Appendix: Summary of Edmonds et al. Paper

THE MODEL: ORIGIN AND STRUCTURE

The Edmonds et al. study used version 4.01 of the Edmonds-Reilly -Barns model (ERB), which projects global energy, and related greenhouse gas emissions, through the end of the next century,[11] The model explicity represents (he energy resource base, plus supply and demand in nine world regions,[12] and includes international trade in fossil fuels, and fuel-specific greenhouse gas emission coefficients.

The model consists of four separate modules representing energy demand, energy supply, energy balance, and greenhouse gas emissions. The energy demand module begins with external assumptions of population and labor productivity growth in each world region, and calculates total final-energy demand in each sector from energy prices (world prices augmented by local energy taxes and tariffs) and incomes. The demand module also maintains a set of energy flow accounts for each region, determining the mix of energy sources used to meet final-energy demands by a S-shaped (logit) function of relative prices.

The supply module represents the fossil fuel resource base of each region in detail, including several grades of resources with increasing extraction costs, and limits on the rate at which production capacity can be expanded. Separate costs and capacity

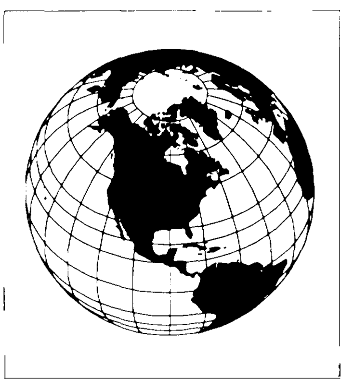

11 A brief description of the Version of the model employed for this project is provided in Edmonds, Barns, Wise and Ton (5). The original model is described in Edmonds and Reilly (2), and major early revisions arc documented in Edmonds et al. (3).

12 The regions modeled are the United States, OECD West (Western Europe and Canada); OECD East (Japan, Australia, and New Zealand): Eastern Europe and the former Soviet Union; China and other centrally planned Asian economies; the Middle East; Africa; Latin American; and South and Southeast Asia.

constraints are presented for nuclear fission in each region, and a nonfossil energy supply (solar electric or nuclear fusion) is assumed available as a "backstop technology "-a technology that can provide essentially any quantity of energy at a constant, high marginal cost, which in this model is also assumed to decline overtime due to technical advance.

The energy balance module represents international trade in liquid, solid, and gaseous fuels, generating world energy prices that yield approximate equilibrium in each global fuel market, Electricity is not traded. The greenhouse-gas emissions module calculates regional and global emissions of greenhouse gases from specific emissions coefficients for each category of fuel.

Given specified input assumptions, the model presents snapshots of the world energy system every 15 years from 1990 to 2095. In each of these years, the model projects each type of energy demand and supply in each region, the price of each energy type, and resultant greenhouse gas emissions. The model investigates the effect of measures to limit greenhouse gas emissions by imposing various forms of regional emissions constraints or taxes, and comparing the resultant projections of energy use and prices to those in a base-case model run. In earlier versions of the model, regional costs of emissions constraints were calculated by a simple, region-specific energy-GNP feedback elasticity; a specified percentage change in world energy prices was assumed to cause a related percentage GNP decrease (for net energy importing regions) or increase (for net energy exporting regions). The authors drew their estimated feedback elasticities from a literature review of estimated energy-GNP interactions. While this mechanism sought to capture macroeconomic effects of changes in energy markets, critics suggested the assumed constant-elasticity relationship was likely invalid for the large energy price changes that would accompany serious emissions constraints. Consequently, in the model version presented to the OTA workshop, the authors modeled regional costs of emission constraints by a partial-equilibrium approach, summing changes in consumer's and producer's surplus in energy markets as the regional economy shifts from the unconstrained to the emissions-constrained equilibrium.

SCENARIOS AND ASSUMPTIONS

The study used three scenarios to specify different plausible rates of emissions growth in the absence of controls. All are based on the same regional population projections, estimated by the World Bank (25) and used in the first report of the Intergovernmental Panel on Climate Change (13). Projected world population reaches 9.5 billion by 2050, and 10.4 billion by 2095.

The three scenarios represent variation in potential economic growth rates by different assumed rates of labor productivity growth. In the first scenario, productivity growth in industrialized nations declines from 1 .5-1.6 percent per year to about 0.9- 1.0 percent over the next century, while China's productivity growth remains around 2.6-2.8 percent, and the rest of the developing world increases from 1.6 to 2.2 percent. The two other scenarios, representing high- and low-growth futures, double and halve all these assumed productivity growth rates.

