

Here's the weird thing: the social cost of carbon today, depends significantly on the year-by-year emissions of carbon in the future, which we obviously don't know. (Because it depends on our own future actions!) It takes some explanation to show why that's true and how it matters.

If you know a bit about climate policy, you know that the SCC — the social cost of carbon — is the amount of harm done by adding an additional ton of CO₂ to the atmosphere. You also know that actually putting a number on the SCC is really hard because of scientific uncertainties about the severity of harm and parameters that economists either don't know at all or can't agree on. But the idea itself seems simple.

Actually, things turn out to be more complicated. That extra ton of carbon is going to be in the atmosphere a long time — centuries, essentially — and so we need to know how much harm it is going to each moment for the entire time. Basically, the temperature increase in proportion to emissions, so the effect of each ton on temperature is effectively a constant. So far, so good.

But — and this is an important “but” — the harmful effects of climate change *aren't* proportional to the temperature change. For example, in one of the leading models, the harmful effects are driven by the square of the temperature change. Models vary, but they all view the relationship between temperature and cost as non-linear (and upward curving, for that matter). What this means is that the harmful effect of an additional ton depends on how much carbon is also in the air each moment that ton is in the atmosphere. An extra ton causes very little harm at times when total carbon is low; much more harm when total carbon is high.

Remember that we want to trace how much harm our extra carbon causes over time. But to do that, we need to know how much other carbon is already in the atmosphere at that future time. And to do *that* we need to know the future trajectory of emissions. That's why the social cost of carbon is not an absolute; it's relative to a future emissions trajectory.

Integrated Assessment Models (IAMs) pick emission scenarios when they are trying to determine the social cost of carbon, but they don't bother to explain why. So it can easily slip by the reader.

Maybe this is interesting, but why is it significant? Here are a few reasons:

1. It introduces another element of uncertainty into calculations of the social cost of carbon, besides the uncertainties about climate impacts from the scientific side and the uncertainties about how those affect the economy from the economic side.

2. It suggests that maybe early carbon emitters should not be held responsible for climate change. Even if they had known as much as we do about climate science, they might not have foreseen that atmospheric emissions would rise so much as to cause significant harm, at least not nearly so quickly.

3. It requires some subtle revision in the standard idea that we should set an emissions tax equal to the marginal cost of the pollutant. Before we impose a carbon tax, the social price of carbon will be one number. But that's not what we want to use as the carbon tax. After we impose a tax, the future emissions trajectory will shift and so will the social cost of carbon. If we pick the right number, the carbon tax will indeed equal the *new* social cost of carbon, but *not* to the old one. The social cost of carbon is a product of the carbon tax, not vice versa.

There are a lot of other problems with figuring out the social cost of carbon, enough so that even some economists think we should take some other approach to setting carbon targets. But this one is particularly interesting because it involves a fundamental aspect of the climate problem, its nonlinearity.