

We don't, and can't, know how much it will cost to tackle climate change

Article by David Roberts taken from Grist (<http://grist.org/>)

How much will it cost humanity to avoid dangerous climate change?

Debates over this contentious topic inevitably focus on a set of economic models that seek to estimate those costs over the course of the next century. The latest such analysis is the [5th IPCC Assessment on Mitigation of Climate Change](#), which has people confidently **asserting** that properly addressing climate change will trim 0.06 percent a year off global GDP growth. Not 0.05 a year. Not 0.07. But 0.06. The much-ballyhooed [New Climate Economy](#) report, out earlier this year, draws many of the same conclusions.

Can we really know the costs of mitigation out through the year 2100 with that level of precision?

I've long had my own opinion on this, but I'm just a dirty hippie blogger. So I've been patiently waiting for somebody with credentials to express it for me, in a fancy peer-reviewed paper.

Well, lucky me. I found the paper. (Thanks to the tweeter who flagged it for me — sorry I can't remember who you are.) And it confirms my considered opinion on the matter, which is: No, we can't know with that level of precision. In fact, we can't know at all. This arguing back and forth over over models is so much posturing, a way of hiding our values and priors behind a sheen of econometrics. We should just drop the facade and debate our values directly.

Let's dig in, shall we?

The models that inform climate economics are called **integrated assessment models** (IAMs), which integrate models of energy markets and land use with greenhouse gas projections. They have come under plenty of criticism in the past; see, for instance, [this paper](#) from economist Frank Ackerman and colleagues, which picks apart many of the assumptions that cause IAMs to overstate costs of climate mitigation and understate the need for action.

The paper I'm talking about — “**The economics of mitigating climate change: What can we know?**”, written by Richard Rosen of the Tellus Institute and Edeltraud Guenther of the Technische Universitaet Dresden — goes beyond that. It questions whether IAMs are up to the job at all.

What IAMs are good for is determining the effects of marginal changes in one or a small set of economic variables. So you run your model of the economy without a \$1 increase in the price of bacon, run it again with the \$1 increase, compare the two, and *voilà*, you have a pretty good idea what the effect of raising bacon prices will be. Crucially, the proposed changes are modeled *ceteris paribus* (all else being equal). So in our example, *only* the price of bacon changes. Other parameters remain the same.

This works great to model the effects of modest tweaks to the economy over modest time spans. But restraining global temperature increase to 2 degrees C or less — right now the generally agreed-upon international goal — will require much more than tweaks. It will take a long time. And all else will not remain equal.

[M]itigation scenarios are not small perturbations on easily forecastable baseline scenarios, where linear ‘first order’ differences would dominate. Rather, mitigation scenarios represent major transformations of the economy relative to baseline scenarios and, thus, represent large and highly non-linear changes that will strongly impact the development of new energy technologies on both the supply and demand sides, as well as other relevant technologies that offset greenhouse gas emissions worldwide.

“Large and highly non-linear changes” are not IAMs’ strong suit. Nor is this level of complexity. Climate IAMs have “hundreds of input parameters, each of which is highly uncertain in the long run.”

Rosen and Guenther’s paper was published before the latest (5th) IPCC report, so it mostly focuses on the IPCC’s Fourth Assessment and the Stern Report (which in turn drew from a few common sources). But the issues it identifies apply regardless. The conclusion — spoiler! — is that **the 50-to-100-year costs of mitigation are unknowable.**

We’ll take a quick look at why, and then we’ll ask what follows from it.

There are several critiques of IAMs in the paper but, recognizing your limited patience for modeling talk, I’ll just hit a few.