Other assumptions and parameters are maintained constant across the three scenarios. These include the size of fossil fuel resources, the size and cost of biomass energy resources, and the economic parameters that define market responses to changes in incomes and energy prices. The income elasticity of energy demand in OECD nations is assumed to be 1.0 through the next century.[13] Income elasticities in other world regions start higher (1.25 in Eastern Europe and the former Soviet Union, and 1.4 in the developing countries), and decline to 1.0 through the next century. The price elasticity of demand for aggregate final energy is assumed to be 0.7 in all regions.

[13] That is a 1 percent increase in income is assumed to generate a 1 percent increase in energy demand, other things remaining equal.

Consequently, all variation between scenarios in factors contributing to greenhouse gas emissions is subsumed into different labor productivity growth rates. The authors argue that this simplification is reasonable, because prior sensitivity analysis with this model has demonstrated the following to be true. While the cost of reducing emissions from any particular baseline level depends on that baseline level (holding emissions to 100 costs more if the unconstrained level was 150 than if it was 120): the cost depends very little on the particular combination of assumptions that generated the baseline level (reducing emissions from a baseline of 120 to 100 costs about the same, regardless of what combination of population and productivity growth, resource availability, and sensitivity to price and income change generated the original baseline of 120) (4).

OTHER IMPORTANT DESIGN ISSUES

In addition to the price and income sensitivity of energy demand, the analysis assumes a non-price-driven, or "autonomous,'' improvement in end-use energy efficiency of 1 percent per year in all regions. The authors argue that this value represents the long-term trend in energy intensity over the past 70 years. This value is controversial, though, with some analysts advocating values both higher and lower than 1 percent, and others arguing that the quantity is misspecified and should, if employed at all, be negative (12, 17,24). After a test of the ability of different values to replicate observed 1987 fuel consumption from historical data, Edmonds et al. reaffirm that within the specification of this model, a value of 1 percent is appropriate.

To represent variation in technological progress, the authors include two mode] runs with widely differing technological assumptions. The first includes optimistic assumptions for improvement in fossil electrical generating efficiency (reaching 55 percent by the year 2020), and solar electric cost (dropping to 5 cents per Kilowatt hour (KwH) by the turn of the century). The second focuses on technology transfer. assuming that some fraction of the price-induced energy efficiency gains realized in regions that control their emissions is transferred to noncontrolling regions without requiring the higher price to elicit it. In effect, some fraction of price-induced efficiency gains are assumed to take the form of innovations that, once discovered, are cost-effective even at lower energy prices. Originally implicitly zero, this fraction is varied from 10 to 100 percent.

Other than the variation represented in the three scenarios for labor productivity growth, the mode1 does not represent uncertainty. Each run of the model generates a single future history. Other input parameters are not varied systematically, and distributions of relevant output variables are not generated.

The model uses a simple carbon-cycle model to show the consequences of different emission paths for atmospheric concentrations, the same model as used in the 1990 IPCC report to calculate global warming potentials (GWPs). This carbon-cycle model makes the common, but controversial, assumption of a "neutral biosphere"-assuming that the unknown total carbon uptake by terrestrial biota is equal to the atinthropogenic source from land-use change, as was roughly true between 1940 and 1980. Computationally, the model splits current-year carbon emissions into three shares, each of which decays in atmospheric concentration with a different time-constant. About 30 percent of emissions decay with a time constant of about 7 years, 34 percent with a constant of71 years, and 36 percent with a constant of 815 years.

POLICIES AND PROTOCOLS STUDIED

The assumptions described thus far define how the mode] represents the basic world energy system and its evolution in the absence of emission controls. The bulk of the work presented to the workshop, though, consisted of imposing various international emission controls and examining how they affect emissions, energy markets, and economies in both participating and nonparticipating regions.

The paper examined five different emission control protocols. All five protocols are expressed as stabilizing carbon emissions, but stabilization is defined as a region's holding its emission constant at then-current levels when the region enters the protocol. That is, a nation joining in 1990 stabilizes emissions at 1990 levels, while a nation joining in 2005 holds them at 2005 levels. While this may be a reasonably accurate prediction of how baseline emission levels would likely be defined politically, the approach differs from recent national emission control pledges, most of which pledge to hold emissions at 1990 levels beginning in 2000 or 2005. Under the paper's definition of stabilization, a nation that delays its accession to a stabilization agreement stabilizes its emissions at a higher level.

