CO2 Concentration Limits, The Costs and Benefits of Control, and The Potential for International Agreement

Stephen C. Peck
Electric Power Research Institute
Palo Alto CA

and

Thomas J. Teisberg
Teisberg Associates
Charlottesville VA

December 31, 1997

Abstract

In this paper, we use the two region CETA-M model to explore some related issues raised by the current interest in CO2 concentration targets as a possible climate change policy objective. First, we identify possible cost and benefit assumptions that would make particular concentration ceilings optimal. Next, we explore the acceptability to both regions of various burden sharing agreements, given particular cost and benefit assumptions. Lastly, we inquire whether the control effort and burden sharing rules may reasonably be negotiated independently.

I. Introduction

Recently, policy analysis has begun to focus on CO2 concentration limits as possible objectives of policy toward climate change. Wigley, Richels, and Edmunds (1996) for example, consider alternative emissions paths designed to achieve one of several possible concentration targets in a given time frame. They argue that some emission paths will have much higher economic costs of achieving a given target, while producing only slightly lower global mean temperature and sea level rises during the period of transition to the ultimate concentration target.

Recent policy discussions also focus on burden sharing rules,
and the theoretical interdependence between burden sharing rules and optimal emissions paths. Chichilnisky and Heal (1994), for example, show that Pareto efficient allocations depend on the

initial distribution of property rights, in an economy with public goods. Chao and Peck (1997) consider this issue in the specific context of climate change. Using a static equilibrium model with emission rights trading between two regions, they conclude that the interdependence between the initial distribution of rights and optimal emissions is quantitatively small.

In this paper, we present an analysis motivated by that of Chao and Peck (1997). In our analysis, however, we use climate change benefit and mitigation cost assumptions that are related to possible CO2 concentration targets, and we use the CETA-M model.[1] CETA-M is a regionalized version of the CETA model[2], and like CETA, it represents the economy and climate in a dynamic context.

We begin our analysis by exploring the implications of a range of possible assumptions about the costs and benefits of controlling emissions. For each set of assumptions, we identify the concentration target that is optimal, given those assumptions. This provides a reasonable way to relate cost and benefit assumptions to possible concentration targets.

Having found cost and benefit assumptions consistent with various concentration targets, we next consider possible international CO2 control systems employing tradeable emission permits or rights. In particular, we focus on finding the allocation of permits to each party to the agreement, such that each party prefers the agreement to no agreement. We also look at the potential acceptability to both parties of a couple of burden sharing rules based on 1990 GDP or 1990 population proportions.

Finally, we address the question of whether the optimal quantity of emission permits can be agreed upon independently of the distribution of those permits. If optimal emissions are independent of the distribution of permits, the negotiation process becomes simpler.

II. The CETA-M Model

The CETA (Carbon Emissions Trajectory Assessment) Model is the foundation for CETA-M. CETA represents world-wide economic growth, energy consumption, energy technology choice, global warming, and global warming costs (costs of damage from and adaptation to higher temperature). Much of the data for CETA is adopted from the base case assumptions of EMF14, the on-going

---

[1] Peck and Teisberg (1997) introduces the CETA-M model and uses it to analyze the costs of alternative proposals for CO2 emission reduction.

[2] See Peck and Teisberg (1992)

Stanford Energy Modeling Forum Global Climate Change study.[3]

CETA-M is similar to CETA in many respects. However, it disaggregates the world into regions, and allows for trade between these regions in multiple goods. Equilibrium is found using an approach based on Negishi welfare weights. The following sections provide more detail.

A. Regions

The EMF14 study disaggregates the world into six regions: the United States (USA), the European Economic Community (EEC), other OECD countries (OOECD), the former Soviet Union (FSU), China, the rest of the world (ROW). For our analysis here, we divide the world into only two regions: the OECD and the Rest of the World. Data for these two regions are obtained by aggregating EMF14 data for the appropriate subregions comprising our two regions.

For each of the two CETA-M regions, our representations of the economy and energy use is essentially the same as that for the world as a whole in the CETA model. Thus regional output depends on exogenous labor input, the endogenous capital stock, and energy use. CO2 emissions depend on the quantity and type of energy used in each region.

B. International Trade

In CETA-M, we allow for international trade in the numeraire good (aggregate output), carbon emission permits (for some policies), and two key energy goods: oil and gas (which are aggregated together as a single commodity in CETA and CETA-M) and synthetic fuels (derived from coal). This choice of energy goods abstracts from the possibility of trade in other energy goods such as coal or electricity, and it ignores some important differences in transportation costs for oil and natural gas. Nevertheless, we think it is a reasonable representation of the most important energy trade flows we would expect to observe over the next century or so.

When international trade in the numeraire good is allowed, the numeraire flows from the region with a lower market rate of interest to the region with a higher market rate of interest, until rates of return are equalized across regions. To prevent unrealistically high capital flows, we follow the approach of Manne and Richels[4] and benchmark our regional utility discount rates (i.e. those used to calculate the present value utility of any given consumption path) so that market rates of return are approximately equal (and equal to 5 percent) for both regions.

---

[3] Energy Modeling Forum (1995).

[4] Manne and Richels (1992), p. 126.

For each traded good and each region, model equations representing the economy and energy use are augmented by an equation requiring regional use of traded goods to equal regional production (or allocation, in the case of carbon permits) plus net imports. In addition, for each traded good, an equation is added to require that the sum of net imports over regions equals zero.

C. Determining the Equilibrium

To determine the competitive equilibrium in CETA-M, we use an approach employing Negishi weights. That is, we specify a problem in which the objective function is a weighted sum of utilities in the two regions; these weights are known as Negishi weights.[5] When this problem is solved for any arbitrary set of weights, the shadow prices of the constraints requiring net imports to sum to zero are the international prices for the corresponding goods. These prices may then be used to calculate a present value trade surplus or deficit for each of the two regions for this model solution. The competitive equilibrium is then found by adjusting the Negishi weights until the present value trade surplus (or deficit) is zero.

With two regions, there is only one independent Negishi weight (the other being completely determined because the weights must sum to one). This makes it simple to adjust the Negishi weight until the present value trade surplus is zero. In fact, we find that the trade surplus is very nearly a linear function of the Negishi weight, which makes it possible to come very close to the equilibrium Negishi weights using two sets of trial weights and interpolating or extrapolating using the results from these trial weights. Repeating the interpolation once or twice produces an even closer approximation to the equilibrium.

D. Regional Warming and Damages in CETA-M

The CETA-M model used for this paper contains a climate change damage representation that is new to this version (as well as all other versions) of the CETA model. We have replaced a damage function based on the globalized damage estimate of Nordhaus (1991) with regionalized damage functions derived from the damage estimates of Fankhauser (1995).

Fankhauser's damage estimates represent "benchmark" damage from a 2.5 degree C. temperature increase, i.e. the temperature increase considered most likely by IPCC (1990) for a doubled CO2 concentration. Since Fankhauser's estimates are presented in 1988 dollars, we have inflated them by 10 percent as a rough

---

[5] Negishi, T. (1972).

adjustment for inflation and growth to 1990.

Fankhauser's damage estimates are presented in categories which may be aggregated into two classes -- market damages and non-market damages. Market damages are those for which market prices can be used directly or indirectly to measure costs; an example is agricultural losses where the prices of crops can be used to value production losses. Non-market damages are those for which there are no market prices to help in valuing damages; an example of non-market damages are health effects (including increased mortality) attributable to climate change.

We aggregated the following of Fankhauser's damage categories into the non-market damage class: wetlands loss (even though fisheries loss is included), ecosystem loss, human life, air pollution, migration, natural hazards (even though this is partly a market damage). The remaining categories are aggregated as market damage: coastal defense, dryland loss, agriculture, forestry, energy, and water.

Table 1 below presents the resulting damage estimates by EMF14 region, together with EMF14 GDP and population numbers. Overall, the market and non-market damages are of approximately equal magnitude, and together come to roughly 1.4 percent of GDP.

