[This document was written by Sam Marks, with help from Eric Neyman. It is a description of the 2022 ACX prediction contest, which Sam and Eric are running.]

It's becoming pretty popular to make annual predictions. Well, "popular" in the sense that everyone *I* read does it, including Scott, the Vox Future Perfect team, and Slow Boring author Matt Yglesias. At the end of the year, all of these authors grade their predictions for *calibration*. You have good calibration if about 60% of the things you assign probability 60% actually happen, 70% of the things you assign probability 70% actually happen, etc.

But calibration isn't everything! For instance, suppose that for every basketball game this year, I picked a team at random and gave a 50% probability that they would win. I would finish out the year with a near-perfect calibration, even better than 538's very calibrated model. But my predictions would have been much less *informative* than 538's: their model would have predicted higher numbers for things that were more likely to happen, and I wouldn't have.

So there is another dimension along which predictions can be good or bad, call it *informativeness*, or perhaps *accuracy*. It means that your predictions actually carry *useful information* in exactly the way that predicting 50% for everything doesn't.

One cool thing about predictive informativeness is that you can compare it from person to person, assuming that they all make predictions from a shared list of events. You do this with what's called a *scoring rule*; these gadgets give a number to each prediction, and when you add up the numbers over many predictions you get a score which reflects the overall quality of your predictions, including calibration and informativeness. (If you'd like, you can read more about scoring rules here.)

For example, Zvi re-predicted the events from Scott's 2021 predictions, and SimonM used a scoring rule to compare their predictions against each other and also against the market here. This wasn't a perfect comparison because Zvi could already see Scott's predictions and could incorporate them into his own. But what if you and Zvi had *both* independently re-predicted from Scott's list? Neither of you could compare yourself to Scott, but you could compare against *each other*.

Long story short, I think I'm better at making predictions than you, dear reader, and now I have a way to show it! Introducing the 2022 ACX predictions contest, which Scott graciously offered to host here, and which I (Sam Marks) will be running with help from Eric Neyman, author of Unexpected Values.

The way the contest works is: you are given a list of events taken from the 2022 predictions of Scott (including his re-predictions of events predicted by Vox/Yglesias); for each event, you predict the probability it happens; then in 2023 we'll grade the predictions according to a scoring rule, and you can see how you did relative to everyone else who participated in the contest.

Here are some more details:

- **The contest deadline is Monday February 14 at noon PST.**
- If you have any issues or questions, submit them using this form. **DON'T email Scott [at] hes-not-the-one-running-the-contest [dot] dont.**
- Don't know much about some topic? For each event, we give the probability predicted by Scott, and if you leave that question blank your prediction will be treated as whatever he predicted. We also give the Vox/Yglesias predictions when they exist, and you can feel free to take all these predictions into account when making your own.
- There are additional rules in the form, like no looking at prediction markets or doing more than 5 minutes of research per question, so make sure to read the rules there before starting.
- There are 74 events to predict, so be ready to put your thinking cap on.
- If you think you're gonna do great in the contest but wish this gave you more than fake internet points, consider betting in a real-money prediction market. Or if you're super psyched about getting fake internet points, consider participating in a no-money prediction website like Metaculus.
- To avoid shaking up prediction markets too much, we won't be posting the full data until at least a month after the contest deadline. That said, Scott might write about interesting things he's noticed in the data before then, we'll see. (Also, Eric and I promise not to make trades based on the contest submissions, at least not until it's all public.)

After you've read all these delicious bullet points, click here to make your predictions and enter the contest!