Lancaster University
Management School

UNIVERSITY OF EXETER | LAND, ENVIRONMENT, ECONOMICS AND POLICY INSTITUTE

CRUCIAL

Climate Risk and Uncertainty Collective Intelligence Aggregation Laboratory

For more information contact Mark Roulston

m.roulston@exeter.ac.uk

July 2022

CONTENTS

EXECUTIVE SUMMARY

What?

Information markets use market-like mechanisms to elicit and aggregate information. Information markets are not designed to transfer ownership or risk but to exploit the proven ability of markets to discover and synthesize information.

Information markets can solve the problems of aggregating diverse climate expertise *and* aligning the incentives of providers. They could become the primary mechanism for obtaining forward-looking expertise about climate-related risks.

Why?

Climate change is recognised as an important source of risk. Regulators now believe that climate-related risks pose, not only a threat to individual firms, but possibly to the stability of the financial system.

In 2017 the G20's Task Force on Climate-Related Financial Disclosures (TCFD) published recommendations that firms should disclose forward-looking information about physical and transition climate-related risks. More than two thousand companies have voluntarily signed-up to these recommendations. Regulators and exchanges in many countries plan to make climate-related disclosures mandatory. These trends have created a demand for forward-looking information about climate, such as emissions pathways and changes in physical climate risks. Companies requiring this information are a nascent source of funding for applied climate research. They are particularly interested in the efficiency with which distributed information can be aggregated and communicated for use in decision making and risk management.

The marketplace for climate information is diverse and fragmented. Relevant expertise exists in public agencies, academia, think tanks, as well as an emerging class of private companies known as climate service providers (CSPs). Given the long horizons involved, few providers have established track records, so it is difficult for users to evaluate their skill. Furthermore, most forecasts are offered with no warranty. This creates a situation in which bad forecasts can drive out good ones; a situation reminiscent of credit ratings for mortgage-backed securities in the run-up to the subprime crisis.

How?

We propose a climate information exchange that hosts markets predicting climate-related variables. These variables will include high-level ones, such as global greenhouse gas concentrations, and more granular ones like sea-level rise in specific locations. Market participants will include individuals and institutions with expertise across the range of disciplines relevant to climate forecasting, including meteorology, climatology, economics, social sciences, and policy. Participants will use on-platform credits representing a claim on a proportion of research funds distributed through the exchange. These funds will come from public, private, and philanthropic sources interested in effective altruism.

The exchange will use AGORA, developed by Hivemind Technologies Ltd. AGORA has been used to run expert information markets for seasonal forecasts, El Niño events, Atlantic hurricane activity, and crop yields. AGORA solves several problems that have afflicted previous attempts to create information markets for climate: an automated market maker guarantees liquidity allowing markets to function in the absence of uninformed "noise" traders. AGORA supports very large outcome spaces, enabling consensus probability distributions to be extracted, not just probabilities for binary events.

Where?

Many of the markets will have horizons of more than a decade. The consensus forecasts generated will be a public good, to be distributed for free. There is a strong case for the exchange to be hosted by a university, which will become a hub for international climate expertise.

TEAM

MARK ROULSTON

Mark led the development of the AGORA prediction market platform while at investment firm Winton Group, where he worked for a decade, and at Hivemind, a technology company spun-out from Winton in 2018. He has a Ph.D in planetary science from Caltech for a thesis on the predictability of El Niño and he continued research on climate predictability at Oxford and Pennsylvania State Universities before working at the U.K. Met Office, prior to joining Winton. At Winton he led the integration of weather and climate information into quantitative trading strategies.

KIM KAIVANTO

Kim Kaivanto is a Senior Lecturer in Economics and the Director of the M.Sc in Money, Banking and Finance at Lancaster University’s Management School (LUMS). Before joining LUMS he held fellowships at the Eitan Berglas School of Economics and Warwick Business School, from where he received his Ph.D. Kim's research interests are theoretical and descriptive models of decision making and behaviour under risk and uncertainty. He has applied his expertise to problems such investor sentiment, security behaviour, civil aerospace R&D support schemes, aviation slot allocation and CO_2 emissions, venture capital, the exploitation of social science research. He has also advised the banking sector on climate risk exposure.

TODD KAPLAN

Todd is a part-time Professor specializing in economic theory and behavioural economics at Exeter University. He also has a position in the Economics Department at the University of Haifa. He has received grants from the Nuffield Foundation, British Academy, iFree Foundation, ESRC, Leverhulme Foundation and the Israeli Science Foundation. His recent research concerns manipulation in prediction markets. He is an Associate Editor of the Journal of Behavioural and Experimental Economics. He received his Ph.D in economics from the University of Minnesota.

BRETT DAY

Brett is an economist working on environmental decision support. He has a Ph.D in economics from University College London and has held faculty positions in the School of Environmental Sciences at the University of East Anglia and the Department of Politics at Exeter University. He is currently a Professor of Economics and Co-Director of the Land, Environment, Economics and Policy Institute at Exeter where he maintains close links with government and business, applying economic methods to environmental management, such as designing real-world “Payment for Ecosystem Service” mechanisms.

1. THE CHALLENGE OF CLIMATE CHANGE

1.1 BACKGROUND

Although climate has changed throughout history, and some of the changes the Earth is currently experiencing have natural causes, most evidence indicates that the dominant cause of rising global temperatures since the industrial revolution is the increase in atmospheric greenhouse gases (GHGs), particularly carbon dioxide emitted by the burning of fossil fuels. Greenhouse gases affect the radiative-transfer properties of the atmosphere, raising the average surface temperature.

1.2 CLIMATE FORECASTING

While there is uncertainty in predictions of the global average temperature, this uncertainty is even larger for more specific predictions, such as regional temperatures and precipitation, and it is larger still when, for example, forecasting future rainfall in the catchment of a particular hydroelectric dam. Furthermore, all predictions are conditional on the future levels of greenhouse gases, which depend on future net emissions, which are themselves uncertain.

Forecasting climate change, given a particular pathway of greenhouse gas emissions, is done using climate models that simulate the behaviour of the ocean-atmosphere system. Building these models requires knowledge of many fields including meteorology, oceanography, chemistry, and even biology. Predicting the likelihood of a particular emissions pathway requires insights into the development and uptake of new technologies as well as possible policy choices across the global economy. The knowledge and expertise required to predict future climate is thus distributed over many disciplines and beyond the capacity of any single research group or organisation.

The Intergovernmental Panel on Climate Change (IPCC) strives to provide an objective consensus view of the science of climate change. IPCC assessment reports, however, are only published once every few years and cannot provide projections of every variable of relevance to decision makers. IPCC projections have been conditioned on Representative Concentration Pathways (RCPs). There is concern that the IPCC's RCP8.5 scenario – the worst considered – has been misconstrued as "business as usual" (Pielke and Ritchie 2020) and that it is "increasingly implausible" (Hausfather and Peters 2020) while others think it is still relevant (Christensen et al. 2018, Schwalm et al. 2020).

Since 1990, the IPCC has published six Assessment Reports. These reports are supported by Coupled Model Intercomparison Projects which are collaborative efforts involving dozens of climate modelling groups around the world. Thanks to these efforts many advances have been made in climate research but mechanisms for synthesising this knowledge into consensus forecasts that are coherent and up to date have lagged—there is an interval of around 5 years between Assessment Reports. Yet decisions affected by future climate change are being made in the present. Those responsible for these decisions, whether they concern mitigation or adaptation, need information that reflects the current consensus views of a broad range of experts.

1.3 VALUE OF FORECASTING

If the more pessimistic predictions of future climate change are correct, civilization faces disruption on a scale not experienced for centuries. Estimating the economic impact of climate change is even more challenging than predicting climate change but most assessments suggest that, globally, investments of $1 to 5 trillion per year might be needed to pay for the mixture of adaptation and mitigation measures required. This is about 1 to 5% of global GDP (Goldman Sachs 2020, Friedmann et al. 2020, Gates 2021).

Forecasting alone cannot solve the problem of a changing climate but it can improve the allocation of resources between different mitigation strategies. For example, predictions of the path of global temperature are a key factor in deciding the time scale on which to decarbonise the economy. Good forecasts can also help optimise adaptation: for example, knowing how much sea-level will rise over the next 50 years may determine whether it is better for a coastal community to improve flood defences or relocate; with an inaccurate forecast they might end up paying for both. Determining the extent to which better climate forecasts can improve resource allocation is difficult but even a 5% improvement in allocation implies a societal pay off to improved forecasts of between $50 and $250 billion per year.

How much should society pay for climate predictions? Society, of course, is not a monolithic entity. Society's spending for climate prediction comes from governments, private companies, and non-profits, with the U.S. leading in all three categories. The U.S. Global Change Research Program budget – an aggregation of research funding for climate change research by ten federal agencies – was $2.4 billion 2019.[1] Assuming, optimistically, that other governments commit a similar proportion of GDP implies that governments spend up to $10 billion per year on climate change research, including forecasting. In the private sector, the global climate change consulting market is estimated to be $5.5 billion, and it is projected to grow to $8 billion by 2026.[2] Of the $730 billion in philanthropic giving in 2019 only 2 per cent was for climate change, and this was mainly directed at mitigation,

[1] https://www.globalchange.gov/about/budget
[2] https://www.coherentmarketinsights.com/market-insight/climate-change-consulting-market-2537

making charitable giving a minor source of funding for climate forecasting, behind governments and the private sector.[3]

1.4 UNDERINVESTMENT IN FORECASTING

When we compare the amount spent on climate forecasting with how much accurate forecasts might save by allowing better allocation of resources it is hard to avoid the conclusion that it is sub-optimal. As well as the total amount being too low, what is spent is arguably misallocated, with too much relative emphasis being placed on physical sciences compared to social sciences which provide insight into the likelihood of achieving emissions reductions and hence the probability of different climate scenarios occurring (Overland and Sovacool 2020).

Climate forecasts produced by governments are a classic public good: non-rival and non-excludable. The tendency for public goods to be under-provided is well documented. It is typically attributed to the absence of a functioning market that allows the monetisation of their supply.

