Free markets. Real solutions.

R STREET POLICY STUDY NO. 248
December 2021

BETTING ON CYBER: OFFERING AN ANALYTICAL FRAMEWORK FOR A CYBERSECURITY CROWD-FORECASTING PLATFORM

By Mary Brooks & Paul Rosenzweig

CONTENTS

Introduction 1
A Primer on the Crowd-Forecasting Technology Space 1
Prediction Markets 2
"Probability Polls" or "Probability Elicitation" 2
Miscellaneous 2
Examining the History of Cybersecurity Crowd-forecasting 2
Previously Asked Questions 2
Specific Challenges Associated with Crowd-forecasting Cybersecurity 3
The Question of Settlement 3
The Question of Expertise 4
Deriving Questions for a Cybersecurity Crowd-forecasting Platform 4
Simple Question 4
Index Questions 5
Conditional Questions 5
Scenario-Based Questions 5
Structuring a Cybersecurity Crowd-forecasting Platform 6
Type of Platform 6
Type of Incentive and Duration 6
Type of Participant 6
Intended Audience 7
Overcoming Moral Hazard 7
The Question of Privileged Information 7
The Question of Platform Manipulation 7
The Question of Settlement 8
Conclusion and Next Steps 8
About the Authors 8

Appendix A 9

INTRODUCTION

Our two previous posts on the Lawfare Blog addressed a series of questions. In our first post, we offered an overview of the current ecosystem of prediction markets and crowd-forecasting platforms, and asked if they might generate useful information for the benefit of government and industry decisionmakers. We concluded there that the answer is most likely yes: while nascent, the science of crowd-forecasting is growing.[1] In a world of increasing uncertainty, the opportunity to more accurately make predictions about future events, explore causality and compare expectations with reality is ever more alluring.[2]

In our second post, we explored the specific ways in which relevant trends in cyber insecurity and vulnerability might be predicted through the tools of crowd forecasting. We offered a rationale for our belief that a cybersecurity-specific platform could be valuable to cybersecurity, and explained how the types of questions asked of the platform could maximize its utility.

In this policy paper, we move from the theoretical to the practical. It is our goal to explain how a cybersecurity crowd-forecasting platform might work, and to make several key decisions to structure it. Essentially, this paper offers a series of answers to our third and final question: what would a cybersecurity-specific crowd-forecasting platform look like, and what might it tell us?

A PRIMER ON THE CROWD-FORECASTING TECHNOLOGY SPACE

Though prediction markets are perhaps the most well-known form of leveraging crowd-forecasting, there are a number of existing techniques and technologies that simi-

1. "How Crowd- Forecasting Might Decrease the Cybersecurity Knowledge Deficit," Lawfare, Dec. 6, 2021. https://www.lawfareblog.com/how-crowd-forecasting-might-decrease-cybersecurity-knowledge-deficit.

2. See Mary Brooks and Paul Rosenzweig, "Let's Bet on the Next Big Policy Crisis—No, Really," The Lawfare Blog, July 13, 2021. https://www.lawfareblog.com/lets-bet-next-big-policy-crisis-no-really.

larly rely on aggregating public wisdom. Below is a brief overview of the options available in the crowd-forecasting space.

Prediction Markets

Similar to other types of financial markets, a prediction market enables people to bet on a commodity. In this case, the commodity is a future event, and participants in the market can bet that the event will or will not occur. Prediction markets rely on betting in a fairly traditional sense, with people putting their money—or their prestige—where their mouth is.[3] Prediction markets that involve money may be classified as betting, and thus are regulated by the Commodity Futures Trading Commission (CFTC). Examples of such markets include PredictIt and Kalshi. Other popular examples, like Polymarket, Augur, or Gnosis, are decentralized and run off blockchain. However, the legality of these decentralized platforms for US-based users is unclear.[4]

"Opinion Polls" or "Probability Elicitation"

There are a variety of non-market options designed to aggregate predictions without the direct exchange of real-world currency. These models are increasingly popular, not in the least because instead of simply placing money on the "yes" or "no" outcome, participants offer more specific answers—for example, putting the odds of a given event at 70 percent—and also explaining the rationale behind their decision.

Perhaps the most widely known probability elicitation platform is Good Judgement Open, which stands as the gold-standard for probabilistic crowd-forecasting, following its success in the Intelligence Advanced Research Projects Activity's (IARPA) forecasting competition.[5] Other well-known examples include Metaculus, Hivemind (HVMD), the Foretell project within the Center for Cybersecurity and Emerging Threats at Georgetown University and the Cosmic Bazaar out of the United Kingdom.

