ISSUES
IN SCIENCE AND TECHNOLOGY

NATIONAL ACADEMIES OF SCIENCES, ENGINEERING, AND MEDICINE
THE UNIVERSITY OF TEXAS AT DALLAS
ARIZONA STATE UNIVERSITY
SUMMER 2017

USA $12.00 / CAN $13.00
7 2>
0 56698 57433 0

M. GRANGER MORGAN
PARTH VAISHNAV
HADI DOWLATABADI
INÊS L. AZEVEDO

Rethinking the Social Cost of Carbon Dioxide

The standard benefit-cost methodology that is used to calculate marginal costs of environmental regulations should not be used for long-lasting greenhouse gases.

There is a very big difference between carbon dioxide and conventional air pollutants. Many of the health and ecological effects of conventional pollutants become apparent in days or a few years. Once emissions cease, conventional pollutants disappear from the atmosphere in just hours or days. Hence it is reasonable to base regulatory policy on an estimate of the damage caused by the emission of an incremental amount of conventional air pollution—that is, on the "marginal damage."

The same is not true for carbon dioxide. A substantial fraction of the carbon dioxide that enters the atmosphere remains there for centuries. Its effects via climate change become apparent only over decades to millennia, and at that point they cannot be reversed by stopping emissions. For this reason, using conventional assessments of marginal damage in benefit-cost analysis to support climate policy fails to consider how little we know about long-term effects of climate change and how these effects should be valued by today's decision makers.

Nevertheless, a number of estimates of the dollar value of the climate change damages associated with the emission of an incremental ton in carbon dioxide emissions have now been made and can be labeled the "social cost of carbon dioxide" (SC-CO_2).

The first serious consideration of using SC-CO_2 by a government agency occurred over a decade ago in the United Kingdom when the Department for Environment, Food and Human Affairs commissioned a pair of studies. In 2006, determining SC-CO_2 was one of the three strategies used by the Stern Review to evaluate the economics of climate change. Although interest in the concept of SC-CO_2 has continued in academic circles, the British researchers Paul Watkiss and Chris Hope have reported that following the Stern Review and the adoption of binding targets for greenhouse gas (GHG) emissions, "the approach to carbon valuation in UK government underwent a major review." Once emission targets were set, the UK government had no further need of SC-CO_2 calculations to justify climate policies.

The United States does not have mandatory GHG emission targets. Since 2009, the federal government, under the direction of the Office of Management and Budget (OMB), has developed and refined official values for SC-CO_2 to be used by government agencies in regulatory decision making. This attempt at rationalization of US policies related to climate change emerged from a legal challenge to a 2006 Final Rule that set Corporate Average Fuel Economy standards for light trucks for model years 2008-2011. In the proposed new standard, the National Highway Traffic Safety Administration discussed the rule's likely effect on carbon dioxide emissions. The rule faced the legal challenge that it failed to monetize the benefits from reducing those climate effects and thus violated President Bill Clinton's 1993 Executive Order 12866 on rulemaking, which, among other things, mandated the use of benefit-cost analysis. In 2008, the United States Court of Appeals for the Ninth Circuit ruled that the highway safety agency's reasoning for not monetizing the benefits of mitigating emissions was arbitrary and capricious. As a consequence, various federal agencies started to comply with the court's ruling by monetizing the costs (or benefits) associated with GHG emissions (or their mitigation) in different ways.

Calculating the social cost of carbon dioxide emissions

The Interagency Working Group on the Social Cost of Carbon constructed its latest SC-CO_2 estimate by performing ten thousand simulations with each of three integrated assessment models, discarding the extreme values found and averaging across the remainder. The average values of SC-CO_2 obtained vary for the year of emissions and the chosen discount rate. For example, if a discount rate of 3% is used, then emissions in 2015 are estimated to have an average value of $36 per ton of carbon dioxide and emissions in 2050 to have an average value of $69. If a discount rate of 2.5% is used, the values are $56 and $95, respectively. All these are in 2007 dollars.

These averages do not make explicit the high uncertainty in SC-CO_2 calculations. The working group concluded that there was a 5% chance that at a 3% discount rate these values would exceed $105 and $212 for emissions in the years 2015 and 2050, respectively. When a cumulative distribution function is fit to the reported results for a ton of carbon dioxide emitted in the year 2020, the 10%, 50%, and 90% points on the resulting cumulative distribution fall at $3, $5, and $20 per ton of carbon dioxide. For a discount rate of 2.5%, the comparable values are $8, $38, and about $150 per ton of carbon dioxide.

