

University of Pennsylvania Law Review

FOUNDED 1852

Formerly
American Law Register

VOL. 155

JUNE 2007

No. 6

ARTICLES

SCIENTIFIC CHALLENGES IN THE ATTRIBUTION OF HARM TO HUMAN INFLUENCE ON CLIMATE

MYLES ALLEN,^a PARDEEP PALL,^b DAITHI STONE,^c PETER STOTT,^d
DAVID FRAME,^e SEUNG-KI MIN,^f TORU NOZAWA^g & SEJI YUKIMOTO^h

^a Atmospheric, Oceanic, and Planetary Physics, Department of Physics, University of Oxford (myles.allen@physics.ox.ac.uk). Dr. Allen was supported by the NOAA/DoE International Detection and Attribution Group; Dr. Pall and Dr. Stone were supported by the UK Department of Environment, Food, and Rural Affairs (DEFRA) contract PB/B38311 with additional support from WWF International and the Tyndall Centre; Dr. Stott was supported by DEFRA contract PECD 7/12/37. Further thanks go to the editors of the University of Pennsylvania Law Review for their hard work and patience.

^b Atmospheric, Oceanic, and Planetary Physics, Department of Physics, University of Oxford.

^c Atmospheric, Oceanic, and Planetary Physics, Department of Physics, University of Oxford; Tyndall Centre for Climate Change Research, Oxford University Centre for the Environment.

^d Met Office, Reading Unit, Department of Meteorology, University of Reading.

^e Oxford University Centre for the Environment.

^f Climate Research Division, Environment Canada.

^g National Institute for Environmental Studies 16, Japan.

^h Meteorological Research Institute, Japan.

The authors of this Article review the current state of the science of attribution of anthropogenic climate change, with particular emphasis on the methodological challenges that are likely to confront any attempt to establish a direct causal link between greenhouse gas emissions and specific damaging weather events. Standard “detection and attribution” analyses, such as those cited by the Intergovernmental Panel on Climate Change (IPCC), are generally sufficient to establish the strength of human influence on large-scale, long-term-average climate, but fall short of quantifying the role of greenhouse gas emissions in almost any conceivable case of actual harm, since nobody is directly exposed to a change in global average temperature alone. The authors argue that it should be possible to agree on a relatively objective approach to quantifying the role of human influence on climate in cases of actual harm. There are, however, a number of questions to be resolved, including: can we apply the concept of Fraction Attributable Risk, developed for population studies in epidemiology, to the analysis of an unprecedented change in a single system such as the world’s climate? Can we rely on computer simulation to address counterfactual questions such as “what would the climate have been like in the absence of twentieth century greenhouse gas emissions,” given that we are working with imperfect simulation models? Due to multiple anthropogenic and natural contributions to changing weather risks, it will always be necessary to apply some kind of principle of ceteris paribus to quantify the role of any particular causal agent, such as greenhouse gas emissions. How is this principle to be applied? These questions are not, in themselves, scientific issues, although how they are to be resolved will have a direct bearing on how and whether climate science can inform specific causal attribution claims. In summary, we need the legal community to ask the scientific community the right questions. It is imperative that these issues be resolved as soon as possible, to avoid having them become entwined in the outcomes of specific cases. Thus, this Article serves as a kind of tutorial, going over some material that many will find familiar in order to place it in the context of attribution.

I. INTRODUCTION AND MOTIVATION: WHY LAWYERS NEED TO UNDERSTAND THIS STUFF

In addressing a crowd of lawyers, the first and overarching challenge for a physicist is to keep the audience awake. This is particularly pertinent when the topic is climate change, because many in the legal community may feel that the methodological details of causal attribution are of academic interest only. When and if cases ever come to court claiming that actual or imminent harm is attributable to human influence on climate, this argument runs, competing sides will invite experts to testify for and against some assertion like “past greenhouse

gas emissions contributed substantially to harmful weather event X,” and cases will be decided on the basis of the credibility and qualifications of the relevant experts. Hence lawyers do not need to trouble themselves with the details; they simply need to line up authoritative expert witnesses.

We would argue that such heavy reliance on expert testimony is both unnecessary and counterproductive. It is unnecessary because it should be possible to agree on a standard method for the attribution of harm to human influence on climate, such that results do not depend, at least as a first approximation, on the prior opinions of the scientists responsible for making the assessment. If the legal community demands it, the meteorological “raw material” for any assessment of harm could be provided by a straightforward extension of routine weather forecasting services, generated by the same neutral government agencies that currently provide forecasts. Expert judgment will naturally still be required for the interpretation of this material in any particular case, but, except in the most marginal situations, this reliance on expert judgment can be minimized if we can agree on an “industry standard” operational approach to the attribution problem, such as the one we advocate here.

Heavy reliance on expert testimony in causal attribution claims is also likely to be counterproductive. First, there is considerable room for confusion over what expert testimony actually means. For example, we would all subscribe to the following two statements: “human influence on climate played a substantial role in causing the European heat wave of 2003”; and “it is impossible in principle to attribute any single weather event to human influence on climate.” To a layperson, these statements appear contradictory. They only make sense if it is understood that “causing” an event in the first statement must be interpreted as “contributing to the risk of” that event occurring, while the second statement is only true if we interpret “attribute” narrowly as “were it not for human influence, this event would not have occurred.” Hence the legal community cannot expect the scientific community to do its job for it: it will always be possible for both sides in a case to persuade experts (and possibly even the same expert) to subscribe to statements that sound as if they support their cause. Expert testimony thus will only be useful if the scientific context is understood.

Overreliance on expert testimony is also likely to have a damaging impact on the scientific community. We do not have to look very far for an example of the damage that can be caused by personalizing sci-

entific debate. Some critics of the Mann et al. reconstruction of Northern Hemisphere temperatures over the past millennium,¹ as cited by the 2001 Scientific Assessment of the Intergovernmental Panel on Climate Change (IPCC),² have made a sustained effort to divert the debate from methodological issues of paleoclimate reconstruction to procedural issues and ad hominem attacks over what source code was available when, how the reconstruction came to be cited by the IPCC, and so forth.³ Scientific research will rapidly become impossible if individual researchers have to operate under the threat of their research software being subpoenaed or if scientific statements must be individually scrutinized for their potential legal implications.

The only way to avoid this (scientists') nightmare scenario is to operationalize the attribution problem now, which must include the development of an agreed approach to incorporating new research into the attribution procedure. Any operational attribution procedure will always tend to be somewhat behind front-line research, whereas the "latest" research results tend, by their very nature, to be relatively unreliable until they have been fully digested by the scientific community. For this reason, weather forecasters are typically conservative about incorporating new research into their forecast models; they generally require some demonstration that an innovation will demonstrably improve forecast skill before making such a change. Demonstrating that an innovation—such as an improved model or a new reconstruction of solar variability—will improve the accuracy of attribution statements is clearly more problematic, however, because we do not have a sequence of cases on which to test out the impact of the innovation as the weather forecasters do.

Popular discussions of the climate issue often give the impression that the whole science of attribution has to return to square one whenever a new feedback process is identified or a new solar reconstruction is published. This would obviously make life impossible for

¹ Michael E. Mann et al., *Global-Scale Temperature Patterns and Climate Forcing over the Past Six Centuries*, 392 NATURE 779 (1998).

² INTERGOVERNMENTAL PANEL ON CLIMATE CHANGE, WORKING GROUP I, CLIMATE CHANGE 2001: THE SCIENTIFIC BASIS 99, 133-36 (J.T. Houghton et al. eds., 2001) [hereinafter SCIENTIFIC BASIS]. This assessment summarized, discussed, and highlighted the reconstruction from Mann et al., *supra* note 1.

³ See, e.g., Letter from Joe Barton and Ed Whitfield, Chairmen, Subcomm. on Oversight Investigations of the H. Comm. on Energy & Commerce, to Michael Mann, June 23, 2006, available at http://republicans.energycommerce.house.gov/108/Letters/062305_Mann.pdf.

the legal community because new research papers that might potentially have some bearing on attribution claims will always be coming out. Nevertheless, we believe it should be possible to agree on a methodology for deciding to what extent new research results refute past attribution claims, and by applying the same standards, to agree on whether or not the current version of the operational attribution system is good enough to be applied to specific questions.

To conclude this introduction, we stress that this review is focusing exclusively on the meteorological attribution problem: to what extent human influence on climate can be “blamed” for observed weather trends and specific weather events such as floods, storms, or heat waves. We will not be addressing the equally important problem of how other factors, such as decisions to build in vulnerable locations or to maintain coastal defenses, as well as offering health services, can contribute to or mitigate actual losses, nor how these losses are themselves related to meteorological variables. We believe that these questions can be considered separately because only in very exceptional cases (such as the carbon released by large-scale regional deforestation) will there be any significant feedback between local decisions affecting vulnerability to weather and global changes affecting the weather itself.

II. “CLASSICAL” DETECTION AND ATTRIBUTION: THE EVIDENCE FOR HUMAN INFLUENCE ON GLOBAL CLIMATE

The aim of this section is to outline the principles behind the standard approach to detecting and attributing human influence on global and regional climate as they are used by the scientific community to provide the basis for successive IPCC Scientific Assessments.⁴ We will avoid methodological details, and instead focus on clarifying the precise question that is being asked in conventional detection and attribution studies and the assumptions underlying their answers in order to highlight issues that might prove contentious if these studies are ever used in any legal action.

Detection and attribution studies have traditionally focused on large-scale, long-term changes in variables such as seasonal or annual mean surface temperature.⁵ Changes in such metrics do relatively lit-

⁴ See SCIENTIFIC BASIS, *supra* note 2, at 695, 700-13 (explaining the history and elements of climate change detection and attribution).

⁵ For examples and reviews of such studies, see Myles R. Allen et al., *Quantifying Anthropogenic Influence on Recent Near-Surface Temperature Change*, 27 SURVS. GEOPHYSICS

tle damage in and of themselves, but these studies provide the overall basis for any attempt to focus on more impact-relevant variables such as sea level, so it is important that they are fully understood.

The scientific context for global-scale detection and attribution studies is illustrated by Figure 1,⁶ which is modeled after the recent analysis by Stott et al.⁷ Panel (a) compares global mean temperatures simulated by a range of climate models driven by a combination of human and natural influences (thin gray lines), with observed changes in global temperature over the twentieth century (heavy black line). Panel (b) compares the same quantity with simulations of a subset of the models shown in panel (a) driven by natural influences alone.⁸ The most important anthropogenic driver of climate change over the period included in these model simulations is the increase in “well-mixed” greenhouse gases (GHGs) (so called because these gases have relatively long lifetimes and mix relatively homogeneously through the atmosphere): predominantly carbon dioxide, methane, nitrous oxide, and chlorofluorocarbons (CFCs).

For a first approximation, the climatic impact of well-mixed GHGs can be evaluated by considering them as a single global driver of climate change, with the contribution of any individual gas at any given time being proportional to its “radiative forcing,” or the change in the overall balance of incoming and outgoing energy evaluated over the whole planet that would result from resetting the concentration of that gas back to its preindustrial value. More generally, the scientific community refers to external drivers of climate change as “climate forcings” or “climate forcing factors” whose magnitudes are conventionally compared with respect to their impact on the global energy budget. There are complications, in that the impact on the energy budget of a given increase in the two most important well-mixed GHGs, carbon dioxide and methane, is affected by the composition of the atmosphere, which is itself changing. The immediate impact on the global energy budget of a ton of carbon dioxide released today is

491 (2006); Peter A. Stott, *Attribution of Regional-Scale Temperature Changes to Anthropogenic and Natural Causes*, GEOPHYSICAL RES. LETTERS, July 16, 2003, at 2-1; and Simon F.B. Tett et al., *Causes of Twentieth-Century Temperature Change Near the Earth's Surface*, 399 NATURE 569 (1999).

⁶ All figures are found in the Appendix, *infra*.

⁷ Peter A. Stott et al., *Observational Constraints on Past Attributable Warming and Predictions of Future Global Warming*, 19 J. CLIMATE 3055 (2006).