Some of these platforms ask relatively simple questions, and some ask more complex scenario-based ones. Some involve playing for prize money, and others for points on a scoreboard. Some are closed to the public while others are broadly open to anyone who signs up. The commonality, amongst such various projects, is that real money is not exchanged on each trade and that participants can offer estimates of the likelihood of an event, rather than a straight yes or no. Beyond that, the platforms rely on different techniques and statistical methods for weighting opinions and expressing outcomes.

There is evidence that these types of probability forecasting polls elicit better results in some respects than pure prediction market alternatives, particularly when talented forecasters are aggregated and trained to be even more accurate, though some reports have found that polling is better in some circumstances, and prediction markets in others.[6]

Miscellaneous

Analogues to crowd-forecasting are not uncommon across industries. One example is collective intelligence: the idea that shared or group intelligence can emerge from collaboration, collective efforts and the competition of many individuals. This most commonly appears in areas of consensus decision making across institutions or entities.

Another analogue can be found in the insurance industry, which relies heavily on digital modeling to make predictions about future trends, run simulations, and understand cybersecurity risk.[7] Similarly, the Centers for Disease Control and Prevention (CDC) recently announced a new "Disease Forecasting Center" to help predict trends in illness.[8] Many of these efforts rely on mathematical or statistical analytics, as opposed to the more human-centric crowd-forecasting model we are recommending. Nevertheless, some augment digital methods with more qualitative crowd-forecasting methods, and vice-versa.

EXAMINING THE HISTORY OF CYBERSECURITY CROWD-FORECASTING

Previously Asked Questions

The idea of applying the techniques of crowd-forecasting to cybersecurity is not a new impulse, but a platform of the type

3. See, e.g., Emile Jacques Servan-Schreiber et al., "Prediction Markets: Does Money Matter?," *Electronic Markets* 14:3 (September 2004), pp. 243--251. https://www.researchgate.net/publication/201169069_Prediction_Markets_Does_Money_Matter.

4. See, e.g. Samuel Haig, "CFTC reportedly investigating decentralized prediction platform Polymarket," CoinTelegraph, Oct. 25, 2021. https://cointelegraph.com/news/cftc-reportedly-investigating-decentralized-prediction-platform-polymarket; Divya Taneja, "CFTC Commissioner: Code Developers May be Accountable for Smart Contracts," Proskauer, Nov. 5, 2018. https://www.blockchainandthelaw.com/2018/11/cftc-commissioner-code-developers-may-be-accountable-for-smart-contracts.

5. "The Aggregative Contingent Estimation Program: Predicting Global Events Through Crowdsourcing," CitizenScience.gov, U.S. General Services Administration, last accessed Nov. 16, 2021. https://www.citizenscience.gov/ace-forecasting/#.

6. See, e.g., Pavel D. Atanasov, et al., "Distilling the Wisdom of Crowds: Prediction Markets vs. Prediction Polls," *Management Science* 63:3 (March 2017). https://www.researchgate.net/publication/281765164_Distilling_the_Wisdom_of_Crowds_Prediction_Markets_vs_Prediction_Polls; Laurie McClellan, "Are Actuaries Superforecasting Material?" Casualty Actuarial Society, Sept. 1, 2016. https://ar.casact.org/are-actuaries-superforecasting-material; J. James Reade and Leighton Vaughan Williams, "Polls to probabilities: Comparing prediction markets and opinion polls," *International Journal of Forecasting* 35:1 (January - March 2019), pp. 336-350. https://www.sciencedirect.com/science/article/abs/pii/S0169207018300633.

7. See, e.g., Maochao Xu and Lei Hua, "Cybersecurity Insurance: Modeling and Pricing," Society of Actuaries, 2017. https://www.soa.org/globalassets/assets/files/research/projects/cybersecurity-insurance-report.pdf.

8. Centers for Disease Control and Prevention, "CDC Stands Up New Disease Forecasting Center," U.S. Department of Health & Human Services, Aug. 18, 2021. https://www.cdc.gov/media/releases/2021/p0818-disease-forecasting-center.html.

we are considering that is dedicated exclusively to cybersecurity does not appear to have been tried. A few previous crowd-forecasting efforts have had a scientific, technological or information security component to them, such as the 2014 IARPA-funded SciCast and Popular Science's Prediction Exchange (PPX). However, from our understanding these are, at best, close cognates for our work.[9]

More relevant are the individual cybersecurity questions that have been asked on existing platforms. Take, for example, the archive of questions that have been asked by Good Judgement Open. Generally, their cybersecurity-related questions have centered on asking about government attribution or sanctions in the wake of a cyberattack, the expected level of internet crime, the amount of disruption caused by a cyberattack, or the likelihood of a cyberattack—often from a given country of origin, during a given period of time either generally or against a particular entity—in addition to the more open-ended question "what questions about cyber should we ask?"[10] Of the questions asked on Good Judgment Open, at least one remains unresolved because of ambiguity over whether the event happened or not.[11]