Private sector provision of climate forecasts is vulnerable to another cause of market failure: *asymmetric information*. Companies that wish to buy climate forecasts from commercial firms have no way of knowing how good the forecasts are likely to be. Unlike short-range weather forecasting, for which users can ascertain forecast skill within a practical time frame, climate forecasting horizons stretch to decades. Buyers of climate forecasts will typically lack the expertise to evaluate the forecasting process itself and so will have no information with which to judge forecast quality. Economist George Akerlof famously explained how, in this type of situation, prices can be suppressed, leading to bad products driving out good products (Akerlof 1970).

The information asymmetry can create a moral hazard in which the information provider neglects rigour in a way that they would not if they were exposed to the risks posed by inaccurate information. Furthermore, consumers of climate information may actively seek forecasts which exaggerate risks or underplay them, as issuers of mortgage-backed securities did in the early 2000s (see box on next page). Even if providers genuinely strive to produce unbiased forecasts, if there is diversity in forecasts available and users can select their forecast this will introduce bias into their risk assessments and disclosures (Skreta et al. 2009).

Providers of climate information products that generate an uncertain value stream over a period of decades can only expect to get paid upfront and unconditionally if they offer large discounts from the expected value of the information. This depresses the returns to investment and innovation in the sector.

[3] https://www.climateworks.org/press-release/funding-trends-climate-change-mitigation-philanthropy/

Addressing structural reasons for underinvestment in climate forecasting should be a priority, as it has the potential to increase the efficiency with which resources are deployed to tackle mitigation and adaptation.

BAD INCENTIVES, BAD INFORMATION: AN EXAMPLE

Credit ratings agencies appeared in the 19th century to provide information about the credit worthiness of railway companies that were proliferating at the time. In the aftermath of bank failures during the 1930s the U.S. Government introduced regulations that meant banks could only hold investment-grade bonds, where "investment-grade" was to be defined by the ratings agencies. Therefore, in a manner not dissimilar to climate risks, regulation helped create a demand for information.

Until the 1970s ratings agencies, such as Standard and Poor, Moody's, and Fitch, produced assessments of the credit risk of bonds which they published in periodicals to which investors subscribed. The investors wanted accurate ratings and, since they were the customer, the agencies had an incentive to provide accurate ratings. The invention of the photocopier made it difficult to sustain this business model, so the ratings agencies switched to charging the issuers of bonds for the ratings that allowed institutional investors to buy their bonds. This created a conflict of interest because the issuers were more interested in getting high ratings than ratings which accurately reflected the risk of default. It has been argued that this conflict of interest was exacerbated when Moody's was spun-out of Dun & Bradstreet in 2000, as executives received stock options making them more concerned about increasing the firm's market share (McLean and Nocera 2010).

With the advent of mortgage securitisation, the issuers of mortgage-backed securities sought triple-A ratings and engaged in "ratings shopping" and pressuring the agencies to get them. Between 2000 and 2007, tens of thousands of mortgage-backed securities were rated, most receiving triple-A ratings even when they contained substantial numbers of subprime loans. When the subprime crisis hit, most of these investment-grade securities were downgraded to junk.

Subsequent investigations partly attributed the subprime crisis, which precipitated the Global Financial Crisis and Great Recession, to the mis-aligned incentives of the credit ratings agencies. A Congressional committee found that, *"The major credit rating agencies played an important – and perhaps decisive – role in enabling (and validating) much of the behaviour and decision making that now appears to have put the broader financial system at risk."* (Congressional Oversight Panel 2009).

This example illustrates the importance of ensuring that the providers of information have the correct incentives for it to be accurate.

2. THE DEMAND FOR CLIMATE FORECASTS

2.1 DECISION MAKING

While the impacts of climate change will play out over decades, there are choices being made today for which decision makers must take a view on future climate. Infrastructure can have a lifespan of decades, so decisions about siting and specifications must take climate into consideration, either explicitly or implicitly.

Sea-level rise is a good example: it is relevant to infrastructure planning yet there is substantial uncertainty among specialists over how much sea-level rise we should expect *for a given amount of warming*; the main source of uncertainty is the contribution of melting ice sheets (see Fig. 1).

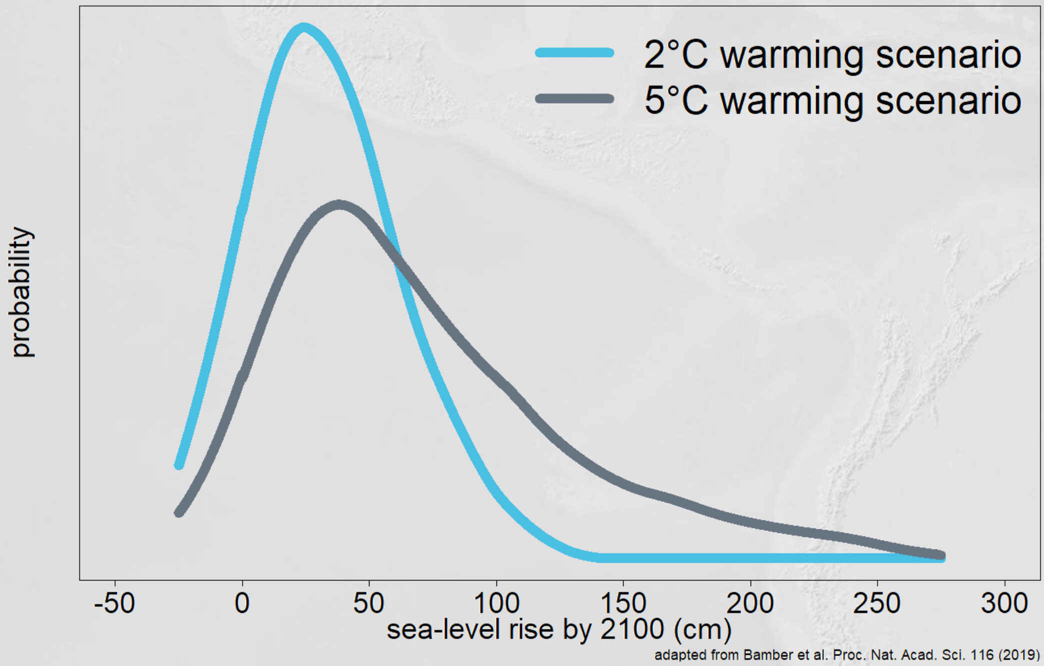

FIGURE 1: Probability distributions for sea-level rise (SLR) by 2100 under two warming scenarios. The distributions were obtained from a structured expert judgement exercise undertaken with 22 experts in the US and UK. While there is consensus that sea levels will rise there is considerable uncertainty in the amount of rise (figure adapted from Bamber et al. 2019).

2.2 DISCLOSURE

In 2017 the Task Force on Climate-related Financial Disclosures (TCFD) produced recommendations for how companies should include decision-useful, forward-looking information on the financial impacts of climate change in their financial filings (Financial Stability Board 2017a). The recommendations have currently been backed by almost 2,000 companies, half of which are financial firms managing more than $150 trillion in assets, including the world's ten largest asset managers and 8 of the 10 largest banks (Economist 2021). While 80% of the top 1,100 global companies have already started disclosing some climate-related information

(Carney 2019) disclosure of the potential impact of changes in climate on finances and strategy is still low (TCFD 2020). Some countries, including Britain, are to make climate disclosures based on the TCFD recommendations mandatory, creating a real need for forward-looking climate information. In March 2022 the U.S. Securities and Exchange Commission published their own framework for climate-related risk disclosures, modelled on the TCFD framework (U.S. Securities and Exchange Commission 2022).

The TCFD recommendations make a distinction between *physical* and *transition* risks associated with climate change. An example of a physical risk would be the decrease in property values caused by an increase in the risk of flooding. A transition risk might be the decline in the value of a coal mine due to carbon taxes or the cost of emission permits. Sometimes liability risks are treated as a distinct third category while sometimes they are regarded as a possible manifestation of physical or transition risks.

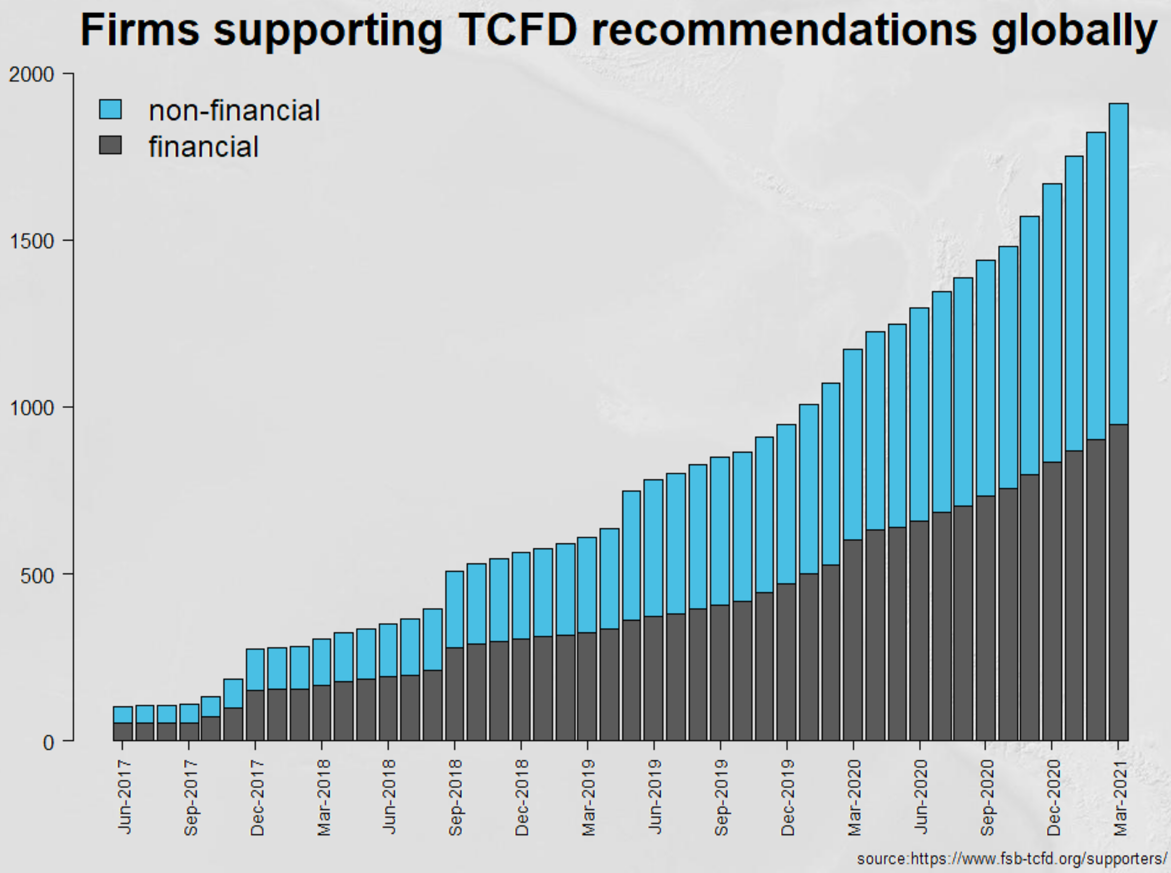

FIGURE 2: The number of firms voluntarily signing-up to support the TCFD recommendations for disclosing climate-related financial risks.