⁸ At the time of writing, not all modeling groups have completed these natural-only simulations.

only about 73% of the impact that ton would have had if it had been released into a preindustrial atmosphere because carbon dioxide levels have already increased by almost 40%. The impact of any methane increase also depends on current carbon dioxide concentrations.⁹

These interactions mean that some agreed interpretation of the *ceteris paribus* principle is necessary even to work out the simplest climatic impact of any anthropogenic change in GHG levels. Within the scientific community, the usual approach to comparing the impact of two GHG emissions (a ton of carbon dioxide versus a ton of methane, both emitted today, for example) is to consider the impact on the global energy budget, integrated over some agreed time scale, of not emitting each one under an agreed, relatively realistic, scenario for the evolution of all other gases. Rather than considering how injecting a given amount of carbon dioxide into the atmosphere would have disturbed the preindustrial climate, we ask how not injecting that amount of carbon dioxide would alter our present-day and projected future climate.

The impacts of the other main anthropogenic drivers of climate change depend much more heavily on location, season, and even local weather conditions, making them much harder to sum up in terms of a global radiative forcing. These drivers include the aerosols that are generated primarily by sulfur emissions from power stations and industrial processes. These have a cooling effect on climate, both directly—by reflecting incoming sunlight—and indirectly—through their influence on cloud amount and cloud properties. Stratospheric ozone depletion (due to CFCs) and tropospheric (near-surface) ozone increase (primarily due to exhaust gases from hot combustion processes, such as vehicle engines) also have an impact on regional climate, although, averaged over the whole globe, their impact appears to be small relative to that of the well-mixed GHGs.¹⁰ Finally, changes in land use, which have been included in only some of the

⁹ For an accessible review, see Keith P. Shine & Piers M. de F. Forster, *The Effect of Human Activity on Radiative Forcing of Climate Change: A Review of Recent Developments*, 20 GLOBAL & PLANETARY CHANGE 205 (1999).

¹⁰ See G.C. Hegerl et al., *Multi-Fingerprint Detection and Attribution of Greenhouse Gas, Greenhouse Gas-Plus-Aerosol and Solar Forced Climate Change*, 13 CLIMATE DYNAMICS 613, 628 (1997) (observing that aerosol-induced cooling is likely to occur primarily in the mid-latitudes of the Northern Hemisphere); V. Ramanathan et al., *Aerosols, Climate, and the Hydrological Cycle*, 294 SCIENCE 2119, 2123 (2001) (suggesting that the cooling effect of sulfate aerosols will manifest itself primarily in the tropics and subtropics, while the effects of GHGs will be most pronounced in the Southern Hemisphere and high latitudes).

model simulations performed to date, can have a highly significant impact on local climate, although their impact on global climate and on regions remote from the land-use change also appears to be relatively small.

The two natural drivers of climate change included in the model simulations shown in Figure 1 are changes in the power output of the sun and aerosols injected into the stratosphere by explosive volcanic eruptions. Solar activity varies naturally over the well-documented eleven-year “sunspot cycle,” but variations on longer time scales are much more controversial.¹¹ The problem is that we only have direct observations of solar output over the past two and a half eleven-year cycles, over which period it is impossible to discern any clear background trend. Efforts have been made to reconstruct solar variations over the past few centuries using, for example, observations of geomagnetic activity and analogs based on observations of sun-like stars. All of these reconstructions suffer from a calibration problem: they tend to be much more confident about the overall shape of recent changes in solar activity (with activity reaching a minimum around the time of the “Maunder minimum” in the seventeenth century) than they are about the size of these changes.¹² Estimates of the change in solar power output between the Maunder minimum and today vary by almost an order of magnitude between studies.¹³ This problem, which also applies to reconstructions of the impact of explosive volcanos, has a direct impact on the techniques used for climate change detection and attribution.¹⁴

¹¹ Compare Judith Lean et al., *Reconstruction of Solar Irradiance Since 1610: Implications for Climate Change*, GEOPHYSICAL RES. LETTERS, Dec. 1, 1995, at 3195 (concluding that approximately half of the global warming observed since 1860, but less than one-third of the warming since 1970, can be attributed to solar activity), with Douglas V. Hoyt & Kenneth H. Schatten, *A Discussion of Plausible Solar Irradiance Variations, 1700-1992*, 98 J. GEOPHYSICAL RES. 18,895, 18,904 (1993) (positing that 71% of the warming that has occurred since 1891 can be attributed to fluctuations in solar irradiance). Judith Lean and David Rind, in *Evaluating Sun-Climate Relationships Since the Little Ice Age*, 61 J. ATMOSPHERIC & SOLAR-TERRESTRIAL PHYSICS 25, 25 (1999), summarize as follows: “while solar variability likely played a dominant role in modulating climate during the Little Ice Age prior to 1850, its influence since 1900 has become an increasingly less significant component of climate change.”

¹² For a review, see Judith Lean, *Living with a Variable Sun*, PHYSICS TODAY, June 2005, at 32, 35-37.

¹³ See *supra* note 11 (comparing the conflicting studies by Lean et al. and Hoyt and Schatten).

¹⁴ See Natalia G. Andronova et al., *Radiative Forcing by Volcanic Aerosols from 1850 to 1994*, 104 J. GEOPHYSICAL RES., 16,807, 16,808 tbl.1 (1999) (noting that recent observations of radiative forcing by Mount Pinatubo, the most intensively studied volcanic

Inspecting Figure 1, we might be tempted to conclude naïvely that around 0.6°C of warming over the past fifty years is unambiguously attributable to human influence on climate. The detection and attribution community has traditionally been more cautious than this, however, not accepting universal agreement among models as sufficient evidence for an attribution claim.¹⁵ The first reason for this caution is that all of these models are driven by similar forcing datasets, and hence might share a common error in, for example, the amplitude of low-frequency solar variations.

The second reason for caution is that, at least until recently, the community has been reluctant to treat the range of responses from available models as spanning the range of responses that could be taking place in the real world. Hence, so this argument goes, if all models display a consistent warming of at least 0.5°C due to anthropogenic influence, this could result from a common error that all models share, and the real-world response (or the response of the next generation of models) could be significantly smaller than 0.5°C. There is considerable debate over the extent to which currently available models span the range of plausible real-world responses. Indeed, given current estimates of uncertainty in both climate forcing factors and response, it might well be argued that the near-universal agreement between models and observations in Figure 1 is evidence that the spread of model-simulated responses is too small to span the current range of uncertainty. To some extent, this is to be expected because each of the modeling groups contributing simulations to the comparison shown in Figure 1 was aiming to provide a “best estimate” of the response to “best estimate” forcing; no systematic effort was made to explore low and high values of either forcing or response.¹⁶

We have tended, as has most of the detection and attribution community, to take a relatively skeptical stance on the issue of whether current models span the range of plausible real-world responses. In a review of this nature, however, it is important to note

eruption as of 1999, yield an uncertainty of greater than 100%); Gabriele C. Hegerl et al., *Detection of Volcanic, Solar and Greenhouse Gas Signals in Paleo-Reconstructions of Northern Hemispheric Temperature*, GEOPHYSICAL RES. LETTERS, Mar. 12, 2003, at 46-1, 46-3 (estimating the uncertainty of their proxy reconstruction of volcanic forcing at approximately 50%).

¹⁵ See the discussion of detection and attribution methods in Mitchell et al., *supra* note 4, at 700-13.

¹⁶ See, e.g., Myles R. Allen et al., *Model Error in Weather and Climate Forecasting*, in PREDICTABILITY OF WEATHER AND CLIMATE 391, 403 fig.15.1 (Tim Palmer & Renate Hagedorn eds., 2006).

that there are other researchers who would regard the agreement between the models displayed in Figure 1 as itself sufficient evidence to attribute a substantial fraction of the warming that has occurred over the twentieth century to the anthropogenic increase in GHG levels. Fortunately for nonspecialists, formal detection and attribution studies that avoid assuming that the models give the correct magnitude of the response to anthropogenic or natural forcing appear to be converging on a very similar conclusion to that suggested by a naïve interpretation of Figure 1. This appears to be a case in which first impressions turn out, in fact, to be remarkably accurate.¹⁷

III. WHAT IS THE “NATURAL” CLIMATE?

Before we discuss how formal detection and attribution studies work, we can use Figure 1 to illustrate a point of principle that the legal community will need to resolve before making use of these studies' results. This is the question of what we mean by *ceteris paribus* in the context of climate change: what is the “natural” climate against which we are comparing present-day and future conditions to assess the extent of human influence? At a global level, the combined effect of solar and volcanic forcing appears to have been to induce a modest (0.1°C or so) cooling over the second half of the twentieth century, so some of the anthropogenic warming to date has been compensating for this natural cooling.¹⁸ Suppose that a plaintiff were to argue that she was being harmed by high temperatures.¹⁹ Is human influence “to blame” for the full 0.6°C of warming or only the 0.5°C that is not compensating for natural cooling? Given that we have no a priori reason to believe that natural forcing would induce cooling rather than warming over this period, one might even argue that human influence can only be blamed for the warming that could not have occurred naturally, giving an even smaller figure.

¹⁷ See Int'l Ad Hoc Detection & Attribution Group, *Detecting and Attributing External Influences on the Climate System: A Review of Recent Advances*, 18 J. CLIMATE 1291, 1298-99 (2005) (reviewing the literature and noting that the Allen study, *supra* note 5, which deliberately excluded global mean temperature information as a sensitivity study, nonetheless found a GHG signal).

¹⁸ See Theodore L. Anderson et al., *Climate Forcing by Aerosols: A Hazy Picture*, 300 SCIENCE 1103, 1103-04 (2003) (suggesting that the impact of aerosol cooling may largely compensate for present greenhouse warming, but concluding that greenhouse warming will exceed aerosol cooling effects in the twenty-first century).

¹⁹ This is a hypothetical example, because nobody is directly harmed by global mean temperature itself, but it serves to illustrate the point of principle.

The problem with this line of argument is that while we are relatively uncertain about the actual magnitude of past natural drivers of climate change, we are even more uncertain about the impact of these drivers, given our limited knowledge of the mechanisms governing low-frequency variability in either solar activity or explosive volcanism. Solar and volcanic activity appear to have caused a cooling over the second half of the twentieth century: how much warming might have occurred if the sun had behaved differently and volcanos had gone off at different times? This would be a very difficult question to answer, so it is to be hoped that the legal community will be satisfied with a much more restrictive interpretation of *ceteris paribus*, under which the impact of any given driver of climate change is measured by how the climate would be different if that driver were absent, assuming all other drivers of climate change evolved as observed.

This is the interpretation of attribution used by the scientific community: when we assess how much warming has been caused by GHGs over the past fifty years, we include in that figure the avoided cooling that would have been caused by other factors had it not been masked by this greenhouse warming. The key reason for this is that the question of how the climate would have been different if the GHG-induced warming had not occurred but all other drivers of climate change evolved as they did is scientifically well posed. The question of how the climate might have been had nothing changed since preindustrial times is not well posed, however, since we have no observations of such a hypothetical stationary climate and very little data on either the preindustrial climate itself or the factors driving it.

The issue of “what is a natural climate” has a direct bearing on the relevance to current attribution claims of paleoclimate reconstructions of the climate of the last millennium. If attribution is evaluated relative to a preindustrial climate, then the questions of what that preindustrial climate was and how much it varied naturally become crucial. The problem is that reconstructions of preindustrial climate and preindustrial solar and volcanic activity are necessarily based on indirect (“proxy”) observations that appear, at least at present, to be much more open to reinterpretation than the direct observations of climate over the past century.²⁰ Instrumental observations are also open to

²⁰ The controversy surrounding the accuracy of such reconstructions has become known as the “hockey stick” debate, reflecting the shape of the graph proposed in Mann et al., *supra* note 1, at 783 fig.5b. For a criticism of Mann’s work, see Stephen McIntyre & Ross McKittrick, *Hockey Sticks, Principal Components, and Spurious Significance*, GEOPHYSICAL RES. LETTERS, Feb. 12, 2005, at L03710, L03710-2 (finding that applying

challenge, as we have seen in the debate over the Microwave Sounding Unit record,²¹ but there is much more information available to resolve contentious issues.

