It's common knowledge— at least, among people who follow these things — that it is hard to estimate the benefits of a proposed climate policy. But it's actually quite difficult to estimate the costs as well.

There are three major problems in estimating the benefits of a proposed climate policy. First, it's hard to estimate the damage associated with any given degree of warming. Second, there is a lot of uncertainty, including the potential for tipping points. Finally, because carbon dioxide stays in the atmosphere for such a long time, the benefits of cutting carbon emissions accrue over many decades or centuries. Economists are not agreed on how to compare these long-delayed future with the costs of emissions today.

Costs seem simpler: they're just the amount of money spent to comply. But there are also serious problems with estimating the cost of a climate policy. First, some policies, such as "zero carbon by 2050", are not specific enough. Without a more concrete plan for achieving the goal, the costs remain unknown. Second, as with weather forecasts, forecasting the future of the energy system is always dicey. For both energy and weather, forecasting gets much worse over longer time periods. In terms of energy forecasts, the uncertainty is partly due to the unpredictability of technology — no one predicted fracking, for instance. Even apart from new technologies, prices aren't always. predictable. The precipitous decline in solar prices is a recent example.

This means that it's very hard to know, for instance, how the energy system would respond over time to a substantial carbon tax. Putting all that together, models of the energy system have a very poor track record in terms of predicting future energy prices, consumption, and supply.

It's for this reason that the Energy Information Agency at U.S. DOE prefers to call its model results "projections" rather than "forecasts." They indicate what will happen if there are no surprises — but there are almost always surprises. You could still try to justify using the model results on the theory that positive surprises are as likely as negative ones. If so, the projection may at least indicate the average we should expect. It's not clear that this assumption is true: for instance, the EIA has fairly consistently overestimated future renewable energy prices.

Even if positive and negative surprises are actually equally likely, using the average is wrong. The reason is that policy is not fixed in place for all time. Our ability to change course in the face of surprises allows us to improve on our original plans.

Here's how this works: Suppose our plan is to go to a zero-emission electricity system by

2040. In 2030, we discover that this will cost much less than we expected. In that case, we can accelerate our plans. This makes the cost-benefit ratio better than expected because costs are lower and benefits will begin sooner. On the other hand, suppose we discover that eliminating the last 10% of carbon emissions from the power system by 2040 will be hugely expensive for unanticipated reasons. Then we're likely to change our plans. We might, for instance, focus instead on cutting emissions from transportation rather than electricity. Or we might extend the deadline or spend our money instead on helping developing countries cut their emissions. So due to our ability to make policy changes, the "good" surprises have an amplified positive effect while the "bad" surprises have a diminished negative effect. So using the "no surprise" scenario as our guide actually overestimates costs.

This doesn't mean it's useless to do projections of costs and benefit. They may be helpful to provide a benchmark or in comparing different policies. They may also give us a sense of how heavily a plan would need "good" surprises in order to be workable. But, especially when we're thinking about long time spans, we should take the numbers with a very large grain of salt.