A search of the archives of another crowd-forecasting company, Metaculus, indicates that the platform has asked four questions on cyber over the past six years, including whether Wikileaks would release more files to the Equation Group archive, a general question on whether a NATO member would invoke Article V, the number of cyberattacks Iran would undertake against U.S. government systems in the first quarter of 2020 and a question about cyber space operations.[12] Notable, cyber-attacks were specifically carved out and excluded from two other questions asking about attacks against a country's security.[13] Of course, these are just a sampling of some of the platforms in existence. It is also our understanding that cyber questions have been asked before in some of the classified intelligence platforms.

SPECIFIC CHALLENGES ASSOCIATED WITH CROWD-FORECASTING CYBERSECURITY

Prior experience provides some lessons for our project. Beyond the issues that are common to all crowd forecasting efforts, we found, unsurprisingly, that there are additional challenges unique to cybersecurity that must be addressed in any successful platform.

The Question of Settlement

First, it is difficult to adjudicate a firm conclusion on some cyber questions. This can be because of secrecy: the cybersecurity industry not only has non-disclosure agreements to attend to, but many operations are conducted in association with governments and happen at the classified or covert level. Cyber-attacks, by their nature, may be difficult to identify, attribute or assign a level of damage to. Furthermore, language used in a question may be imprecise. As used colloquially, cyber-attacks can refer to a breach for espionage purposes or to the takedown of critical physical infrastructure. In short, it is often challenging to determine what happened and thus to determine what the answer to the prediction question truly is.

9. See, e.g., Kevin Kelly, "Popular Science Prediction Exchange," The Long Now Foundation, June 18, 2007. https://blog.longnow.org/02007/06/18/popular-science-prediction-exchange; Kathryn Blackmond Laskey, et. al, "Combinatorial Prediction Markets for Fusing Information from Distributed Experts and Models," Conference: 18th International Conference on Information Fusion, July 2015. https://www.researchgate.net/publication/285232308_Combinatorial_Prediction_Markets_for_Fusing_Information_from_Distributed_Experts_and_Models.

10. See, e.g., "Before 26 December 2020, will the Australian government accuse the Chinese government, by name, of being behind the "copy-paste compromises" cyberattacks?," Good Judgment Open, Dec. 26, 2020. https://www.gjopen.com/questions/1670-before-26-december-2020-will-the-australian-government-accuse-the-chinese-government-by-name-of-being-behind-the-copy-paste-compromises-cyberattacks; "How many total complaints of suspected Internet crime will the FBI's Internet Crime Complaint Center (IC3) report for 2019?," Good Judgment Open, Jan. 1, 2020. https://www.gjopen.com/questions/1393-how-many-total-complaints-of-suspected-internet-crime-will-the-fbi-s-internet-crime-complaint-center-ic3-report-for-2019; "Between 15 November 2019 and 21 August 2020, will the New York Stock Exchange (NYSE) attribute a cross-market trading halt to a cyberattack?," Good Judgment Open, Aug. 22, 2020. https://www.gjopen.com/questions/1397-between-15-november-2019-and-21-august-2020-will-the-new-york-stock-exchange-nyse-attribute-a-cross-market-trading-halt-to-a-cyberattack; "Will North Korea execute and/or sponsor a low or higher level cyber attack against networks owned by a US entity between 19 September 2017 and 30 November 2017?," Good Judgment Open, Nov. 30, 2017. https://www.gjopen.com/questions/608-will-north-korea-execute-and-or-sponsor-a-low-or-higher-level-cyber-attack-against-networks-owned-by-a-us-entity-between-19-september-2017-and-30-november-2017; "What questions about cyber security should GJ ask in 2017 and how should we ask them?," Good Judgment Open, March 8, 2017. https://www.gjopen.com/questions/408-open-what-questions-about-cyber-security-should-gj-ask-in-2017-and-how-should-we-ask-them.

11. "Will China execute and/or sponsor a low or higher level cyber attack against networks owned by a US entity between 19 September 2017 and 30 November 2017?," Good Judgment Open, Nov. 30, 2017. https://www.gjopen.com/questions/607-will-china-execute-and-or-sponsor-a-low-or-higher-level-cyber-attack-against-networks-owned-by-a-us-entity-between-19-september-2017-and-30-november-2017.

