ENERGY, ENVIRONMENT & NATURAL RESOURCES

CLIMATE CHANGE REGULATION AND PREDICTION MARKETS

An impartial party—the market—can weigh in on the climate's future.

BY SHI-LING HSU

Much of what we believe about the future of climate change relies on complicated climate models run on supercomputers. Climate models break the earth's atmosphere down into three-dimensional cells representing blocks of space. The most advanced climate models with the strongest resolution use cells that are 100 kilometers square by 1 kilometer high. Given those dimensions, the best climate models contain about 1.7 billion cells, each with a basket of physical, chemical, and biological characteristics that govern the climate within it. In addition to the complexity of the relationships among cells, there are the biotic feedbacks to consider. Changing temperatures give rise to changing biological and oceanographic conditions, which themselves may alter the physical and chemical properties of each cell.

Climate modelers readily concede limits to current models because much can happen within a cell. Clouds, for example, are often smaller than a cell, so models have only crudely modeled the behavior of clouds. Uncertainty and large error bars are chronic in the business of climate projection.

Layered on top of that literature is another set of educated guesses about the economic effects of a changing climate. What are the damages from the many aspects of climate change? As economist Robert Pindyck has pointed out, at least there is historical data to anchor the climate models; there is nothing analogous in the way of baselines or history to inform economists about the economic harms resulting from higher temperatures or other climate effects. Yet, the economic models and climate models are combined into "integrated assessment models" that project both climate changes and economic effects together, as interrelated phenomena. The integrated assessment models are fundamental in the formation of climate policy throughout the world.

Despite the uncertainties, modeling work is important as it represents the current best attempt to understand the earth's climate. Even so-called "climate skeptics" acknowledge that climate modeling is, in the words of Roy Spencer of the University of Alabama, Huntsville, "an absolute essential part of climate research."

Over time, the accuracy of the models will improve. More importantly, even at these early stages, the models produce some useful information. Even with very conservative model parameterizations, almost all of the models suggest at least a small carbon price, on the order of $20 per ton of carbon dioxide–equivalent. But as a foundation for climate policy, it is safe to say that the whole effort has failed to move large parts of the American electorate. The inability of highly accomplished and talented climate scientists, economists, and modelers to deliver a credible message appears to be a problem for at least the near-term future.

Prediction market / There is an alternative, complementary means of projecting a climate-changed future: climate prediction markets. A prediction market is a market of contingent contracts that pay off in accordance with the outcome of some future event. It is essentially a market for *bets*. As the bets are bought and sold, the prices for those contracts reflect market expectations for the probability of the specified outcome.

Experiences with prediction markets to date have largely vindicated market enthusiasts. Given the right conditions, markets do an excellent job of organizing and sorting information and compiling it all into a tractable result: a price. Prediction markets have largely done an excellent job of predicting. One of the most hallowed and reliable prediction markets is that of the presiden-

SHI-LING HSU is the Larson Professor at the Florida State University College of Law.

SASHA RADOSAVLJEVIC/ISTOCKPHOTO

tial election markets run by the Iowa Electronic Markets (IEM) at the University of Iowa, which have been administered for every U.S. presidential election since 1988. For most of 2012, prices on traded shares of "Obama" and "Romney" were actively traded and widely believed to reflect credible, spin-free, objective estimates of the probability of each candidate winning. The performance of IEM prediction markets for political elections has been remarkably accurate in predicting not only outcomes, but margins of victory—often outperforming exit polls.

VIRTUES OF A MARKET

For the problem of climate change, a prediction market offers the possibility that a market could harness the vast amounts of climate science in ways that even the most sophisticated climate models do not. In outperforming exit polls following elections, and outperforming marketing studies projecting product line sales, prediction markets have demonstrated ways of processing information that escape some very sophisticated algorithms. Markets seem to be able to do some things that even supercomputing cannot.

More importantly for the problem of climate change, a prediction market could impose some discipline on the processing and interpretation of climate science. Both sides in the climate science debate—the concerned and the skeptical—complain that the other side manipulates its science. Too often, climate science is used the way a drunk uses a lamppost: for support, rather than for illumination. In the current environment, beliefs about climate change are too intertwined with a variety of economic and professional interests, such that virtually no one can make an assertion about climate change without being accused of having *some* interest—economic, professional, or psychic—in convincing others. But in prediction markets, it is simply too costly to sustain a disingenuous position. It is harder to put your money where your mouth is when you do not *truly* believe what you are saying, particularly when market prices are providing constant feedback. What a prediction market can do is scrub out the ideological taint of a climate message and deliver a prediction that is not just believable by political ideologues, but also by self-interested traders.