Figure 2, from Hegerl et al.,²² illustrates how much disagreement exists about the earth's climate during preindustrial times. It shows a number of reconstructions of temperatures in the Northern Hemisphere over the past millennium relative to their mid-twentieth-century values. There are clearly pronounced differences between the reconstructions regarding how cold temperatures were in the so-called "Little Ice Age" around the seventeenth century. Likewise, there is even less agreement about whether temperatures were warmer than today during the "Medieval Warm Period" around the eleventh century. This level of disagreement is unsurprising given that all of these reconstructions are attempting to estimate a global or hemispheric temperature from a limited number of indirect observations, such as tree-ring widths, that provide, at best, an indication of local temperatures over a particular season.

The situation is even more problematic when we consider that harm, in almost all instances, will depend on centennial climatic trends as well as on absolute climate, and also on variables other than large-scale temperatures. If a plaintiff were to argue that she had been dispossessed because an increase in heat-wave risk had rendered her village uninhabitable, she would have to explain why she was there in the first place. After all, there are locations which have always been effectively uninhabitable because of heat waves. The issue is thus not the absolute climate, but the climate relative to that which a plaintiff might reasonably have expected to experience when she invested in the region. Estimating the range of past centennial temperature trends may prove even more difficult than estimating past temperatures themselves and very few studies have even begun to address the problem of reconstructing centennial fluctuations in other variables, such as precipitation.²³

Mann's methodology to red noise often produces similar, hockey-stick-shaped graphs).

²¹ For a comprehensive discussion of this debate and how it appears to have been resolved, see U.S. CLIMATE CHANGE SCI. PROGRAM, TEMPERATURE TRENDS IN THE LOWER ATMOSPHERE: STEPS FOR UNDERSTANDING AND RECONCILING DIFFERENCES 56-58 (2006), available at <http://www.climate-science.gov/Library/sap/sap1-1/finalreport/sap1-1-final-all.pdf>.

²² Gabriele C. Hegerl et al., *Climate Sensitivity Constrained by Temperature Reconstructions over the Past Seven Centuries*, 440 NATURE 1029, 1030 fig.2 (2006).

²³ For additional information, see the discussion of attribution results in non-temperature variables (or scarcity thereof) in chapter 9 of IPCC, WORKING GROUP I,

Hence, if the question of attribution of current temperature trends were dependent on our knowledge of climatic trends prior to the Medieval Warm Period, it is unlikely to be resolved in the foreseeable future. Fortunately, it does not, provided that the legal community accepts the restrictive definition of *ceteris paribus* used by the scientific detection and attribution community. As far as we can tell, most of these preindustrial fluctuations in climate were externally driven. Figure 2, for example, compares these paleoclimate reconstructions with simulated hemispheric temperatures from a simple climate model driven by estimated volcanic activity, solar variability, and anthropogenic influences over the past millennium. Cold periods in the seventeenth and early nineteenth centuries are reproduced by the model, in both cases largely driven by an increase in volcanic activity, while the Medieval Warm Period appears to have been caused by a temporary reduction in volcanic activity, possibly with some contribution from solar variability.

Given that these large-scale warming and cooling trends observed in Figure 2 appear to be externally driven, if the second half of the twentieth century had been characterized by a substantial increase in solar activity and a decrease in volcanic activity then we might have seen a natural warming trend comparable to that observed over the past fifty years. But this is irrelevant to the attribution of recent warming trends, because we can be reasonably confident that neither natural forcing occurred. The mid-twentieth century was a relatively quiescent period for volcanic activity followed by three major eruptions—Agung in 1963, El Chichón in 1982, and Pinatubo in 1991—giving a significant net volcanic cooling trend over the period from 1950 to 2000. On time scales longer than the eleven-year solar cycle, solar output has been relatively constant throughout this period, so the net effect of natural drivers has been to induce the cooling trend observed in the “natural” simulations in Figure 1.

Even if the ambiguities of paleoclimate reconstruction could be resolved, there is also the question of arbitrariness: should we define “preindustrial climate” as that of the past 700 years—in which case all experts agree that current global temperatures are probably unprecedented—the past 1000 years (in which case there would be more disagreement)—or the past 5000 years (which very likely contain periods warmer than the present day due to changes in the configuration of

the Earth's orbit). One might argue that the conditions in the seventeenth century are more relevant because predictable factors affecting climate, specifically orbital forcing—were then close to their present-day values. But this suggests that the real question in which we are interested is not, in fact, preindustrial climate at all, but the climate that would have prevailed today in the absence of human influence. If we require orbital forcing to take near-present-day values, then why not volcanic forcing and so on, which brings us back to the more restrictive interpretation of *ceteris paribus* used by the detection and attribution community.

We highlight this issue for two reasons. First, the debate over reconstructions of the climate of the past millennium is frequently cited as evidence that important issues regarding the attribution of human influence on climate remain unresolved. In fact, it has relatively little relevance to the question of causal attribution of changes over the past half century because much stronger evidence is available for this period in the form of instrumental observations and detailed model simulations. If future research were to conclude that the eleventh century was indeed warmer than the present day, the simplest explanation would be that the net impact of increased solar and reduced volcanic activity on eleventh century temperatures was larger than our current best estimates of these climate drivers indicate. Since these estimates are already highly uncertain, such a revision would be unsurprising and would have very little impact on our estimates of the impact of natural drivers of climate on recent decades, for which we have much more information. Millennial temperature reconstructions are useful as a means of checking model-simulated climate variability on centennial time scales, a test which current models typically pass, but their role is supplementary and is not the primary focus of attention.

The second reason we highlight this issue is that this is an interdisciplinary Article, and we recognize that the legal community might take a different view from our own on the appropriate reference point to use when defining climate change. The advantage of defining changes relative to the climate of the sixteenth through nineteenth centuries, for example, is that we would then be focusing on “reference conditions” that actually occurred, even though they can only be observed indirectly. The alternative that the scientific community uses, is to define our reference conditions as the climate that would have occurred in the early twenty-first century in the absence of specific human influences. However, these conditions can only be ex-

plored through computer simulation. Hence, this is the first question that we would like the legal community to resolve: what is the appropriate baseline against which to quantify human influence on climate? If it is acceptable to define this baseline or reference climate as that which would have occurred either today or in the future in the absence of human influence, the entire “hockey stick” debate becomes largely irrelevant; the fact that volcanos may have been anomalously quiet in the eleventh century, resulting in a temporary upturn in temperatures, plays no direct role in our understanding of the causes of climate change over the past half century. The *ceteris paribus* issue will come up again in the discussion of quantifying the role of human influence in cases of actual harm.

IV. QUANTIFYING HUMAN INFLUENCE

Allowing for uncertainty in model-simulated responses to both human and natural external drivers, how do formal detection and attribution studies assess how much human influence has contributed to the warming observed in Figure 1? The standard approach, first set out by Hasselmann,²⁴ and implemented by Santer et al.²⁵ and Hegerl et al.,²⁶ is generally known as “optimal fingerprinting.” It is essentially nothing more than a form of linear regression and assumes that models provide the correct *pattern* of response in both space and time to a given external driver, such as increasing GHGs or explosive volcanic activity, while allowing for the possibility that they may err on the *magnitude* of the response.²⁷

²⁴ See generally K. Hasselmann, *Multi-Pattern Fingerprint Method for Detection and Attribution of Climate Change*, 13 CLIMATE DYNAMICS 601, 610 (1997) (extending the “optimal fingerprint” method to multiple competing explanations of climate change); K. Hasselmann, *Optimal Fingerprints for the Detection of Time-Dependent Climate Change*, 6 J. CLIMATE 1957 (1993) (proposing the “optimal fingerprint” method).

²⁵ See Benjamin D. Santer et al., *Ocean Variability and Its Influence on the Detectability of Greenhouse Warming Signals*, 100 J. GEOPHYSICAL RES. (OCEANS) 10,693-10,725 (1995) (applying Hasselmann’s method to ocean model temperatures).

²⁶ See Gabriele C. Hegerl et al., *Detecting Greenhouse-Gas-Induced Climate Change with an Optimal Fingerprint Method*, 9 J. CLIMATE 2281 (1996) (applying Hasselmann’s method to surface temperatures).

²⁷ See Stephen S. Leroy, *Detecting Climate Signals: Some Bayesian Aspects*, 11 J. CLIMATE 640, 643 (1998) (demonstrating how the fingerprinting approach fits into conventional Bayesian estimation theory); M.R. Allen & S.F.B. Tett, *Checking for Model Consistency in Optimal Fingerprinting*, 15 CLIMATE DYNAMICS 419 (1999) (arguing that fingerprinting should only be used in conjunction with checks on model-simulated internal variability to avoid spurious detection claims).

There are sound physical reasons to believe that model-simulated response patterns are more likely to be correct than simulated response magnitudes. The timing of explosive volcanic eruptions and their relative magnitude in terms of potential climatic impact, for example, are relatively well known, while the absolute amount of reflective material injected into the stratosphere, how long it stayed there, and the net cooling effect on climate are much more ambiguous, even for recent volcanos.²⁸ Likewise, different reconstructions of past solar activity agree on the overall timing of increases and decreases in solar output while differing substantially on their magnitude.²⁹

The spatial pattern of response to global climatic drivers such as solar activity and increasingly well-mixed GHGs is primarily determined by the different heat capacities of land and ocean, and therefore causes land temperatures to respond to both warming and cooling much faster than sea surface temperatures. The positive feedback effect of melting snow and ice also causes high-latitude temperatures to change more than tropical temperatures. Again, these are well-understood processes that are replicated in all plausible climate models.

Responses can also be differentiated by their vertical structure: an increase in well-mixed GHGs tends to cause warming in the troposphere and cooling in the lower stratosphere. By contrast, an increase in solar activity causes warming over the full depth of the atmosphere away from the poles. Volcanos warm the stratosphere and cool the surface and troposphere. Again, all of these features of the responses to different drivers are well understood and generally accepted across the scientific community.

The spatial patterns of response to spatially heterogeneous drivers such as anthropogenic and volcanic aerosols are less well established, although it is generally accepted that anthropogenic aerosols have a much larger impact on temperatures in the Northern Hemisphere, where most aerosol sources are located, than in the Southern Hemisphere. The key uncertainty in the spatial pattern of response to volcanic eruptions is due to the dynamic response of the atmosphere,

²⁸ See, e.g., Alan Robock, *Volcanic Eruptions and Climate*, 38 REV. GEOPHYSICS 191, 214 (2000) (discussing the ability of current climate models to simulate the response to large volcanic eruptions).

²⁹ Hoyt & Schatten, *supra* note 11, at 18,897-18,901 figs.1-7; Lean et al., *supra* note 11, at 3196-97 figs.1-4.

which may not be accurately simulated by relatively low-resolution climate models.³⁰

The time history of anthropogenic aerosol influence on climate is much better understood than its spatial structure and, crucially, is distinct from the time history of the influence of anthropogenic GHGs, even though both GHG emissions and aerosol emissions are closely related to industrialization. The reason is that the dominant GHG, carbon dioxide, accumulates in the atmosphere, so its influence depends on cumulative emissions over the past century or so. In contrast, anthropogenic aerosols are “washed out” by precipitation on a time scale of only a few days, so their influence depends on current emissions rather than on cumulative emissions. Over the past few decades, improved power generation systems have also substantially reduced the aerosol burden that accompanies a given level of carbon dioxide emissions.³¹

The net effect of these factors, all of which are relatively noncontingent, is that aerosol cooling increased rapidly until the mid-1970s, more than offsetting greenhouse warming in the Northern Hemisphere and inducing a modest hemispheric cooling trend through the mid-twentieth century.³² There has not been much of an overall trend in aerosol-induced cooling since the 1970s, which has hence ceased to mask the greenhouse warming trend. In the Southern Hemisphere, the warming trend due to GHGs has been proceeding since the early twentieth century, with an apparent interruption in the 1940s that may be an artifact of the paucity of wartime observations. No other combination of natural or anthropogenic factors can explain the observed Northern Hemisphere cooling until the mid-1970s and subsequent rapid warming or the observed Southern Hemisphere warming since the 1940s.

The fact that the responses to these different factors are distinct allows us to estimate the magnitude of their contributions to recent temperature trends without relying on the magnitude of model-simulated responses. We do this by what is, in effect, a least-squares fit between observations and model simulations. If there were only a single factor that could explain the observed climate change, this

³⁰ Georgiy Stenchikov et al., *Arctic Oscillation Response to Volcanic Eruptions in the IPCC AR4 Climate Models*, J. GEOPHYSICAL RES. (ATMOSPHERES), No. D7, Apr. 11, 2006, at 15.