Table 1

1990 Climate Change Damage Estimate
Derived from Fankhauser (1995)
(Billion $)

EMF14 Region	Non-Mkt Dam.(D_{NM})	Market Dam.(D_M)	EMF14 GDP	EMF14 POP
USA	33.11	33.995520	250	
EEC	35.64	34.325710	244	
OOECD	32.12	29.264970	259	
FSU	8.91	11.111310	289	
CHINA	7.26	11.111330	1134	
ROW	41.03	18.593110	2976	

Our next step is to "explain" statistically the cross sectional variation in market and non-market damages in terms of GDP and population. It is reasonable to expect that market damages would be linearly related to GDP, and indeed we find that the data are consistent with this expectation:

$$D_M = \alpha + \beta_1 \cdot GDP \qquad \mathbf{1}$$

where:

$\alpha = 3.573223$
(SE = .97596)
$\beta_1 = .005327$
(SE = .000237)
$R^2 = .992083$

For non-market damages, it seems reasonable to suppose that:

$$D_{NM} = f(y) \cdot POP \qquad \mathbf{2}$$

Here y is income per capita, and f(y) may be interpreted as an amount per person that represents willingness to pay to avoid non-market damage. Intuitively, it is plausible that the function f(y) might be non-linear in income per capita. However, after a little experimentation, we concluded that the following linear relationship best fits the data:

$$\frac{D_{NM}}{POP} = \alpha + \beta_1 \cdot \frac{GDP}{POP} \qquad \mathbf{3}$$

The regression results for the above equation were:

$\alpha = .003705$
$\beta_1 = .006017$
(SE = .000200)
$R^2 = .995566$

Multiplying both sides of equation (3) by POP, D_{NM} is seen to be a linear function of population and GDP (with no constant term).

Equations (2) and (3) provide functional relationships between income, population, and benchmark damages. These may used to produce regionalized (and time varying) benchmark damages from the projected future regional populations and incomes. Benchmark damages, however, only indicate the damages at a certain temperature increase, 2.5 degrees C., in this case. To get actual estimated damages, it is necessary first to have regional temperature changes and then to specify the functional relationship between temperature change and estimated damages. We next describe the procedures we use to accomplish these last steps.

In CETA-M we explicitly track global mean temperature. To go from this to regional temperature, we assume there is a regional temperature differential relative to the global mean. This differential is developed from some regional temperature results presented in IPCC (1990). Specifically, Figure 5.3 (p. 140) shows climate sensitivity by latitude and month of year. Roughly speaking, the figure suggests that if latitude exceeds 45° north or south, the temperature change is significantly different. In the north high latitude, it's warmer (than average) for roughly half the year (fall and winter); in the south high latitude it's warmer all year. However, since there is relatively little inhabited land south of latitude 45°, we ignore the south and focus on the north.

Although the fall and winter temperature increase ranges from 4° C. to 12° C. north of latitude 45°, the more populated land areas are close to latitude 45°. Thus, we assume that north of latitude 45° is characterized as having +5° C. for half the year and +3° C. for half the year, or an average of +4° C for the whole year. Below latitude 45°, on the other hand, might reasonably be characterized as having +3° C. throughout the year. Thus, in an obviously rough way, we assume that the temperature rise above latitude 45° is 1.33 times the temperature rise below latitude 45°.

If high latitude warming is 1.33 times low latitude warming, and global mean temperature is the average of the high and low latitude warming, then high latitude warming must be 1.14 times the global mean, while low latitude warming must be .86 times the global mean.

Having characterized high latitude regions as having 1.33 times the warming of low latitude regions, we next need to decide which of the EMF14 regions (see Table 1 above) should be treated as high latitude regions. Again, in a rough way, we assume that the EC and FSU are reasonably identified as high latitude, while all other regions are low latitude. The most troubling aspect of this decision is the placement of the OOECD, consisting of Australia, Japan, Canada, and New Zealand, in the low latitude category. While Canada is unambiguously high latitude, Japan is the most important country economically and in terms of population, and it is predominantly low latitude, by our definition. Thus, we include OOECD in the low latitude group.

To summarize, then, we specify regional temperature change by assuming that warming in the EU and FSU is 1.14 times the global mean temperature rise, while warming in the other regions is .86 times the global mean temperature rise.

Finally, it is necessary to assume a relationship between regional temperature rise and regional damages, when temperature rise is something other than the 2.5 degrees C. benchmark. We assume that actual damages are a quadratic function of regional temperature rise which passes through the benchmark damage amount when the temperature rise is 2.5 degrees C.

III. Optimal Concentration Targets for Alternative Assumptions

A. Alternative Benefit and Cost Assumptions

The benefits of emission control are the climate change damage costs avoided. We consider two possible levels of climate change damage. For our low case, we assume that damages are those derived from Fankhauser's benchmark estimates (see the discussion in the preceding section). Then, since non-market damage estimates are inherently more speculative than market damage estimates, we increase Fankhauser's non-market damages to obtain a high case assumption about climate change damages. We increase non-market damages by a factor of 3; this increases total damages by approximately a factor of 2.

There are many parameters that affect the costs of emission control. However, a key parameter is the cost of the non-electric backstop technology. The backstop technologies in CETA-M are admittedly speculative future technologies that are characterized as providing carbon-free energy in practically unlimited amounts, but at relatively high cost, and not until some later date. The electric backstop technology might be photovoltaics or some form of advanced nuclear power. The non-electric backstop technology might be hydrogen produced by electrolysis, where the required electricity is obtained using the electric backstop technology. Because the non-electric sector is large and relatively carbon intensive, the cost of the non-electric backstop technology is a key parameter of the CETA-M model.

We consider two alternative costs of the non-electric backstop technology. Our high cost is $80 per barrel of oil equivalent, while our low cost is $60 per barrel.

The above alternative assumptions about benefits and costs of control define four cases, which we identify as follows:

1xND/HC: Fankhauser's damages and backstop cost of $80/barrel.
3xND/HC: Fankhauser's non-market damages increased by a factor of three and backstop cost of $80/barrel.

1xND/LC: Fankhauser's damages and backstop cost of $60/barrel.

3xND/LC: Fankhauser's non-market damages increased by a factor of three and backstop cost of $60/barrel.

As we show in the next section, the 1xND/HC set of assumptions results in the lowest optimal amount of emission control, since the benefits of control are low and the cost of control is high. Conversely, the 3xND/LC set of assumptions results in the greatest optimal amount of emission control.

B. Concentration Targets

For each of the four sets of assumptions specified in the preceding section, an optimal path of emissions can be determined. This is done by including the regional damages from warming in the CETA-M model, and then finding the Negishi equilibrium in which both the OECD and the ROW satisfy their intertemporal budget constraints. The results are shown in Figure 1.

In Figure 1 emissions are highest in the 1xND/HC case; the peak concentration level in this case is around 1100 PPM, or nearly four times the pre-industrial concentration. When the non-market damages are increased by a factor of three (3xND/HC), the peak concentration level falls to roughly 900 PPM. When the non-electric backstop cost is reduced to $60 per barrel (1xND/LC), the peak concentration falls further to about 700 PPM. Finally, when both damage is increased and cost reduced (3xND/LC), the peak concentration is a little over 500 PPM, or about twice the pre-industrial level. We interpret the peak concentrations as optimal concentration targets for the underlying assumptions about costs and benefits of emission control.

The United Nations Framework Convention on Climate Change sets an objective of stabilizing greenhouse gas concentrations "at a level that would prevent dangerous anthropogenic interference with the climate system." As yet, there has been no attempt to determine what this concentration level might be. Indeed, most of the policy proposals for controlling climate change have been expressed in terms of emissions rather than concentrations. For example, proposals have been advanced to stabilize emissions at 1990 levels, or to stabilize them until some date such as 2020, and then reduce them by 20 percent.

It is interesting that neither the proposal to stabilize emissions nor the proposal to cut them by 20 percent causes CO2 concentrations to stabilize, at least not in the time frame out to 2250 (which is the terminal period in the CETA model). These proposals do tend to result in concentrations in the 500-550 PPM range by 2150, the approximate time by which the concentration trajectories in Figure 1 have all reached or passed their peaks. So in a very loose sense, the policy discussions to date are most consistent with the optimal concentration for the 3xND/LC case assumptions.

IV. International Emission Control Using Tradeable Permits

When benefits of emission control are not considered, it is only possible to analyze the costs of control and how these might be shared among regions participating in an agreement to limit emissions. When both benefits and costs of control are considered, it is possible to analyze the net gains from an emission control agreement (relative to no agreement) and how these gains are shared among regions participating in the agreement. In this section, we look at the net gains and the distribution of these gains, for the two most extreme sets of assumptions considered in the preceding section -- the 1xND/HC case and the 3xND/LC case.

A. A Tradeable Permit System

A relatively simple way to implement an agreement to control emissions is through a system of tradeable emission permits. Under such a system, emissions permits are issued for each time period in a number equal to the optimal emissions for that time period. Then, to emit a unit of CO2 into the atmosphere, a region would be required to use a permit (either one it was issued, or one it had purchased from the other region).