A prediction market for climate change would serve a third vitally important function: it would help organize the vast, disparate, and disorganized sources of climate research. Climate science emanates from an enormous number of fields of study, institutions, and agencies. There are probably no major research universities in the world without at least some climate change research. Perhaps that is appropriate; every ecosystem in every corner of the world will be affected by climate change, and may have something to teach us about potential feedback effects. Understanding the climate is thus a gargantuan task, and making sense of this onslaught of information is a huge challenge.

This is where a prediction market can be useful: markets organize large quantities of information. In markets, traders weight the importance of different pieces of information. This kind of organization has proven extremely daunting given the onslaught of climate information at varying time scales. With climate science coming from so many quarters and drawing on information that is local in many ways, prediction markets are a singularly effective way to process the otherwise intractably numerous sources of climate science.

A MARKET PROPOSAL

What would a prediction market for climate change look like? In my past work, I have proposed a tax-and-cap-and-trade program to set up a prediction market. That proposal had two parts: (i) a small carbon tax that is indexed annually to a basket of climate outcomes, and (ii) a cap-and-trade program for a limited number of permits to emit in lieu of paying the tax. I proposed indexing the carbon tax to:

- global mean temperature
- days of unusually high or low temperatures
- extreme rainfall and drought events
- sea level
- ocean acidity
- hurricanes above a certain intensity level

All of those phenomena are thought (not uncontroversially) to be among the potential and anticipated effects of climate change. Absent a successful geo-engineering effort, they are outcomes that are nonmanipulable. All of them are directly relatable to significant damages, though adaptation efforts may alleviate some of the damages. (For example, developed countries such as the United States could clearly do a better job of protecting their most vulnerable populations from heat waves.) All of those climatic events are routinely monitored internationally, so that even in remote parts of the planet, weather anomalies are susceptible to measurement and counting.

If, in any given year, some combination of conditions produces an increase in the index, the carbon tax would increase. If they produce a decrease, then the carbon tax would decrease. Importantly, to smooth out the fluctuations inherent in climate conditions, the index should be based on a moving average of climate outcomes over a number of years.

The prediction market of this program, however, is the cap-and-trade program. Permits would be unitary exemptions from a future carbon tax: an emitter can either pay the carbon tax or surrender an emissions permit to emit in the specific vintage year. Because of this link between the carbon tax and the permit market, the trading price of the permits should reflect market expectations of what the carbon tax will be in the future and, concomitantly, expectations of future climate outcomes. The idea is to link the price of tradable permits to future climate outcomes so that a market is created in which accurate and credible information about future climate conditions is input into the price of permits. The market for tradable permits to emit in the future is essentially a prediction market for climate outcomes. If permits to emit in the future are issued, then the market price of those permits should reflect credible evaluations of future climate conditions. If traders truly fear that temperatures will rise, they will bid up the price of future permits; if not, then the market price for future permits will remain low. Others have proposed simpler prediction markets: a series of futures contracts traded on an existing private exchange, or operating on something like the Iowa Electronic Markets project, so as to avoid the need for congressional approval. There may be a number of design schemes, but the important point is to create a *market for bets* on events occurring in the future.

In a prediction market for climate outcomes, it is important to choose climate indices that are nonmanipulable and uncontroversially and reliably measured. In a prediction market for climate outcomes, indices should be focused on direct climatological effects such as temperature or droughts, and not secondary, indirect effects such as forest fires or flooding, which are sometimes the product of a mix of climate events and public policy. If expensive and numerous bets are to be placed on the likelihood of various climate outcomes, it is essential that the outcomes be universally understood. I chose the six indices listed above (Arctic sea ice loss could also be added) because they are susceptible to measurement. Scholars have differing opinions about the reliability of this information, but those differences appear to be tractable.

Effects of change / Beyond the need for projecting climate outcomes, however, there exists perhaps an even greater need to predict damages from climate change. In setting climate policy, it is helpful to have a sense of the likelihood or extent of climate conditions such as temperature, drought, or flooding. But several steps would still be needed to translate that into an estimate of the *economic* seriousness of the problem. The social cost of carbon is the estimate of the external damages of greenhouse gas emissions. And to re-emphasize Pindyck's point, current estimates of economic damages stand on even more tenuous footing.