12. Will Wikileaks release a significant augmentation to the Equation Group cyberespionage archive?," Metaculus, Sept. 1, 2016. https://www.metaculus.com/questions/316/will-wikileaks-release-a-significant-augmentation-to-the-equation-group-cyberespionage-archive; "Will any NATO member invoke Article 4 or Article 5 before 8 September 2018?," Metaculus, Sept. 9, 2018. https://www.metaculus.com/questions/710/will-any-nato-member-invoke-article-4-or-article-5-before-8-september-2018; "How many cyberattacks by Iran against US Govt. systems in Q1 2020?," Metaculus, Oct. 11, 2020. https://www.metaculus.com/questions/3453/how-many-cyberattacks-by-iran-against-us-govt-systems-in-q1-2020; "ASAT Weapons Tests and Space Debris by 2023," Metaculus, last accessed Nov. 16, 2021. https://www.metaculus.com/questions/7644/asat-weapons-tests-and-space-debris-by-2023.

13. "Will Iran execute or be targeted in a national military attack between 6 June 2019 and 5 October 2019?," Metaculus, August 25, 2019. https://www.metaculus.com/questions/2926/will-iran-execute-or-be-targeted-in-a-national-military-attack-between-6-june-2019-and-5-october-2019.

An example will perhaps make the challenge of settlement easier to understand. Let's say the question "will the Chinese government launch a cyberattack against a U.S. government system in calendar year 2015" was asked on January 1, 2015.

Famously, in June 2015, it was revealed that there had been a massive breach of security clearance files stored by the Office of Personnel Management, on the Department of the Interior's servers. This was a significant hit to U.S. security. But would it satisfy our question, above?

The answer is more complicated than it perhaps first appears.

Take the phrase "Chinese government." Attribution is difficult to confirm and often overtly political. If the attack originated in China, was it undertaken by the government? Or were they criminals? Or perhaps government workers moonlighting with the skills learned in their day jobs?

And whose attribution matters? The Obama administration never officially attributed the breach to the Chinese, though then-Director of National Intelligence James Clapper said publicly "you have to kind of salute the Chinese for what they did [in OPM]."14 Three years later, in a press conference, then National Security Advisor John Bolton directly blamed China.[15]

Similarly, take "cyberattack"—is espionage an attack? Does there need to be physical damage to a system or network to consider it as an "attack," or is it simply, "malicious cyber activity?"

Finally, take "in calendar year 2015." The breach of OPM's servers was discovered in April 2015 and revealed to the public in June 2015. But forensic evidence of the breach indicates that the original intrusion may have occurred in 2014 or even earlier.[16] So, the attack was "launched" before 2015.

Despite the obvious challenges here, we do not believe these issues are insurmountable. There may be some questions that are simply unresolvable. Others may need to be reopened or nullified following the discovery of new information—or at the very least, delayed in adjudication for some time. But some of the problems can be addressed by speaking with greater precision. This challenge makes it clear that there is a critical need to carefully structure our questions.

The Question of Expertise

Crowd-forecasting platforms frequently ask questions that, while difficult to predict, are not inherently niche or challenging to work out. For example, while a former pro football player turned sports analyst may be particularly well-positioned to predict who will win the next Super Bowl, it is highly plausible that an outside fan of the sport will guess just as well, because there only so much information that impacts a football game, and much of it is publicly accessible.

Our observation is that when cyber questions have previously been publicly asked, relatively few people have offered an opinion. This could be for any number of reasons, but it may be because there are fewer people comfortable with answering the question, fewer people interested in the issue or high levels of secrecy and classification. This lack of response will, we suspect, be especially true of questions that are important to the technical side of the domain, whereas geopolitical questions may be more accessible.

It might seem intuitive that the more experts you have on an issue, the better your outcomes. This is not always true.

For example, in a true prediction market, you need a mix of skillful and unskillful bettors in order for the market to work. If all bettors are always right, there will be little profit in being right. You need people who can guess wrongly more often in order to pay out the people who guess correctly more often.

This is less true of a non-market crowd-forecasting platform. One thing that is important and at the heart of crowd-forecasting is that you need enough people to nullify the wrong answers of the so-called "idiot respondents"—i.e. the people that are widely off. As such, the platform may benefit from having only experts—indeed, that is the very purpose of the Good Judgement project: to identify talented "superforecasters" and then to train them to make them even better.

Of course, "expert" does not necessarily mean a skillful bettor. Part of the premise of crowd forecasting is that so-called experts are often wrong—they may rely on group-think, be trapped in older paradigms, or rely on information or theories that are accepted in their industry but that are flawed. If this is the case, there is a reason to believe that including non-elites and non-experts would be advantageous.

DERIVING QUESTIONS FOR A CYBERSECURITY CROWD-FORECASTING PLATFORM

When using a crowd-forecasting platform the questions we ask and their associated outcomes must be 1) clearly-defined, 2) circumscribed/mutually exclusive, 3) time-constrained and 4) knowable (that is, capable of being resolved). However, questions that meet these criteria can be structured very differently, depending on the information we wish to elicit.