The allocation of permits among regions participating in the agreement determines how the net benefits of controlling emissions are distributed among the participating regions. If region 1 is granted more permits, it will be in a position to buy fewer permits or sell more permits, thereby improving its welfare; conversely, region 2 would see a reduction in its welfare, since it would be able to sell fewer permits or would need to buy more permits.

B. The Bargaining Ranges

To be willing to participate in an agreement, each region would need to achieve an improvement in its welfare under the agreement, relative to its welfare in the absence of an agreement. Generally, if the overall benefits of emission control are large, there will be a fairly wide range of emission permit allocations that will improve the welfare of both regions; conversely, if the overall benefits of control are small, the range of permit allocations that would improve both regions' welfares is small. We refer to this range of allocations as the "bargaining range." Outside this range, the agreement cannot be concluded; within this range, the allocation between the regions is the result of the relative bargaining abilities of the regions.

We illustrate the bargaining range using the CETA-M model with two regions: the OECD and the ROW. In doing so, we effectively assume that each region acts as a cohesive entity; this ignores the very real divisions that would exist within various blocks of countries within the OECD or the ROW. In addition, we assume that in the absence of an international agreement to control emissions, there would be no control at all. In fact, a non-cooperative solution involving two regions would be characterized by more than zero emission control. Of course, in reality, there will be many more than two regions involved in negotiating an international control agreement, and the non-cooperative solution for this larger number of regions might in fact be quite close to the zero control solution we posit in our analysis.

Figure 2 illustrates the bargaining range in the 3xND/LC case, where the benefits of controlling emissions are greatest and the bargaining range is largest. In Figure 2, the axes measure the welfares (i.e. present value utilities) of the OECD and ROW regions. The point labelled "No Control" indicates the welfare levels achieved by the two regions absent an agreement. The line labelled "Utility Frontier" represents the highest welfare levels achievable with optimal emissions control. The ends of the bargaining range are found by moving vertically or horizontally from the No Control point to the Utility Frontier. At each end, the welfare of one of the regions is the same as that in the no control situation, while the welfare of the other region is maximally improved relative to the no control situation. Between the two ends of the bargaining range, the welfare of both regions is higher than in the no control situation.

In the context of an emission control system employing tradeable emission permits, different points within the bargaining range can be achieved by changing the allocation of permits between the two regions. A regional permit allocation is necessarily a time path of permit grants. The time path could be specified as a simple constant share of total permits in each time period, or it could be related to other benchmarks such as regional GDP shares, regional population shares, or regional optimal emissions.

Some permit allocation time paths could be favorable to a region in the early years, and unfavorable to that region in the later years, raising the possibility that agreements made now might be broken in the future by the region with an unfavorable future permit allocation. In this paper, we assume that an agreement made today will continue to be honored in the future, even if it later becomes unfavorable to one of the regions. In future work, we intend to give further consideration to the possible incentives to break agreements, and the permit allocation rules that are likely to minimize such incentives.

If regions will remain bound by an agreement made today, it is straightforward to calculate the welfares (i.e. present value utilities) obtained by the two regions under any specific permit allocation rule. Then these welfares can be compared to the welfares at the ends of the bargaining range to see if the allocation rule produces an agreement within the bargaining range.

Perhaps the simplest emission permit allocation time path involves annually giving each region a fixed percentage of the total quantity of permits for that year. With this kind of time path of emission permits, we find that the end of the bargaining range most favorable to the ROW may be achieved by annually allocating 66.5 percent of the total permits to the ROW, while the end of the range most favorable to the OECD may be achieved by annually allocating of 41.0 percent of the total permits to the ROW. These fixed annual allocations provide one simple way to characterize the ends of the bargaining range in Figure 2.

Another simple way to understand the range of permit allocations characterizing the bargaining range is in terms of the value of the emission permits. In the 3xND/LC case illustrated in Figure 2, the total present value of the permits is about $3,200 billion, and the present value of the permits going to the OECD changes by about $800 billion dollars from one end of the bargaining range to the other. This change in permit allocation value also represents a measure of the overall efficiency gain achieved by a policy of optimally controlling emissions in this case.

Figure 3 illustrates the bargaining range in the 1xND/HC case, where the benefits of controlling emissions are the smallest and the bargaining range is the smallest. The bargaining range here is so small as to be difficult to make out in the figure. The ends of the bargaining range here are achieved by an annual ROW permit allocation that ranges from 57.0 percent to 59.5 percent -- a swing of less than three percentage points. This swing represents a change in permit value of about $60 billion. This small change in the value of the allocation corresponds to a small overall efficiency gain from a policy of optimally controlling emissions in the 1xND/HC case. It is not surprising that the efficiency gain is small in this case, since benefits of controlling emissions are low and costs of controlling emissions are high, relative to the 3xND/LC case.

C. Burden Sharing Rules

In this section, we consider some alternative possible rules for allocating permits. We retain the assumption that rules for allocating permits which are agreed to at the present time remain binding on the regions over future decades and centuries.

Various ideas about how to distribute emission permits have been suggested. Permits might be allocated among regions in proportion to their base period populations or GDPs. For some period of time, the developing countries might be given permits equal in number to their expected emissions in the absence of control. Combining these notions, the ROW might be given permits equal to uncontrolled emissions for a period of time, after which there would be a transition to a permit allocation based on either population or GDP.[6] We will consider two allocation rules of this type. In both, the ROW is given permits equal to its uncontrolled emissions until 2030; then there is a transition, between 2030 and 2050, to a permit allocation based on 1990 population proportions in one case, and based on 1990 GDP proportions in the other.

Figures 4 and 5 illustrate the ROW permit allocation paths that result from application of the above allocation rules, under 1xND/HC case assumptions; Figures 6 and 7 illustrate these paths for the 3xND/LC case assumptions. In all four figures, the ROW permit allocation path follows the ROW uncontrolled emissions path until 2030. Then there are transitions either to 1990 GDP-based shares or 1990 population-based shares of the optimal emissions path for each set of assumptions. Since the optimal emissions path is much higher for the 1xND/HC case assumptions, the post-2050 ROW permit allocations are much higher in Figures 4 and 5 than in Figures 6 and 7 (note the change of vertical scale).

Figure 8 shows, for 1xND/HC case assumptions, the two rule-based permit allocation paths together with the paths that result from annually assigning 59.5% or 57.0% of permits to the ROW (i.e. the constant annual allocation percentages that characterize the ends of the bargaining range in the 1xND/HC case). Note that the paths characterizing the ends of the bargaining range are very close together. The GDP-based permit allocation path is initially close to these paths, but starting in 2030 it diverges below them. The population-based allocation path is also initially close to these paths, but starting in 2030 this path diverges above them. Based on a visual inspection of Figure 8, it seems doubtful that the welfares resulting from application of the GDP-based rule or the population-based rule would lie within the bargaining range.

Figure 9 shows, for 3xND/LC case assumptions, the two rule-based allocation paths together with the paths that result from annually assigning 66.5% or 41.0% of permits to the ROW (i.e. the constant annual allocation percentages that characterize the ends of the bargaining range in the 3xND/LC case). Here both the GDP-based allocation path and the population-based allocation path lie within the paths characterizing the bargaining range until 2030. Starting in 2030, however, the population-based allocation path diverges above these paths; and starting in 2050, the GDP-based allocation path diverges below them.

---

[6] A similar burden sharing rule based on population is analyzed in Manne and Richels (1997).

Allocation paths that lie partly but not entirely between the paths characterizing the bargaining range may or may not produce welfares that lie in the bargaining range; to determine whether they do, we need to calculate directly the welfares (i.e. present value utilities) of the two regions that result when permits are allocated according to the specific rules under consideration. We have done these calculations for the 3xND/LC case assumptions where there is some chance that welfares will lie within the bargaining range, as well as for the 1xND/HC case assumptions where the chance that welfares will lie within the bargaining range is remote. The results of these calculations are shown in Figures 10 and 11, for the 1xND/HC case and 3xND/LC case, respectively.

Figure 10 shows the 1xND/HC case bargaining range (Figure 3), with additional points plotted to indicate the welfares of the two regions resulting from the rule-based permit allocation paths shown in Figures 4 and 5. Not surprisingly, the rule-based allocations of Figures 4 and 5 lie outside this bargaining range, confirming the visual impression obtained from Figure 8.