To address this, a climate prediction market might go beyond climate conditions and be extended to predicting the economic effects associated with climate change. Again, there is a need for nonmanipulable and reliably measured indices that represent economic effects. Importantly, it need not be the case that the indices measure damages attributable to climate change. As with a prediction market for climate outcomes, the point is not to quantify the effect of climate change, but to induce traders to bring to market information about the likelihood of some future state. A prediction market for climate outcomes is just meant to project future temperature, not temperature increases from climate change; project the number of hurricanes in the future, not just those attributable to climate change. A prediction market imparts valuable information because it communicates market expectations of future states, not why those states occur. In other words, in establishing a prediction market for economic damages related to climate change, it does not matter what the baseline level of damages is absent climate change. Just as a prediction market for climate outcomes can usefully project future temperatures without knowing what the baseline temperature would be without climate change, a prediction market can usefully project the economic damages from hurricanes without knowing what "normal" hurricane damages would be in the absence of climate change. If hurricane damages are projected to be very similar to current damages from hurricanes, that is valuable information.

A prediction market imparts valuable information because it communicates market expectations about future states, not why those states occur. It does not matter what the baseline level of damage is absent climate change.

An index of damages need not be a comprehensive measure of damages from climate change. To restate, I do not propose prediction markets as a substitute for other efforts to project the future; I propose prediction markets as an added input into climate policy. A prediction market could, for example, help economists develop a functional form for damage models. That itself would be an important contribution. A prediction market might just place a more credible dollar figure in front of a polity, to provide a more tangible indicator of the importance (or unimportance) of climate change in one respect (say, hurricane damages). The strength of prediction markets is in harnessing vast, disparate pieces of information. Knitting together the far-flung pieces of information to produce a number—a dollar figure—is a paradigmatic use of prediction markets.

What would a prediction market for climate-related damages look like? Just as a prediction market for climate science might project the number of hurricanes in the future, a prediction market for damages might project the amount of damages from hurricanes in the future. In fact, a paired prediction market of that sort might provide economic modelers with some guidance for developing damage functions. Imagine three related prediction markets:

- global mean temperature
- the number of hurricanes of a defined intensity
- economic damages flowing from the hurricanes

Data mined from prediction markets could help economic modelers with determining functional form. It could serve as an input into the integrated assessment modeling exercise, which is currently based on mostly educated guessing.

Data / As noted above, it is important for a prediction market for climate science to have well-defined, easily understood indices—doubly so for a prediction market for damages for events associated with climate change. The need to delineate the boundaries of damages is crucial because such an index forms the basis for market trading and predictions. Every trader must have the same conception of how damages are measured in order for the prediction market to work. For this reason, it would be necessary to have some fairly detailed criteria for the specific kinds of damages that would form the basis of the index.

In the United States, the National Oceanic and Atmospheric Administration compiles data on fatalities, injuries, property damage, and crop damage for seven categories of extreme weather events. That information is derived from a database called Storm Data, which is itself a compilation of reports from the National Weather Service's 122 field offices, aggregating economic and crop-loss damages. Those records are required to be certified and are used from time to time in court proceedings as evidence of the severity and losses of an extreme weather event. By themselves, they could serve as the basis of an index for a prediction market. But with 122 field offices, each collecting data in possibly idiosyncratic ways, data integrity could be a concern.

Alternatively, one could cobble together data from several reliable data sources to produce a single damage figure. Returning to the example of hurricanes, there are three highly destructive effects of hurricanes: flooding, high winds, and storm surge. Data on damages for those three types of costs of hurricane events come from four reliable sources:

- the Federal Emergency Management Agency's presidential disaster declaration assistance program, which delivers aid to a variety of victims of hurricanes following a presidential disaster declaration
- the FEMA National Flood Insurance Program, which insures coastal commercial and residential properties at risk of flooding
- the private consortium Insurance Services Office Property Claim Services (PCS), which tracks claims and payouts from private property insurance policies following a variety of short-term extreme weather events
- the U.S. Department of Agriculture, which administers a crop loss insurance program and covers a wide variety of extreme weather events affecting crops

All of the data are uncontroversially supplied and accepted. Of course, they could become more controversial if the value of some securities in a prediction market hinged on their magnitude. But the nature of those data, based as they are on large numbers of private transactions, enjoys some insulation from manipulation. The data are, in fact, widely regarded by industry and private actors as reliably measured and high-quality information.