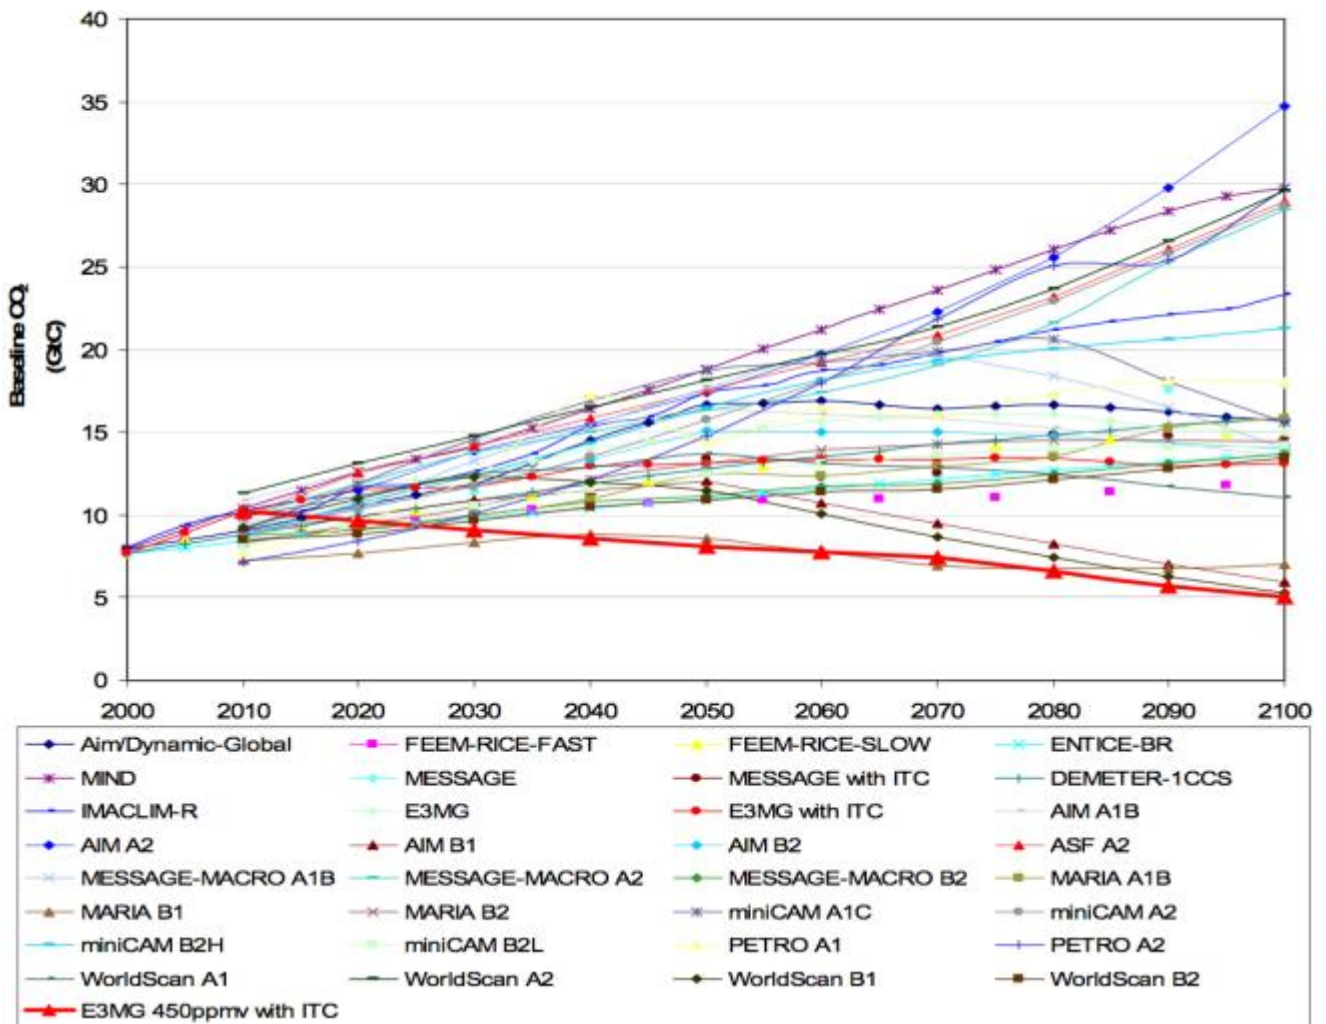
First, there's the question, compared to what? To calculate the net costs or benefits of climate mitigation policies, one must compare the carbon they reduce to a world in which the policies weren't passed — a "baseline scenario" for CO2 emissions.