Figure 11 shows the 3xND/LC case bargaining range (Figure 2), again with additional points plotted to indicate the welfares of the two regions resulting from the permit allocation paths shown in Figures 6 and 7. As may be seen, both allocations do lie within the bargaining range in this 3xND/LC case. Not surprisingly, the population-based allocation (Figure 6) is more favorable to the ROW, while the GDP-based allocation (Figure 7) is more favorable to the OECD. Evidently, the departures of the permit allocation paths under consideration from the range of paths defining the ends of the bargaining range do not prevent the welfares from lying in the bargaining range. This is true because the paths under consideration are within the paths defining the bargaining range until several decades into the future -- by that time, the present value effects of subsequent departures from the paths defining the bargaining range are small.

D. Separability of Optimal Emissions and Permit Allocation

The allocation of permits, within the bargaining range, affects the wealth of the two regions. In principle, a change in the relative wealth of the two regions could cause the optimal total amount of emissions to be different. This could happen because a region's marginal willingness to pay to avoid intangible damages depends on its overall welfare level.

To determine whether this is a practical consideration for the two region model we use here, we examine optimal emissions at both ends of the bargaining range. These optimal emissions are found by maximizing the Negishi weighted sum of utilities, subject to constraints that require one or the other regional welfare to equal its value in the no control situation. We perform this experiment in the 3xND/LC case, where there is the greatest welfare shift between ends of the bargaining range.

Figure 12 shows the optimal emissions paths at the OECD's and the ROW's ends of the bargaining range in the 3xND/LC case. As is apparent, the two optimal emissions paths are virtually indistinguishable, implying that optimal emissions do not depend on the end of the bargaining range to which a negotiation ultimately leads. The reason for this is that the swing in the present value of the permit allocation between ends of the bargaining range (i.e. $800 billion) is not large relative to the present value of consumption over this time frame (about $500,000 billion for the OECD and $260,000 billion for the ROW).

V. Summary and Conclusions

In this paper, we began by considering a range of possible assumptions about the costs and benefits of CO2 emission control. For each of these assumptions, we identified the maximum concentration reached as the optimal concentration target for those assumptions. The purpose here was to illustrate possible cost and benefit assumptions that are consistent with alternative possible concentration targets that might be agreed to in future negotiations. We found that for 1xND/HC case assumptions, the optimal concentration target is very high, around 1100 PPM. At the other end of the spectrum, with high emission control benefits and low emission control costs (i.e. the 3xND/LC case), the optimal concentration target is in the 500-550 PPM range.

Next, we considered international emissions control using tradeable emission permits. First, for the 1xND/HC case and the 3xND/LC case, we identified the range of welfare outcomes, achieveable by controlling emissions, that are preferred by one or both regions over the no control situation -- we referred to this range as the bargaining range. We found that the bargaining range is very small for 1xND/HC case assumptions, while it is reasonably large for 3xND/LC case assumptions.

We next found the annual emission permit allocations (as fractions of optimal emissions) that would produce welfares at the ends of the bargaining ranges. Also, we analyzed a couple of possible permit allocation rules based on 1990 GDP or 1990 population proportions to see if these would produce welfares that lie within the bargaining ranges. The particular rules we tested did lie in the bargaining range under 3xND/LC case assumptions, but not under 1xND/HC case assumptions. However, in obtaining these results, we assumed that agreements made at the present time would be honored, even if they became disadvantageous to one of the parties over the coming decades or centuries. Under a more realistic assumption that agreements will be broken if the incentive to do so becomes strong, it is possible that neither the GDP-based nor the population-based allocation rules would actually be in the bargaining range; in our future research we intend to explore this issue more fully.

Finally, we considered the question of whether negotiations about the total quantity of emission permits can be undertaken separately from negotiations about the allocation of the permits between regions. We found that this is indeed possible, since the optimal amount of emissions is virtually unaffected by the allocation of emission permits between regions.

Regarding the likelihood of actually concluding an international agreement, our results are mixed. For 1xND/HC case assumptions, the range of acceptable agreements is very small and there is little incentive to reach agreement. For 3xND/LC case assumptions, the range of acceptable agreements is larger and there is substantial incentive to reach agreement, making an agreement more likely under these assumptions. However, even in this case the total amount of money at stake, i.e. the $3.2 trillion present value of all permits, is so large that the negotiating parties may be led into endless self-interested posturing that ultimately frustrates an agreement. On the bright side, however, and for either set of assumptions, it appears that the total quantity of emission permits can be negotiated separately from the allocation of those permits, which tends to simpify the negotiation process somewhat.

Acknowledgements

We thank Richard Richels and Samuel Fankhauser for helpful comments and suggestions.

Disclaimer

This paper does not represent the position of EPRI or of its members.

References

Chao, Hung-Po and Stephen Peck (1997), "Optimal Environmental Control and Distribution of Cost Burden for Global Climate Change," Journal of International and Comparative Economics, forthcoming.

Chilchilnisky, Graciela and Geoffrey Heal (1994), "Who Should Abate Carbon Emission? An International Viewpoint," Economics Letters.

Energy Modeling Forum (1995), "Second Round Study Design for EMF14: Integrated Assessment of Climate Change" Draft, Stanford University Energy Modeling Forum, Stanford CA, March 3.

Fankhauser, Samuel (1995), Valuing Climate Change: The Economics of the Greenhouse, Earthscan Publications Ltd., London.

Intergovernmental Panel on Climate Change (1990), Climate Change: The IPCC Scientific Assessment, Cambridge University Press, Cambridge.

Manne, Alan S. and Richard G. Richels (1992), Buying Greenhouse Insurance: The Economic Costs of CO2 Emission Limits, MIT Press, Cambridge, MA.

Manne, Alan S. and Richard G. Richels (1997), "On Stabilizing CO2 Concentrations - Cost-Effective Emission Reduction Strategies," draft, Electric Power Research Institute, Palo Alto CA, February 25.

Negishi, T. (1972), General Equilibrium Theory and International Trade, North-Holland, Amsterdam.

Nordhaus, William D. (1991), "To Slow or Not to Slow: The Economics of the Greenhouse Effect," The Economic Journal, Vol. 101, pp. 920-937.

Peck, Stephen C. and Thomas J. Teisberg (1992), "CETA: A Model for Carbon Emissions Trajectory Assessment," The Energy Journal, Vol. 13, No. 1.

Peck, Stephen C. and Thomas J. Teisberg (1997), "International CO2 Emissions Targets and Timetables: An Analysis of the AOSIS Proposal," Environmental Modeling and Assessment, Vol. 1, No. 4, February.

Wigley, T.M.L., R. Richels, and J.A. Edmunds (1996), "Economic and Environmental Choices in the Stabilization of Atmospheric CO2 Concentrations," Nature, Vol. 379, No. 18, January.