This is not to say that the data are care-free. For one thing, "scaling up" that kind of data measurement to a global scale may be problematic. Also, PCS data measure only private insurance payouts, and they can be a partial function of the percentage of properties in an area that are insured. So if, over time, the percentage of uninsured properties increases, then PCS data would exhibit a downward bias in the form of payouts. To fix that, climate damage researchers Adam Smith and Richard Katz employed a weighting method in which PCS payouts were weighted by the inverse of the percentage of eligible properties insured. In other words, to obtain an accurate measure of the damages redressed by private insurance payouts, they adjusted PCS data by dividing by the fraction of eligible properties insured. For a prediction market, that kind of adjustment would be important to insure intertemporal consistency.

A prediction market for the economic effects of climate change must also confront an endogeneity problem: it is possible that prediction market activity could reflect more than simple expectations of damage trends. It is possible that prediction markets could predict some adaptation to destructive extreme weather events related to climate change. For example, a prediction market might reasonably reflect an expectation that government policy will make coastal areas more robust in the face of hurricane damages. But that kind of endogeneity could be a benefit. There is nothing to say that a damage function must capture the *pure* destructive effect of climate change, unaffected by adaptive capacities of future societies. In fact, uncertainty about the ability of future generations to adapt to climate change remains an awkward, if critically important, issue. While a prediction market for damages related to climate change might be designed to capture only the destructive capacities of an inflexible society, the more relevant question for current policymakers is how much to invest today to protect future generations. A prediction market that can help forecast the adaptive capacity of future generations should be an important input into the decisionmaking process.

There are a number of uncertainties that could be endogenized in a damages model. Will there be a successful geo-engineering effort to directly capture carbon dioxide from the ambient air? Will future crops be genetically modified to better tolerate drought? Will water management schemes be able to cope with much more sporadic and perhaps much less manageable precipitation? Will there be some benefits to climate change? Those are all questions that wrap up climate outcomes with human responses, making them particularly challenging to model. That is why integrated assessment modeling poses such gargantuan challenges. A prediction market can cut through some of the complexity, even if it poses the possibility that it might obscure some specific processes. In the end, what matters most is that humankind has an objective, if imperfect, projection of damages from certain types of destructive climatic conditions.

CONCLUSION

If one can accept that some events are related to climate change—like droughts, extreme rainfall, and hurricanes—then a prediction market could be developed for damages from those types of events. One need not accept that the events are attributable to climate change in order to accept that there could be more or less of those damages in the future. The point of a prediction market for damages related to climate change would be to have an impartial party—the market—weigh in on whether the future holds more or less or the same amount of those types of destructive events. It is not necessary to have a prediction market that accurately predicts the damages attributable to climate change; the attribution problem would be too intractable to form the basis of an asset that is actively traded in a prediction market. It is enough for a prediction market to help us decide if people really believe in an upward trend in these kinds of events. If such an upward trend were coupled with an upward trend in climate outcomes traded on a parallel prediction market, then some relationship between climate change and damages can be inferred.

How exactly do markets work? How does information travel from one market participant to another, what form does that information take, and how does it get translated into prices? Nobody knows. As economist Maureen O'Hara has quipped, "While markets appear to work in practice, we are not sure they work in theory." The advantage of having prediction markets in the realm of climate change and climate policy is that markets are at least more trusted than most sources of information bearing on climate change these days. The virtue of having markets make pronouncements about the future climate and future damages is that there can be no accusation of bias and self-dealing. Anything that can raise the level of public discussion over climate change would be an improvement. R

READINGS

- "A Prediction Market for Climate Outcomes," by Shi-Ling Hsu. *University of Colorado Law Review*, Vol. 83 (2011).
- "Climate Change Policy: What Do the Models Tell Us?" by Robert S. Pindyck. NBER Working Paper no. 19244, July 2013.
- "Energy and Climate Change: a Climate Prediction Market," by Michael Vandenbergh, Jonathan Gilligan, and Kaitlin Toner. *UCLA Law Review*, forthcoming.
- "Prediction Markets," by Justin Wolfers and Eric Zitzewitz. *Journal of Economic Perspectives*, Vol. 18 (2004).
- *Predictocracy: Market Mechanisms for Public and Private Decision Making*, by Michael Abramowicz. Yale University Press, 2008.
- "U.S. Billion-Dollar Weather and Climate Disasters: Data Sources, Trends, Accuracy, and Biases," by Adam B. Smith and Richard W. Katz. *Natural Hazards*, Vol. 67, No. 2 (2013).

You'll find:
• Stimulating articles
• Authoritative blogs
• Education news
• Audio & Video
• Searchable archives