³¹ Meinrat O. Andreae et al., *Strong Present-Day Aerosol Cooling Implies a Hot Future*, 435 NATURE 1187, 1188-89 (2005).

³² See fig.3, *infra*.

analysis would consist simply of making a scatter plot of observations against model simulations and fitting a straight line through the points. In that case, the test of whether that external influence was detectable would simply depend on the uncertainty in the slope of the line: if the scatter were consistent (at some confidence level) with a line of zero slope, then we would conclude that external influence is not detectable (at that confidence level).

With multiple external drivers contributing to observed climate change, the situation is only slightly more complicated. Instead of fitting a straight line to a two-dimensional scatter plot, we fit a flat surface through a cloud of points, where each “horizontal” dimension of the cloud represents the expected response to a different external driver and the “vertical” dimension represents the observations. This is easy to visualize if there are only two external drivers considered, so the cloud is three-dimensional. Including more drivers makes the problem harder to visualize and tends to increase the uncertainties, but the principle is unchanged. This procedure is illustrated in Figure 3 using a subset of the models shown in Figure 1 for which simulations of the twentieth century are available. The effects of the various external drivers—specifically, well-mixed GHGs, other anthropogenic influences (predominantly sulfate aerosol cooling), and natural influences (solar and volcanic activity)—are displayed separately. The left panels show observed and model-simulated time series of Northern and Southern Hemisphere annual mean surface temperatures plotted against time, while the right panels show the same data plotted against each other.

The top panels (a) and (b) show the combined influence of all external drivers, as shown in Figure 1 (a). The fit between models and data is very good, as illustrated by the tight scatter plot in panel (b), resulting in only a small uncertainty in the best-fit line. Hence we are very confident that most of the observed large-scale changes in surface temperature over the twentieth century are externally driven. This was a surprising result for many in the climate research community when it was first pointed out by Peter A. Stott and his coauthors in 2000.³³ At the time of the first IPCC Scientific Assessment in 1990, for example, the prevailing view was that most observed large-scale temperature variability was generated internally by chaotic fluctuations in the atmosphere-ocean system, with external drivers making only a

³³ Peter A. Stott et al., *External Control of 20th Century Temperature by Natural and Anthropogenic Forcings*, 290 SCIENCE 2133, 2135-36 (2000).

small, barely detectable, contribution.³⁴ In fact, it emerges that large-scale temperature changes on time scales longer than a few years are largely controlled by external factors, both anthropogenic and natural. This has important implications for the issue of whether harm is foreseeable. The impact of these external drivers on large-scale climate appears to be much more predictable than the climate research community would have expected in the early 1990s.

Panels (c) and (d) show the influence of the increase in well-mixed GHGs, with the best fit contributions from the other two factors—anthropogenic aerosols and natural forcing—removed. In geometric terms, this is equivalent to viewing a four-dimensional scatter plot in which observed temperatures are plotted in the vertical and the simulated response to GHGs, anthropogenic aerosols and natural drivers are plotted on the other three “horizontal” dimensions from an angle at which the contribution of the two non-greenhouse factors becomes invisible. Panels (e) and (f) show the influence of anthropogenic aerosols, and (g) and (h) show the influence of natural factors and solar and volcanic activity. In each case, the best-fit line through the scatter plots is consistent with the diagonal corresponding to the model response that is equal to the observed response; and on the basis of near-surface temperature changes, there is no reason to suppose the models are not responding correctly to all three of these external drivers.

There are greater uncertainties about the contributions of individual drivers of climate change than about the total change shown in the top panel. This is to be expected, because the main uncertainty in our interpretation of recent climatic trends is not in the total change, but in how much greenhouse warming may have been masked by changes induced by other factors. Nevertheless, a formal uncertainty analysis based on any of these models taken alone or using all of these models combined together indicates—as one would intuitively expect simply from viewing the scatter plots in Figure 3—that we would be very unlikely to have obtained this level of agreement between models and data purely by chance. In the conventional language of detection and attribution, the hypotheses of no response to greenhouse, sulfate,

³⁴ Although the IPCC's 1990 Assessment reported that the rise in atmospheric carbon dioxide concentration “is beyond any doubt due to human activities,” they pointedly did not attribute any of the observed temperature rise to human influence at that time. See R.T. Watson et al., *Greenhouse Gases and Aerosols*, in CLIMATE CHANGE: THE IPCC SCIENTIFIC ASSESSMENT 1, 17 (J.T. Houghton et al. eds., 1990).

or natural forcing can all be rejected at a high confidence level—specifically, with a less than 5% chance that the rejection is in error.

Our aim here is to convey the basic principles of the climate change detection and attribution procedure in order to emphasize its fundamental simplicity. A four-dimensional scatter plot may seem rather difficult to visualize, but this is a standard, classical statistical procedure, and it is important to stress that there is no simpler way of addressing this problem without significantly biasing the outcome. This procedure, known as multiple regression, is precisely what is used in epidemiological studies in which the confounding influences of age, wealth, employment status, and so forth are taken into account when quantifying the impact of, for example, smoking on health.³⁵

The basic principle of plotting model simulations against observations and fitting a straight line or flat plane through the resulting scatter of points is very simple. The precise method used to estimate the slope of the line and uncertainties therein can become quite complicated, but key results do not depend on the methodological details. It is much more important to ensure that the comparison is set up in a sensible way that will not automatically introduce bias into the results. Highly misleading results can be obtained if inappropriate inputs are made to the comparison, as illustrated by Figure 4.

For example, ostensibly simpler procedures have been used to argue for a dominant role of solar forcing in recent climate change. If one takes some estimates of past solar output and scales them to match observed temperature changes, an impressive fit can be obtained, which appears to leave little room for a substantial human influence. The problem with this procedure is that results are heavily influenced by the arbitrary decision to consider solar influence before anthropogenic causes. There is no reason to expect that solar influence over climate is greater than human influence: if anything, our understanding of the physical magnitude of these two drivers points to the opposite expectation. If one were to begin by fitting observed temperature changes to estimated human influence, an equally impressive fit, which leaves apparently little room for a substantial solar influence, can be obtained.³⁶ The only nonarbitrary way of comparing

³⁵ See generally Sander Greenland & Kenneth J. Rothman, *Measures of Effect and Measures of Association*, in MODERN EPIDEMIOLOGY 47 (Kenneth J. Rothman & Sander Greenland eds., 1998) (explaining basic concepts necessary to estimate the effects of associations in epidemiological studies).

³⁶ For a clear and accessible discussion of this long-standing problem, see W.J. Ingram, *Detection and Attribution of Climate Change, and Understanding Solar Influence on*

the evidence is to consider all candidate explanations simultaneously, in a multiple regression.

A great advantage of framing the problem of climate change detection as a multiple regression is that we can examine the residuals of regression—the variations in the observations that are “left over” after we have extracted the estimated responses to known climate forcings—to assess whether there is evidence for any “missing forcing.”³⁷ This is particularly pertinent to the ongoing debates over whether land-use changes, particularly urbanization effects, may have contributed to the apparent warming trend in the observations, and over whether (still hypothetical) interactions between solar activity and global cloudiness could have contributed to the observed warming. In general, detection and attribution studies have found that regression residuals are consistent with what we would expect from internally generated climate variability, hence providing no evidence that either of these missing forcings is necessary to explain the observations. Some studies have found tentative evidence that models underestimate the response to solar forcing, but the discrepancy is too small to have a substantial impact on the estimated response to greenhouse forcing.

Detection and attribution studies cannot, of course, entirely exclude the possibility that an unknown missing forcing that happens to give a pattern and time history of response very similar to, for example, anthropogenic GHG increase might affect their results. It remains essential, therefore, to continue to explore any such possibility for which there is a sound physical basis. This certainly applies to land use changes, which for well-understood reasons are known to have an impact on uncorrected local temperature records.³⁸ Considerable effort has gone into minimizing the impact of urbanization on the observations shown in Figures 1 and 3 through quality control of the input data. If urbanization were contributing a substantial fraction of the observed warming trend, then various consequences would follow, such as differential trends under low and high wind speed conditions.

Climate, 125 SPACE SCI. REVS. 199, 206-09 (2006).

³⁷ See generally Allen & Tett, *supra* note 27 (observing that such a missing forcing should be detected as a deficiency in the detection model and that the absence of any evidence that the model is deficient therefore provides some (necessarily incomplete) reassurance that key forcings have been accounted for).

³⁸ See, e.g., Richard A. Betts, *Biogeophysical Impacts of Land Use on Present-Day Climate: Near-Surface Temperature Change and Radiative Forcing*, 1 ATMOSPHERIC SCI. LETTERS (2001) (explaining that land uses such as deforestation and vegetation affect climate).

These consequences are not observed, so it continues to be reasonable to assume that these data are approximately unbiased.³⁹ Likewise, studies of the impact of other potential drivers of global climate change, such as “black carbon” (soot) aerosols, have found their global impact to be generally much smaller than the impact of increasing GHGs.

There will always be proponents of a particular climate-forcing mechanism prepared to argue that attribution claims should be disregarded because their favored mechanism has not, in their view, been adequately accounted for. Without wishing to forestall research into these very important questions, it is important that the legal community reach agreement on what constitutes an adequate explanation of recent climate trends. Within the scientific community, an explanation that is both physically coherent and consistent with the available data—meaning the data that provide no indication that anything is missing from either forcings or response—is generally considered adequate. On these terms, the combination of greenhouse warming, anthropogenic aerosol cooling, and solar and volcanic activity provides an adequate explanation of large-scale temperature changes over the course of the twentieth century. Other factors have certainly contributed, but based on current evidence their impact appears to be relatively small.

It is interesting to note that critics of recent detection and attribution studies have focused exclusively on the inputs used—both observations and model simulations—rather than the attribution methods themselves; this is in stark contrast to the debate over reconstructions of the climate of the past millennium.⁴⁰ This probably reflects the underlying simplicity of the standard attribution methodology, but it is helpful to note that in this area, at least, the science seems to be relatively uncontroversial.

We refer readers to the technical literature, particularly the 2007 report of the IPCC and the references therein, for the details of such formal analyses.⁴¹ There are technical issues that need to be resolved that are not obvious from Figure 3: for example, how many effectively

³⁹ David E. Parker, *Large-Scale Warming Is Not Urban*, 432 NATURE 290, 290 (2004).

⁴⁰ See, e.g., Climate Science: Roger Pielke Sr. Research Group Weblog, Main Conclusions, <http://climatesci.colorado.edu/main-conclusions/> (last visited Apr. 5, 2007) (arguing that recent climate change reports “have overstated the role of the radiative effect of anthropogenic increase of CO₂”).

⁴¹ See Hegerl et al., *supra* note 23. Note that the report’s “Summary for Policymakers” is available at <http://www.ipcc.ch/SPM2feb07.pdf>.

independent data points do we have and how large a discrepancy should we expect between models and data due to internal climate variability and uncertainty in the model-simulated responses to different drivers of climate change? These important issues are the subject of ongoing research,⁴² but the strength of the observed signal has enabled recent attribution studies to produce robust conclusions. In summary, most of the global warming over the past fifty years is very likely to have been due to the observed (anthropogenic) increase in GHG levels.⁴³ The warming from the 1920s to the 1940s is most likely explained by a combination of increasing solar activity and a temporary reduction in volcanic activity, possibly augmented by increasing GHGs.⁴⁴ The lack of warming from the 1940s to the 1970s is likely to have been caused by compensation between warming induced by increasing GHGs and cooling due to anthropogenic aerosols.⁴⁵ Generally, large-scale changes over the twentieth century can be explained by current models driven by current best estimates of natural and anthropogenic drivers of climate change. No substantial “missing” forcing, such as land-use changes or exotic sun-climate interactions, is needed to explain the observed record. Quantitative comparison of current models with the observed record provides no evidence that any models are systematically overestimating the response to increasing GHGs: in fact, one or two models may be underestimating it.⁴⁶

We stress that we are simply providing a summary of recent conclusions and expressing these conclusions as simply as we can without attempting to use the more constrained language of governmental and intergovernmental assessments. The key point is that we are considerably more confident about the causes of climate change over the past fifty years than we are about the causes of the early-twentieth-century warming, or of any change prior to the twentieth century. This is partly because of the paucity of Southern Hemisphere observations in this earlier period and partly because of the irresolvable un-

⁴² Examples of further study of these issues include Myles R. Allen et al., *Quantifying the Uncertainty in Forecasts of Anthropogenic Climate Change*, 407 NATURE 617 (2000); Gabriele C. Hegerl et al., *Effect of Observational Sampling Error on the Detection of Anthropogenic Climate Change*, 14 J. CLIMATE 198 (2001); and Gabriele C. Hegerl & Myles R. Allen, *Origins of Model-Data Discrepancies in Optimal Fingerprinting*, 15 J. CLIMATE 1348 (2002).