Figure 1: Optimal CO2 Concentrations - Four Cases

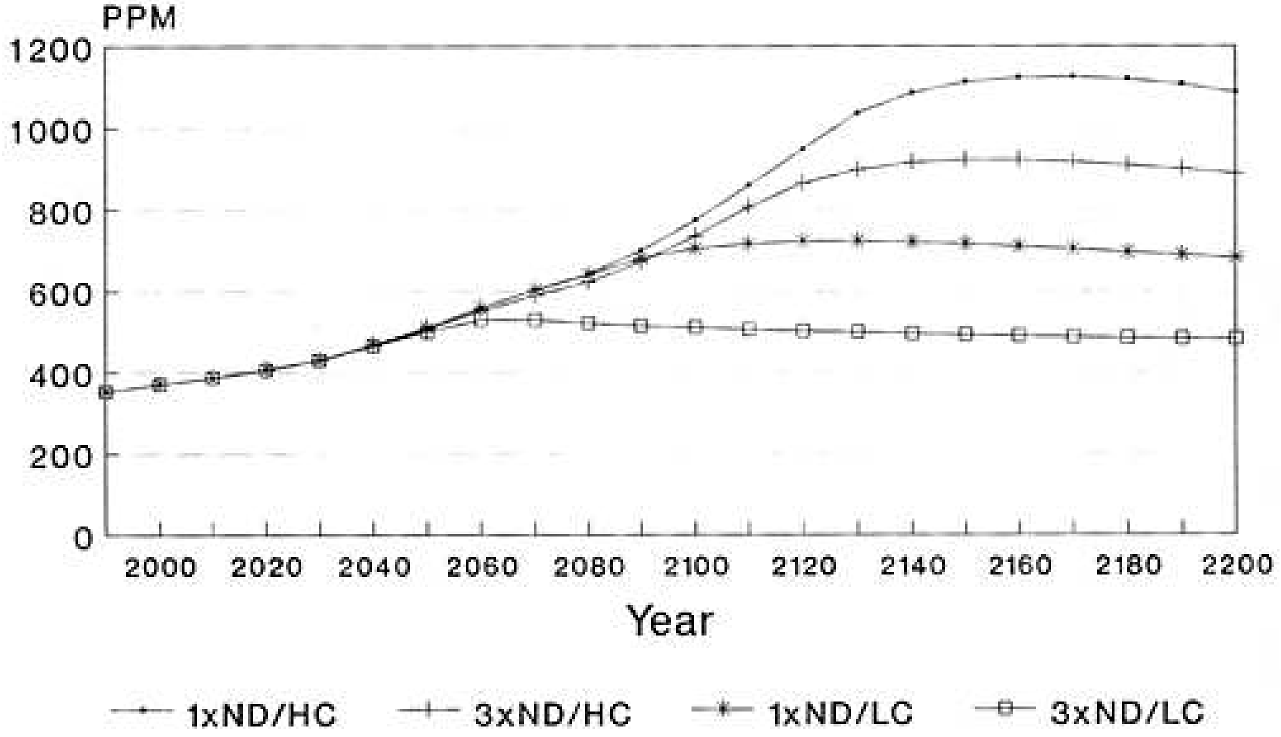

Figure 2: Bargaining Range in the 3xND/LC Case

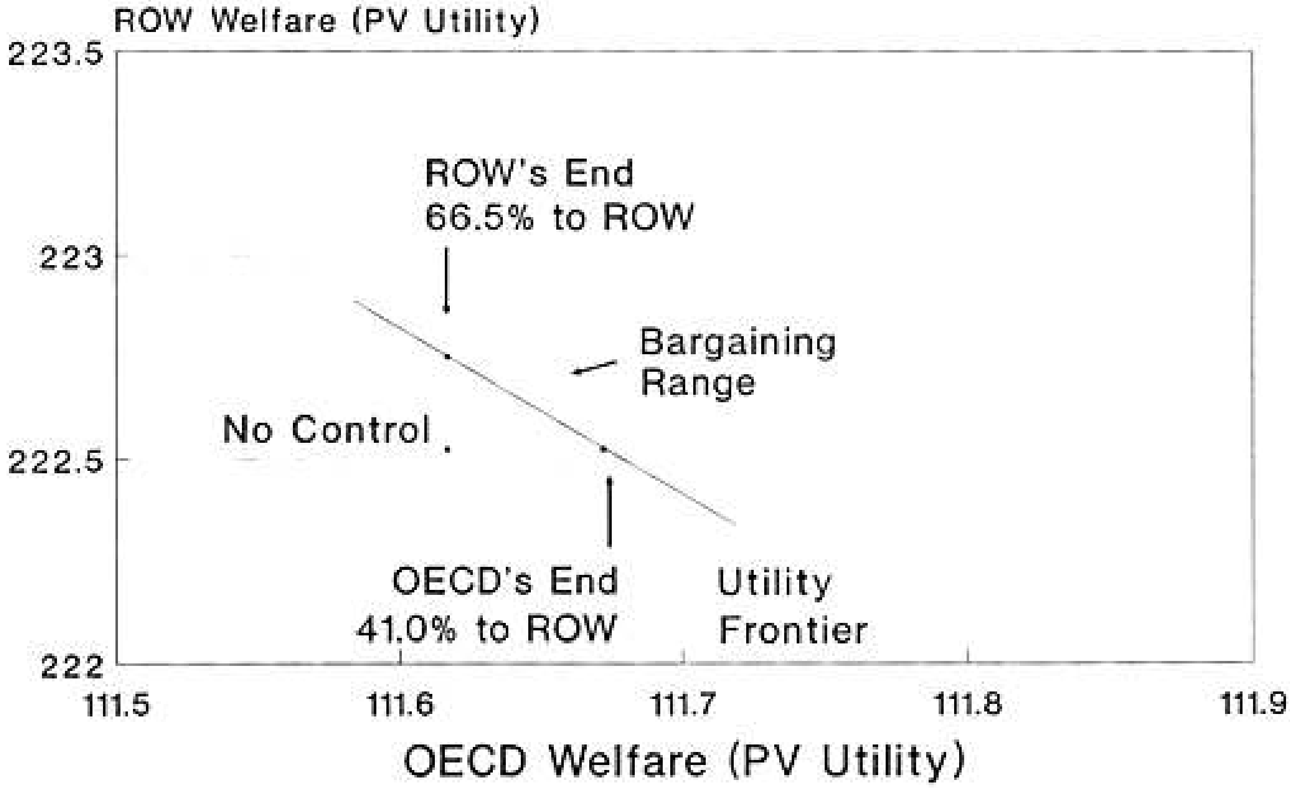

Figure 3: Bargaining Range in the 1xND/HC Case

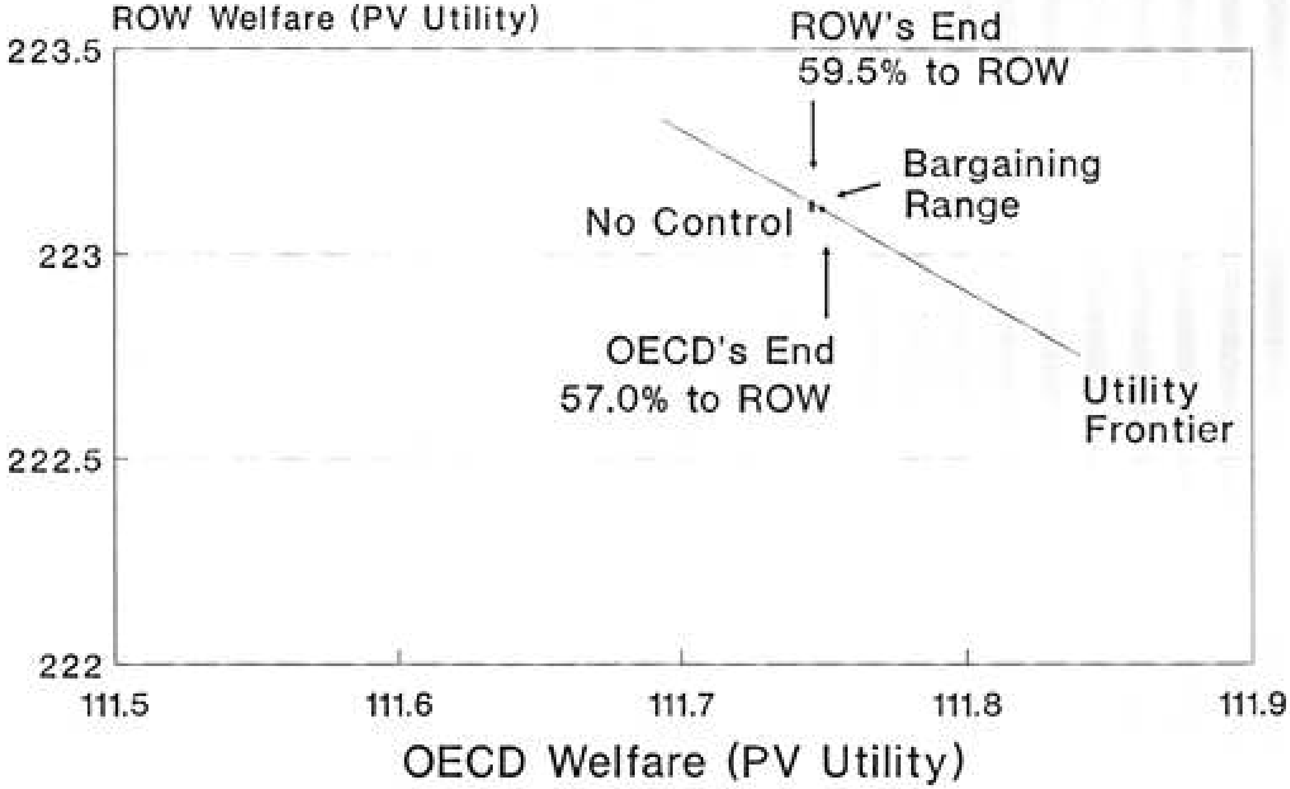