Given this definition of stabilization, the paper's five protocols differ only in the years that emissions are stabilized in different world regions. In the first protocol, all nations stabilize at 1990 levels in 1990 (the first year represented in the model); a second protocol delays worldwide participation to 2005. Variants of these two protocols have staggered participation. in which the OECD stabilizes emissions in the specified year, while Eastern Europe and the former Soviet Union (EEFSU), China, and the rest of the developing world delay their stabilization by 15, 30, and 45 years respectively.[14] A final protocol variant delays OECD stabilization to 2020, with other regions staggered at the same intervals.

In a separate examination, the paper considers the effect of protocols that only ever achieve limited participation. A "Big Three" protocol includes only the three largest coal-bearing regions —OECD, EEFSU, and China—all stabilizing emissions from 1990. Separate analyses were also conducted of the same protocol with only "Big Two" (OECD and EEFSU), and OECD-only participation.

Each protocol is examined under three forms of implementation: a uniform carbon tax; regional emission targets; and tradable permits. Under a carbon tax, all participating regions tax carbon emissions at the same level, so as to stabilize their total emissions joint] y.[15] Taxes are in fact the only policy that the model can represent directly; other measures are represented by surrogate taxes, whose effects are equivalent to the specified measure under the assumption of competitive en–ergy markets in equilibrium. For example, under separate regional emission targets, each region imposes a tax high enough that it meets its own target. Under this system, taxes—and hence marginal abatement costs-are unequal between regions, so total compliance costs are higher than under a uniform tax. Since the model only represents multination regions, though (except the United States), this system does assume joint reduction among the nations in each region, so is less restrictive and more efficient than separate national targets.

A tradable permit system is also modeled indirectly by imposing a uniform tax on participating regions, and assuming regions will trade permits to reach the efficient distribution of emissions from whatever starting point the initial distribution is defined. All trades are assumed to be at the competitive price, equal to the marginal cost of emission reduction and the uniform tax rate, equivalent to assuming no market power in the market for emission permits. Under these assumptions, the paper examined the consequences of six different rules for distributing permits. Whatever rule applies, permits are redistributed according the rule in each 15-year modeling period. The six distribution rules are as follows:

[14] For example, in the variant of the first protocol, OECD controls begin in 1990, EEFSU in 2005, China in 2020, and the rest of the world in 2035.

[15] The revenues from carbon taxes are assumed to be retained within the tax ing region, and recycled in some manner that does not affect the rate of capital formation. Consequently, the potential gains available from shifting the total tax burden away from investment-deterring taxes on capital toward carbon taxes are not represented in the model.

- "Grandfathered emissions:" permits are distributed to continue emitting at 1990 levels.
- Equal per capita: permits are distributed in proportion to adult population each year (10).
- Equal per GDP: permits are distributed in proportion to regional GDP in each period.
- "GDP-adjusted grandfathered emissions:" an original distribution by 1990 emissions is adjusted over time based on regional differences in GDP growth.
- "No harm (o developing nations:" developing nations receive enough permits that their revenue from permit sales restores their GDP to that of the noncontrol case. The allocation between OECD and EEFSU is in proportion to grandfathered emissions.
- "NO harm to non-OECD:" as above, except that EEFSU is also given enough permits to restore its unconstrained GDP.

RESULTS

The model's three uncontrolled scenarios generate time paths of world energy and carbon emissions that roughly span the range of estimates in the literature. World primary energy consumption grows from 350 Exajoules (EJ) in 1990 to 750, 1300, and 1980 EJ in the year 2095 under the low, medium, and high-growth scenarios, while global fossil carbon emissions grow from about 6 Petagrams (Pg) in 1990 to 11, 20, and 32 Pg, respectively. Under all scenarios, primary energy is increasingly dominated by coal, with oil and gas contributions peaking in the first half of the next century, then declining. Nonfossil primary en– ergy provides shares up to about 30 percent. Regionally, the developing countries, especially China, account for an increasing share of both primary energy and carbon emissions under all three scenarios, with total carbon emissions from the industrialized countries flat or declining through the century under both the low and medium-growth scenarios.

Applying the carbon-cycle model to these emission scenarios gives the range of atmospheric carbon concentrations that result. Atmospheric CO^2 concentration passes 550 parts per million, roughly double the pre-industrial value. in approximately the years 2095,2070, and 2060 under the high, medium, and low-growth scenarios. In all three cases, concentration is increasing at the end of the modeling period.