⁴³ See generally PHYSICAL SCIENCE BASIS, *supra* note 23, and references therein.

⁴⁴ *Id.* at 10-11.

⁴⁵ *Id.*

⁴⁶ See Stott et al., *supra* note 7, at 3067 fig.9; PHYSICAL SCIENCE BASIS, *supra* note 23, § 9.4.1.4 fig.9.9.

certainty in the magnitude of external drivers of climate change prior to the advent of direct observations.

Hence, the question of whether most of the warming observed over the full twentieth century or since preindustrial times is attributable to human influence is still open to dispute because of the uncertainty in the magnitude of natural influence before 1950. Nevertheless, this does not detract from the fact that we do understand what has been causing changes over more recent decades. This knowledge about recent climate change is most relevant to quantifying the impact of GHG emissions to date and to predicting future climate change.

Failure to recognize these (admittedly subtle) distinctions between attribution of warming over the past fifty, one hundred, or one thousand years can lead to considerable energy being wasted on issues, such as the temperature of the Medieval Warm Period, that are largely irrelevant to what is happening now and what will happen in the next couple of decades. For example, in the amicus curiae brief filed by climate scientists in *Massachusetts v. EPA*, David Battisti and his coauthors posited that “[h]uman activities likely caused most of the approximately 0.6°C (1.1°F) rise over the 20th century.”⁴⁷ Although this statement is supported by the majority of detection and attribution studies analyzing the causes of climate change over the twentieth century, it refers to the full twentieth century rather than the past fifty years, and it goes beyond the conclusions of the 2001 IPCC Reports⁴⁸ and the U.S. NAS/NRC Report,⁴⁹ which stated that “the changes observed over the last several decades are likely mostly due to human activities.” In a counter-brief filed in support of the respondents in the same case, Sallie Baliunas and other climate scientists argued that Battisti’s statement is “scientifically wrong” because the origins of the early-century warming are still not fully understood and are of comparable magnitude to the more recent warming.⁵⁰

⁴⁷ Brief for David Battisti et al. as Amici Curiae Supporting Petitioners at 12, *Massachusetts v. EPA*, 126 S. Ct. 2960 (2007) (No. 05-1120), 2006 WL 2563377.

⁴⁸ See generally PHYSICAL SCIENCE BASIS, *supra* note 2, at 695-738 (discussing in depth the attribution of causes of climate change).

⁴⁹ COMM. ON THE SCI. OF CLIMATE CHANGE, NAT’L RESEARCH COUNCIL, CLIMATE CHANGE SCIENCE: AN ANALYSIS OF SOME KEY QUESTIONS 1 (2001).

⁵⁰ Brief for Sallie Baliunas et al. as Amici Curiae Supporting Respondents at 14, *Massachusetts v. EPA*, 126 S. Ct. 2960 (2007) (No. 05-1120), available at http://www.aei.org/research/nri/publications/pubID.25076,projectID.22/pub_detail.asp.

Since current models can, in fact, explain the early-century warming, whether it was greater or less than the late-century warming is entirely irrelevant to the question of the human contribution to the more recent warming. For instance, if the Krakatoa volcano had happened to erupt a couple of decades later, further depressing temperatures prior to the 1920-1940 period, we might have seen an even greater warming in that period than we have seen since the 1970s, but this would not alter either our understanding of what is taking place today or our predictions for the coming decades. But the exchange between the climate scientists in *Massachusetts v. EPA* illustrates the importance of meticulous phrasing of attribution statements and highlights the importance of focusing on the past half century and not being sidetracked into what might or might not have happened earlier. Notably, although Battisti and his coauthors were eminent climate scientists, they did not include any of the specialists in quantitative detection and attribution that advised the NAS/NRC and IPCC panels.⁵¹ This, perhaps, illustrates the increasing specialization of climate science: not all climate scientists, however eminent, will necessarily be attuned to the niceties of the attribution question.

The only question to which the magnitude of the early-century and pre-twentieth-century warming periods might be considered relevant is the issue of whether human societies and natural ecosystems can or should be expected to cope with climate changes of this magnitude. This is a more subtle question beyond the scope of the attribution issue that we address in this Article. We simply draw attention to two points. First, the fact that climatic changes have occurred naturally does not necessarily mean that they were benign. These past natural climate fluctuations may well have caused considerable loss and suffering that could be very difficult to quantify with the historical records available today. Second, many prospective plaintiffs will be more concerned with the climate they are likely to have to contend with over the coming decades rather than with the changes that have already occurred. The key reason the warming trend since the 1970s is qualitatively different from, and more damaging than, previous natural warming periods is that the trend is expected to be the begin-

⁵¹ Brief for David Battisti et al., *supra* note 47, at 1. We understand that the intention of the drafters of the brief was simply to reproduce the IPCC and NAS/NRC statements, and that this subtle inconsistency only arose through the editing process, so it may seem rather uncharitable to belabor the point, but it does serve to illustrate the increasingly technical nature of attribution statements.

ning of a sustained change rather than a temporary upturn in global temperatures.

V. IMPLICATIONS OF PAST ATTRIBUTABLE WARMING

The relatively tight scatter of the points in panels (c) and (d) of Figure 3 indicates that the response to past GHG emissions is known to a relatively high degree of certainty. The models agree with the observations in that the points fall along the diagonal; in contrast, a model that predicted substantially more or less greenhouse warming than displayed in Figure 3 would agree significantly less well. Thus, we can be reasonably confident that current models neither overpredict nor underpredict the response to rising GHGs by a substantial margin.

What are the implications? The warming attributable to GHGs over the past fifty years is closely related to the warming predicted, assuming (as seems overwhelmingly likely) that the current increase in GHG levels continues. The crucial property of the climate system that relates these two periods is the “Transient Climate Response” (TCR), which is defined as the warming we should expect at the time of carbon dioxide doubling—around year seventy—if carbon dioxide levels were to increase at 1% per year, starting with a climate in equilibrium. TCR is a theoretical measure of climate system properties that is much more relevant to forecasts of change over the coming decades than the much more widely known “Equilibrium Climate Sensitivity” (ECS) because TCR determines current trends, rather than the long-term equilibrium response to a hypothetical stabilization of atmospheric GHG levels that may take centuries to materialize. This is essentially because TCR focuses on the dynamical response, and hence takes account of the thermal inertia of the system, while the equilibrium sensitivity—since it is a long-run, static parameter—does not. For this very reason, the equilibrium sensitivity is not well constrained by recent data.⁵²

The reason TCR is closely related to the warming attributable to carbon dioxide increase to date is that recent carbon dioxide changes are relatively close to an idealized scenario in which carbon dioxide levels began to increase from their preindustrial values at 0.45% per

⁵² See D.J. Frame et al., *Constraining Climate Forecasts: The Role of Prior Assumptions*, 32 GEOPHYSICAL RES. LETTERS, May 6, 2005, at 1 (discussing how “universal consensus on long-term equilibrium warming consistent with any given stabilisation level for greenhouse gases may prove impossible to achieve”).

year beginning around 1935.⁵³ The response to such a compound increase in carbon dioxide scales almost exactly with the rate of increase, so after seventy years of an approximately 0.45% per year increase, we should expect the warming to date attributable to past carbon dioxide increase to be approaching half the theoretical TCR.⁵⁴ When earlier carbon dioxide increases are also taken into account, the total carbon-dioxide-induced warming to date is slightly greater than this, or slightly over the warming that should be expected after about seventy years of a 0.5% per year increasing carbon dioxide scenario.

The “warming to date” referred to here is the total warming attributable to past carbon dioxide emissions, meaning the difference between today’s expected global temperature and the global temperature we would expect, *ceteris paribus*, had this increase not occurred. Other measures of attributable warming are also possible, such as the warming since 1900 or the linear warming trend over the twentieth century.⁵⁵ All of these are closely related to the total attributable warming, so this argument applies to them all, although the conversion factors will be different.

What actually is estimated in detection and attribution studies is the warming due to the observed increase in well-mixed GHGs, including the impact of methane, nitrous oxide, and CFCs. In terms of the carbon dioxide concentration that would have the equivalent impact on the global energy budget, this is closer to a 0.6% per year increase beginning in 1940, so—again, allowing for changes prior to 1940—we can expect the warming to date attributable to past well-mixed GHGs to be around 70% of the theoretical TCR.

Hence, because the climate system has already been subjected to a ramp-up in carbon dioxide levels very close to that which determines TCR, TCR is no longer a purely theoretical, model-based quantity, but a property of the world that we can, in principle, observe. Estimating climate system properties from observations must, however, be done with care. In particular, it is necessary to allow for the impact of multiple natural and anthropogenic influences on the climate system.

⁵³ See *infra* fig.4(a).

⁵⁴ See *infra* fig.4(b).

⁵⁵ Linear warming trends over the twentieth century have been used as a standard measure of attributable change since the publication of Simon F.B. Tett’s article in 1999. See Tett et al., *supra* note 5. The problem with this measure is that the changes we are observing accelerate over the century, so fitting a straight line to them tends to understate the magnitude of the change that has actually occurred.

A graphic example of the pitfalls of an oversimplistic comparison of models with data is given by Figure 5.⁵⁶ The thin solid lines show global temperatures in a number of climate models driven by a 1% per year increase in carbon dioxide; the dots show observed temperatures over the past thirty years, and the dashed line shows an extrapolation of a best-fit line through these observed temperatures. The scientists who created the graph argue that the observed rate of increase in carbon dioxide over the thirty-year period was slightly under 0.5% per year and that extrapolation of a 0.5% per year rate of increase is more credible than a 1% per year scenario because “the change in atmospheric concentration shows no evidence for an imminent doubling from its rate for the last three decades.”⁵⁷ Therefore, they argued, their extrapolated trend of 1.8°C per century represents a more credible forecast of future warming than these model simulations.

If it were true that a 0.5% per year increase in carbon dioxide gives a 1.8°C per century rate of global warming, this would imply that all but one of the models shown in Figure 5 are *underestimating* the TCR of the real climate system. If the real world were to warm by 1.3°C over eighty years in response to a 0.5% per year carbon dioxide increase, as the dashed line in Figure 5 indicates, then we would expect it to warm by double that amount in response to a 1% per year increase. This would be higher than all but one of the model predictions of the response to this scenario, and that model is no longer in use. Indeed, if this comparison between model simulations and observations were taken at face value, the fact that the models’ response to a 1% carbon dioxide increase appears to match the real world’s response to a 0.5% increase over the past thirty years suggests that the models are underestimating the true TCR of the climate system by a factor of approximately two.⁵⁸

In fact, such an alarming conclusion is completely unjustified, as is the diametrically opposite conclusion that Baliunas and her coauthors drew from this figure, which is that current models are currently overpredicting the real-world response.⁵⁹ The reason this conclusion is

⁵⁶ Figure 5 is reproduced exactly from the Baliunas brief. See Brief for Baliunas et al., *supra* note 50, at 13 fig.2.

⁵⁷ *Id.* at 12.

⁵⁸ We realize that this point has previously been made by Halldor Björnsson. See Posting of Halldor Björnsson to RealClimate, <http://www.realclimate.org/index.php/archives/2006/12/further-comment-on-the-supreme-court-briefs/> (Dec. 7, 2006 14:22 EST).