Simple Questions

Simple or binary questions are questions that stand alone and have a yes or no answer. For example, will the Russia government criminally charge a Russian-based non-state ransomware operator for launching ransomware operations outside of Russia on or before December 31, 2022?

14. David Welna, "In Data Breach, Reluctance to Point the Finger at China," *National Public Radio*, July 2, 2015. https://www.npr.org/sections/parallels/2015/07/02/419458637/in-data-breach-reluctance-to-point-the-finger-at-china.

15. "National Security Adviser John Bolton on Cyber Strategy (AUDIO ONLY)," C-SPAN, Sept. 20, 2018. https://www.c-span.org/video/?451807-1/national-security-adviser-bolton-briefs-cyber-strategy-audio-only.

16. "The OPM Data Breach: How the Government Jeopardized Our National Security for More than a Generation," House Committee on Oversight and Government Reform, Sept. 7, 2016. https://republicans-oversight.house.gov/report/opm-data-breach-government-jeopardized-national-security-generation.

Index Questions

Index or spread questions take the form of multiple-choice questions. For example, when will Congress pass the 2023 National Defense Authorization Bill: before October 31, 2022; between November 1 and December 31, 2022; on or after January 1, 2023? Or never?

Conditional Questions

Conditional questions are traditional if-then questions, designed to determine causality. These types of questions introduce some complications, and appear to be best answered in two parts, with the first question working as a simple question and the second as a conditional question.

A simple question looks like this: will a cryptocurrency operator or employee be charged or fined by the United States for facilitating a ransomware payment to a country subject to OFAC sanctions in 2022?

A conditional question would state: if a cryptocurrency operator or employee is charged or fined for facilitating a ransomware payment to a country subject to OFAC sanctions in 2022, will the total imputed amount of ransomware payments decrease in the six months following the public announcement as compared to the six months prior to the announcement?

There are risks with conditional questions, the most significant of which is that if your premise does not occur, then you have no conditional second question to ask and the question becomes defunct. However, if the premise does occur, it can be a valuable set-up: rather than requiring the question-maker to impute causality on their own, you can ask the platform its opinion.

Scenario-Based Questions

Scenario-based questions are questions that build upon each other to produce trend lines and offer insight into significantly more ambiguous and open-ended questions of the type that policymakers may be most interested in.

If an investor is simply treating the crowd-forecasting tool as a hedging mechanism to offset any losses in the event of a cyberattack, then a straightforward question asked in a public prediction market is likely to suit their needs. By contrast, if a policymaker is looking to take advantage of crowd-forecasting capacity, they're likely interested in questions that are much less straightforward and resolvable by the terms of a binary question in a prediction market.[17]

Georgetown's Foretell program provides some excellent examples of the latter sort of scenario questions. Their methodology identifies a research scenario that asks a specific question about what the world will look like in a few years.[18] Researchers work backwards from that query to define "predictors" of possible scenarios and identify the metrics that drive the predictors. The metrics questions are recognizable as questions of the sort that more traditional prediction markets have historically asked.[19]

Consider this scenario outline as a way of illustrating how this process might work if modified, simplified and applied to the type of questions we may want to ask:

Predictor One: Will Putin shift his attitudes toward cyber-crime emanating from Russian territory in a way that reduces the ability of cyber criminals to operate?

Metrics for Predictor One:

- Will at least one cyber-criminal in Russia face legal penalties directly connected to digital intrusions into non-Cyrillic speaking countries before December 31, 2022?
- Will Putin verbally promise to XYZ (where XYZ is some possible anti-ransomware action, such as cooperating with American law enforcement)?
    - » If Putin makes promise XYZ will he take action that authoritative outside observers (in the case above, the U.S. government) acknowledge publicly as having satisfied the verbal promise?

Predictor Two: Will the U.S. government and private industry endeavor to harden their cybersecurity posture in 2022?

Metrics for Predictor Two:

- Will at least one major American insurance company (defined by its market capitalization) decide to stop binding coverage for ransomware payments before December 31, 2022?
- Will the Federal government enact a know-your-customer law for cryptocurrencies before July 31, 2022?
- Will the government add money (in excess of $ABC million) for federal anti-ransomware software to be deployed on federal systems in the 2023 NDAA?
- Will the government mandate compliance with XX security measures by YY date?
- Will the Department of Defense pass a final rule to implement the Cybersecurity Maturity Model Certification by January 31, 2022?