In an effort to impose consistency across agencies, an Interagency Working Group on the Social Cost of Carbon (IWG) was formed in 2009. The IWG was charged with producing an estimate of the marginal benefits of carbon dioxide mitigation. The group produced its first recommendations in 2010 and subsequently published updates in 2013, 2015, and 2016. The resulting SC-CO_2 estimates are intended to provide a yardstick to assess whether government policies for mitigation of carbon dioxide emissions yield net benefits and allow for different alternatives to be ranked in terms of efficacy and effect. SC-CO_2 values are also now widely used outside of government when analysts address technology and policy alternatives that influence the release of GHGs to the atmosphere.

In 2015, the IWG asked the US National Academies to review the SC-CO_2 with the objective of guiding future revisions. In early 2017, the study committee released a detailed report that makes recommendations on the choice of models and damage functions, climate science modeling assumptions, socioeconomic and emissions scenarios, presentation of uncertainty, and temporal discounting. Shortly after the report's release, President Trump signed on March 28, 2017, an executive order on "Promoting Energy Independence and Economic Growth," which disbanded the IWG and withdrew all of its reports "as no longer representative of governmental policy." We will return to these two developments later.

The machinery behind the curtain

Conceptually, the IWG computes the social cost of carbon dioxide by running an integrated assessment model (IAM) to assess the present value of the future monetized consequences of climate change. Present value is obtained by using a technique called exponential discounting. Then an additional ton of carbon dioxide is added, and the model is run again. The difference between the two present values is computed and taken to be the SC-CO_2. Depending on the assumptions it made, the IWG has estimated values for the SC-CO_2 that fall between a few tens of dollars per ton to over $100 per ton.

Four things are needed to compute the SC-CO_2: a reasonable projection of how future global emissions of GHGs are likely to evolve; a model that estimates how those future emissions of GHGs will change the climate; a model of all the consequences of that climate; and a way to assign monetary values to all those consequences (at least partly so that qualitatively disparate damages may be combined).

Since the early 1990s many researchers have developed increasingly elaborate models of climate change, its dynamics, and its impacts. Some models have tried to integrate across all key elements, from demographics and economics through climate change and effects, in order to deliver a coherent, albeit less detailed, system for policy analysis. There is of course uncertainty about both how future GHG emissions and land use will evolve and how the climate will change as a result. There is even greater uncertainty about the consequences of these changes and how they should be valued. Thus, it is not surprising that Paul Watkiss and Thomas Downing, also a British researcher, reported in a 2008 review that estimates of SC-CO_2 "span at least three orders of magnitude, reflecting uncertainties in choices of key parameters/variables." In 2014, Robert Pindyck, an economist at the Massachusetts Institute of Technology, wrote: "IAMs are of little or no value in evaluating alternative climate change policies and estimating SCC [social cost of carbon]. On the contrary, an IAM-based analysis suggests a level of knowledge and precision that is non-existent, and allows the modeler to obtain almost any desired result."

Meaningful quantitative valuation not possible

In the 1990s, two of us (Dowlatabadi and Morgan) led the development of one of the first integrated assessment models, called the Integrated Climate Assessment Model (ICAM). This model was designed with the express purpose of reflecting key uncertainties (in model structure, parametrization, and valuation) with internally coherent projections of drivers, dynamics, and impacts of and interventions for climate change mitigation, adaptation, and geoengineering. Our experience mirrored Pindyck's conclusion that IAMs cannot produce quantitative estimates on which policy should be based. However, we believe well designed and internally consistent IAMs can produce useful qualitative insights about alternative climate policies. After a decade of work on ICAM, we chose to end further development for two reasons: we could not produce trajectories that were internally consistent within ICAM and also matched those produced by the Intergovernmental Panel on Climate Change (IPCC), and when we included structural uncertainties, it became possible to produce almost any outcome. We were also concerned that quantitative results from integrated assessment models such as ours were being used without an adequate discussion of the vast uncertainties. Unfortunately, false precision from IAMs is being used in the generation of quantitative "answers" that have come to serve as an inappropriate foundation for public policies.

As noted above, GHG concentrations are cumulative. How emissions will evolve in the future is unclear and will obviously depend on myriad social choices. In IPCC's baseline scenario, the Earth is projected to run out of economically recoverable oil and gas by the 2050s, with coal returning as the dominant primary source of liquid and gaseous energies. However, renewable energy sources such as solar and wind power are now more economical than fossil energy in many parts of the world. In other parts, coal is being eschewed because of concern about air pollution. Hence, the range of likely future GHG emissions spans the gamut from the gloomy return to coal of the IPCC baseline to far lower figures.