⁵⁹ Brief for Baliunas et al., *supra* note 50, at 13.

unjustified is that this is not a like-for-like comparison between the models and observations. The models in Figure 5 were started from equilibrium conditions at the beginning of the experiment, reducing their initial warming trends, whereas the real world was already responding to past GHG emissions by the 1970s. The models were only driven by an increase in carbon dioxide, whereas the real world was responding to the combined influence of increasing carbon dioxide, other GHGs, anthropogenic aerosols, solar variability, and volcanic activity. Studies that do take into account the influence of these other drivers, as well as the time scales of the response and uncertainties therein, tend to conclude that most current models simulate approximately the correct TCR. These studies typically allow for a TCR as high as that indicated by the naïve extrapolation of recent trends shown in Figure 5, but typically assign it less than a 10% chance of occurrence.⁶⁰

Systematic comparisons between model simulations and observed climate change can be used to constrain climate forecasts, but it is essential to base these on a like-for-like comparison in which models are driven by all the principal causes of climate change that are likely to have affected the real world. Such studies must allow for uncertainty in the influence of these multiple drivers on past climate trends through the use of some kind of multiple regression procedure to avoid arbitrarily favoring one driver over another. They must also allow for the fact that carbon dioxide is not the only driver of climate change affected by human decisions and that the balance between these drivers is likely to change in the future.

A 1% per year compound increase in carbon dioxide over the twenty-first century would give a level of 1000 parts per million (ppm) by the year 2100, corresponding to the upper end of the range of increases considered by the IPCC in 2001. The frequently made assertion⁶¹ that climate scenarios typically assume a 1% per year increase is not, in fact, correct. The standard “medium” scenario used by the modeling groups contributing to IPCC’s 2007 report, denoted “A1B,” leads to the equivalent, when other GHGs are taken into account, of a carbon dioxide level of 850 ppm by the year 2100.⁶² Much lower scenarios have also been considered. The 0.5% per year increase sug-

⁶⁰ See, e.g., Stott et al., *supra* note 7, at 3055-69.

⁶¹ Brief for Sallie Baliunas et al., *supra* note 50, at 11.

⁶² For descriptions of the IPCC’s special emissions scenarios, see PHYSICAL SCIENCE BASIS, *supra* note 23, at 1, 12.

gested by Baliunas and her coauthors would lead to slightly over 600 ppm by the year 2100. This is slightly lower than the IPCC's "A1T scenario," which is relatively optimistic about future technological development encouraging a rapid shift away from fossil fuels, or its "B2 scenario," which is relatively pessimistic about future economic development and global economic convergence. The IPCC projects a warming over the twenty-first century of approximately 2°C-3.5°C under these scenarios.⁶³ There are a number of reasons that these figures are higher than the 1.8°C suggested by Figure 5. Most importantly, reduced use of fossil fuels is accompanied by a substantial reduction in the sulfur emissions that cause aerosol-induced cooling in the A1T scenario, while continued reliance on low-technology food production methods causes a substantial increase in methane emissions in the B2 scenario.⁶⁴

The text of the Baliunas brief implies a much higher level of agreement with the mainstream climate modeling community than Figure 5 suggests: in particular, they endorse the prediction of a 0.5°C-1.0°C warming over the coming fifty years in response to a sustained equivalent of a 0.5% per year increase in carbon dioxide.⁶⁵ This corresponds to a range for TCR of 1.5°C-3.0°C, which approximately spans the 1.3°C-2.6°C range of TCRs of current models used by the IPCC.⁶⁶

There is a surprisingly strong consensus regarding TCR,⁶⁷ which includes authors of both the Battisti and the Baliunas briefs and is potentially very useful for a court in establishing potential liability. This is because the warming attributable to either carbon dioxide or total GHG emissions to date is closely related to TCR, as both represent responses to a driver that is relatively close to the idealized scenario upon which TCR is defined. Carbon-dioxide-induced warming to date is equal to approximately half the TCR (hence 0.75°C-1.5°C) while total warming due to well-mixed GHGs is 40%-50% higher. Any attempt to quantify the damage done by past carbon dioxide emis-

⁶³ *Id.* at 13.

⁶⁴ *Id.* at 18.

⁶⁵ This prediction was first put forth by James Hansen in 2001. James E. Hansen & Makiko Sato, *Trends of Measured Climate Forcing Agents*, 98 PROC. NAT'L ACAD. OF SCI. OF THE U.S. 14778, 14,778-83 (2001). Hansen and his coauthors suggest that this prediction could be achieved by a relatively rapid reduction in methane emissions.

⁶⁶ See PHYSICAL SCIENCE BASIS, *supra* note 23, at 723.

⁶⁷ D.J. Frame et al., *Alternatives to Stabilization Scenarios*, 33 GEOPHYSICAL RES. LETTERS, July 26, 2006, at L14707-1.

sions, either now or over the next couple of decades, will depend, *inter alia*, on warming to date attributable to carbon dioxide. The fact that we have such a high level of agreement, at least on the lower bound, means that the debate on the reality of a carbon-dioxide-induced warming, comparable in magnitude to the total warming over the past fifty years, is essentially over.

VI. PROBABILISTIC EVENT ATTRIBUTION: ADDRESSING CAUSE AND EFFECT IN A CHAOTIC SYSTEM

The previous section argued that there is not only strong empirical evidence, but also a high level of consensus that most if not all of the warming observed over the past fifty years is attributable to the anthropogenic increase in GHG levels. Indeed, quantitative estimates suggest that the total greenhouse-induced warming may be significantly greater than the net observed warming, with a substantial fraction being offset by cooling due to other factors, notably anthropogenic aerosols and volcanic activity. The strength of the evidence for human influence on global climate is not, however, likely to be enough for a plaintiff to argue that she has suffered, or is likely to suffer, actual harm due to GHG emissions.

For the purpose of this Article, we will restrict our attention to actual harms due to past or present injury or loss and immediate potential harms, by which we mean harms which may manifest themselves within the next decade or two. The reason for this focus is that quantifying this class of harm depends on understanding present-day climate and current climate trends. Crucially, it does not depend on future carbon dioxide emission rates, which depend on decisions outside the control of any conceivable class of defendants, making them highly contentious. The reason is simple: the climate responds to the current level of carbon dioxide, not its rate of increase, so any foreseeable change in the carbon dioxide emission rate will take at least a couple of decades to have a significant impact on climate. We can safely consider the climate of the current quarter century to be a result of natural factors (some of which may happen in the future, such as a volcanic eruption in the next few years) and human decisions that have already been made. Although we will continue to experience climate change over this period, decisions made now or in the future are unlikely to have much impact on climate prior to 2025.

The importance of this restriction is again illustrated by the amici curiae exchange cited above.⁶⁸ Baliunas et al. argue that the warming expected by the year 2100 under their proposed 0.5%-per-year scenario for the increase of carbon dioxide (which, starting from today's levels, would lead to levels reaching approximately 600 ppm by 2100) would be lower than the warming expected under, for example, the IPCC A1B "medium" scenario in which GHG levels reach 850 ppm carbon dioxide by 2100.⁶⁹ This may be so, but it is entirely irrelevant to the climate of the next couple of decades because a small acceleration in the rate of carbon dioxide increase, as predicted by the IPCC in response to rising emissions, would take at least twenty years to have any discernible impact on climate.⁷⁰ Hence, if a court were to decide that it is interested solely in impacts over the current quarter century, it would not need to concern itself with the details of future emissions scenarios and their associated assumptions of global population and economic growth, technological development, and so forth. This would avoid a vast and highly contentious area of debate.

This restriction is not intended to downplay the importance of more distant impacts of climate change nor to presume that these are not the proper concern of the courts. Our point is simply that the science required to support or contest claims about the likelihood, for example, of a massive sea level rise due to the collapse of the Greenland Ice Sheet in the twenty-second century is very different from the science required to evaluate claims about much more mundane issues like increased drought, storm, or flood risk today.

The key difference between long-term, catastrophic impacts of GHG increases and more mundane short-term impacts is that we might, in some instances, be able to say with confidence that some of these long-term impacts would not have occurred in the absence of human influence on climate. Our descendants might, in the twenty-second century, be in a position to say that the Greenland Ice Sheet would have been preserved if we had not interfered with the climate, and could straightforwardly attribute a catastrophic decline in ice volume to elevated GHG levels. Our current understanding of present

⁶⁸ See *supra* notes 47 & 50.

⁶⁹ See Brief for Baliunas et al., *supra* note 50, at 11, 12 n.10 (arguing that computer models using a 1% per year carbon dioxide increase are likely to overstate global warming by a factor of two); PHYSICAL SCIENCE BASIS, *supra* note 23, at 12 (outlining the IPCC scenarios).

⁷⁰ See, e.g., Peter A. Stott & J.A. Kettleborough, *Origins and Estimates of Uncertainty in Predictions of Twenty-First Century Temperature Rise*, 416 NATURE 723, 725 (2002).

climate trends, however, suggests that we are unlikely to see specific events that simply could not have occurred in the absence of human influence on climate within the next couple of decades. Local temperature records continue to be broken every year, but this simply means that the temperature in question has reached a level not seen for a hundred years (or whenever that record began). It does not mean that human influence is necessarily to blame, since natural climate fluctuations can also cause weather records to be broken. Most of the damage attributable to climate change on this time scale is likely to be due to events, such as floods and heat waves, which might have occurred naturally, but which have been made more likely by human influence on climate. Thus, although focusing on this shorter time scale simplifies our task, in that we do not have to consider future emission scenarios, it does mean we have to take a more sophisticated approach to attribution than the simple but-for test that might be applicable to the Greenland Ice Sheet problem.

It is still widely assumed, even in scientific circles, that if an event might have occurred naturally, it cannot be attributed to climate change, and hence that the most that can be said is something along the lines of “this is the kind of event that we might expect to become more frequent under climate change.” We believe a quantitative approach to causal attribution is possible, but it needs to be framed in a probabilistic framework. Some of the authors of this Article have proposed an approach to this problem that we will refer to as “probabilistic event attribution,” which focuses on attributing changes in the risk of an event occurring to external drivers of climate rather than attempting to dissect the event itself.⁷¹ If we are talking about a global temperature trend, it makes sense to say that $x\%$ of the trend is attributable to carbon dioxide increase, $y\%$ to other GHGs, and so on. If a flash flood occurs due to heavy rain, it makes no sense to say that this many millimeters of rain were due to factor A while this many millimeters were due to factor B. In a chaotic system, events are self-reinforcing, so the various external drivers of weather do not simply add different amounts of rain. We have to consider the weather system that occurred as a single, indivisible event and ask how various causal factors may have contributed to the risk of such an event occurring.

⁷¹ Myles Allen, *Liability for Climate Change*, 421 NATURE 891, 891-92 (2003); Dáithí A. Stone & Myles R. Allen, *The End-to-End Attribution Problem: From Emissions to Impacts*, 71 CLIMATIC CHANGE 303, 303-04 (2005).

The simplest analogy is with a loaded die. If a die is loaded to double the odds on a six, and it comes up six, then there is a clear sense in which some of the risk of that six occurring can be attributed to the loading. It makes no sense to say that three of the six dots on the face of the die are due to the loading, and three are due to chance, but this is what we would be doing if we were to try to dissect a weather event into a component due to human influence and a component due to natural variability.⁷²

The die example is simplified by the fact that rolling a die is a discrete event and, provided it is rolled properly, nothing apart from the loading has any detectable impact on the odds of the die coming up six. In the case of weather, we have to deal with ambiguity in when exactly a weather event begins and ends, and the fact that many factors will have contributed to the sequence of events that led up to it. This complicates the definition of what we mean by “the risk of an event”: do we mean the probability of that event occurring given the information available the day before it occurred, or the month before, or the year before? This issue matters considerably for quantifying external contributions to risk: the odds of a storm tomorrow are almost completely determined by the state of the weather today, with almost no role for external influences like elevated GHG levels.

If a die rolls forever, can we coherently talk about the influence of loading on the odds of a particular outcome? Fortunately, because of the highly chaotic nature of weather, the die does not roll forever, in that the impact of these precursor events becomes completely unpredictable after a relatively short time. Most mid-latitude variability displays predictability time scales of only a few days. The most predictable phenomenon to have a substantial impact on global weather is probably El Niño, which is only predictable a year or so in advance. The quest for longer-term predictions continues, but so far, despite decades of searching, predictable phenomena on time scales beyond a few months account for only a small fraction of the changes in the weather variables that tend to affect people’s lives. (Most predictability resides, for example, in the deep oceans.) This is particularly true once variations due to external factors are taken into account. Hence, we could argue that the roll of the weather die that resulted in a particular flood began at most a year or two before the flood occurred:

⁷² We do not wish to belabor this point more than is necessary, but this seems to be a common source of confusion among nonspecialists.

the state of the atmosphere and ocean prior to that time will have had only a negligible impact on the odds of the flood occurring.