17. Adam Siegel, "Tracking the Outcome of Strategic Questions with Crowd Forecasting," Cultivate Labs, Oct. 7, 2021. https://www.cultivatelabs.com/posts/tracking-the-outcome-of-strategic-questions-with-crowd-forecasting.

18. Michael Page et al., "Future Indices: How Crowdforecasting Can Inform the Big Picture," Center for Security and Emerging Technology at Georgetown University, Oct. 19, 2020. https://cset.georgetown.edu/publication/future-indices.

19. "Issue Campaign: What is the future of the DoD-Silicon Valley relationship?," Center for Security and Emerging Technology at Georgetown University, last accessed Nov. 16, 2021. https://www.cset-foretell.com/issue-campaigns-dod-sv.

These are, of course, examples of the types of questions we might ask on a crowd-forecasting platform. As part of our research we have surveyed a host of experts to solicit their suggestions as to additional questions that might be asked. A summary of some of these questions can be found in Appendix A.

STRUCTURING A CYBERSECURITY CROWD-FORECASTING PLATFORM

With that background in mind, we now turn to the meat of the problem: what, precisely, should a cybersecurity crowd-forecasting platform look like? In this section we describe the mechanics of how we would propose to create a separate, specialized cybersecurity crowd-forecasting platform.

Type of Platform

Intention: We believe it would be best to create a tool that relies on probability elicitation from experts and semi-experts, rather than on a traditional prediction market.

Explanation: Traditional real-money prediction markets are overseen by the CFTC and are subject to a relatively extensive approval process. They require significant cash reserves to preserve liquidity. An alternative, of course, is a smaller-scale prediction market that caps the amount of money that an individual can wager, such as PredictIt.

However, there remain several other reasons why non-money probability elicitation would likely work better for cybersecurity prediction purposes. For one thing, it may be easier to get a broad range of participants to join, whereas a prediction market would require the market to either draw from those with the willingness to put their own money in the market, or for us to subsidize the market with external funds. Secondly, relying on probability elicitation would also enable participants to answer the question of why they believe a certain event is or is not likely, rather than simply requiring them to offer a straight yes or no response.

Thus, both for ease of implementation and for possibly superior results we have concluded that it would be reasonable to begin research with a probability elicitation platform.

Next Steps: There are a number of companies that offer crowd-forecasting software as a service. These companies manage the mechanics of the crowd-forecasting effort and offer optional support for question-writing and participant-recruiting. We intend to solicit their services.

Type of Incentive and Duration

Intention: Create an elicitation platform that is interval-defined for recurring time periods in which participants will be graded within a ranking system and earn cash prizes.

Explanation: In a real-money prediction market, the incentive to play and to do well is baked into the system. Winners make money and losers lose money. But non-real-money crowd-forecasting options also have incentives. Some play for prestige. Some play for cash prizes or in-kind prizes. We propose to combine both.

Note that the incentive cannot only be a token prize, but should be substantial enough to offset the amount of time and effort that participants will be expected to invest in the project.

Next Steps: In this structure it seems necessary to offer a monetary reward at specific intervals. For example, the platform could offer cash prizes to the top three scorers every six months, at which time scores will reset. There could also be further incentivization for more detailed participation in a particular area by creating a specific subset of questions with its own pot of money. A public leaderboard updated either in real-time or on a periodic basis would also work to show who is top-performing.

Type of Participant

Intention: Cultivate an invited group of participants with a broad range of expertise in cyber and information security as well as in surrounding areas of expertise such as geopolitics, finance, sociology, and legal and regulatory systems.

Explanation: Because cybersecurity is interdisciplinary, it seems best to consult a wide range of individuals, including government officials, cyber technical experts, geopolitical researchers, financial analysts, CISOs, marketers, industrial control systems operators, IT professionals and more. These individuals should have some experience with information systems or cybersecurity, but not all share the same view on a given issue.

The target users of the crowd-forecasting tool will be best cultivated purely through an invitation and referrals participant group (at least initially). We would need this group to be sufficient in number to generate aggregate knowledge, though there does not appear to be a specific minimum number of participants necessary.

Notably, we propose to allow for anonymity on the public-facing side of the platform. Users can self-identify by their real names or under online handles. However, an invitation and referral only program will necessarily mean that operators of the platform know who is playing in the platform in the aggregate, even if there is not a database that connects real names to login emails and online handles. Similarly, any

individuals receiving a monetary prize would need to file taxes—although refusing prize money is also an option.

Next Steps: Cyber and information security is a robust and highly-networked industry. There are existing groups of potential participants we could recruit from. In the interest of creating a slightly more controlled environment for beta testing, the platform would be offered to professionals associated with known groups (professionals associations, networking groups, loose affinity groups, etc.) and their referrals, with the expectation that it could be expanded at a later time.

