analysis translating emission flow impacts into cost and benefit impacts that can be appropriately discounted using traditional financial methods. If the use of physical carbon discounting as a policy "shortcut" is unavoidable, however, the process will only be meaningful if an effort is made to select a discount rate that captures the relevant underlying impact and cost dynamics: the decay rate of the gas in the atmosphere, some representation of the relationship between stock and warming damage, and an economic, or social, discounting term. It is therefore a non-trivial exercise to select an appropriate carbon discount rate, and not one that can extrapolate directly from the range of numbers generally used for economic discounting.

Fundamental to the selection of an appropriate physical discount rate (or rates) will be recognition that the project horizon and the impact horizon are distinct time horizons that must be defined independently of each other and of the carbon discount rate. It is not appropriate to lump the two together into a single analytical time horizon that attempts to simultaneously capture the physical dynamic of emissions impacts over time and the policy or project dynamic of changing emissions from a parcel of land over time. For example, selection of a single 100-year time horizon in an attempt to capture the lasting impacts of carbon in the atmosphere will greatly over-estimate the likely period over which carbon benefits from the land conversion will continue to accrue. And while use of a non-zero discount rate can make selection of the amortization period less significant in influencing results, it does not preclude the need to select a period that is theoretically justifiable. In other words, while use of a non-zero discount rate may make the analytical results emerging from a 100-year analysis look more similar to those of a 30-year analysis (by discounting the latter-year results), use of a larger discount rate is not sufficient justification for using the 100-year rather than the 30-year horizon. The only way to define a theoretically justifiable analysis period is to break out the relevant analytical horizons and define them separately.

The time analysis variables are critical in influencing the magnitude of the net carbon benefit (or loss) estimates associated with an activity or project in any GHG accounting framework. In the context of evaluating the net GHG emissions associated with land-use conversion for biofuel production, use of a higher discount rate will tend to produce more cautionary results about the benefits of land conversion by discounting the latter year land-conversion benefits of displaced gasoline use but leaving the early-year conversion costs relatively intact. Use of a single, longer amortization period tends to have the opposite effect—increasing net benefits relative to net costs by including more years of benefits in the analysis.

The relevance of physical carbon discounting is not limited to the biofuels sector; similar issues are being debated in the forestry sector, for instance, to address issues of how to value temporary carbon storage. As diverse as the potential applications of physical carbon discounting are, those applications share many elements, such as the dual time horizons and discounting structures, sensitivity to common assumptions about carbon residence time in the atmosphere, and cost estimates for warming impacts. Further development of a vocabulary and a common analytical structure for evaluating scenarios involving carbon discounting will therefore be broadly applicable across sectors.

Given the ongoing debates about discounting in the world of environmental project evaluation, regulators and rule-makers must develop stronger justification and greater transparency in the selection of the discounting figures that are used to support regulatory accounting efforts. If the practice of applying carbon discount rates in GHG accounting is adopted into the GHG accounting methodologies attached to the current low-carbon regulatory policies, it is critical that selection of an appropriate discount rate be awarded the attention that it requires and that it capture the underlying atmospheric damage dynamics that it is supposed to reflect. It is not appropriate to rely on a simple transferral of the monetary discount rates, or of the social discount rates often used in longer-term or intergenerational analyses, because the factors that are implicitly represented by a physical discounting term go far beyond those related to the changing value of money over time.

REFERENCES

Delucchi, M.A. 2008. "Important Issues in Life-cycle Analysis of CO_2-Equivalent Greenhouse-Gas Emissions from Biofuels." Presented at "Workshop on Measuring and Modeling the Lifecycle GHG Impacts of Transportation Fuels," Berkeley, CA, July 2008. Available online at http://edf.org/fuels_modeling_workshop.

Farrell, Alexander E., Richard Plevin, Brian Turner, Andrew Jones, Michael O'Hare, and Daniel Kammen. 2006. "Ethanol Can Contribute to Energy and Environmental Goals." *Science* 311: 506-508.

Fearnside, P.M., D.A. Lashof, P. Moura-Costa. 2000. "Accounting for Time in Mitigating Global Warming Through Land-Use Change and Forestry." *Mitigation and Adaptation Strategies for Global Change* 5(3): 239-270.

Guo, J., C. Hepburn, R.S. J. Tol, D. Anthoff. 2006. "Discounting and the social cost of carbon: a closer look at uncertainty." *Envir. Sci. & Pol.* 9(3): 205-216.

Hill, Jason, Erik Nelson, David Tilman, Stephen Polasky, and Douglas Tiffany. 2006. "Environmental, economic, and energetic costs and benefits of biodiesel and ethanol biofuels." *Proc. of the Natl. Acad. of Sci.* 103(30):11206-11210.

Howarth, R.B. 2005. "Against High Discount Rates." *Advances in the Economics of Environmental Research* 5: 103-124.

Nepstad, D.C. C.M. Stickler, B. Soares- Filho, and F. Merry. 2008. "Interactions among Amazon land use, forests and climate: prospects for a near-term forest tipping point" Phil. Trans. R. Soc. B, DOI: 10.1098/rstb.2007.0026.

Nordhaus, W. 2007. "Critical Assumptions in the Stern Review on Climate Change." *Science* 317: 201-202.

O'Hare, M., R.J. Plevin, J.I. Martin, A.D. Jones, A. Kendall, and E. Hopson. 2009. "Proper accounting for time increases crop-based biofuels' greenhouse gas deficit versus petroleum." Environ. Res. Lett. 4: 024001 (7 pp.).