In the context of probabilistic event attribution, we define the risk of an event as the probability of that event occurring in an atmosphere-ocean system identical to our own, driven by the same external drivers (natural and anthropogenic), and subject to a small perturbation far enough in advance of the event that further advancing it has no impact on the odds of the event occurring. With a rigorous definition of the risk of an event, we can then ask counterfactual questions about how that risk might change if a particular driver of climate change, such as past GHG emissions, had not occurred.

Crucially, such questions can only be addressed through computer simulation. Every sequence of meteorological events is effectively unprecedented, and the events that cause damage are, almost by definition, exceptional, so there is little scope for examining the historical record to determine empirically how external factors affect their risk of occurrence. The point is frequently made in the debate over human influence on extreme weather risk that the period of human influence remains too short to detect any significant trend in event frequency. This may well be true for many classes of extreme weather events, but it misses the point that risks may nevertheless be changing, even though we cannot rely merely on observation to quantify the magnitude of that change. Through computer simulation, we combine our observations with physical understanding of what is taking place.

If we had a perfect simulation model of the climate system and complete knowledge of how external drivers have varied since preindustrial times, then the problem of assessing the contribution of human influence on climate to the probability of occurrence of a particular weather event would be straightforward. We would simply perform two “ensemble” experiments, each comprising multiple runs of our simulation model in which initial conditions are disturbed to yield an independent trajectory through the chaotic atmosphere-ocean system. The requirement of independence means, in effect, that simulations are started far enough in advance of the season in which the event occurred for their starting conditions to have only a negligible impact on the odds of the event occurring in any member of the ensemble. There can be no hard and fast rule about how far in advance this needs to be (because it will be affected by the nature of the event under investigation), but it is straightforward to test whether

the simulations are long enough simply by varying the start date and examining whether it has any significant impact on results.

One ensemble would include, and the other would exclude, a specific human influence, such as elevated GHG levels, allowing us to compare the probability of, say, summer temperatures in 2003 exceeding a particular threshold between the two ensembles. Suppose P_1 is the probability of the threshold being exceeded in our simulated present day climate, and P_0 is the corresponding probability in a climate with human influence removed. Then the so-called Fraction Attributable Risk ($FAR = 1 - (P_0/P_1)$) is a measure of the proportion of current risk that can be attributed to human influence.⁷³

Even before we introduce scientific uncertainty, a decision needs to be made about how to treat other drivers of climate change. Under the interpretation of *ceteris paribus* suggested above, the relevant comparison is between present-day risk and risk in a climate in which external factors (like volcanic eruptions), with the sole exception of human influence, evolve exactly as they have up until now. This seems to be the most natural line to take, and also that which is most consistent with epidemiological practice. If a smoker exercises regularly, then the contribution of smoking to her risk of high blood pressure is generally assessed relative to non-smokers who take a similar level of exercise, rather than relative to the population as a whole. Hence, in our view, the fact that a volcano might have gone off in 2002, significantly reducing the risk of the 2003 European heat wave, is immaterial: no such volcano erupted.

No climate model is perfect now, nor will one ever be, so it is essential to have a methodology for attributing risk that is applicable to imperfect simulation models. In principle, the approach we can take is very simple. Instead of a single pair of ensembles, with and without human influence on climate, we can generate many pairs of ensembles from a range of possible models, with each model weighted by its relative likelihood of providing an accurate representation of the variable in question.

Let us consider a specific example. The summer of 2003 brought sustained, exceptionally high temperatures to large areas of Europe, with some regions of France experiencing average temperatures over 10°C higher than the average for 2000, 2001, 2002, and 2004, over the first twenty days of August. This heat wave, and the accompanying drought, are estimated to have caused over \$10 billion in continent-

⁷³ Stone & Allen, *supra* note 71, at 305.

wide agricultural losses, almost entirely uninsured, and over \$1.6 billion in losses from forest fires, primarily in Portugal.⁷⁴ Even more seriously, the heat wave has been blamed for up to 35,000 excess deaths across the continent, most of which occurred during the most intense period in early August.⁷⁵

In a 2004 article,⁷⁶ we explored the extent to which this heat wave could be attributed to human influence on climate using this probabilistic event attribution framework. We focused on the statistics of three-month-averaged (June, July, and August) temperature anomalies over a large area of southern Europe that had previously been identified as a climatically coherent region.⁷⁷ Standard definitions of “summer” (June through August standard conditions) and “Southern Europe”⁷⁸ were used, even though the heat wave itself was most intense in much smaller regions and over a shorter time (the first two weeks of August). This was done because, if the definition of the threshold used to define a particular extreme event is tailored too closely to the details of the event that actually occurred, we will inevitably find that every event appears to be highly improbable. Given sufficient detail, every meteorological situation is unprecedented. For the same reason, we defined a heat wave event as an area-averaged and seasonally-averaged temperature anomaly exceeding a threshold of 1.6°C above the average for 1961-1990, even though the actual temperature anomaly in 2003 was 2.3°C, since 1.6°C would exceed the temperature of the previous warmest summer on record—2001. Avoiding this so-called “selection bias” requires that the definition of the threshold should be independent of what actually occurred, or there should be a natural selection threshold before the event takes place. While it is plausible that we might have posed the question, “What are the chances of a record-breaking heat wave this summer?” early in 2003, it is less plausible that we would have specifically asked, “What are the chances of temperature anomalies exceeding 2.2°C?”

⁷⁴ Emily Black et al., *Factors Contributing to the Summer 2003 European Heatwave*, 59 WEATHER 217, 217-23 (2004); Christoph Schär & Gerd Jendritzky, *Hot News from Summer 2003*, 432 NATURE 559, 559 box 1 (2004).

⁷⁵ Schär and Jendritzky, *supra* note 74, at 559 box 1.

⁷⁶ Peter A. Stott, D.A. Stone & M.R. Allen, *Human Contribution to the European Heatwave of 2003*, 432 NATURE 610 (2004).

⁷⁷ *Id.* at 610.

⁷⁸ The region used was one of the “climatically coherent” regions technically defined as “Mediterranean” by F. Giorgi & L.O. Mearns, *Probability of Regional Climate Change Based on the Reliability Ensemble Averaging (REA) Method*, 30 GEOPHYSICAL RES. LETTERS, June 24, 2003, at 31-3.

It would be of considerable interest for future attribution studies if an effort were made now to catalogue and define damaging weather events whose likelihood might be affected by human influence on climate before they actually occur. In many cases, risks are relatively well known from catastrophic loss modeling in the insurance industry, which also has considerable experience in precise specification of what an event (such as a hurricane of Category 4 or greater striking a major Gulf Coast city) exactly entails. It would be a major undertaking to draw all this information together. However, doing so would significantly simplify and potentially strengthen future attribution claims. If an event is predefined and then it occurs, we completely avoid the problem of assessing how much our attribution results may be biased by the possibility that we may be tailoring our definition of the event to the knowledge of what actually occurred.

The statistics on summer temperature variability about the expected “background” level were found to be simulated adequately by the climate model used, and were found to conform to a standard model for the behavior of extreme values, displaying no evidence of changing systematically over time.⁷⁹ This last condition allowed a very significant simplification of the problem because it meant that the impact of human influence on the risk of summer temperatures exceeding a particular threshold is simply determined by the characteristics of this distribution and the human impact on expected summer temperatures. The impact of any external driver on expected summer temperatures is straightforward to define because of the deterministic, predictable relationship between external drivers and climate.

We obtained a range of magnitudes for human influence on European summer climate using precisely the multiple regression approach described above: in effect, scaling the response to anthropogenic drivers up and down until the scaled version of the model was no longer consistent with observations of European summer temperatures over the twentieth century. This scaling approach provides a method of synthesizing a set of possible models without explicitly running them. It is justified by the observation that temperature responses to external drivers differ between models primarily in magnitude and much less in the temporal evolution or spatial pattern of the

⁷⁹ This condition is unlikely to be satisfied on smaller scales to those considered in the Stott et al. study. See, e.g., Christoph Schär et al., *The Role of Increasing Temperature Variability in European Summer Heatwaves*, 427 NATURE 332 (2004); Sonia I. Seneviratne et al., *Land-Atmosphere Coupling and Climate Change in Europe*, 443 NATURE 205 (2006).

response, something that is less true for other variables like precipitation.⁸⁰ At the same time, uncertainty in the response to natural drivers, including fluctuations in solar and volcanic activity, was taken into account using a similar approach: scaling the model-simulated response to these other drivers up and down and allowing the scaling factor to vary along with the scaling on the anthropogenic response. This approach indicates an anthropogenic summer warming of around $0.5^{\circ}\text{C} \pm 0.2^{\circ}\text{C}$, with relatively little warming due to natural influences.

Given that year-to-year variability about this expected summer temperature appears to be stationary, computing the possible impact of such an uncertain change in expected mean temperatures on the probability of a given temperature threshold being exceeded is straightforward. We simply subtracted possible values of the anthropogenic summer warming from a simulation of late-twentieth-century conditions with all external drivers imposed to simulate a set of possible models with a range of different responses to anthropogenic influence. Crucially, this approach assumes that, before any comparison is made with observations, all values of anthropogenic or natural warming are equally likely, so that there is an equal chance of greenhouse-induced cooling or warming. Although such an outcome is highly unlikely on physical grounds, it is here ruled out by the record of observed summer temperatures, not by any model simulation. The model is used to provide an estimate of the spatial and temporal form of the response, not its magnitude or sign.

Having synthesized distributions of summer temperatures with and without anthropogenic influence for a range of “possible models” of varying degrees of likelihood, we can then compute the risk of summer temperatures exceeding any given threshold. Results are shown in Figure 6. The fuzzy line labeled “With human influence” shows how the “return period” (the inverse of the odds of an event occurring in a particular year) for summer regionally- and seasonally-averaged temperatures exceeding a particular threshold varies with

⁸⁰ See generally H. Douville, *Detection-Attribution of Global Warming at the Regional Scale: How To Deal with Precipitation Variability?*, 33 GEOPHYSICAL RES. LETTERS, Jan. 17, 2006, at L02701-1 (discussing the effect of the “natural variability of precipitation” on our ability to attribute changes in surface temperature to human activity); Gabriele C. Hegerl et al., *Detectability of Anthropogenic Changes in Annual Temperature and Precipitation Extremes*, 17 J. CLIMATE 3683, 3689-90 (2004); F. Hugo Lambert et al., *Attribution Studies of Observed Land Precipitation Changes with Nine Coupled Models*, 32 GEOPHYSICAL RES. LETTERS, Sep. 21, 2005, at L18704-1 (noting that precipitation is subject to greater variability than temperature).

the threshold considered assuming key external drivers of climate change take values similar to those that obtained in 2003. On this analysis, temperature exceeding the 1.6°C threshold was approximately a one-in-one-hundred-year event, or, rather, there was only a 1% chance of it happening in that particular year. The spread or the fuzziness of the line is an indication of the uncertainty in what the return time of this event actually was.

The fuzzy line labeled as “Without human interference” shows how the return period varies with threshold, assuming the impact of human influence has been removed and all other factors are held near 2003 values. In the absence of human influence, temperatures exceeding the 1.6°C threshold would have been closer to a one-in-one-thousand year event, so the impact of human influence has been to increase the risk of such an event occurring by a factor of four to ten. Note that in the 2004 article,⁸¹ the statistics of these extremes were modeled with a Gaussian distribution, giving a more conservative (smaller) estimate of the increase in risk. We believe it is both more realistic and consistent with standard practice to model these distributions with an extreme value distribution.⁸² With the limited sample of model simulations available (and even more limited observations), we cannot distinguish between different representations, so it is appropriate to use the most flexible representation, which is the generalized extreme value distribution.

