altruists do not. Tolerant agents have lucrative exchanges with outsiders; parochial agents do not. A high proportion of parochials in groups restricts trading opportunities for all.

Among the four possible combinations of traits, TN is the most profitable. These self-interested traders profit both from contact with outsiders and from the donations made by altruists. The most costly combination is PA. These generous warriors make donations and also risk their lives to protect noncombatants and conquer new territory for the group's offspring. Individual selection favors the T and N alleles over the P and A alleles. Victory in war favors groups with more PA types over those with fewer. The other two trait combinations are PN bullies, who are both hostile and selfish, and TA philanthropists, who both trade and donate to others.

In each generation, groups are randomly paired. What happens next depends on the proportions of tolerant types and warriors in the paired groups. If two highly tolerant groups are paired, tolerant members reap the benefits of trade. If the proportion of tolerant types drops below a strong majority in either group, however, the likelihood of peaceful trade plummets. Instead, the groups have either an unproductive standoff or a war. If both groups have the same numbers of warriors, a standoff results. War becomes increasingly likely the greater the imbalance of power, and wars end in a victory or a draw. Some proportion of warriors are killed regardless of outcome. In a victory, however, many civilians on the losing side are also killed, and offspring from a postwar baby boom among the victors migrate into the conquered territory.

The societies that evolve are stable in two conditions: when either selfish traders (TN) or generous warriors (PA) are the dominant type. A few PN bullies and even fewer TA philanthropists can coexist within trader or warrior regimes. The trading regime is peaceful. Standoffs and wars are more common in the warrior regime, but even infrequent war—10 to 20% of encounters—can maintain high levels of parochial altruism. Similar findings for the impact of intermittent war on the evolution of heroism (6) suggest that war need not be "constant" to act as a powerful selective force.

The convergence of altruism and parochialism in Choi and Bowles' simulation is consistent with links between the two found in behavioral studies. Selfish choices in social dilemma experiments, for example, diminish markedly when the game is embedded in an intergroup context (7). The boost in altruism caused by awareness of an outgroup is also more marked among women than men (8),

consistent with war exerting stronger selective pressure on males as warriors. Interestingly, altruism levels for women, although relatively unaffected by intergroup hostility, were still high. It appears that the relative importance of alternative evolutionary pathways to altruism may differ for men and women.

A full accounting of such pathways must include cultural evolution. In other work, Bowles and colleagues show how norms can support altruism by promoting conformity (9). In the current simulation, warrior-rich groups enforce a trading ban. However, this norm is predetermined. An obvious extension would be to allow norms to evolve. Can pro-trade norms outcompete more isolationist parochial norms? Do norms that punish cowards naturally coevolve with war and altruism?

The simulation findings suggest that one legacy of war is an inherent tension between tolerance and altruism. Cross-cultural studies, however, provide grounds for optimism. In one study, people from 15 small-scale societies played a donation game (10). Average generosity correlated with the amount of market exchange and economic cooperation typical in the society. By adding mutable norms to the simulation, the poten-

tial viability of societies of tolerant altruists could be further explored.

A better understanding of how our impulses to give, to trade, and to attack outsiders are intertwined should help in the quest to promote pro-social behavior while keeping the sharp end of altruism sheathed.

References and Notes
1. J.-K. Choi, S. Bowles, *Science* **318**, 636 (2007).
2. W. D. Hamilton, *J. Theor. Biol.* **7**, 1 (1964).
3. R. L. Trivers, *Q. Rev. Biol.* **46**, 35 (1971).
4. E. Sober, D. S. Wilson, *Unto Others: The Evolution and Psychology of Unselfish Behavior* (Harvard Univ. Press, Cambridge, MA, 1998).
5. L. H. Keeley, *War Before Civilization; The Myth of the Peaceful Savage* (Oxford Univ. Press, New York, 1997).
6. H. Arrow, O. Smirnov, J. Orbell, D. Kennett, in *Conflict in Organizational Teams: New Directions in Theory and Practice*, L. Thompson, K. Behfar, Eds. (Northwestern Univ. Press, Evanston, IL, 2007), pp. 113–142.
7. G. Bornstein, I. Erev, *Int. J. Conflict Manag.* **5**, 271 (1994).
8. M. Van Vugt, D. De Cremer, D. P. Janssen, *Psychol. Sci.* **18**, 19 (2007).
9. S. Bowles, J.-K. Choi, A. Hopfensitz, *J. Theor. Biol.* **223**, 135 (2003).
10. J. Henrich *et al.*, *Am. Econ. Rev.* **91**, 73 (2001).
11. I thank J. Orbell, K. Henry, and B. H. Rogers for helpful comments, and O. Smirnov and D. Kennett for earlier discussions on altruism and war.