What should that baseline scenario look like? Well, what will the global economy look like over the next 50 to 100 years without more climate mitigation? Uh ... I don't know. No one knows. How could anyone know? It's ridiculous.

Unchecked climate change may cause economic dislocations that substantially alter the course of GDP. "But integrated assessment modelers never model feedback between the amount of climate change and economic growth," write the authors, "and would have an extremely difficult time doing so if they tried," given the uncertainty of climate impacts. Or what if there's an oil crisis, or financial crisis? Who the hell knows what's going to happen over a century?

Sure enough, IAMs have widely varying baseline scenarios, reflecting different assumptions about the pace and nature of growth. The chart below is from a **meta-analysis of IAMs** done by Barker et al., commissioned by the Stern Report and relied upon by the IPCC as well. It shows the range of baseline scenarios in the models it considered:

Baseline CO₂ Emissions (GtC)



“Mitigation Costs of GHG with ITC: a meta-analysis,” Barker et al.

What’s a 30 gigaton difference among friends?

Since the cost differences between baseline and mitigation scenarios within a particular model tend to be fairly small, differences in baseline can substantially complicate any effort to do inter-model comparisons. Just averaging different baselines — throwing a bunch of contestable assumptions into a pot and stirring — seems unlikely to produce accurate results.

“Because forecasting the future of the energy economy for the next 50–100 years is impossible (not just difficult),” they conclude, “there is no valid baseline emissions scenario to which the costs of a mitigation scenario can be compared.”

Second, and related, to calculate the net costs and benefits of climate mitigation, one needs a grasp of the climate damages that will be avoided by such mitigation. Yet most IAMs “do not include any estimates of the likely future damage due to climate change at all.”

Because no reasonable damage estimates have been incorporated into most IAM results, and because there is tremendous uncertainty inherent in even trying to make monetized estimates of damages that have yet to happen on a global scale, it is very hard to estimate, even roughly, the net benefits of mitigating climate change in the long run.

The IPCC’s Fourth Assessment, chapter three, says, “Due to considerable uncertainties and difficulties in quantifying non-market damages, it is difficult to estimate SCC [social cost of carbon] with confidence. Results depend on a large number of normative and empirical assumptions that are not known with any certainty.” Yet quantifying the social cost of carbon is the only way to know the monetary value of avoiding climate change!

Chapter three also says, “the risk of climate feedbacks is generally not included in the ... analysis.” But climate feedbacks, the vicious cycles whereby the effects of warming in turn accelerate warming, are the greatest danger of climate change, precisely the changes that have the potential to substantially affect GDP.

Cumulatively, “these differential impacts [of climate change] on metrics such as GDP could be very substantial and could even be larger than the impacts on GDP of attempting to simply mitigate climate change.” Catch that? The damages avoided by climate mitigation could be larger than the cost of mitigation, yet they are invisible in most IAMs. It’s crazy.

Third and finally, there’s the question of technology costs. To calculate the total net costs and benefits of climate mitigation, one needs some estimate of the cost of investment in various mitigation technologies. Add up those investments and you get the *gross* cost of mitigation. Then subtract the value of the displaced carbon energy to get the *net* cost of mitigation.

So, what is the cost of an increment of solar power in, say, 2060? Ha ha. Who the frack knows that? How about biomass, or CSS, or electric, self-driving cars, or so-called **smart cities**? How will advances in some of these technologies affect the cost of

others? How will changes in consumption patterns or land use affect their costs? No one knows.

Not surprisingly, IAM models have different assumptions about technology costs. So even for IAM runs with similar CO₂ baselines, cumulative CO₂ reductions can vary widely.

The technical criticism here is that the Barker meta-analysis did not factor these differing cost assumptions in as independent variables when doing its regression analysis on IAM results. (Yeah, yeah, I know, zzzzz.)

But the larger criticism, again, is of the sheer presumption:

Can we, and should we, pretend to be able to forecast the future economic costs of each technology, in each scenario, with any reasonable accuracy? Or, even more challenging, can we accurately forecast small cost differences between two very different possible future scenarios involving very different mixes of technologies? If we cannot, can we get a reasonable approximation of the likely net costs or benefits of mitigating climate change, or not?