Under the first protocol, emissions are stabilized immediately at 1990 levels and held there through the century. Achieving this stabilization through a common global carbon tax requires a tax level that is initially modest, about $40 per ton in 2005, and grows to $400 to $500 by the end of the century. (These figures are for the medium-growth scenario. Required taxes in the high- and low-growth scenarios are roughly double and half these values.) Through the century, this common tax Progressively redistributes emissions from the industrialized to the developing countries. While the bulk of the early cost burden is borne by OECD nations, later in the century costs are increasingly borne by the LDCs, especially China: the tax redistributes emissions toward China, but not enough to keep its costs down. In terms of GDP, costs of emission control remain below 1 percent in all regions except EEFSU (where they reach 2 percent by 2095) and China (where they reach 3 percent). Under the high-growth scenario, costs of stabilizing emissions are higher in absolute terms, but are not in all cases higher as a fraction of GDP (since GDP is also higher). Under this protocol, per capita emissions converge to the range of 0.5 to 1.0 tons in all regions except the United States, whose emissions decline from about 5.5 to about 3.0 tons per person over the century.

Broadly speaking, the high and increasing tax rates required to stabilize emissions are predictable consequences of a few model assumptions. Energy demand is exponentially increasing, with no technological miracle available to provide large quantities of low-cost, low-emission energy. Because the carbon tax is recycled in a way that does not stimulate investment, it provides no offsetting macroeconomic benefit. Under these conditions, holding emissions constant becomes increasingly costly over time, and requires increasingly high marginal tax rates.

The same global stabilization goal yields substantially different consequences when implemented by separate stabilization targets in each region rather than by a common global tax. Because marginal abatement costs are no longer equalized across regions, world abatement cost is higher—in fact, more than double the cost under a common tax. Cost differences vary sharply across regions, though, with some regions paying less under uniform targets than under a common tax. EEFSU costs are zero under uniform targets through the first part of the century; because unconstrained emissions do not surpass 1990 levels, a regional stabilization constraint is not binding. OECD costs start higher under uniform targets, then become lower, because the uniform tax eventually requires OECD nations to reduce their emissions below 1990 levels.

When stabilization is realized by tradable carbon emission permits, total world costs and the distribution of emissions are the same as under a uniform tax (by assumption), but the various rules for distributing emission permits yield wide variations in the distribution of costs. Since regions are assumed to buy and sell permits at marginal cost to move from the initial endowment to the efficient inter-regional distribution of emissions, a distribution of permits uniquely determines a distribution of costs.

Under ● *grandfathered emission s," the distribution of permits remains at 1990 emission levels through the next century. Under this system there is little trading in the early years, but as LDC growth outstrips the OECD, LDCs buy permits in increasing quantities. Since the OECD and EEFSU earn money by selling permits, their total cost burden is very small (in fact, negative in EEFSU), while transfers from the LDCs to the industrial countries due to trading reach $1.2 trillion annually by the end of the century. These transfers clearly render this scheme infeasible.

Equal per capita distribution of permits mostly reverses the direction of trading and transfers. LDCS receive the most permits, and their share increases further over time due to their higher rates of population growth. Most LDCS sell excess permits while OECD and EEFSU buy them, but China provides a surprising anomaly. China's share of perm its does not grow fast enough to meet the demand driven by its rapid economic growth, so it turns from selling to buying permits in midcentury. In effect, China's projected rapid economic growth and abundant coal resources cause its demand for emissions to exceed its entitlement, which grows only with its modest population growth. By the end of the century, China buys $500 billion of permits annually and the industrial countries buy another $500 billion.

Two other allocation systems are based on regional GDP. In the first, permits are distributed according to current GDP. Under this system trades are initially small, while late in the century China's more coal-intensive resource base leads them to buy permits from the OECD. The second scheme begins with grandfathered emissions, but modifies future shares in proportion to relative GDP growth. Under this scheme, distributions closely track the determinants of furture emissios, so inter-regional trade is small. Total transfers from trading remain at a few tens of billions through the first half of the century, growing to $200 billion by 2095. Most trade is from EEFSU to OECD.

Two final allocation schemes investigate a "hold-harmless" rule. These schemes give enough permits to specified regions that their revenue from selling excess permits precisely offsets the costs of reducing emissions, leaving them as well off as under an unconstrained growth path. In the first of these schemes, only the developing countries are held harmless; in the second, the EEFSU countries are added, in effect making the OECD nations bear the entire world's emission abatement cost.