Intended Audience

Intention: Our goal is to ask various question types (binary, index, conditional and modified scenarios) that cover a wide range of topics (technical/incident, industry and geopolitical/policy) in order to generate information that could be applicable to decisionmakers both inside and outside of the government.

Explanation: Our original intent, when proposing this type of information-generation and aggregation platform, was to offer it specifically to policymakers and government decisionmakers. However, in private discussions with others who have tried to champion crowd-forecasting and prediction markets for government decisionmakers, it became clear that many years-long efforts to create cyber or political platforms within the government were not always well-supported—and suffered from challenges like insufficient buy-in at a high level and insufficient funding.

And in recent discussions with current government officials and those who work with them, there is a sense of skepticism that seems borne of information overload. Several are not convinced that more information will be beneficial—though the intent of the platform is actually to decrease and diminish noise. Others express uncertainty around when the information might be taken up into the policy-process. As has been the case for the past two decades, the lion's share of interest remains with more experimental and innovative sectors of the government such as the intelligence community, where the National Intelligence Council and Central Intelligence Agency are already both at work.[20]

To some degree, we are optimistic that if we build a successful crowd-forecasting model the government will eventually accept it. Yet, we also recognize the intrinsic institutional lethargy of government. With this in mind, we primarily aim to demonstrate the efficaciousness of a crowd-forecasting model, and that suggests that we need to widen the scope of our intended audience.

Next Steps: Offer a prediction platform that focuses on accuracy and reliability with broad-based insights on a wide range of interests connected to cybersecurity. Identify interested parties in the private and public sectors to introduce the concept and, eventually, expand to include offerings of interest to government policy makers.

OVERCOMING MORAL HAZARD

In this final section, we look at some of the problems of moral hazard, which might attend this work. As we make clear, while we are cognizant of these possibilities, we believe that they can be mitigated.

The Question of Privileged Information

Elicitation platforms are designed to reveal and allow the discovery of collective aggregated knowledge. This offers an advantage to those with additional or insider knowledge. Unlike in financial markets, where taking advantage of privileged information is often illegal, this may be a boon to the platform. Our interest, after all, is consolidating good information regardless of the source.

However, in the context of our proposed cybersecurity platform, this differential privileged knowledge could become a problem if some participants have access to classified information. By definition, our crowd-forecasting platform would be unclassified. This is advantageous as it offers an opportunity to incorporate robust open source intelligence. However, we will need to be cautious in selecting questions and participants to avoid offering the impression that classified information could be reverse-engineered by analyzing platform outcomes.

The Question of Platform Manipulation

The ability to manipulate the market or platform is of course not unique to crowd-forecasting. Consider, for example, the GameStop stock surge.[21] However, if the intent of the platform is to generate accurate data fit for decision making, we have a higher standard to meet, because manipulation could impugn the credibility of the platform. If someone suspected that decisionmakers were relying purely on a platform for decision making (something that should never occur, as crowd-forecasting should be considered only one tool among many), they might have an incentive to attempt to shift the results in a certain direction.

20. Perry World House, "How to See the Future: Forecasting and Global Policy," University of Pennsylvania, Sept. 28, 2021. https://global.upenn.edu/perryworldhouse/how-see-future-forecasting-and-global-policy.

21. Chris Young, "GameStop's Reddit-Fueled Stock Surge Explained in Simple Terms," Interesting Engineering, Jan. 28, 2021. https://interestingengineering.com/gamestops-reddit-fueled-stock-surge-explained-in-simple-terms.

Indeed, the manipulation risk goes both ways: while people could manipulate the platform, the platform could also manipulate people. For example, a question could name an individual person or company, e.g. will XYZ software be successfully exploited by a ransomware attack in 2022? If platform participants believe the likelihood is high, then criminal actors may focus efforts on that software specifically, out of the belief that participants have knowledge of software insecurity. This seems to invite legal or at least ethical consequences.

Part of this risk will be mitigated by a carefully-selected participant body. It may also be mitigated by a prize/prestige incentive, rather than a traditional betting market. One final measure is structuring the questions to minimize the threat of incentivizing compromise or inviting retributive litigation—for instance, not publicly questioning the vulnerability of an individual company's software.

The Question of Settlement

A political outcome is obvious. For example, if the Associated Press (AP) calls an election for President Joe Biden, we pay off on contracts that say Joe Biden won the 2020 election. But there are occasions in which an unexpected third outcome might occur, when the reports are mixed, or when there is a delay or reversal in external reporting or evidence. An excellent example would be if the AP called the election for former President Donald J. Trump, realized it was mistaken, and reversed it. This is not an impossible scenario—as the iconic 1948 *Chicago Tribune* headline "Dewey Defeats Truman" attests.[22]

The same uncertainty of settlement may occur in our cyber platform. Indeed, it may be even more prevalent because, as we have already noted, some settlement questions might prove impossible to resolve. There are two ways we propose to handle this.