The TCFD hope that the increase in transparency will increase the efficiency of markets and the stability of economies. The disclosure of climate risks by listed companies could act as a transmission mechanism for climate forecasts to influence the real economy: for example, a lender, worried about its "climate value-at-risk", denies loans to a developer to build in an area predicted to become vulnerable to flooding, thus preventing the ill-judged development from going ahead.

Recently the Network for Greening the Financial System (NGFS) reported that, to comply with disclosure recommendations, stakeholders need more forward-looking climate information, such as emissions pathways (NGFS 2021).

2.3 CLIMATE SCENARIO ANALYSIS

The TCFD advocates the use of *climate scenario analysis* by companies to stress-test their businesses and investment portfolios under different scenarios for future climate and mitigation pathways. The scenarios can be based on those provided by organisations like the IPCC or the NGFS. These scenarios are high-level and do not contain predictions of every variable a company may need to evaluate its exposure to climate risks. The Bank of England is currently conducting a Biennial Exploratory Scenario exercise focusing on climate risk. The Bank considers that it is the responsibility of participating institutions to perform *scenario expansions* to model input variables that they require but which are not part of the high-level scenarios (Bank of England 2019).

The scenario most used by financial firms that have already conducted scenario analysis is RCP8.5.[4] Although the physical impacts of the RCP scenarios do not diverge substantially until after 2050 this choice is interesting, given the lack of consensus among experts as to the likelihood of this scenario mentioned in section 1.2, as well as the lower transition risks implied by RCP8.5.

While the IPCC provides predictions conditioned on different emissions and concentration scenarios in their most recent summary for policymakers for their 6th Assessment Report they state, *"…the feasibility or likelihood of individual scenarios is not part of the assessment."* (IPCC 2021).

2.4 CLIMATE SERVICE PROVIDERS: A NEW INDUSTRY

Because few companies have the in-house capacity to conduct scenario expansions there is a burgeoning industry of *Climate Service Providers* (CSPs). Companies such as Jupiter Intelligence, founded in 2017, Four Twenty Seven, founded in 2012 and acquired by Moody's in 2019[5] and Trucost, acquired by S&P in 2016. Trucost provides analytics on the carbon intensity of firms as well as projections of physical climate risks to which they are exposed. Catastrophe modelling firms, that have traditionally provided natural disaster risk assessments and loss modelling for the

[4] https://climate.garp.org/insight/climate-scenario-analysis-stress-testing-the-future/

[5] Moody's now uses Four Twenty Seven's climate projections to provide information about physical climate risks over to properties included in commercial mortgage-backed securities and real estate collateralized loan obligations. Moody's includes this information in their presales reports for these products (Moody's press release 11 August 2020).

reinsurance industry based on historical statistics, are also entering the space with forward-looking projections. There are many smaller boutique firms as well as data and analytics companies and consultancies that have created climate units (e.g. KPMG, Rhodium Group etc.). In Europe, for-profit companies now constitute around 30 per cent of climate information providers, academia is home to 40 per cent, the public sector about 20 per cent and non-profits 7 per cent (Cortekar et al. 2020).

The proliferation of CSPs is welcome but presents a problem for users of climate information: how should they select a forecast? The accuracy of climate predictions cannot be verified for years, so this cannot be the basis for a decision. Some CSPs use proprietary models for which the details are not disclosed to clients while, even if this isn't the case, clients are unlikely to have the in-house expertise required to critically evaluate models. In a market where experts cannot be differentiated from charlatans there is no return to expertise. One observer has noted that, *"In many ways, the marketplace is operating on blind faith in model validation and the professional competence in applying those models,"* with some, "*products and services that may be operating outside of the bounds of scientific merit*" (Keenan 2019). End users of climate disclosures will also not know whether scenario expansions used by two different companies are climatologically consistent. If they are not, their disclosures will not be comparable.

2.5 PRECISION VERSUS ACCURACY

Several CSPs offer high-resolution climate forecasts which can be combined with information about the location of a corporation's assets to provide a score measuring the physical climate risks to which the corporation might be exposed. These climate forecasts can be at a much higher resolution than those produced by global climate models and are produced by *downscaling* the output of global models. Downscaling has been widely used in weather forecasting since the method of *Model Output Statistics* (MOS) was introduced in the 1970s (Glahn and Lowry 1972). MOS essentially uses the output of a numerical weather prediction model as predictors in a statistical model for observations. Initial versions used linear models, but now non-linear models and machine-learning algorithms are also used, although the underlying principle remains the same.

Because they work on much shorter time horizons, weather forecasters can assume that the relationship between model variables and observations are relatively stationary. Stationarity assumptions are more uncomfortable for climate forecasters to make because the climate is changing. For example, while there is more than one statistical relationship between sea surface temperatures and hurricane activity that fits the historical data well, the choice of relationship has a profound impact for forecasting hurricane activity in a warmer world (Vecchi et al. 2007). It is possible that predictions from global climate models cannot be refined on scales below several hundreds of kilometers, although this scale may improve with future research (Zhang et al. 2016). For example, the UKCP18 probabilistic projections at a resolution of 25km have been criticized for not being a true reflection of the uncertainty in future climate at those scales (Baldissera Pacchetti et al. 2021).

There is a danger that users will be seduced by high-resolution, granular data sets without questioning whether the precision of such products is justified. A recent study of half-a-dozen scores from different providers purporting to measure the exposure of individual corporations to physical climate risk showed very little correlation between the scores, even when they claimed to use similar methodologies (Hain et al. 2021). Such results raise questions about how much information such scores really contain about real-world risk.

Companies that view climate-related financial disclosures as a box-ticking exercise might not worry about the accuracy of climate forecasts. These companies only need them to be plausible enough to satisfy regulators and shareholders; once that criterion is met, they may decide on price, precision – even if it's spurious – or by choosing projections that flatter their disclosures. However, for companies wanting to make a good faith assessment of their climate-related risk exposures information asymmetry is a serious problem. These companies need to have confidence in the forecasts they use in the disclosure process, and they need other stakeholders to have confidence in them too. Providers of climate forecasts also need a way to credibly signal the confidence they have in their forecasting technology if they are to avoid being pushed out of the market by less principled competitors.

2.6 TRACKING PROGRESS TO NET ZERO

"Net zero" refers to the aim of cutting emissions of greenhouse gases to a low enough level that residual emissions can be absorbed, either by natural sinks or carbon-capture technologies (IEA 2021). There is a proliferation of companies and countries making net zero "pledges" and publishing plans indicating how they hope to achieve this goal. Some countries, such as the U.K., have enshrined these pledges in law. However, making pledges and plans, or even passing laws, is one thing, actually achieving net zero is something else. At the moment, experts can express opinions about the feasibility of national net zero plans but there is no forum to synthesise their views into a unified consensus.

Information markets for global concentrations of greenhouse gases, as well as markets to predict national and regional emissions, would provide evolving assessments of the credibility of national net zero commitments.

3. INFORMATION MARKETS

3.1 WHAT IS AN INFORMATION MARKET?

Fraedrich Hayek described market prices as a "marvel" for information discovery (Hayek 1945). Information discovery, however, is a side-effect of most markets whose primary purpose is to transfer ownership, hedge risk, or allocate capital (Arrow et al. 2008). In contrast, information markets (also known as prediction markets) are designed specifically for information discovery. These markets are typically built around contingent contracts that pay out 1.00 if a specified event occurs. Participants in the market trade these contracts and a price emerges, between 0 and 1. This price can be interpreted as a consensus probability of the event occurring. One of the longest-running information markets is the Iowa Electronic Markets (IEM) that was started by the University of Iowa in 1988. IEM focusses on the outcomes of political elections (Berg and Rietz 2006).

In the IEM, participants trade contracts through a *continuous double auction* (CDA) which matches every buyer with a seller. While the market operator is not exposed to any financial risk, this design can give unsatisfactory results if there are many possible outcomes or a small number of participants. Ironically, if there is a strong consensus there will be little trading, inhibiting the ability of the market to reveal this consensus. The underlying problem is that, in a "balanced budget" market, the rewards that accrue to informed participants *must* come from the losses of other participants. To function properly, these markets must attract uninformed traders who will lose money and indirectly subsidise information discovery. The necessity of attracting participants, some of whom are certain to lose money, raises both practical and ethical concerns.

3.2 ALTERNATIVE MARKET DESIGNS

Information markets using a CDA have been successful for sports, entertainment, and politics[6] but have struggled when applied to more specialized topics for which it is harder to attract uninformed traders. Even when used for popular subjects, CDA markets only perform well when there are a small number of possible outcomes. In 2011 an experimental information market for global mean temperature was run on the InTrade platform, using a CDA. This market suffered from low liquidity, resulting in high bid-ask spreads, making inference of probabilities difficult (Boslough 2011).