Figure 4: ROW Permits - Trans. to POP-Based Alloc. in 1xND/HC Case

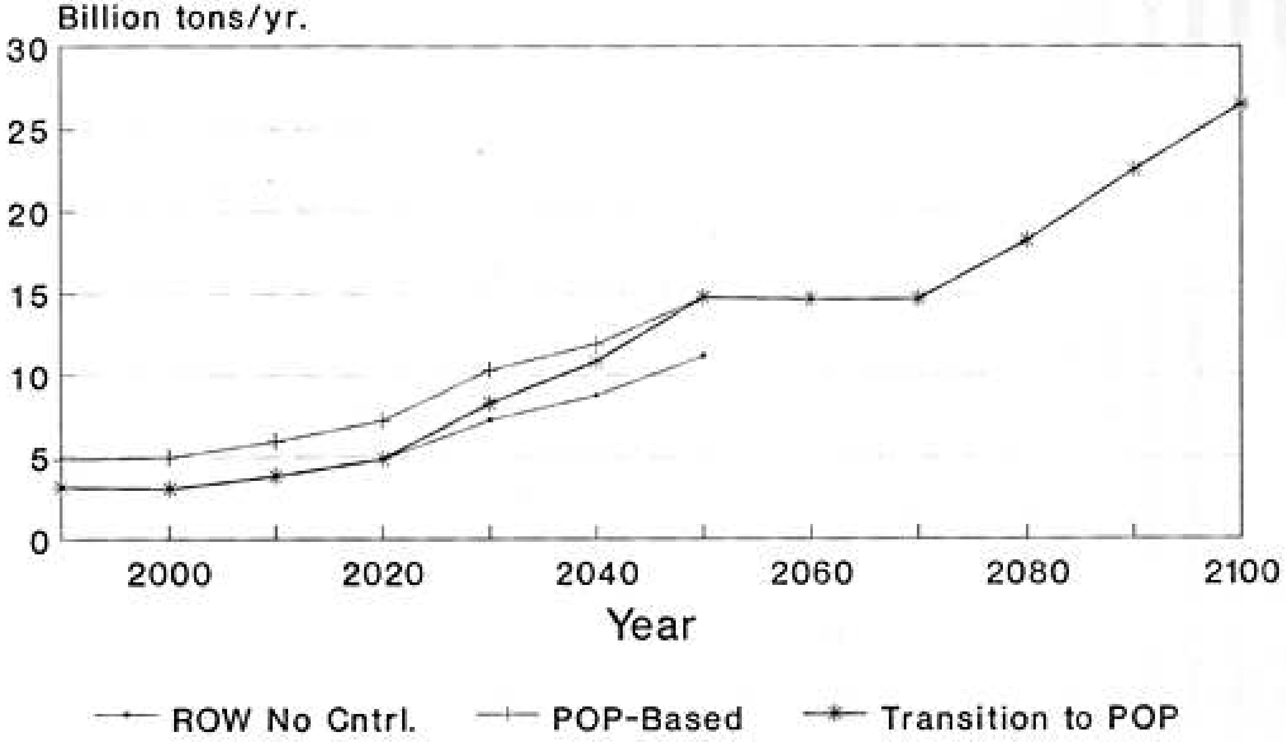

Figure 5: ROW Permits - Trans. to GDP-Based Alloc. in 1xND/HC Case

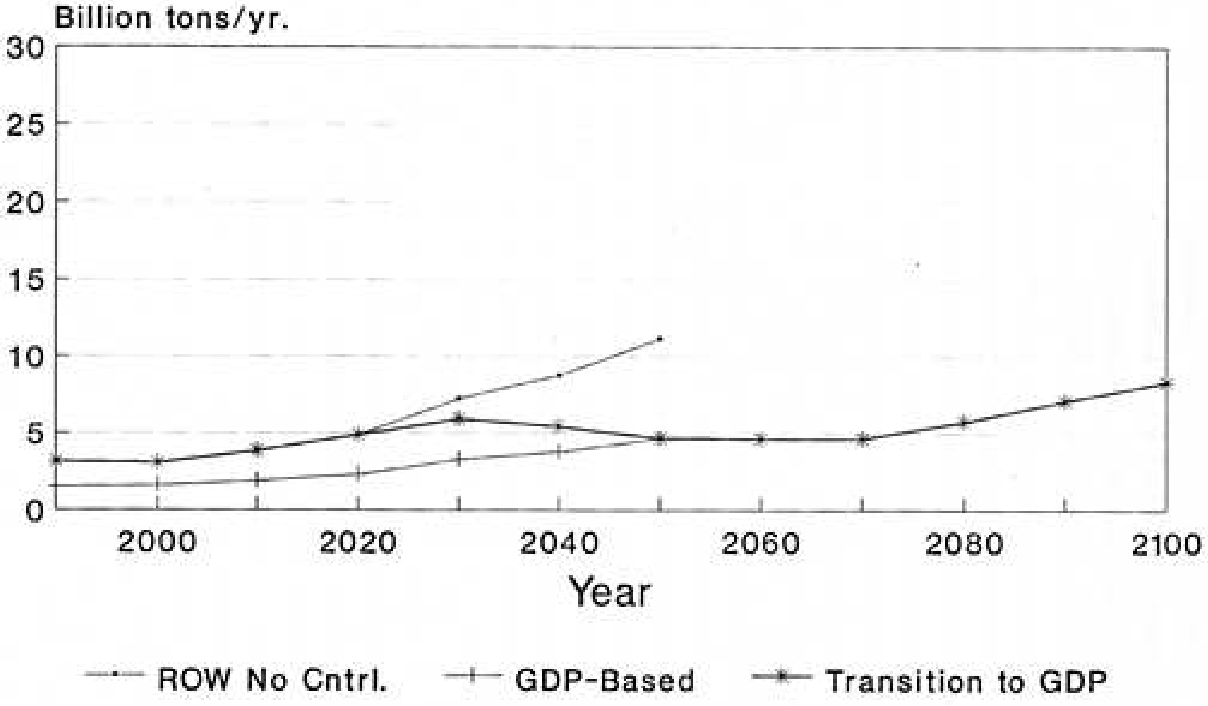

Figure 6: ROW Permits - Trans. to POP-Based Alloc. in 3xND/LC Case

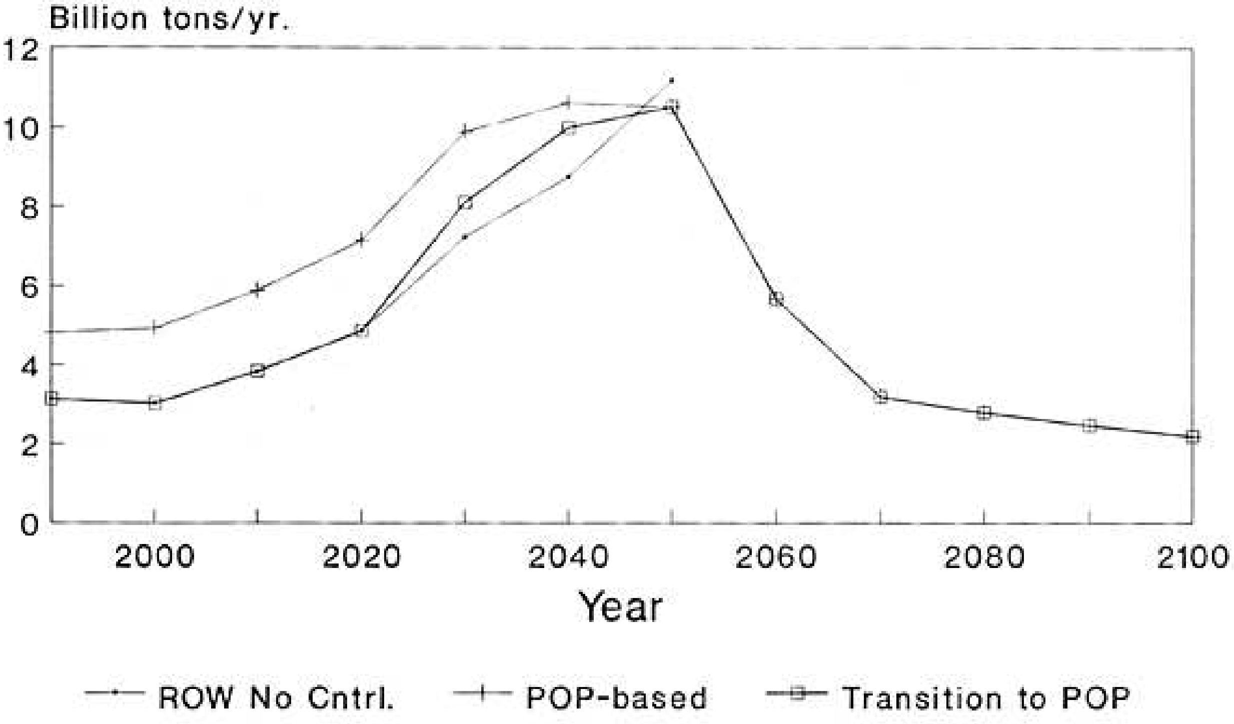