The consequences of these schemes change sharply over time, as the LDC economies grow. Initially, it requires only small transfers to LDCS to make them whole, so the allocation of permits is close to the "grandfathered" scheme and transfers are only about $15 billion in 2005. But the LDC losses that must be compensated grow rapidly through the century, so that by 2080 even giving LDCS all the permits in the world fails to yield transfers large enough. Consequently, this scheme eventually requires that the industrial countries be allotted "negative emission permits" and be required to buy their way out of the hole before they begin purchasing emission permits. Total transfers to LDCs reach $1 trillion annually by 2080, and $1.7 trillion by 2095. When EEFSU countries are also "held-harmless," the negative allocation to OECD appears sooner (in 2065) and grows larger; associated transfers reach $2 trillion in 2095.

A separate series of model runs examined the effect of protocols in which only a limited set of nations participate. In particular, a "Big Three" protocol examined the effect of controls by only OECD, EEFSU, and China, which together hold 96 percent of world coal resources. These regions stabilize their emissions by a tax levied half at the point of combustion and half at the point of extraction. This case exhibited a number of remarkable results. Global emissions grew modestly through the first half of the century through increases in nonparticipating regions, but declined abruptly near the end of the century as nonpm-ticipants exhausted their coal. Through the century, progressive exhaustion first of conventional oil and gas, and later of nonparticipants' coal, drive up world energy prices, eventually bringing involuntary price-driven emission reductions even in countries that did not intend to do so.

The ‘Big Three" is the only scenario modeled in which world emissions drop below 1990 levels, and the only one in which atmospheric carbon stabilizes by the end of the next century, reaching about 500 parts per million (ppm). Two subsequent runs tested the same form of controls in the "Big Two" (OECD and EEFSU) and the "Big One" (OECD only). These agreements were much less effective at reducing global emissions, indicating that the strong results of this scenario depend on participation by nations holding the great bulk of the world's coal. The OECD-only protocol, for example, yielded global emissions that differed by only a few percent from the reference case.

The involuntary reductions in nonparticipating nations that the "Big Three" protocol brings represent an interesting reversal of the well-known "offshore effect," by which controls in some countries induce adjustment in nonparticipating countries. The standard offshore effect dilutes the effect of emission control measures that a subset of countries enact through taxes or controls on *consumption.* When such consumption measures reduce world energy demand and hence prices, participating countries' emissions reductions are partly offset by price-induced emission increases elsewhere. The "Big Three" scenario illustrates that when participants tax fossil fuel *production,* hence reducing their energy exports or turning themselves from exporters to importers, world energy prices can increase and so cause price-induced emission reductions in nonparticipating nations.

A final set of model runs explored the effect of technology development and diffusion. One run used optimistic assumptions of improved fossil generation efficiency and solar-electric cost. With these assumptions, unconstrained global emissions were 25 percent lower than in the reference case (due to a 20 percent reduction in primary energy demand and a doubling of primary solar energy), while the cost of stabilizing world emissions at 1990 levels was cut by half. A further series of runs examined technology diffusion. varying the fraction of price-induced technical efficiency gains in participating regions that become available to nonparticipating regions for free. The benefit from such diffusion is necessarily temporary, lasting only as long as some regions participate in a protocol while others do not. Under the most extreme assumptions, in which all price-induced efficiency gains are transferred, world emissions in some years can be 25 percent below the reference case. These temporary reductions yield atmospheric concentrations in year 2100 at most 2 or 3 percent below the reference case, a significant reduction, though smaller than that generated by the optimistic technological] development scenario.

The following are the major results of the Edmonds et al. project, as presented in the text of the paper and summarized by Edmonds for workshop participants.

- No simple system for distributing emission control obligations and resultant costs is likely to remain acceptable through the next century, so any system of international controls that bites must be designed flexibly enough to allow renegotiation as conditions change.
- The savings available from mechanisms that distribute in ternational abatement efficiently are large, of the order of trillions of dollars, and these savings may be particularly large for developing countries.
- Several simple, plausible schemes for allocating tradable emission permits can have paradoxical effects, and may become infeasible as time passes. For example, in the latter part of the century, China becomes a net loser under equal per capita distribution.
- Not all countries need to participate. The few nations with the largest coal resources can limit cumulative global carbon emissions themselves by controlling or taxing production. To control atmospheric carbon concentration, *only coal matters;* there is not enough oil and gas in the world to make a large difference.
- Modest variation in assumed rates of change of technical end-use efficiency, and rates of diffusion of technical progress among nations, make large differences in the cost and effectiveness of treaties. For example, plausible assumptions of accelerated technology development and deployment can cut total cost of stabilizing world emissions in half.
- Initial delays of 10 or 20 years in implementing emission stabilization have little effect on ultimate atmospheric carbon concentrations.