Breakthroughs in aviation followed the realisation that aircraft do not have to emulate birds; progress in information markets has been made once it was realised that they do not have to mimic traditional futures markets. Because the purpose of

[6] The type of politics for which information markets have been successful is arguably "politics as entertainment": Who will win elections? Who will have to resign? Information markets for whether particular policies will be enacted, are generally much less popular.

information markets is information discovery, rather than the transfer of assets, there is no need for them to use a CDA. Instead, an automated market maker, always willing to buy and sell contracts, can provide a minimum level of liquidity. This market maker can be directly subsidised by a market sponsor to provide incentives for informed participants to take part, even in the absence of noise traders. The logarithmic market scoring rule (LMSR) market maker is particularly elegant (Hanson 2003): it ensures that the aggregate gains of market participants — and hence the cost to the market maker — is equal to the logarithmic scoring rule for probability forecasts (Good 1952) which is a direct measure of the information content of the forecast (Roulston and Smith 2002). The amount the sponsor must pay for the forecast generated by the market is therefore directly linked to its accuracy (see box on next page).

As well as removing the need to attract uninformed participants, subsidised automated market makers can facilitate information markets with very large outcome spaces. A market to predict the opening date of Carnegie Mellon University's computer science building used an LMSR market maker to support a space of 365 possible days (Othman and Sandholm 2013) and, more recently, markets to simultaneously predict temperature and rainfall in a space partitioned into more than 5,000 outcomes have succeeded using the same approach. Large outcome spaces allow continuous variables to be partitioned into intervals, enabling detailed probability distributions to be produced.

Information markets, such as the IEM, that make participants put in money are legally classed as gambling or futures markets in most jurisdictions. The IEM operates under a "no-action" letter from the Commodity Futures Trading Commission which permits it to run the market as an academic exercise with tight constraints on participants and their financial stakes. However, once it is understood that the participants in an information market should be a source of information rather than revenue, the rationale for requiring them to pay to take part is weakened. Removing this requirement allows information markets to be established in jurisdictions that do not allow gambling, as well as removing regulatory hurdles in jurisdictions that do. In this form, an information market becomes a mechanism for distributing funding in which the amount received by each participant is tied to the improvement they contribute to the forecast.

PROPER SCORING RULES AND INCENTIVE COMPATIBILITY

Probability forecasts can be scored using scoring rules. If a forecaster optimizes the expected value of a score by giving their true belief about a probability, then the score is said to be *proper*. This condition is also known as *incentive compatibility.*

The Brier score (Brier 1950) is an example of a proper scoring rule and is widely used in weather forecasting. If there are m outcomes and f_i is the predicted probability of outcome i occurring, and the actual outcome that occurs is outcome k then the Brier score for the forecast is given by

$$BS = \sum_{i=1}^{m}(f_i - \delta_{ik})^2 \quad \text{where } \delta_{ik} = 1 \text{ if } i = k \text{ and } \delta_{ik} = 0 \text{ if } i \neq k$$

As defined, Brier score is negatively oriented: smaller values are better and the best possible score of zero is achieved if the actual outcomes was predicted with a probability of 1.0.

Another proper scoring rule is the logarithmic score (Good 1952). This score is given by

$$LS = -\log f_k$$

The logarithmic rule is also negatively oriented and represents the information deficit of someone in possession of the forecast. If the actual outcome has been predicted with a probability of $f_k = 1$ then the logarithmic score is zero: the holder of the forecast has no information deficit because they know what will happen with no uncertainty. However, if $f_i = 1$ and outcome i does not occur the logarithmic score is undefined. This is an important difference between the Brier score and the logarithmic score: the logarithmic score severely punishes unjustified certainty!

Not all scoring rules are proper. For example, the (positively oriented) linear score, which is just f_k, seems superficially sensible as it assigns a higher score when the probability assigned to the actual outcome is higher. However, under this score, the maximum expected score is achieved if the forecaster reports that $f_i = 1$ for the outcome they consider most likely even if they believe the probability of it occurring is less than 1.0. Furthermore, nonlinear functions of proper scores are not themselves proper scores. Setting targets on proper scoring rules means they are no longer proper as the forecaster can increase the probability of achieving the target by reporting probabilities that differ from their true beliefs (Roulston 2007). Forecasting tournaments scored using proper scoring rules can also suffer from a similar distortion of incentives (Kilgour and Gerchack 2004, Witkowski 2022).

3.3 ACCURACY OF INFORMATION MARKETS

Information markets produce probabilistic forecasts. A key characteristic that probabilistic forecasts should possess is *calibration*: the frequency of events should match the predicted probability of their occurrence. Figure 3 is a calibration plot of all election markets run by the IEM from 1988 until 2006. The diagonal line represents calibrated forecasts for which the event frequency equals the predicted probability. It is seen that the IEM markets are generally well calibrated (Wolfers and Zitewitz 2006). Furthermore, the IEM markets to predict vote share were closer to the actual share than opinion polls in 76% of elections (Berg et al. 2008).

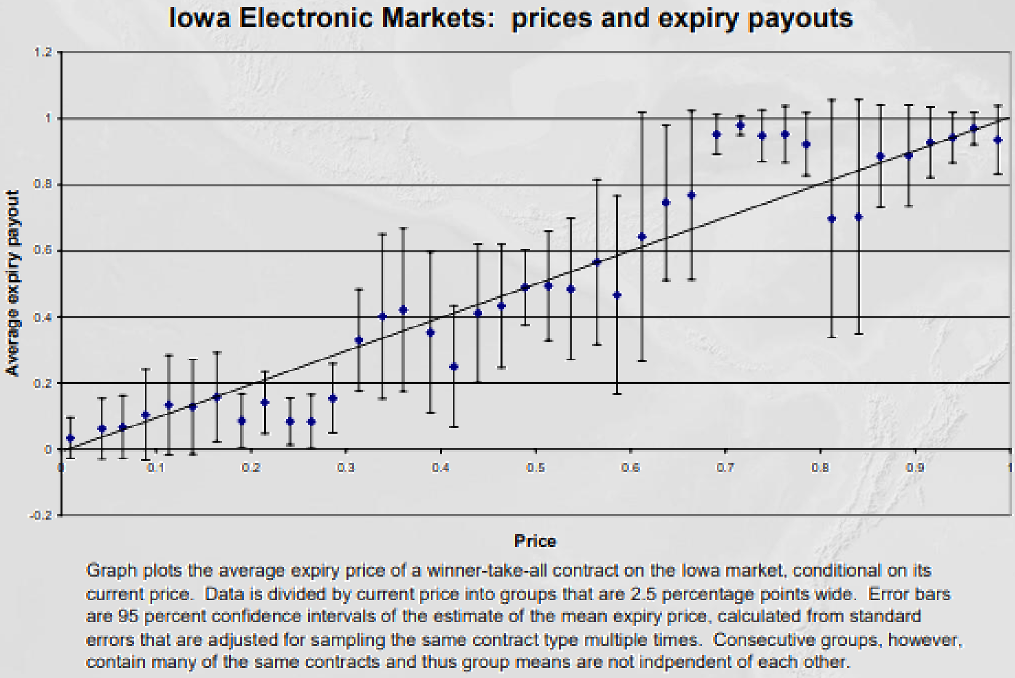

FIGURE 3: The calibration of political election forecasts from the Iowa Electronic Markets (Wolfers and Zitewitz 2006).

Predicting elections, however, is not the same as predicting climate-related variables: the required expertise is different, and elections have a small number of possible winners, whereas we would like to predict distributions of climate variables by partitioning their values into large numbers of possible outcomes. The concept of calibration can still be applied by asking whether the actual value of a variable falls below the 10th percentile in 10 per cent of forecasts, below the 20th percentile in 20% of forecasts and so on. Figure 4 is a calibration plot for 23 climate-related markets run on Hivemind's AGORA information market platform. The predicted variables were U.K. monthly temperatures and rainfall, monthly average sea surface temperature anomalies in the NINO3.4 region of the tropical Pacific and the number of Atlantic hurricanes occurring during the 2020 hurricane season.[7] The plot indicates that the information market forecasts were consistent with well-calibrated predictions.