The first option is delayed adjudication. For questions where it is appropriate, we could impose a freeze option on the resolution when they hit the deadline. Thereafter we would make a final determination some time (hypothetically we can say six months) after the close of the contract. This would slow down the pace of payouts and awards, but ultimately would support the integrity of the platform and would minimize the number of questions that had to be retroactively re-determined. We also recognize that this is an experiment and that it is possible that some questions will simply have to be canceled and ruled null.

The second option we are considering is an arbitration mechanism. While we do not anticipate needing this, it may become necessary if a significant percentage of questions run into settlement issues. In this case, the settlement decision—or decision to nullify the question—could be determined by a judge or panel jury. To be candid, if this is our result, then our beta test of a platform might wind up proving that cyber-specific crowd forecasting is too difficult to achieve.

CONCLUSION

We have identified a knowledge gap in the field of cybersecurity, and posited that crowd-forecasting methods may help us close this gap. We have also explored how to formulate questions that could move the needle on cyber issues and offered a conceptual proposal for creating a cyber-specific crowd forecasting platform. Finally, we have laid out some key guidelines and considerations for creating a platform of this sort.

There are many reasons to be optimistic about the potential of cyber crowd-forecasting—and it would appear that several leaders within the Biden-Harris administration are similarly optimistic. A recent policy report co-chaired by now-Director of National Intelligence Avril Haines recommended that the intelligence community (IC) embrace and "leverage the wisdom of crowds," noting that the long-running IC prediction market has been found to be "more accurate than traditional analytic methods."[23]

A brave new world of prediction forecasting for cybersecurity lies ahead. Only time will tell if it works. For now, we embrace the possibility.

ABOUT THE AUTHORS

Mary Brooks is a resident fellow for Cybersecurity and Emerging Threats at R Street Institute.

Paul Rosenzweig is the founder of a homeland security consulting company, Red Branch Consulting, and a senior advisor to The Chertoff Group. Rosenzweig formerly served as deputy assistant secretary for policy in the Department of Homeland Security.

22. Tim Jones, "Dewey defeats Truman: The most famous wrong call in electoral history," *The Chicago Tribune*. Oct. 31, 2020. https://www.chicagotribune.com/featured/sns-dewey-defeats-truman-1942-20201031-5kkw5lpdavejpf4mx5k2pr7trm-story.html.

23. Technology and Intelligence Task Force, "Maintaining the Intelligence Edge: Reimagining and Reinventing Intelligence through Innovation," Center for Strategic and International Studies, January 2021, p. 16. https://csis-website-prod.s3.amazonaws.com/s3fs-public/publication/210113_Intelligence_Edge.pdf.

APPENDIX A

In the spirit of crowd-forecasting, we asked several private groups of cyber experts a very broad question: whether they were aware of questions that, if answered in a way that we could generally rely on (that is, that are at least somewhat correct), could reasonably help us improve cybersecurity. Note that some of the questions have been edited and would in many cases require further specificity in order to be viable on a crowd-forecasting platform.

Geopolitics / Government	Will the U.S. government require software manufacturers to provide warranties by June 30, 2022? Will the US implement SOX-style executive accountability for corporate cybersecurity before December 31, 2024? Will ransomware payments become illegal by December 31, 2022? Will Congress pass a data security and/or data privacy law before June 30, 2022? Will there be U.S. criminal charges threatened or filed against a security researcher for overstepping boundaries about pen-testing or the Digital Millennium Copyright Act (DMCA) before December 31, 2022? Will we see a "gift box" of information on REvil accessed by government authorities before December 31, 2021?
Industry / Technology	Will a [US] insurance company announce a policy of refusing to pay ransomware-related losses before December 31, 2022? Will we see a public report of cyber vendor's key researchers being *physically* targeted by bad actors/ criminal organizations by December 31, 2022? Will there be a public doxxing of a ransomware crew in the next 12 months?
Technical / Incident	Will the volume / effect of distributed denial of service (DDoS) attacks on American networks exceed the available ISP / MSP capacity to stop them—resulting in disruption of service to XX thousand or more customers on YY or more occasions in 2022? Will a major cloud service go down (in the sense of failing to service more than XX percentage of its customers for longer than eight hours) more than once in the next year? Will there be a publicly reported incident involving ransomware on a space-based asset in the next 12 months? Will there be a "fire-sale" event [massive cyberattack crippling transportation, finance, and utilities) affecting a major city or country?