Figure 7: ROW Permits - Trans. to GDP-Based Alloc. in 3xND/LC Case

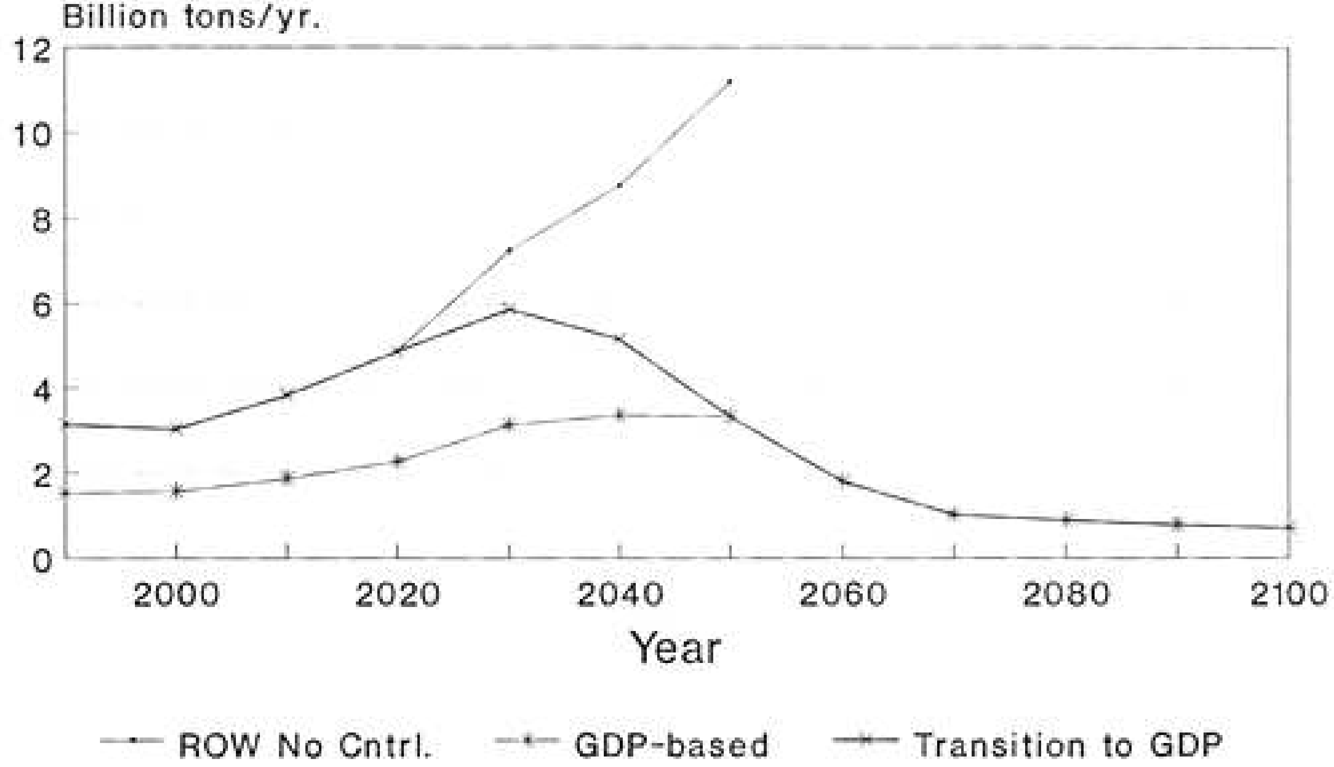

Figure 8: Rule-Based Permit Paths vs. B.R. Paths - 1xND/HC Case

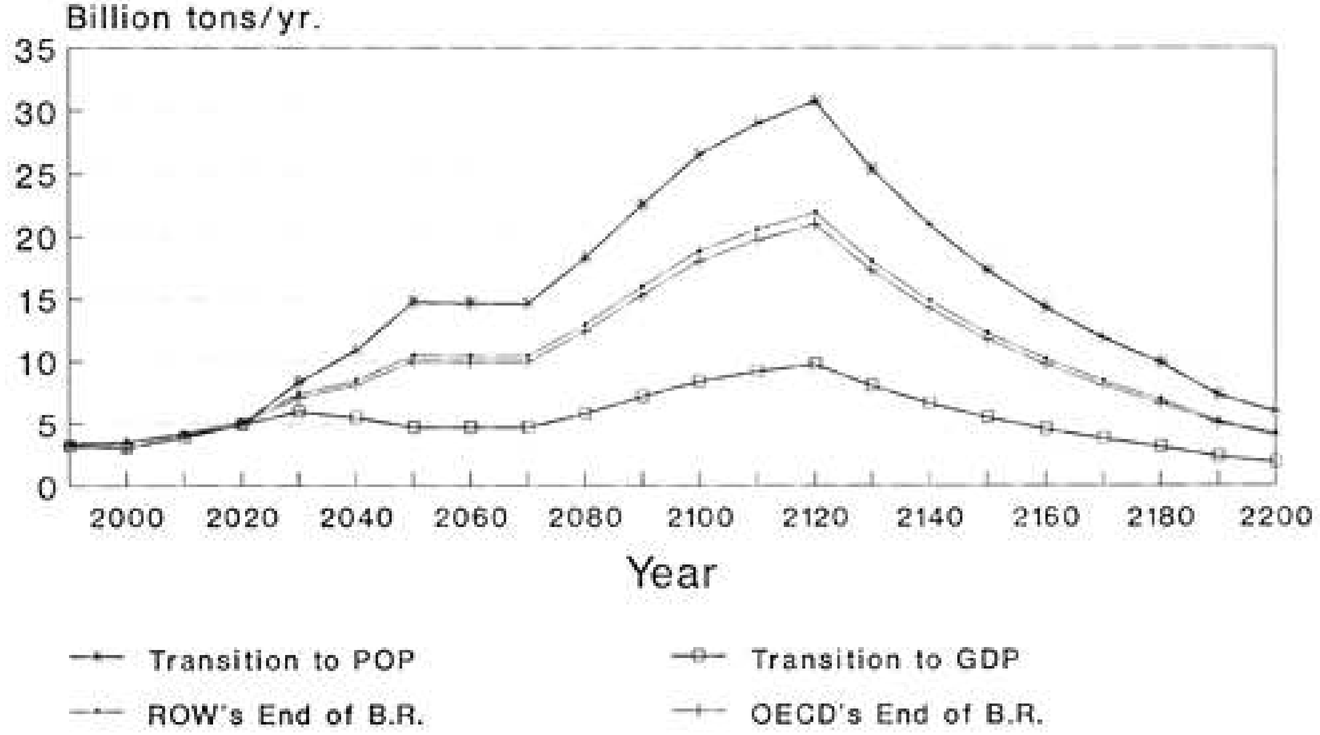

Figure 9: Rule-Based Permit Paths vs. B.R. Paths - 3xND/LC Case

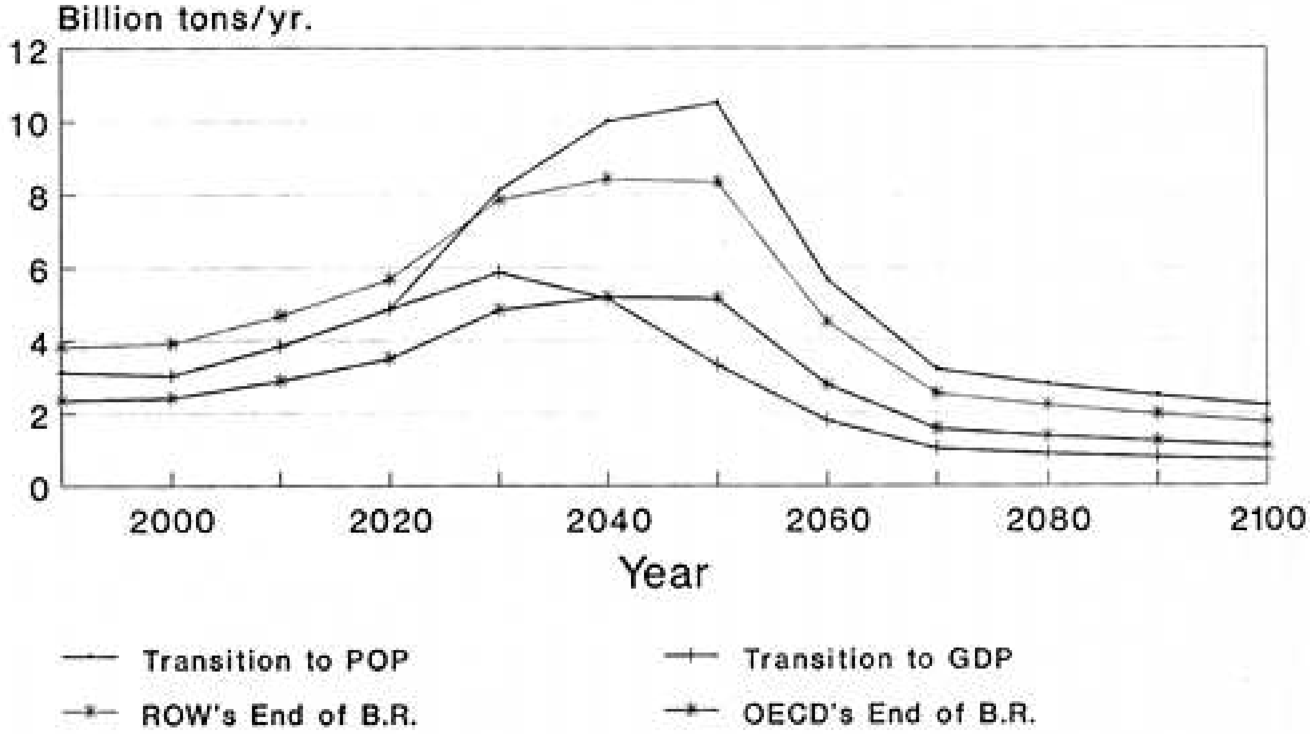