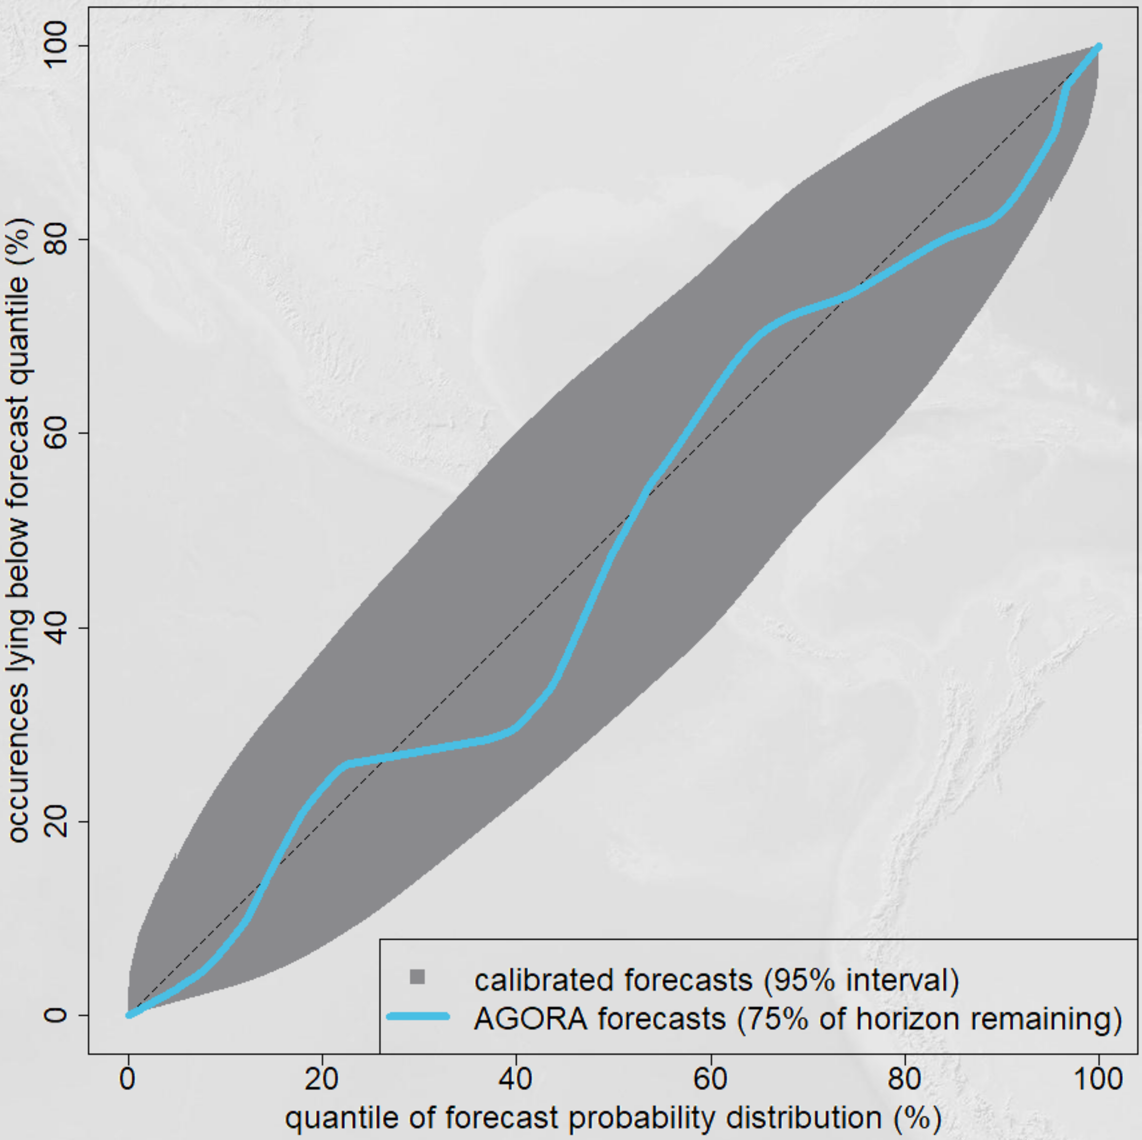

FIGURE 4: A calibration plot of 23 climate-related information markets run on Hivemind's AGORA platform. Markets were for monthly U.K. temperature and rainfall, tropical Pacific sea-surface temperatures, and seasonal hurricane activity. The grey envelope denotes the range of calibration curves that would be expected for 23 perfectly calibrated forecasts. The total forecast horizon ranged from 2 to 9 months across markets and this curve is for when 75% of the forecast horizon remained for each market.

While there is plenty of evidence that information markets can be efficient mechanisms for eliciting and aggregating information, it is not possible to know whether they will be effective in any given application without trying them. One

[7] The hurricane markets were run as part of Hivemind's participation in the Lloyd's Lab InsurTech accelerator.

concern is that calibration issues may arise with long-range information markets if participants are not compensated for the time value of money (Page and Clemen 2013).

3.4 MANIPULATION OF INFORMATION MARKETS

Once concern with information markets, particularly those addressing a politically polarizing topic such as climate change, is whether they are vulnerable to manipulation. Specifically, to what extent could a participant who is prepared to lose money distort prices in the market to either exaggerate or downplay the risk of climate change? The participant's motive might be that the forecast produced by the market may be used to make a decision that would be far more consequential for them: e.g., a fossil fuel firm wanting to create the impression that climate change is not a major concern to keep carbon taxes low, or perhaps a civil engineering firm wanting a higher estimate of sea-level rise to obtain a lucrative contract for building flood defences.

Theoretical studies have shown that, at least in some circumstances, the presence of manipulators can enhance the accuracy of information markets by providing higher returns to informed participants (Hanson and Oprea 2009). In a laboratory study the presence of manipulators has been shown to distort prices, but not to an extent that decision makers are no longer better off using the prices. The laboratory study indicated that a bigger problem might be the reduction of trust decision makers have in the market forecast if they suspect the presence of manipulators (Choo et al. 2022).

While it is important to understand how manipulators can distort forecasts produced by information markets, and what steps can be taken to minimise their impact, the vulnerability of these markets to manipulation should be compared with the status quo, or practical alternatives, rather than an ideal situation. The existing process for creating a consensus view on future climate, which includes the IPCC process as well as the media and political discourse that surrounds it, is also prone to manipulation by parties who are willing to spend money (Oreskes and Conway 2011).

4. A CLIMATE INFORMATION EXCHANGE

4.1 OUTLINE

We propose to establish an exchange dedicated to running climate-related information markets designed to elicit and aggregate the judgments and models of diverse experts from relevant disciplines. The exchange will address the problem of information asymmetry by linking the funding of forecast providers to the accuracy of the information they provide. The benchmark predictions generated by the exchange will increase transparency in decision making and the disclosure process for climate-related risks.

The exchange will host individual information markets for specified variables related to climate risks (physical and transition) at specified future time horizons.

Using information markets for climate prediction has been suggested by others (e.g. Hsu 2011, Vandenbergh et al. 2013, Lucas and Mormann 2018, Aliakbari and McKitrick 2018) but no practical market has been established. The reasons for this include regulation and the long horizons involved. We will address the regulatory issues by not making participants pay to take part and tackle the horizon problem by situating the exchange within a university and using a segregated fund and full collateralisation to mitigate counterparty risk.

4.2 MARKET DESIGN

The information markets will be hosted on the AGORA platform originally prototyped by Winton Group (Roulston et al. 2016, Roulston 2017). AGORA was further developed by Hivemind Technologies after it spun-out of Winton in 2018. AGORA uses the LMSR market making algorithm (Hanson 2003, see box) enabling it to support markets with large numbers of outcomes and small or large numbers of participants.

AGORA allows participants to define contracts covering one or more possible outcomes and then buy and sell these contracts via the market maker using on-platform credits. The automated market maker adjusts prices in response to the trading activity. When the actual outcome becomes known every contract containing this outcome is worth 1.00 credit while all other contracts are worth nothing. If participants do not believe they have relevant information they have an incentive to not trade, thus not degrading the consensus forecast. This arrangement matches a previously suggested way of rewarding experts to screen out uninformed ones (Sandroni 2014), except the market also aggregates expertise as well as incentivising it.

Since AGORA can support markets with large numbers of outcomes, continuous quantities – such as temperature – can be partitioned into intervals allowing probability distributions to be inferred. This contrasts with many prediction market platforms that focus on binary markets predicting whether a specified event will occur or not.

AGORA also supports multi-dimensional outcome spaces. This means that markets can jointly predict atmospheric greenhouse gas concentrations *and* global temperatures. Having an implied two-dimensional probability in GHG concentration-temperature space is powerful because it solves the "circularity problem" (Sumner and Jackson 2008): suppose that only temperatures are predicted, if the market is forecasting only a modest rise in temperatures this could be because the market believes there is a low sensitivity to GHG concentrations, or it could mean the market anticipates drastic action to curb concentrations. If these two possibilities cannot be differentiated the information generated by the market is of limited use when deciding on mitigation policies. Having predictions of temperature *conditioned* on GHG concentrations allows these possibilities to be distinguished resulting in a more valuable information product.

AGORA has been used to host several climate-related markets with horizons out to 10 months. These include markets for monthly sea surface temperature anomalies in the tropical Pacific (a key diagnostic for El Niño events), markets for the number of hurricanes and landfalling hurricanes during the Atlantic hurricane season (Fig. 5), and markets simultaneously predicting monthly temperatures and rainfall in the United Kingdom (see Fig. 6 and Fig. 7). The temperature-rainfall markets demonstrated the ability to extract joint probability distributions from participants. The market structure shown in Fig. 6 could just as easily be for atmospheric carbon dioxide concentration and global temperature anomaly for a specified future year.

LOGARITHMIC MARKET SCORING RULE (LMSR)

The logarithmic market scoring rule (LMSR) market maker was developed by Hanson (2003). It is an example of a subsidised market maker based on a proper scoring rule. The market maker uses an algorithm to price outcomes based on how many contracts for that outcome it has already sold, relative to other outcomes.

Let there be m outcomes and let the market maker's exposure to outcome i be q_i. The pricing rule is based on a cost function given by:

$$C(\boldsymbol{q}) = b \log\left(\sum_{i=1}^{m} e^{q_i/b}\right)$$

where b is a *liquidity parameter* that determines how much the market maker will change prices in response to a given change in exposure and the maximum loss the market maker can suffer.

The prevailing prices the market maker assigns to each outcome, p_j, are given by the slope of the cost function:

$$p_j = \frac{\partial C}{\partial q_j} = \frac{e^{q_j/b}}{\sum_{i=1}^{m} e^{q_i/b}}$$

Note that the prices are normalised so that $\sum_{j=1}^{m} p_j = 1$ and that as the exposure of one outcome becomes large (relative to other outcomes) its price tends toward 1.00.

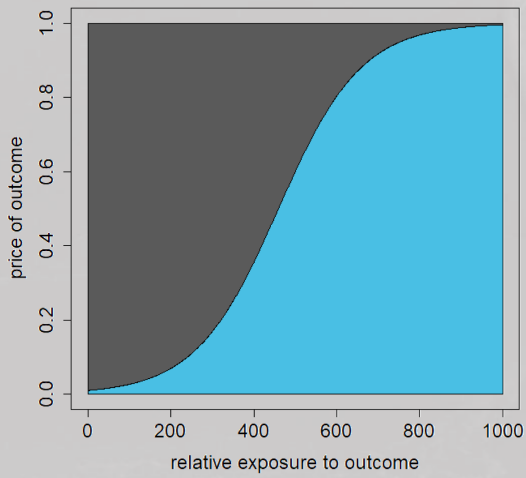

The blue area under the curve represents the amount the market maker will raise selling contracts up to a given exposure. If that outcome occurs, the market maker will pay out 1.00 for each contract, so the grey area represents its net loss. The size of the grey area, and hence the net loss, is bounded.