We know that the response of the climate system to changes in radiative forcing (the heat energy added to the atmosphere as a result of increasing GHG concentrations) is nonlinear. Geologic evidence indicates that the Earth has several quasi-equilibrium climate states. The feedbacks that have blessed the planet with a stable "climate optimum" for the past ten thousand years are uncertain in magnitude and operate over limited perturbations. Beyond that range of perturbations climate system dynamics may tip to a very different climate state. Nobody can adequately assess the probability and consequences of such climate transitions. If and when such transitions occur, many resulting changes will *not* be marginal.

Even if we knew all the consequences of changing climate, the idea that one can find an optimal global policy makes little sense given the uneven distribution of costs and benefits around the world and among different stakeholders. Many of these changes will not be marginal in nature. Although side-payments are sometimes proposed, the practicality of such payments is based on the idea that the costs borne by the losers can be meaningfully monetized, the cost of compensating them adequately estimated, and the compensation actually paid. Even in the simple case of inundation through sea level rise, experience with displaced populations from places such as the Bikini Islands and Diego Garcia suggests that compensation of the "value of lost real estate" does not begin to make up for the loss experienced by the affected peoples. The inhabitants of these communities were moved previously during the Cold War. Their resultant high suicide rates, short life expectancy, and broken social structures make it clear that they have failed to "adapt" to their new locations, even after half a century. The problem grows only more complex when other damages are considered for valuation.

Climate change and its effects will vary by location, ecosystem, and socioeconomic context. The responses of social, economic, and ecological systems are also likely to be nonlinear, with some entering protracted periods of unstable chaos while others undergo rapid transition to conditions fundamentally, not marginally, different from today. We neither know how to characterize such effects or how they will be valued across different cultures, societies, and future generations. Indeed, monetizing, combining, and discounting these heterogeneous and contextual effects as a single global monetary metric displays a hubris that has been roundly condemned by ethicists and decision analysts.

As noted above, it is possible that the climate system and a number of social, political, economic, and ecological systems can undergo transitions to other states that are not reversible, at least on timescales relevant to human affairs. Whereas some of the changes could be global or hemispheric in nature (such as dramatic shifts in the El Niño-Southern Oscillation, the Meridional Overturning Circulation, and the Indian monsoon), some will be quite local (such as a long-term change in circulatory patterns that makes local rain-fed agriculture possible or impossible in some regions). Tipping points related to effects likely also display a wide range of scales. The range of changed climate patterns and states, along with the range of changed effects, considerably complicates the issue of what constitutes catastrophic change. A change that is viewed as minor by some may be viewed as catastrophic by others.

As we have argued above, nonmarginal effects cannot be translated into marginal damage "costs." The nonmarginal effects may be local, not in the market, incalculable, and not amenable to compensation. As such, the local damage function can be almost infinite. For example, an ecosystem may be eliminated or a traditional way of life that depends on an ecosystem may disappear. Impact studies incorporate such damages and evaluate them "at the margin," then aggregate them to form a damage function used to calculate the SC-CO_2. But these figures mask the inadequacy of financial compensation for the subjective damages being incurred.

Global damage functions may combine losses of amenities, such as higher air conditioning costs in the US Southwest (which may be large in economic terms but marginal in nature), with losses of natural or human patrimony (to which a small economic value may attributed, but which are nonmarginal and irreversible). None of the suggested approaches to equity weighting or discounting schemes address such nonmarginal damages.

The ability to choose appropriate policy given uncertainty in cost and benefits has been one of the greatest theoretical achievements of resource economics. When the marginal damages are much shallower than marginal costs of mitigation, it is appropriate to use price mechanisms, such as a carbon tax. When uncertain marginal damages are likely steeper than the marginal cost of mitigation, it is appropriate to cap emissions. This is what the United Kingdom (and the European Union) did following the Stern review. In the United States, the absence of GHG emission targets signals the philosophical divide across the Atlantic as well as the continued reliance on the SC-CO_2.