ATMOSPHERE

Call Off the Quest

Myles R. Allen and David J. Frame

Knowledge of the long-term response of Earth's climate to a doubling of atmospheric carbon doixide may be less useful for policy-makers than commonly assumed.

Over the past 30 years, the climate research community has made valiant efforts to answer the "climate sensitivity" question: What is the long-term equilibrium warming response to a doubling of atmospheric carbon dioxide? Earlier this year, the Intergovernmental Panel on Climate Change (1) concluded that this sensitivity is likely to be in the range of 2° to 4.5°C, with a 1-in-3 chance that it is outside that range. The lower bound of 2°C is slightly higher than the 1.6°C proposed in the 1970s (2); progress on the upper bound has been minimal.

On page 629 of this issue, Roe and Baker (3) explain why. The fundamental problem is that the properties of the climate system that

M. R. Allen is in the Department of Physics, University of Oxford, Oxford OX1 3PU, UK. D. J. Frame is in the Oxford University Centre for the Environment, Oxford OX1 3QY, UK. E-mail: myles.allen@physics.ox.ac.uk; dframe@atm.ox.ac.uk

we can observe now do not distinguish between a climate sensitivity, S, of 4°C and $S > 6$°C. In a sense, this should be obvious: Once the world has warmed by 4°C, conditions will be so different from anything we can observe today (and still more different from the last ice age) that it is inherently hard to say when the warming will stop. Roe and Baker formalize the problem by showing how a symmetric constraint on the strength of the feedback parameter f (which determines how much energy is radiated to space per degree of surface warming) gives a strongly asymmetric constraint on S. The reason is simple: As f approaches 1, S approaches infinity. Roe and Baker illustrate the point with the information provided by recent analyses of observed climate change, atmospheric feedbacks, and "perturbed physics" experiments in which uncertain parameters are varied in climate models.

It might be objected that some models that displayed high sensitivities in perturbed physics experiments also poorly reproduce the energy budget at the top of the atmosphere (4) and hence perform poorly in short-term climate forecasts (5). Likewise, the fact that direct studies of atmospheric feedbacks provide only a weak constraint on S does not mean that no stronger constraint is possible. But these objections miss Roe and Baker's main point: The fact that uncertainties in climate processes add up to give an approximately Gaussian uncertainty in f means that there are innumerable ways of generating a climate model with f close to unity and hence a very high S. Ruling all of these out requires us to find observable quantities that are consistently related to S in all physically plausible climate models, and to show that observa-

stabilization concentration of carbon dioxide (or equivalent) for which the risk of dangerous warming is acceptably low. Fortunately, we do not need to.

To understand why, consider two scenarios for carbon dioxide–induced warming, based on large numbers of runs of a simple climate model constrained by recent temperature observations (6–10). In the first scenario, carbon dioxide concentrations are stabilized at 450 ppm from 2100 onward. If S turns out to be close to our current best estimate, then achieving this concentration target gives an eventual equilibrium warming of 2°C (see the figure, left, dashed line). But S is uncertain; thus, even if we stabilize at 450 ppm, we cannot rule out much more than 2°C of eventual warming, as shown by the shaded plume. Notice that observed tem-

a low (and hence expensive) C2K is much better constrained by data than is the risk of a high (and hence dangerous) climate sensitivity. This is because C2K, like f, scales approximately with things we can observe, and hence is not subject to the problems that bedevil efforts to constrain sensitivity. The uncertainties in how the available policy levers translate into global emissions, and how emissions translate into concentrations through the carbon cycle, are so large that uncertainty in the final concentration we are aiming for in 2200 is probably the least of our worries—provided we resist the temptation to fix a concentration target early on. Once fixed, it may be politically impossible to reduce it.

The temperature response to this adaptive-stabilization scenario (see the figure, right) is much better constrained because it depends on current trends, not on S. If S turns out to be toward the upper end of the current uncertainty range, we may never find out what it is: Some models with $S = 4$°C are effectively indistinguishable from others with $S = 6$°C under this scenario. But provided our descendants have the sense to adapt their policies to the emerging climate change signal, they probably won't care.