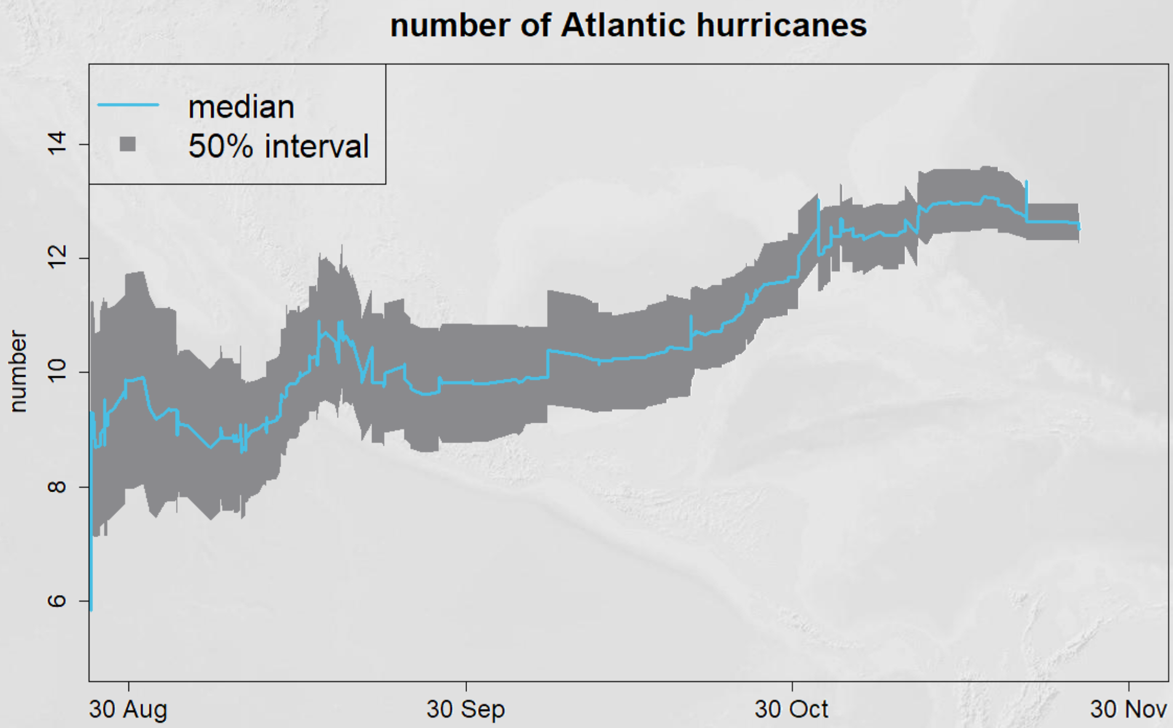

FIGURE 5: The evolution of the number of Atlantic hurricanes (Category 1 or higher) during the 2020 hurricane season in an AGORA market run during the Lloyd's Lab InsurTech accelerator. Historically the average number of hurricanes during a season has been 6. When the market opened, at the end of August, there had already been two and the market immediately predicted 9. By the end of the season there had been 13 (although the National Hurricane Centre upgraded Tropical Storm Gamma to a hurricane in its post-season reanalysis released in April 2021). The market produced a complete distribution for the number of hurricanes, from which the median and 50% interval were inferred.

AGORA allows participants to trade either through the user interface or using the API. The API enables programmatic trading and allows participants to integrate their own models directly into their trading strategies enhancing efficient information discovery and aggregation. The API can also be used by end-users of the information generated by the markets to access it in real-time.

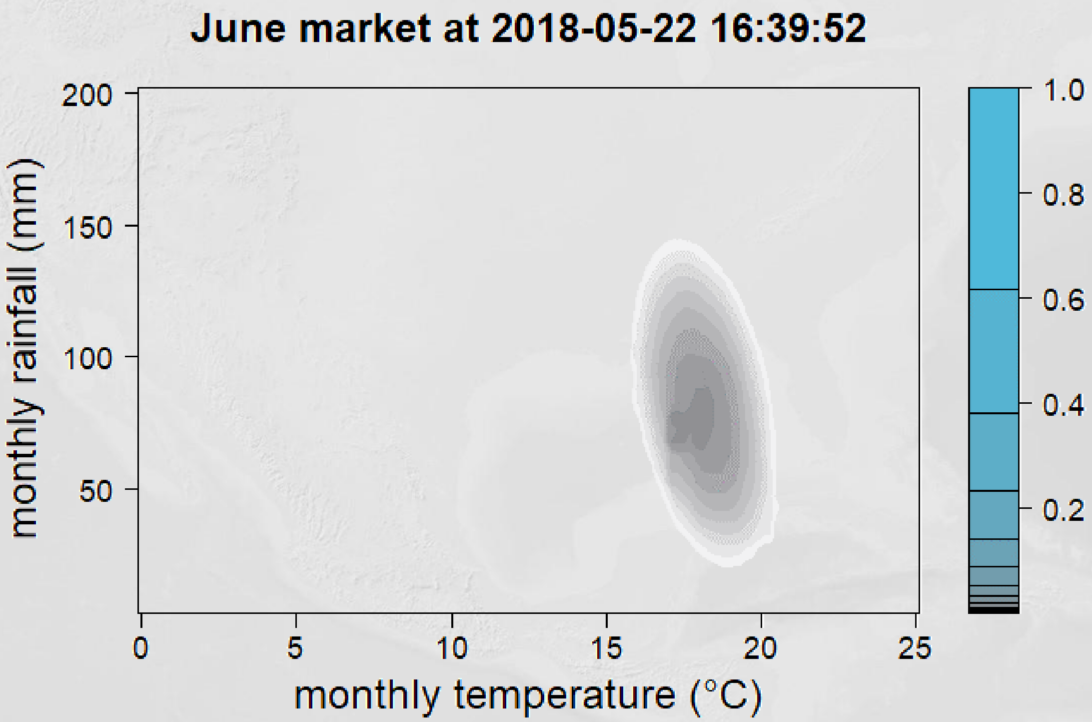

FIGURE 6: The prices on 22nd May 2018 in a market simultaneously predicting monthly temperature and rainfall for the United Kingdom in June 2018. The outcome space is divided into 5,207 outcomes each covering 0.2°C in temperature and 5mm in rainfall. The participants were 24 groups from British universities with expertise in meteorology, climate science and statistics. The prices can be interpreted as a joint probability distribution. The actual values for June were 19.9°C and 34.7mm (2018 tied for the second warmest June on record in the UK). Notice the negative correlation between rainfall and temperature (hotter summers tend to be drier) which emerged from participant trading.