There are two flaws to carbon pricing based on net marginal cost: the first is the assumption that the

Choi + Shine Architects: Infrastructure

Choi + Shine Architects is an international architecture and design studio run by Boston-based husband and wife team Jin Choi and Thomas Shine. They draw inspiration for their innovative and experimental practice from unusual places: Lego blocks, a silk blouse, and sea urchins, to name a few, as well as their varied backgrounds in sociology, fine art, and architectural history and theory (Choi), and engineering and ceramics (Shine). Many of their projects are not inherently architectural, and their work ranges from product design to commercial and residential architecture to artwork and infrastructure design. They have proposed several designs to make new and existing infrastructure aesthetically appealing and whimsical, transforming electrical towers into monumental sculptures in projects that include *The Land of Giants, Mantis Tower, Bamboo Tower, Centipede Tower,* and *Swords Tower.*

Their infrastructure project, *The Land of Giants,* winner of a design competition sponsored by Landsnet, a power transmission company in Iceland, transforms mundane electrical pylons into giant figures on the Icelandic landscape. Making only minor alterations to well established steel-framed towers, they propose creating iconic towers that will become monuments. Seeing the pylon figures will become an unforgettable experience, elevating the towers to something much more than merely a functional design of necessity. The project has won several awards, including the Boston Society of Architects Design Award for Best Unrealized Architecture (2010).

Adding aesthetic appeal to functional infrastructure, Choi + Shine's work also often contains a narrative element that has the potential to make viewers think more about their relationship to the environment and the natural resources they consume.

CHOI + SHINE ARCHITECTS, Concept proposals (this page and opposite) for *The Land of Giants*, 2008.

damages are marginal and shallow (clearly not so, given the above discussions), and the second is the assumption that mitigation is costly and steep. This, too, is not an assertion that is supported by evidence of energy supply choices of the past decade. Fossil fuels are being rejected for air pollution and energy security reasons. They are also facing stiff economic competition from renewables. In fact, with the strictest of emission caps, negative emissions can be achieved through the capture of carbon dioxide from the free atmosphere. As one can infer from the recent National Academies' report *Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration*, this option is available at a fixed, not rising, marginal cost within the range of calculated SC-CO_2.

An alternative quantitative approach

For decades both Republican and Democratic administrations have issued executive orders requiring quantitative benefit-cost assessments of major federal regulations. The Interagency Working Group on the Social Cost of Carbon used a complex process to do this (that the National Academies' recommendations would make even more complex)—and it did produce numbers. But we do not believe these numbers are meaningful, even as they have given agencies values they can plug into their benefit-cost analyses.

The executive orders requiring benefit-cost assessments contain language that says if it is not feasible or appropriate to quantify costs, other approaches can be adopted. However, neither federal agencies nor the courts have demonstrated much willingness to adopt such alternatives, even for assessing civil rights laws such as the Americans with Disabilities Act. Although it is tempting to say that in addressing climate change, OMB should abandon the search for dollar values and adopt some other strategy, such a proposal is not likely to succeed, and in today's political climate it could further contribute to retrograde policy developments.

In place of using the SC-CO_2, we believe that a more defensible method can be based on identifying and avoiding climate change thresholds: temperatures or GHG concentration levels at which damages are likely to become unacceptable. Such an approach meets conditions that damage estimates cannot. The great advantages of such an approach are that the costs involved in achieving different emissions reduction levels are in the marketplace and in metrics that are universally accepted; the costs of emission reductions can be covered through side payments or technology transfers or both; and, since the goal is to transition to an entirely new energy system, the marginal costs may even start falling as the policy progresses. In such systems, richer countries can act as early adopters and drive down the cost of technology, allowing later adopters to make the transition at a lower cost. Early adopters can even see themselves as part of a social movement and view their expenditures not as a cost but as an expression of their commitment to civic responsibility.

In place of developing policies based on a SC-CO_2, the European Union has adopted a strategy of setting a cap on member nations' emissions of carbon dioxide and other greenhouse gases. As outlined below, once such a cap has been established, it is perfectly feasible to back out a dollar cost for eliminating each incremental addition of carbon dioxide or other GHG to the atmosphere.

In the climate negotiations in Paris, which led to the Paris Agreement, most of the world's political decision makers reached the conclusion that the consequences of global temperature change above 2 degrees Centigrade (2°C) were unacceptable. Hence, their pledge to limit climate change to 2°C (3.6 degrees Fahrenheit) or less. (It must be noted, however, that on June 1, 2017, President Trump announced that he planned to withdraw the United States from the climate accord.) In a series of studies, the Norwegian climate researcher Glen Peters and his colleagues have estimated how much more carbon dioxide can be added to the atmosphere before the average temperature of the planet rises by 2°C. A similar calculation of remaining "atmospheric capacity" can be done for any temperature increase. Although such a calculation involves uncertainty, it has a much narrower range of uncertainty and is more defensible than SC-CO_2 calculations.