The fact that these return period plots are straight lines on a log scale is a simple consequence of the shape of the exponential tail of the extreme value distributions used. This would also apply to a broad range of other distributions. The fact that these lines are parallel is a consequence of our assumption that the impact of human influence is to shift the mean of the distribution without significantly altering its shape. This would not necessarily be the case in a more complex problem.⁸³ But in this case it clearly simplifies matters, because the ratio P_0/P_1 which determines FAR is independent of the threshold used or possible modeling errors that might shift the distributions up or down. Since the absolute value of the return period for an event will always be poorly known, this is a very useful property.

⁸¹ Stott, Stone, & Allen, *supra* note 76, at 612.

⁸² See STUART COLES, AN INTRODUCTION TO STATISTICAL MODELING OF EXTREME VALUES 105-23 (2001) (providing a general introduction to modeling extreme-value distributions).

⁸³ For example, it would probably not be the case for temperature extremes on smaller scales.

This analysis indicates that the most likely value for the fraction of risk attributable to past human influence on climate is around 0.85—or, equivalently, that the overall summer warming caused by past human influence on climate has increased the risk of summer temperatures exceeding the 1.6°C threshold by over a factor of six. There is, however, a high level of uncertainty in this estimate of FAR, as indicated by the spread of the fuzzy distributions in Figure 6. Most of this spread arises from uncertainty in the size of human influence on expected European summer temperatures, with some contribution from uncertainty in the estimated parameters of the statistical model used to infer the probability of rare events from a limited ensemble. Averaging over possible models, weighted by their relative likelihood, indicates a mean likelihood-weighted FAR of 0.75, while an FAR of less than 0.5 can be ruled out as unlikely at the 10% level. Hence, the 2004 article concluded that it is very likely that human influence on climate increased the risk of the 2003 heat wave by a factor of at least two, with the most likely increase in risk considerably greater than two.⁸⁴

It is important to reemphasize that the conclusions of the 2004 article pertained to large-area seasonally averaged temperatures,⁸⁵ which are still relatively far removed from the scale of the extreme weather events that caused the bulk of the damage over the heat wave. In principle, extending the analysis to smaller spatial scales and shorter time scales is quite feasible. The essential question is the same: how does the probability of this event differ between an ensemble simulation of present-day climate and an equivalent simulation excluding past human influence, allowing for uncertainty in the formulation of the simulation model by averaging over possible models weighted by some measure of their relative likelihood? The assumptions used to address this question would, however, need to change. In particular, the influence on nonlinear feedbacks on smaller scales would rule out the assumption of stationary variability about a changing background climate, necessitating explicit simulation with ensembles of higher resolution models.

Despite these caveats, the conclusions of the 2004 article are of interest in that the overall estimate for attributable risk is not inconsistent with the possibility of liability. David A. Grossman argues that, in the context of United States tort law, “plaintiffs . . . must show that,

⁸⁴ Stott et al., *supra* note 76, at 612-13.

⁸⁵ *Id.* at 611.

more probably than not, their individual injuries were caused by the risk factor in question, as opposed to any other cause. This has sometimes been translated to a requirement of a relative risk of at least two.”⁸⁶ If the risk factor in question were human influence on area-averaged summer temperatures in southern Europe in 2003, it would appear that this threshold has already been passed.

VII. OUTSTANDING QUESTIONS FOR THE LEGAL COMMUNITY

We conclude this review by reiterating a range of questions that we would like to encourage the legal community to consider. These are not scientific questions, but they bear on the nature of the science required to support any causal attribution of harm to human influence on climate and may, in some cases, determine whether such science is even possible.

First, what is “natural” climate? In quantifying the extent of human influence on climate- and weather-related risks, what is the appropriate reference climate? Is it a “preindustrial” climate? If so, which preindustrial climate? Is it relevant that natural factors might—had things like the timing of volcanic eruptions evolved differently—have caused a warming as large as that observed over the past half century? Or is the relevant reference climate simply the climate that would have occurred in the present day, all other things being equal, in the absence of human influence? As a scientific community, we hope that the legal community can be satisfied with the second option, even though it can only be explored through computer simulation, because the problem of characterizing the “climate that might have been” if we had not interfered with it is considerably better posed than the problem of characterizing preindustrial climate, as evidenced by the “Hockey Stick” debate over reconstructions of the climate of the last millennium.

Second, are we primarily interested in impacts in the current quarter century, or over longer time scales? The advantage of focusing on the current quarter century is that impacts on this time scale are unlikely to be affected by current or future decisions regarding carbon dioxide emissions (although they may be affected by other human and natural influences on climate). This simplifies the attribution of responsibility for impacts but, of course, makes it difficult to

⁸⁶ David A. Grossman, *Warming Up to a Not-So-Radical Idea: Tort-Based Climate Change Litigation*, 28 COLUM. J. ENVTL. L. 1, 23 (2003) (footnote omitted).

design effective measures of injunctive relief. The disadvantage from an attribution perspective is that few impacts on this time scale will be events that could not have occurred in the absence of human influence, making some form of probabilistic approach to attribution essential.

Third, if we are concerned with long-term future impacts, how should we deal with uncertainty in future emissions? The exchange between the amici curiae⁸⁷ in *Massachusetts v. EPA* illustrates that it is possible to argue over likely future emission rates and consequent carbon dioxide levels more or less indefinitely. Most of the factors on which these emission rates depend, such as future global population trends or the economic performance of large developing nations, would lie outside the control of most prospective defendants, complicating any link between defendants' decisions and future climate impacts.

Fourth, if we are concerned with current or near-term future impacts, most will be due to extreme weather events that might have occurred naturally, but whose risk of occurrence may have been increased by climate change. Can we treat changing risks as if they are real quantities, despite the fact that these risks must be estimated from computer simulations that cannot, because we only have a single realization of the observed record, be directly validated against observations?

Finally, if we are prepared to base an assessment of liability in terms of attributable changes in risk, is the overall approach to quantifying Fraction Attributable Risk described above⁸⁸ the best way to go forward? It was developed within the scientific community with reference to the epidemiological literature, but with no particular regard to the needs of the legal community, so there is no reason to suppose it is the optimal approach.

As we stressed at the outset of this Article, if these questions can be resolved, it is to be hoped that attribution scientists can then play a relatively noncontentious, mundane, and technical role in forthcoming cases in which the attribution of harm to human influence on climate is an issue. This would allow the courts to focus on other matters such as standing and negligence that are closer to their traditional remit, and would also avoid polarizing the scientific community any further than it is already.

⁸⁷ See *supra* notes 47 & 50.

⁸⁸ See *supra* note 73 and accompanying text.

APPENDIX

Figure 1




Figure 1: (a) Observed global temperatures over 1901-2005 from near-surface temperature observations (thick black line) and as simulated by a range of climate models driven by the combination of increasing GHGs, changes in anthropogenic aerosols, solar variability, and volcanic activity (thin gray lines). Average of model simulations shown as thick grey line. (b) Same as in (a), but removing the impact of human influence from the model simulations.

Figure 2⁸⁹




Figure 2: Reconstructions of Northern Hemisphere temperatures over the past millennium from “proxy” observations (primarily tree-ring records) from various sources (grey bands, thin lines, and dotted lines), compared with instrumental observations (thick black line) and a simulation using a simple

⁸⁹ Hegerl et al., *supra* note 22, at 1029 fig.1. Figure reproduced with permission. Monochrome version kindly provided by G. Hegerl.

climate model (thick grey solid line) driven by the combination of natural and anthropogenic forcings.

Figure 3

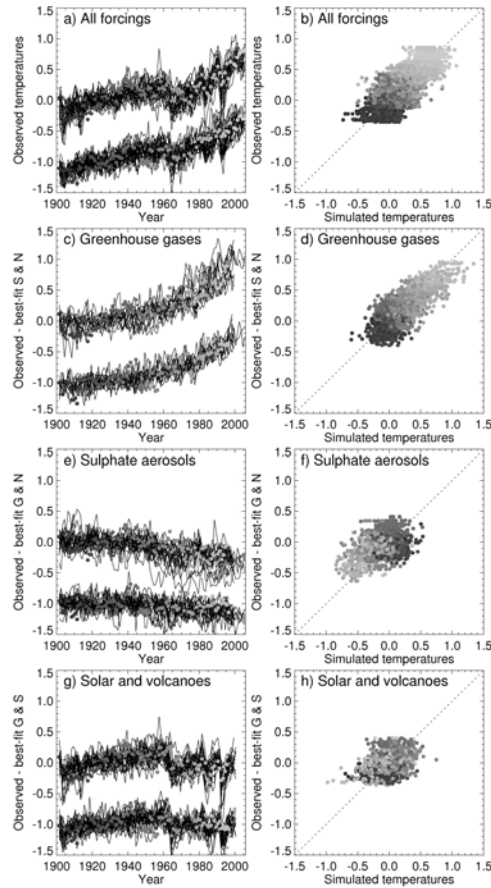


Figure 3: A simple attribution analysis, comparing model simulations with observed temperature changes over the twentieth century. (a) Grey dots: Observed Northern and Southern Hemisphere area-averaged near-surface temperature anomalies during the period 1901-2005 relative to average temperatures between 1900-1940. Grayscale indicates time, with lighter being more recent. Black lines: Corresponding simulated temperatures from six of the models shown in Figure 1 driven by the combination of GHG increase, anthropogenic sulfate aerosols, and natural (solar and volcanic) variability. Southern Hemisphere points are offset by 1°C . (b) Same data, plotting model simulations (horizontal) against observations (vertical). Grayscale indicates time, as

in panel (a). (c) Grey dots: Observed temperature anomalies after removing the best-fit contribution from sulfate and natural forcing. Best-fit is obtained from a three-way, least-squares multiple linear regression between the observations and model-simulated responses to GHGs, sulfate, and natural forcing, obtained from simulations in which drivers are prescribed separately (ensemble means smoothed with a five-point running mean). Black lines: Simulated temperatures from three models driven by GHGs alone. (d) Simulated greenhouse response versus observed temperatures after removing best-fit sulfate and natural contributions. Regression fits are obtained for the models separately, hence allowing the models to make different errors in the magnitudes of their responses. Fitted points are plotted separately in panel (d) and averaged together before being removed from the observation in panel (c). (e) and (f): same as in (c) and (d), but showing the response to anthropogenic sulfates. (g) and (h): the response to natural (solar and volcanic) variability. Formal uncertainty analysis of regression slopes requires a more sophisticated treatment. The aim of this Figure is simply to illustrate the strength of the greenhouse signal over the aerosol and natural contributions (which, although weaker, are still detectable). The fact that the dots in panel (d) lie along the leading diagonal indicates that these models are neither overestimating nor underestimating the response to GHG increase.

Figure 4




Figure 4: (a) A comparison of observed carbon dioxide increase since 1750 with an idealized scenario in which carbon dioxide levels increase at 0.45% per year for seventy years from 1935-2005. Note that carbon dioxide levels since the 1970s are very similar between the real and idealized scenarios. (b) A comparison of the response of a simple climate model to both observed carbon dioxide increases (solid line) and the idealized scenario (dashed line). By the year 2000, carbon-dioxide-induced warming is very close to what is expected under this idealized scenario, the impact of the earlier “ramp-up” having dissi-

pated. Hence, estimates of carbon-dioxide-induced warming to date are very close to half the theoretical Transient Climate Response.

Figure 5⁹⁰




Figure 5: An oversimplistic comparison of models with data. Observed near-surface temperature anomalies 1975-2005 (dots) plotted against model-simulated responses to an imposed 1% per year increase in carbon dioxide (thin lines). Thick line is the mean model-simulated response and thick dashed line shows an extrapolated trend fitted to these observations. The agreement between the models driven by a 1% per year increase and the observations, driven by an average of slightly under 0.5% per year over this period, is coincidental and should not be taken to imply that the models are systematically underestimating the response to a given rate of carbon dioxide increase.

⁹⁰ Brief for Baliunas et al., *supra* note 50, at 13 fig.2.

Figure 6




Figure 6: Change in risk of European summer temperature anomalies exceeding a range of temperature thresholds. The thick line labeled “With human influence” shows the best estimate of the risk under present-day conditions, with dotted lines showing the range of uncertainty in these return periods. The thick line labeled “Without human influence” shows the best estimate of the risk under conditions similar to the present day except with human influence (the impact of increased GHG levels and sulfate aerosols) removed, with dotted lines showing the range of uncertainty.