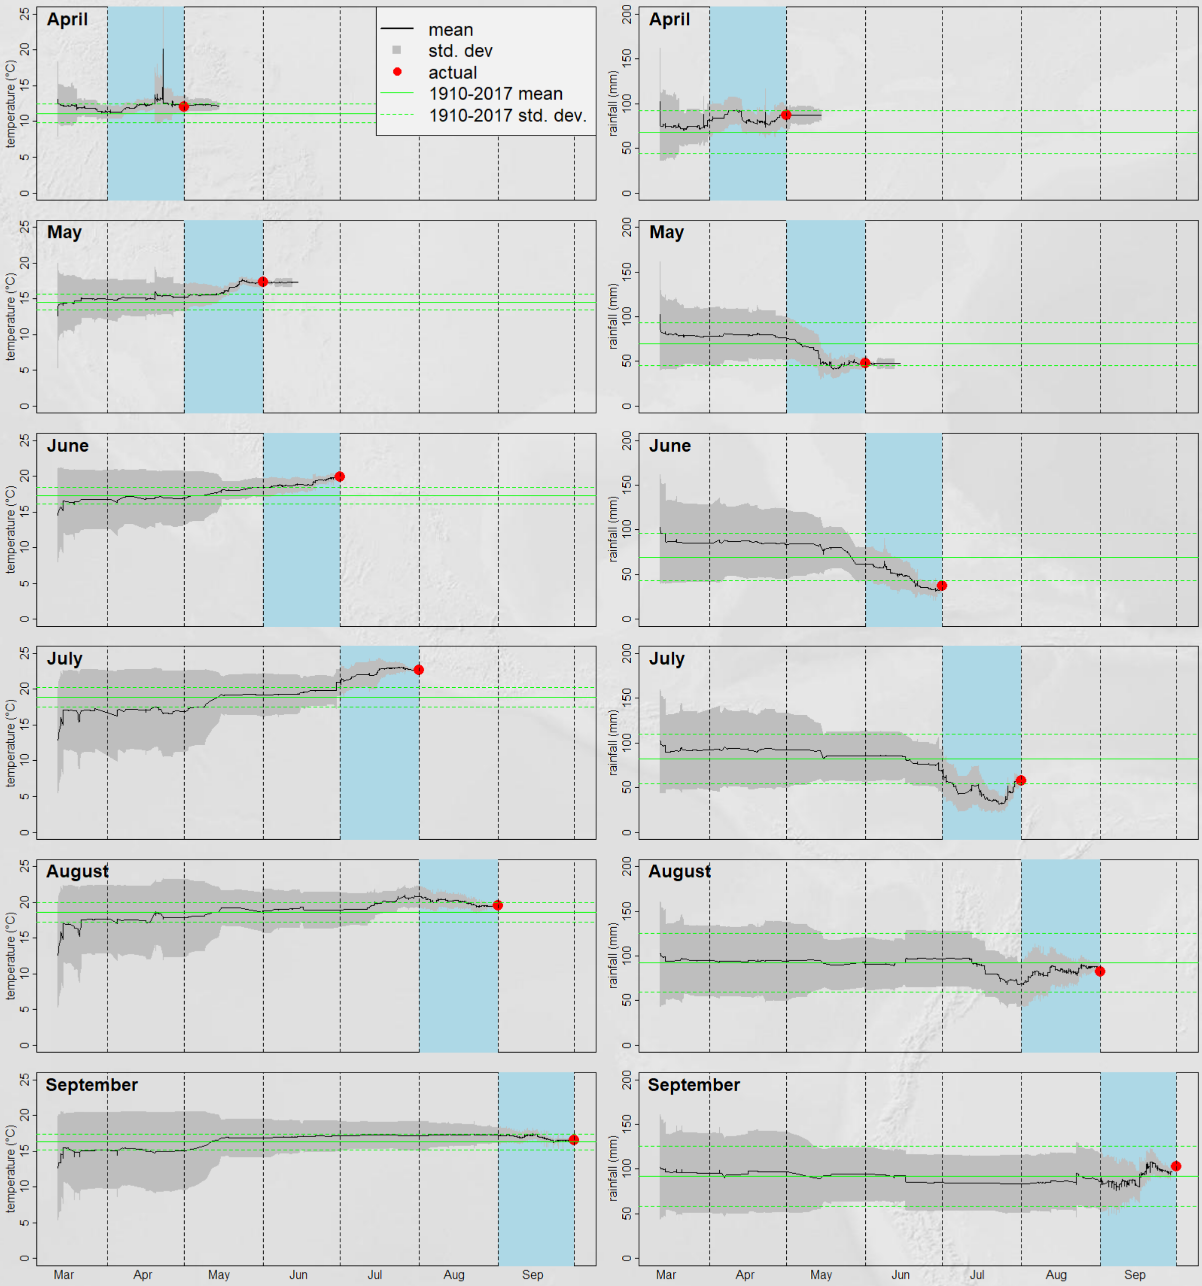

FIGURE 7: The evolution of the marginal price distributions for UK average daily high temperature and monthly rainfall for the markets covering April to September 2018. All the markets opened on 12th March 2018. The markets for April and May were closed for trading on 15th May and 15th June respectively, after the Met Office had published the official figures (the "actual" denoted by the red dots). The markets for June to September were closed on the last day of the relevant month, before publication of the official number. The green lines show the mean and standard deviation for the relevant month calculated using 1910-2017 data.

4.3 CREDITS SYSTEM

AGORA participants trade using on-platform credits without participants having to "pay-to-play". Previous AGORA markets have rewarded participants with cash compensation either through a tournament structure, in which those who accumulate the most credits receive fixed cash pay outs, or through a lottery system, in which the probability of a participant winning a cash payout is proportional to the number of credits they accumulate. While both arrangements worked well, the tournament structure can, in theory, distort behaviour by motivating people to accumulate a certain number of credits rather than maximize their expected credits. The lottery system does not suffer from this problem, but while it incentivises people to behave in a risk neutral way, *given they have chosen to participate,* the random element of the reward may deter people from participating in the first place.

To overcome the drawbacks of tournament and lottery structures, the climate information exchange will use on-platform credits that represent claim on a proportion of funding from the users of information. Each quarter, a share of funding will be deposited in a segregated trust account.[8] Participants will receive an endowment of credits. The market makers in each market will also be allocated enough credits to cover their maximum possible losses (which are strictly constrained by the LMSR market making algorithm). The prevailing cash value of these credits will be the funds in the segregated account divided by the total number of outstanding credits.

When new credits are created, whether to endow participants or the market makers of newly created markets, cash will be deposited in the segregated fund at the prevailing credit price, thus leaving the cash value of credits unchanged. Money may also be deposited in the fund without creating new credits causing the value of each outstanding credit to appreciate. Under this arrangement the cash value of credits will never decrease. Initially, the cash value of the credits may be quite low. Nevertheless, participants will have an incentive, even in these early stages, because of the anticipation that the value of credits gained through trading will increase substantially as funding from sponsors rises.

Credits are transferable between markets. This encourages participants to self-select and focus their attention on those markets where they believe they have an informational advantage. Due to the long-range horizons of some of the markets endowments of credits to participants will be "locked-up" for multi-year periods encouraging them to participate in these longer-range markets.

The credit system allows the exchange to be a mechanism for distributing money from the users of forward-looking climate information to the providers of that information in a way that directly ties the compensation of providers to the accuracy and timeliness of the information that they provide. This will serve as an additional funding stream for applied climate research and align the incentives of providers to overcome the flaws in the marketplace for climate information described in Section 1.

[8] The share deposited should be determined by a stable and transparent policy but allowing some discretion.

4.4 PARTICIPATION

The climate information exchange will host markets for physical risks (e.g. temperature and sea level rise) as well as indicators of transition risk (e.g. GHG emissions and uptake of clean technologies). Therefore, we will recruit participating members with a wide range of expertise in social sciences and policy as well as physical climatology. Both individuals and institutions will be eligible for membership. Institutions would include think tanks, academia, and private companies. We have received letters of interest from institutions who would be keen on becoming participating members.

Over one hundred people with expertise in climatology, meteorology, statistics and machine learning from universities and the private sector have already participated in AGORA markets. They have validated the market design and demonstrated the effectiveness with which it is able to elicit and aggregate the expertise of participants.

4.5 MARKETS: CORE AND CUSTOM

Markets on the exchange will be either *core* or *custom*. Core markets will be funded from a general funding pool. The information they generate will be publicly available. Core markets will typically provide high-level forecasts of wide interest: e.g., global carbon dioxide concentrations and temperature anomalies in future years.

Custom markets will be bespoke. The subsidy for the market maker for these markets will be provided by sponsors interested in forecasting a specific climate-related risks, e.g., sea-level rise in a specific location. To what extent market sponsors have exclusive access to the information a market generates will need to be decided. In some cases, the market sponsor may have no objections to public dissemination (possibly with the information being co-branded) but there may be potential sponsors who want information from their market to be private.

4.6 MARKET HORIZONS

Timescales relevant for climate risks are intrinsically decadal so many markets will be on horizons of 20 or 30 years and, ultimately, even longer than this. However, there are reasons to run additional markets on shorter time scales of 1 or 2 years. Firstly, information about short-range climate risks, e.g., hurricane activity or the intensity of the Asian monsoon, is potentially useful and shorter-range markets will help engage participants and allow them to familiarise themselves with the market design. In addition, near term markets provide a baseline to compare with longer range markets: e.g., a market generating a probability distribution for hurricane numbers next season can be compared with an otherwise identical market for 20 years hence, implying a consensus view on how participants expect hurricane activity to evolve.

4.7 EXAMPLE MARKETS

The table below provides some examples of information markets that could be hosted on the exchange. For longer horizons there will be separate markets targeting specific years at intervals of 1, 5, or 10 years.

MARKET	HORIZON	DESCRIPTION
CO_2 emissions	1-10 years	Annual global and regional carbon dioxide emissions.
CO_2 concentration	10-30 years	Annual average of carbon dioxide concentration (as measured at Mauna Loa observatory)
CH_4 concentration	10-30 years	Annual and global average of methane concentration according to NOAA Global Monitoring Lab.
global temperature anomaly	10-30 years	Annually averaged global temperature anomaly relative to a defined baseline, according to a specified index such as HadCRUT.
Joint CO_2-global temperature	10-30 years	Two-dimensional outcome space to generate predictions of global warming contingent on CO_2 concentrations.
global sea-level rise	10-30 years	Global average (eustatic) sea-level rise according to recognized source: e.g., Univ. of Hawaii Sea Level Center.
Arctic sea ice	1-10 years	Minimum annual sea ice as reported by the National Snow & Ice Data Center.
El Niño	1-12 months	Monthly value of sea surface temperature anomaly in NINO3.4 region of tropical Pacific.
Atlantic hurricane activity	1-20 years	Markets for number of hurricanes per season, number of severe hurricanes (CAT3+)
All India rainfall index	1-20 years	All India rainfall index during the summer monsoon period.
Crop yields	1-20 years	Crop yields in key producing regions.

4.8 INAUGURAL MARKETS

The exchange will open with a group of inaugural core markets for predicting annually averaged global carbon dioxide concentrations and global temperature anomalies. There will be a market for each year for the next 20 years. Each market will have a two-dimensional outcome space defined by partitioning the annual average carbon dioxide concentration (as reported by the Scripps Mauna Loa observatory) and the global average temperature anomaly (as defined by the most recent version of the Hadley CRU temperature anomaly time series).[9] The market outcome space would be analogous to that for temperature and rainfall shown in Fig. 6.

The choice of carbon dioxide and temperature as the inaugural set of markets is partly driven by the desire to provide relative likelihoods of different scenarios, a data gap already mentioned. The two-dimensional outcome space is dictated by the need to solve the circularity problem highlighted by Sumner and Jackson (2008). Because understanding future carbon dioxide emissions and concentrations requires synthesising views about national economic and energy policies while predicting global temperature anomalies conditioned on carbon dioxide concentrations is the central challenge of physical climate modelling the two-dimensional market will be an excellent showcase for the ability of information markets to aggregate judgment and modelling across a large array of diverse disciplines.

The near-term markets (less than five years) won't provide much information about long term climate change, but they should capture knowledge of phenomena like El Niño-Southern oscillation and volcanic eruptions and incorporate them into a complete "forward curve" for global temperature. The settlement of near-term markets should also stimulate interest and promote engagement.

The opening of the markets will be phased, with the near-term markets being opened first and then augmented with longer-range markets over a period of months.

The information product of the market—consensus probability forecasts of carbon dioxide concentration and temperature anomalies and the relation between them—will be of wide interest. If the participants include a large enough number of recognized international experts in relevant fields the consensus forecasts will provide a valuable reference point in any discussions of climate change.

4.8 AGORA INTELLECTUAL PROPERTY

The intellectual property and codebase for the AGORA platform was donated by Hivemind Technologies Ltd. to the University of Lancaster in July 2022 for use in the climate prediction market project.

[9] As with weather derivative contracts, contingency arrangements specifying alternate measures of carbon dioxide concentration and global temperature anomaly to be used to settle markets if the primary measures are, for some reason, unavailable will be specified in the prospectus for market participants.