Um, not.

And yet IPCC's Fourth Assessment conclusion, duly delivered to policymakers, is that climate mitigation will cost, on average, "0.12 percentage points [of global GDP] per year for 50 years." The authors raise an eyebrow.

[G]iven all the uncertainties and variability in the economic results of the IAMs ... the claimed high degree of accuracy in GDP loss projections is highly implausible. After all, economists cannot usually forecast the GDP of a *single* country for *one year* into the future with such a high accuracy, never mind for the entire world for 50 years, or more.

So you're saying it might be 0.13 percent a year?

It's a little silly. Given the "huge uncertainties that exist for each input parameter" of IAMs, and the uncertainties they omit entirely, the overall costs of mitigation are "quite unknowable to any reasonable degree of accuracy." (There's a lot more in the paper to support this, including a whole big thing about how IAMs unduly constrain energy-efficiency investments, but I sense patience exhausted.)

Interestingly, the aforementioned Barker has coauthored a more recent paper with the aforementioned Ackerman — "[Pushing the boundaries of climate economics](#):"

critical issues to consider in climate policy analysis” — rounding up many of these critiques of IAMs. They write that ...

... baseline assumptions employed in modeling studies are often arbitrary and inconsistent with each other, particularly when projections are taken off the shelf from different sources. A complete model of climate policy costs and impacts should, in theory, make some of these data endogenous; climate damages can affect the rate of (business as usual) growth of per capita incomes; climate policies can change the price of oil.

Unfortunately, making these feedback effects endogenous (internal to the model) quickly makes a model complex, far more complex than most IAMs and possibly intractably complex. They write bluntly, “Economies are highly complex non-linear systems and it is impossible to accurately predict their future evolution.” Impossible indeed. And thus the costs of climate mitigation over the 50-to-100-year horizon are unknowable.

Rosen and Guenther write in their conclusion:

For the reasons cited [in the paper], not only do we not know the approximate *magnitude* of the net benefits or costs of mitigating climate change to any specific level of future global temperature increase over the next 50–100 years, but we also cannot even claim to know the *sign* of the mitigation impacts on GWP, or national GDPs, or any other economic metric commonly computed.

Given that ignorance, “the IPCC and other scientific bodies should no longer report attempts at calculating the net economic impacts of mitigating climate change to the public in their reports.” Basically, in terms of sound policymaking, no numbers are better than wild-ass guess numbers misleadingly presented as reliable. So economists should just stop pretending those long-term costs are knowable. What do we do, then, without credible estimates of how much it might cost to tackle climate change over the course of this century? Does it *matter* that we don’t have credible estimates? Here Rosen and Guenther are quite bold:

Our answer is “no,” because humanity would be wise to mitigate climate change as quickly as possible without being constrained by existing economic systems and institutions, or risk making the world uninhabitable. This conclusion is clear from a strictly physical and ecological perspective, independent of previously projected economic trade-offs over the long run, and it is well-documented in the climate change literature. As climate scientists constantly remind us, even if the world successfully implemented a

substantial mitigation program today, a much warmer world is already built into the physical climate system. And since we can never know what the cost of a hypothetical reference case would be, and since we must proceed with a robust mitigation scenario, we will never be able to determine the net economic benefits of mitigating climate change, *even in hindsight*.

In short, the physical sciences have given us reason to act. Rather than twiddle knobs to produce largely fantastical long-term cost-benefit analyses, we should simply act. We should imagine the end state we want to achieve and then use modeling to determine the most cost-effective way to proceed toward it in the short-term. That is what models are good for: tools, not masters.

Climate change thrusts upon us the question of what kind of world we want, what we care about, what kind of ancestors we want to be. Long-term modeling attempts to give those decisions a scientific sheen, but it's an illusion. Economics cannot make those decisions for us. They are up to us, and no matter what decisions we make, we are leaping into the unknown.

David Roberts: About Me



David Roberts writes about energy and climate change and joined Vox (from Grist) in April 2015 as environment and climate reporter.

©David Roberts and AECB (October 2015)

Reprinted with permission from David Roberts