Independent of how the remaining atmospheric capacity is allocated among emitting parties, if the planet is going to hold warming below catastrophic levels, the United States, the European Union, China, and all other major emitters will need to reduce their emissions of long-lived GHGs by 80-90% in the next two or three decades. Although two or three decades is a very long time for many firms making investment decisions in a market economy, it is almost instantaneous in terms of institutional change and the turnover of long-lived physical infrastructure. This means that the prospects of holding the amount of warming below a level of 2°C looks increasingly unlikely.

Writing in *Nature* in 2009, the British climate scientist Miles Allen and his colleagues have observed that "either we specify a temperature or concentration target and accept substantial uncertainty in the emissions required to achieve it or we specify emissions and accept even more uncertainty in the temperature response." One of us (Dowlatabadi) has argued that a target based on atmospheric GHG concentration involves less uncertainty and is more easily implemented.

But these are details. Either way, the path is the same: set some target; estimate an "emissions reduction supply curve"; and from that estimate either an ultimate cost to achieve all of the needed reduction or a per-ton cost that evolves rapidly over time as the world transitions away from a fossil economy. By an emissions reduction supply curve, we mean a plot of cost as a function of the amount of emission reduction. Such a curve starts out negative (there are ways to reduce some emissions while also saving money—for example, with improved energy efficiency and conservation) and then rises, at least for a while, as deeper reductions are required. Over time, technological innovation and managerial experience might slow the rise in cost or even eliminate it.

If the target is specified as staying below some specific average increase in average global temperature, we will also need a plot of how temperature change will be related to emissions (call it the "warming curve"). Then, combining the emission reduction supply curve with the warming curve will allow one to compute the needed amount of reduction in emissions and hence an average cost per ton of carbon dioxide to stay below a certain temperature change. Both these curves involve some uncertainty, so the resulting cost would actually be a probability distribution. The Office of Management and Budget could use that distribution in its benefit-cost analyses, or it could specify how risk averse it wants to be and choose a single cost point on that distribution.

An alternative way to specify a target is as a desired emission trajectory over time—the kinds of curves that the Intergovernmental Panel on Climate Change and many others in the climate community have produced in abundance. Having chosen such a trajectory, one can use the emission-reduction supply curve to compute a cost per ton (roughly the equivalent of an emission tax) that would be required to stay on that curve. Of course, if the carbon dioxide that is already in the atmosphere were allocated to the different nations that have emitted it and the level of economic development that each has achieved is considered, an argument could be made that over the next several decades different blocs of nations should undertake different amounts of reduction as a function of time.

Such strategies do not involve an estimate of the marginal damage arising from each emitted ton of carbon dioxide. Rather than applying future discounting, which makes even the most catastrophic outcome appear small if it falls far enough in the future, they would simply depend on a scientifically informed normative judgment that there is a point beyond which more climate change runs too high a chance of producing catastrophic damage to the planet's peoples and ecosystems. There are, of course, uncertainties associated with such approaches, but we stress that such frameworks involve far less uncertainty than the SC-CO_2 approach, since they do not require the series of assumptions that must be made to coerce a range of disparate damages into a single global monetary metric.

Clearly it is essential that research, development, and deployment for energy technologies be continued to drive down the cost of low- and zero-carbon energy technologies. For the next few years the US Department of Energy may reduce its support for such work, but many states, as well as other nations, will push forward, as will firms that adopt a longer view of likely future market demand. If costs of low-carbon energy technology continue to fall, the net benefits provided by such technologies will increase.

National Academies v. Trump

In a remarkable display of confidence that further refinement of models and methods will make it possible to meaningfully forecast and quantify the consequences of future climate change, the National Academies' report on the social costs of carbon dioxide endorses the IWG's basic approach. It recommends that rather than use existing integrated assessment models, a new and improved IAM should be constructed. It argues for "the creation of an integrated modular SC-CO_2 framework that provides transparent articulation of the inputs, outputs, uncertainties, and linkages among the different steps of SC-CO_2 estimation." It calls for an improved treatment of "interactions and feedbacks among the modules of the SC-CO_2 framework if they are found to significantly affect SC-CO_2 estimates" and argues to extend assessment "far enough in the future to provide inputs for estimation of the vast majority of discounted climate damages." Last, it calls for greater use of statistical techniques and greater use of expert elicitation to quantify key uncertainties.