Figure 10: Rule-Based Allocations Are Not in 1xND/HC Barg. Range

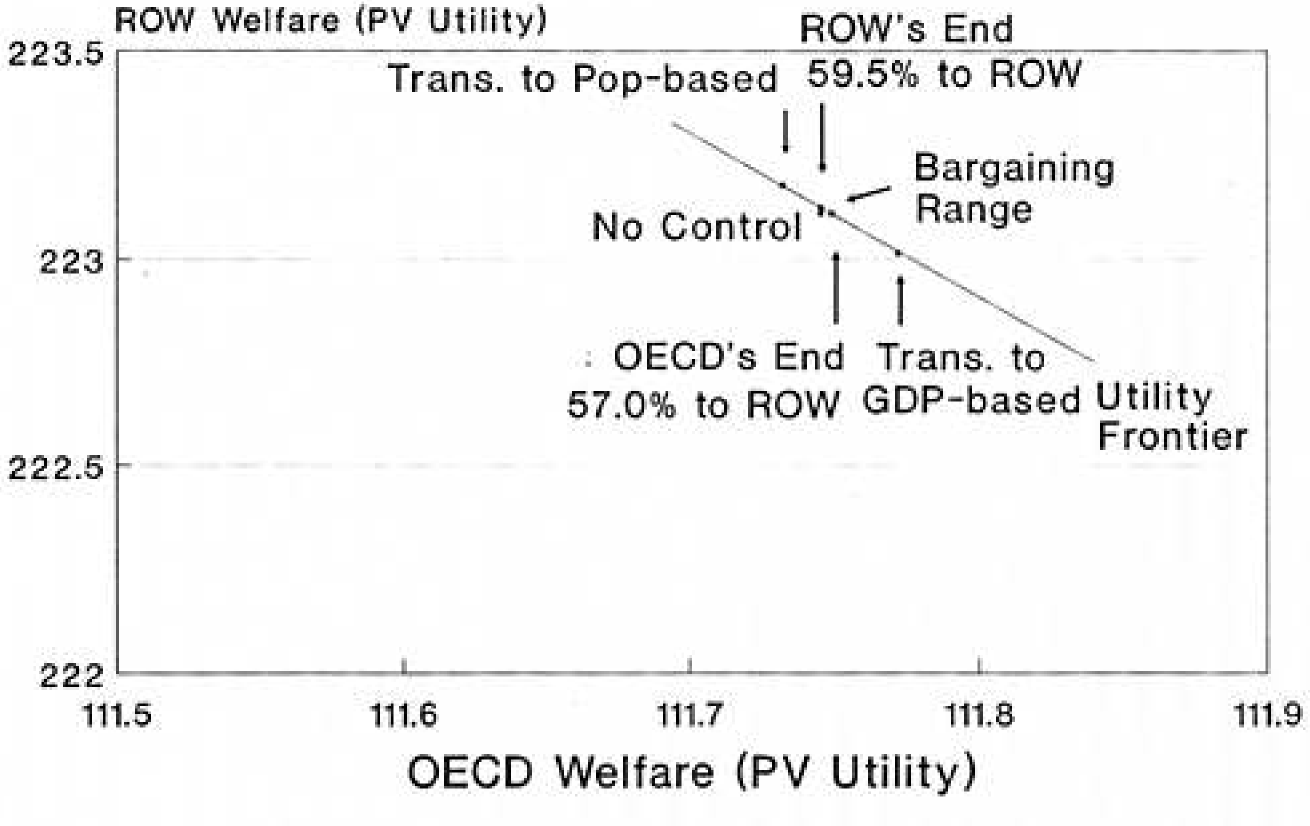

Figure 11: Rule-Based Allocations Are in 3xND/LC Bargaining Range

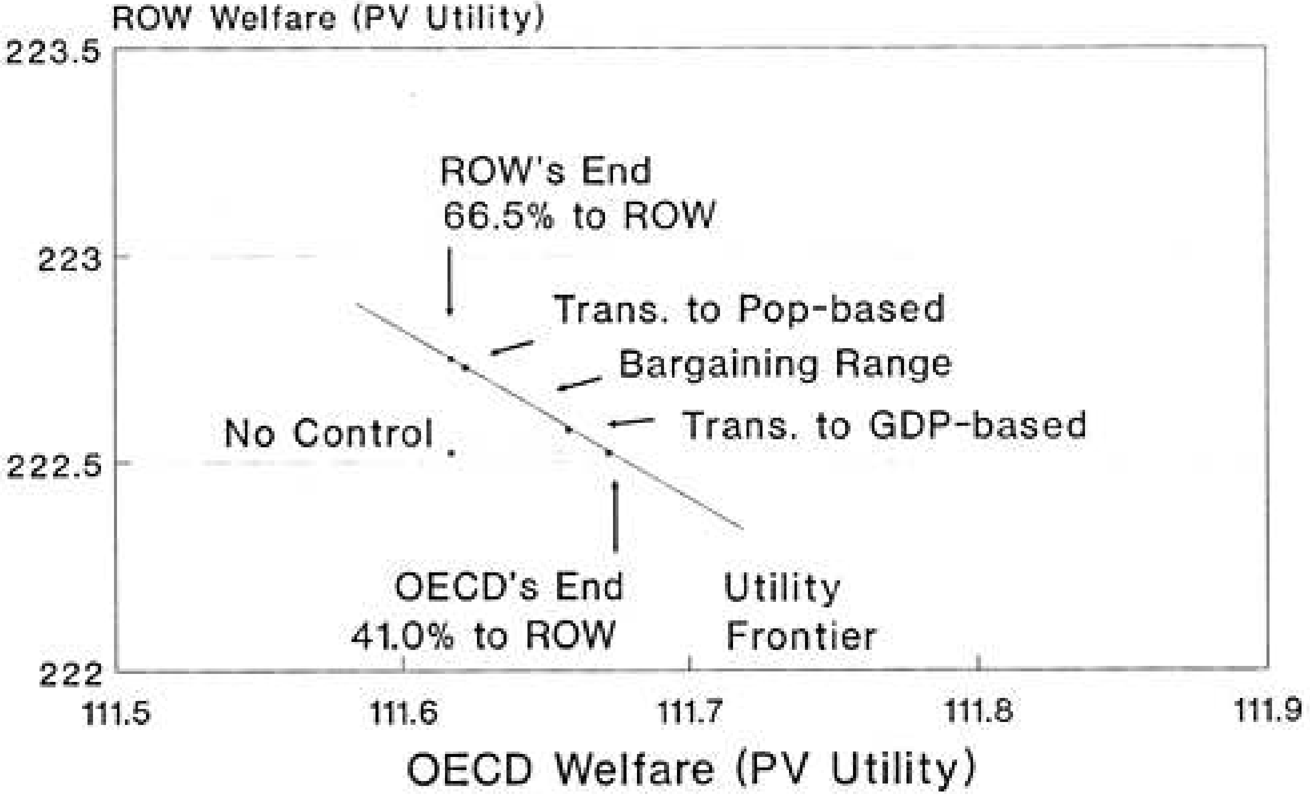

Figure 12: Optimal Emissions at Ends of Barg. Range - 3xND/LC Case