An upper bound on the climate sensitivity has become the holy grail of climate research. As Roe and Baker point out, it is inherently hard to find. It promises lasting fame and happiness to the finder, but it may not exist and turns out not to be very useful if you do find it. Time to call off the quest.
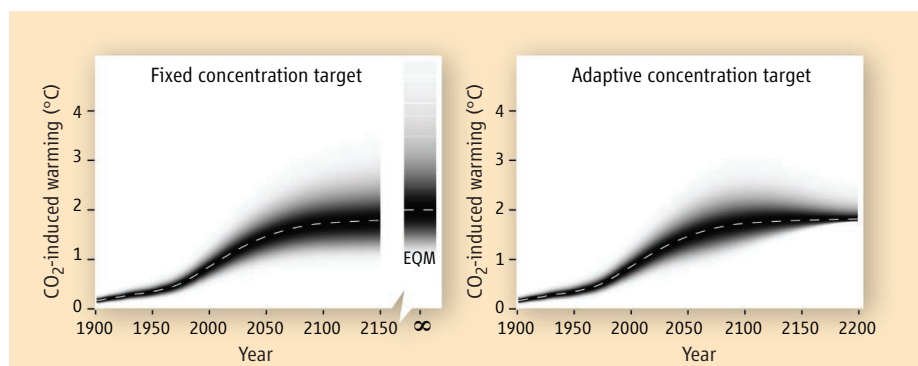
Carbon dioxide–induced warming under two scenarios simulated by an ensemble of simple climate models. (**Left**) CO_2 levels are stabilized in 2100 at 450 ppm; (**right**) the stabilization target is recomputed in 2050. Shading denotes the likelihood of a particular simulation based on goodness-of-fit to observations of recent surface and subsurface-ocean temperature trends (7, 8). Simulations are plotted in order of increasing likelihood, so worse-fitting models are obscured. The bar labeled "EQM" shows the models' likelihood against their long-term equilibrium warming at 450 ppm. How these likelihoods are translated into forecast probabilities is controversial, and the more asymmetric the likelihood function, the greater the scope for controversy.

tions of these quantities are inconsistent with a high S. Despite much searching, such observations remain elusive.

There are even more fundamental problems. Roe and Baker equate observational uncertainty in f with the probability distribution for f. This means that they implicitly assume all values of f to be equally likely before they begin. If, instead, they initially assumed all values of S to be equally likely, they would obtain an even higher upper bound. This sensitivity of the results to prior assumptions shows that the real problem with the upper bound on climate sensitivity is not that it is high (in which case we could hope that more data will bring it down), but that it is controversial: Opaque decisions about statistical methods, which no data can ever resolve, have a substantial impact on headline results.

All this would be very bad news if avoiding dangerous anthropogenic interference in the climate system required us to specify today a

perature trends provide a much stronger constraint on forecast warming even 50 years after stabilization than on the long-term equilibrium response (shown by the bar labeled EQM). Hence, if the true climate sensitivity really is as high as 5°C, the only way our descendants will find that out is if they stubbornly hold greenhouse gas concentrations constant for centuries at our target stabilization level.

In reality, of course, our descendants will revise their targets in light of the climate changes they actually observe. Suppose that, in 2050, they simply divide our 450-ppm target forcing by the fraction by which the observed carbon dioxide–induced warming trend between 2000 and 2050 over- or underestimates our current best-guess forecast (11). They then recompute concentration paths to stabilize at this revised level in 2200.

The long-term carbon dioxide concentration consistent with a 2°C warming (which we call C2K) is currently uncertain, but the risk of

References and Notes

1. Intergovernmental Panel on Climate Change (IPCC), *Climate Change 2007: The Physical Science Basis*, S. Solomon et al., Eds. (Cambridge Univ. Press, Cambridge, 2007).
2. J. Charney et al., *Carbon Dioxide and Climate: A Scientific Assessment* (National Academy of Sciences, Washington, DC, 1979).
3. G. H. Roe, M. B. Baker, *Science* **318**, 629 (2007).
4. B. M. Sanderson, C. Piani, W. J. Ingram, D. A. Stone, M. R. Allen, *Clim. Dyn.*, 10.1007/s00382-007-0280-7 (2007).
5. M. Rodwell, T. N. Palmer, *Q. J. R. Meteorol. Soc.* **113**, 118 (2007).
6. The model is a simple energy balance atmosphere coupled to a diffusive ocean as used in (7), with data constraints updated as in (8). In such simple models, feedbacks are assumed to be independent of the climate state, despite evidence to the contrary (9, 10).
7. D. J. Frame et al., *Geophys. Res. Lett.* **32**, 10.1029/2004GL022241 (2005).
8. D. J. Frame et al., *Geophys. Res. Lett.* **33**, 10.1029/2006GL025801 (2006).
9. C. A. Senior, J. F. B. Mitchell, *Geophys. Res. Lett.* **27**, 2685 (2000).
10. G. Boer, B. Yu, *Clim. Dyn.* **21**, 167 (2003).
11. Note that this does not require our descendants to "discover" the true S, unlike the similar "learning" scenario of Yohe et al. (12).
12. G. Yohe, N. Andronova, M. Schlesinger, *Science* **306**, 416 (2004).
13. We thank the James Martin School and the Tyndall Centre for support.