Given President Trump's recent Executive Order disbanding the IWG, the US government is not likely to undertake such an effort in the next few years, nor is it likely to actively support efforts to reduce GHG emissions. However, there are others talking about using private funding to continue work in refining the SC-CO_2 framework. Given the many needs facing the US climate research communities, we do not believe that it makes sense to invest scarce funds in further refining the SC-CO_2. Yet even as other approaches than the SC-CO_2 should be pursued to guide policy, the present inaction by the US federal government is a grave mistake. The window to limit warming to anything like 2°C is rapidly closing.

Although President Clinton's Executive Order 12866 required that "agencies should assess all costs and benefits," it does recognize that "some costs and benefits are difficult to quantify." In such cases, it requires that agencies act based "upon a reasoned determination that the benefits of the intended regulation justify its costs." Congress has not enshrined the current SC-CO_2 approach in statute. The requirement that some monetary value be assigned to greenhouse gas emissions stems from the Ninth Circuit Court of Appeals' decision, which does not specify a method by which this value should be arrived at, asserting only that "the value of carbon emissions reduction is certainly not zero." The courts are therefore likely to give the executive branch, including agencies, considerable latitude in determining how to implement Executive Order 12866.

As an alternative to further refining the SC-CO_2 framework, we believe a group of private foundations, corporations, and others—perhaps in collaboration with several supportive state governments (and the federal government if in the future it could be persuaded to participate)—should undertake a serious cooperative effort in policy-focused research designed to develop cost estimates derived from the creation of a cap on future warming or future US greenhouse gas emissions. Choosing caps is inherently normative. Hence, as a second phase of such an effort, we believe that a high-level national commission should be assembled, made up of thoughtful citizens including but not limited to climate scientists, economists, ecologists, technology experts, and ethical leaders. Building on the policy analytic work just outlined, this commission should be charged with reviewing the scientific evidence on the consequences of climate change and the costs of emission controls, as well as with developing recommendations for the choice of caps for the United States. We believe that the courts could be persuaded that policies that emerge from such an exercise are based on "a reasoned determination that the benefits of the intended regulation justif[ies] its costs."

There is an urgent need to take serious action now to reduce emissions of carbon dioxide and other greenhouse gases. At the federal level, the United States may not make much progress on reducing its emissions of carbon dioxide in the next few years, but we should not let those years be wasted. Many states and cities are taking action now. Within a few years the federal government (and the Office of Management and Budget) may once again become serious about controlling emissions. When that happens, we should have already laid the foundations for a system that OMB and others can use to drive emissions reductions that is more defensible than the SC-CO_2 framework. Emitting a ton of carbon dioxide to the atmosphere causes damage. We may not be able to defensibly monetize the damage done by each ton of emissions, but the evidence is clear that it is high and growing higher with each passing year.

M. Granger Morgan, Parth Vaishnav, and Inês L. Azevedo are on the faculty of the Department of Engineering and Public Policy at Carnegie Mellon University. Hadi Dowlatabadi is on the faculty of the Institute for Resources Environment and Sustainability at the University of British Columbia in Vancouver, Canada.

Recommended reading

Report of the Committee on Assessing Approaches to Updating the Social Cost of Carbon, *Valuing Climate Damages: Updating Estimation of the Social Cost of Carbon Dioxide*, Board on Environmental Change and Society (Washington, DC: National Academies Press, 2017).

H. Dowlatabadi, "Bumping against a gas ceiling," *Climatic Change* 46, no. 3 (2000): 391-407.

Interagency Working Group on Social Cost of Carbon, *Technical Update of the Social Cost of Carbon for Regulatory Impact Analysis Under Executive Order 12866* (2016).

National Research Council, *Climate Intervention: Carbon Dioxide Removal and Reliable Sequestration* (Washington, DC: National Academies Press, 2015).

G. P. Peters, R. M. Andrew, S. Solomon, and P. Friedlingstein, "Measuring a fair and ambitious climate agreement using cumulative emissions," *Environmental Research Letters* 10, no. 10 (2015).

Presidential Executive Order on Promoting Energy Independence and Economic Growth (March 28, 2017), available online: https://www.whitehouse.gov/the-press-office/2017/03/28/presidential-executive-order-promoting-energy-independence-and-economi-1.