Renewable Fuels Agency. 2008. "The Gallagher Review of the Indirect Effects of Biofuel Production." 92 pp. RFA, UK. Available Online at http://www.dft.gov.uk/rfa/_db/_documents/Report_of_the_Gallagher_review.pdf.

Richards, K.R. 1997. "The Time Value of Carbon in Bottom-Up Studies." *Critical Reviews in Environmental Science and Technology* 27(special): S279-s292.

Righelato, R., and D.V. Spracklen. 2007. "Carbon Mitigation by Biofuels or by Saving and Restoring Forests?" *Science* 317(5840): 902.

Searchinger, T., R. Heimlich, R. A. Houghton, F. Dong, A. Elobeid, J. Fabiosa, S. Tokgoz, D. Hayes, and T.-H. Yu. 2008. "Use of U.S. Croplands for Biofuels Increases Greenhouse Gases Through Emissions from Land Use Change." *Science* 319 (5867): 1238 – 1240.

Stavins, R.N. and K.R. Richards. 2005. "The Cost of U.S. Forest-Based Carbon Sequestration." 40 pp. Prepared for the Pew Center on Global Climate Change. Available online at http://www.pewclimate.org/docUploads/Sequest_Final.pdf.

Stern, N. and C. Taylor. 2007. "Climate Change: Risk, Ethics, and the Stern Review." *Science* 317: 203-204.

Zah, R., H. Boni, M. Gauch, R. Hischier, M. Lehmann, and P. Wager. 2007. "Life Cycle Assessment of Energy Productions: Environmental Assessment of Biofuels Executive Summary." EMPA, Switzerland. Available online at http://www.bioenergywiki.net/images/8/80/Empa_Bioenergie_ExecSumm_engl.pdf.

ABOUT THE AUTHOR

Liz Marshall is a senior economist at the World Resources Institute, 10 G. St, NE, Suite 800, Washington, DC 20002. Ph: (202) 729-7719 Email: lmarshall@wri.org.

ACKNOWLEDGMENTS

WRI would like to thank the U.S. Department of Energy, U.S. Department of Agriculture, National Science Foundation, Doris Duke Charitable Foundation, Robertson Foundation, and BP Foundation for support of its climate and biofuels programs.

ABOUT WRI

The World Resources Institute is an environmental think tank that goes beyond research to find practical ways to protect the earth and improve people's lives. Our mission is to move human society in ways that protect the Earth's environment and its capacity to provide for the needs and aspirations of current and future generations.

NOTES

1. Other analyses have yielded more moderate figures for indirect impacts by changing Searchinger's assumptions about crop yields, advances in conversion technology, and distribution of land types that are ultimately converted.
2. Net emissions are the total emissions associated with producing and using the fuel minus the avoided emissions associated with not producing and using an energy-equivalent amount of fossil fuel.
3. Land for expanding biofuel production can come from a number of sources, including existing cultivated land, retired or abandoned cropland, unmanaged grasslands, or forest. The carbon implications of the land-use conversion differ widely depending on the pre-existing land use. To illustrate the concepts of GHG accounting introduced in this paper, we present a hypothetical case of forest conversion.
4. Carbon losses may be deferred when biomass carbon is stored in wood products rather than released upon land conversion. The discussion that follows can be modified to accommodate that case, but we don't specifically address it in the general framework that follows. Current carbon quantification best practice conventions for land-use change assume that all biomass counts as an emission when it is harvested, in part because the questions addressed in this paper about the "value" of deferring carbon emissions by storing them in wood products have not yet been resolved.
5. "Philosophers and economists have conducted vigorous debates about how to apply discount rates in areas as diverse as economic growth, climate change, energy, nuclear waste, major infrastructure projects, hurricane levees, and reparations for slavery," explains William Nordhaus (2007). Nordhaus has first-hand experience in such debates; he was involved in a high-profile dispute with Nicholas Stern about the Stern Review's use of a near-zero discount rate in its analysis of the economics of climate change (Stern and Taylor, 2007).
6. As in the quote from Richards (1997) above, this paper uses the term "carbon emissions" synonymously with "carbon dioxide emission."
7. In a recent report on biofuels and GHG accounting, O'Hare et al. (2009) refer to the project horizon as the "production period" and the impact horizon as the "analytic horizon."
8. Once "current values" are discounted, they are called "present values." Current values are the values that would be current at the time of emission, while present values are those values discounted back to the present.
9. The discount rate structure applied to the impact horizon, however, must be identical for all units of emissions over the project horizon. The structure itself can be quite sophisticated, involving declining discount rates over time for instance, but it must be identically applied to all units of emissions. If a non-identical discount structure is applied to the emissions, it will result in changing current value estimates of damages, and this conclusion no longer applies.
10. Analogously, the method will favor those projects whose displaced emissions occur early in the project horizon.
11. There is a third way in which time can enter policy analyses for GHG reductions, and that is through the specification of target dates for achievement of an objective. California's Low Carbon Fuel Standard, for instance, calls for a 10 percent reduction in the average carbon intensity of California's transport fuels by 2020. While such formulations may imply that we are not concerned about impacts beyond 2020, and that a fixed impact horizon truncated at 2020 is therefore appropriate, a closer examination of the quantification methodology and purpose will usually show that is not the case.
12. This is in some ways analogous to the fact that differing half-lives mean that, among a set of GHGs, relative global warming potentials (GWPs) will vary depending on the analysis horizon chosen.