5. INFORMATION MARKETS AND ACADEMIA

5.1 CLIMATE FORECASTS ARE A PUBLIC GOOD

Many of the envisaged information markets will produce consensus forecasts of value to many parties. Forecasts of carbon dioxide and other greenhouse gas concentrations and global temperature anomalies are of wide interest. In conversations with financial firms about climate risk, the paucity of information about the relative likelihood of different scenarios has been one of the first data gaps they highlight. A sequence of probabilistic forecasts for atmospheric concentrations of carbon dioxide would allow relative likelihoods to be assigned to the IPCC or other scenarios.

While information about carbon emission pathways might be of considerable interest and value, it is also difficult for a provider of the information to capture this value. The provider is likely to find itself in the same position as the credit ratings agencies after the invention of the Xerox machine (see Box "Bad incentives, bad information: An example" in Section 1). Once the information is produced it can be easily copied and distributed by others, it is "non-excludable" and "non-rival" as economists would say. With some information products the market can be segmented based on the freshness of the data: you can get share prices delayed by 15 minutes for free, but trading firms pay substantial amounts for a "low latency" live feed. Unfortunately, while there may be situations when consensus forecasts of future climate change rapidly (for example, the announcement of a new carbon tax), in general, climate forecasts evolve slowly enough that freshness cannot command a premium price.

The non-excludable and non-rival nature of the forecast information that will be generated by the exchange mean it is a classic example of a "public good". Such goods are typically provided by governments, philanthropists (including crowdsourced philanthropy), or consortia of users. This aspect of the information exchange suggests a non-commercial model is appropriate.

One of the functions of universities is to discover information and disseminate it. The channels typically used are academic journals and conferences. The mechanics of an information market are obviously different to a journal, but they serve the same purpose: the elicitation, aggregation and communication of information and expertise. Academia's mission to discover and disseminate information for the benefit of society provides a justification for a climate information exchange to be hosted by university.

5.2 LONGEVITY AND GOVERNANCE

Another argument for a university hosting a climate information exchange are the long horizons of some of the markets. As mentioned, some markets will not settle for more than 20 years after they first open. While participants will be able to trade in and out of the market there still needs to be an expectation that the exchange will be around when the actual value of the target variable is known and the market settles. There is likely to be more confidence in this longevity if the market is hosted by an academic institution than if it is hosted by a commercial entity, particularly one of very recent origin. British universities are subject to the Freedom of Information Act would ensure enhanced transparency.

Even though participants are likely to have more confidence in the longevity and stability of a university there should still be a contingency plan for winding-up markets before their scheduled settlement date. Circumstances may mean the institution can no longer host the markets and, if an alternative operator cannot be found, markets will have to be closed and settled before the forecast can be verified.

Markets can be settled prematurely by cashing everyone out at prevailing prices. Under an LMSR market maker in which market subsidies were deposited in a segregated fund when the market was initialised, and participant endowments were similarly deposited, the contents of the fund will be able to cover the necessary pay outs. The existence and design of this procedure for prematurely settling markets should be known to participants.

5.3 HOST INSTITUTION BENEFITS

The aim is to attract a population of market participants representing international expertise across all disciplines relevant to climate-related risks, including physical sciences, economics and other social sciences, politics and policy studies and technological innovation. This will raise the profile of the hosting institution among researchers in all these fields.

The information exchange could serve as a nexus for interdisciplinary research. The necessity of aggregating expertise across a wide range of disciplines is integral to the project. By providing an apparatus and reason for bringing diverse researchers together the information exchange could stimulate cross-cutting activities such as workshops and research projects.

It is hoped that the forecasts generated by the climate information markets will become a standard reference point in discussions of climate change. For example, journalists could refer to the relative likelihood of different IPCC scenarios implied by carbon dioxide concentration markets. Coupled with a good branding strategy for the information products, this should raise the international recognition of the host institution among the public, in the context of climate change research.

The information exchange will generate forecasts of climate-related risks, but it will also generate data about the behaviour of markets. It will facilitate research into market design which can hopefully feedback and inform the development of the exchange platform.

6. SUMMARY

As climate change becomes a more recognized source of risk by companies, shareholders, and regulators there will be an increase in the demand for forward-looking information concerning both physical climate risks and transition risks caused by the decarbonization of economies. In addition, public and private sector organisations will increasingly need forecasts of climate-risks, that are credible and transparent, to assist their decision making and risk management.

The marketplace for forward-looking climate information is fragmented and suffers from information asymmetries. These asymmetries introduce moral hazard, making it vulnerable to market failure, in which bad products crowd out good products.

Information markets solve the problems of fragmentation and asymmetric information because they are tools for eliciting and aggregating distributed knowledge and expertise and provide incentives for accuracy. Participation in information markets will allow providers to signal the confidence that they have in their predictions and get compensated for skill.

We propose establishing a climate information exchange that will use information markets to produce consensus forecasts of climate-related risks through the participating of relevant experts from academia, think tanks and private companies.

The exchange will serve as a mechanism for the consumers of climate research to fund that research in a way that is directly tied to its effectiveness. It could become an important source of new funding for institutions engaged in climate-related research helping to mitigate the problem of underinvestment in such research.

The exchange should become a hub for international expertise in an array of disciplines relevant to forecasting climate risks, including physical sciences, economics, social sciences, and policy.

REFERENCES

Akerlof, George A. "The market for “lemons”: Quality uncertainty and the market mechanism." Uncertainty in economics. Academic Press, 1978. 235-251.

Aliakbari, Elmira, and Ross McKitrick. "Information aggregation in a prediction market for climate outcomes." Energy Economics 74 (2018): 97-106.

Arrow, Kenneth J., et al. "The promise of prediction markets." (2008).

Baldissera Pacchetti, Marina, et al. "Assessing the quality of state-of-the-art regional climate information: the case of the UK Climate Projections 2018." Climatic Change 168.1 (2021): 1-25.

Bamber, Jonathan L., et al. "Ice sheet contributions to future sea-level rise from structured expert judgment." Proceedings of the National Academy of Sciences 116.23 (2019): 11195-11200.

Bank of England. “The 2021 biennial exploratory scenario on the financial risks from climate change”. Discussion Paper. December (2019).

Berg, Joyce E., and Thomas A. Rietz. "The Iowa electronic markets: Stylized facts and open issues." Information Markets: A New Way of Making Decisions (2006): 142-169.

Berg, Joyce, et al. "Results from a dozen years of election futures markets research." Handbook of experimental economics results 1 (2008): 742-751.

Boslough, Mark. Comparison of Climate Forecasts: Expert Opinions vs. Prediction Markets. No. SAND2011-8232C. Sandia National Lab. (SNL-NM), Albuquerque, NM (United States), 2011.

Brier, Glenn W. "Verification of forecasts expressed in terms of probability." Monthly weather review 78.1 (1950): 1-3.

Budish, Eric, Peter Cramton, and John Shim. "The high-frequency trading arms race: Frequent batch auctions as a market design response." The Quarterly Journal of Economics 130.4 (2015): 1547-1621.

Carney, Mark. "TCFD: strengthening the foundations of sustainable finance." Proceedings of the TCFD Summit (2019).

Choo, Lawrence, Todd R. Kaplan, and Ro’I. Zultan. "Manipulation and (mis) trust in prediction markets." Management Science (2022).

Christensen, Peter, Kenneth Gillingham, and William Nordhaus. "Uncertainty in forecasts of long-run economic growth." Proceedings of the National Academy of Sciences 115.21 (2018): 5409-5414.

Congressional Oversight Panel. “Modernizing the American Financial Regulatory System: Recommendations for improving Oversight, protecting consumers, and ensuring stability.” Special Report on Regulatory Reform. January (2009).

Cortekar, Jörg, Matthias Themessl, and Katja Lamich. "Systematic analysis of EU-based climate service providers." Climate Services 17 (2020): 100125.

Deloitte. “The predictive power of stress tests to tackle climate change.” 2020.

Economist, “Telling all”. The Economist. March 13th (2021).

Financial Stability Board. "Recommendations of the task force on climate-related financial disclosures." Financial Stability Board, Basel, Switzerland (2017a).

Financial Stability Board. "Technical supplement: The Use of Scenario Analysis in Disclosure of Climate-Related Risks and Opportunities" Financial Stability Board, Basel, Switzerland (2017b).

Flaherty, J. C., G. Gourgey, and S. Natarajan. "Five lessons learned: risk management after the crisis." The European financial review 30 (2013).

Friedmann, Julio, et al. "Levelized Cost of Carbon Abatement: An Improved Cost-Assessment Methodology for a Net-Zero Emissions World." Columbia University SIPA Center on Global Energy Policy: New York, NY, USA (2020).

Gates, Bill. How to Avoid a Climate Disaster: The Solutions We Have and the Breakthroughs We Need. Knopf, 2021.

Glahn, Harry R., and Dale A. Lowry. "The use of model output statistics (MOS) in objective weather forecasting." Journal of Applied Meteorology and Climatology 11.8 (1972): 1203-1211.

Gneiting, Tilmann, and Adrian E. Raftery. "Strictly proper scoring rules, prediction, and estimation." Journal of the American statistical Association 102.477 (2007): 359-378.

Goldman Sachs. “Carbonomics: Innovation, Deflation and Affordable De-carbonization”, Equity Research October 13 (2020).

Good, I.J. “Rational decisions.” Journal of the Royal Statistical Society Series B 14.1 (1952): 107-114.

Hain, Linda Isabella, Julian F. Kölbel, and Markus Leippold. "Let's get physical: Comparing metrics of physical climate risk." Available at SSRN 3829831 (2021).

Hanson, Robin. "Combinatorial information market design." Information Systems Frontiers 5.1 (2003): 107-119.

Hanson, Robin, and Ryan Oprea. "A manipulator can aid prediction market accuracy." Economica 76.302 (2009): 304-314.

Hausfather, Zeke, and Glen P. Peters. "RCP8. 5 is a problematic scenario for near-term emissions." Proceedings of the National Academy of Sciences 117.45 (2020): 27791-27792.

Hayek, Friedrich August. "The use of knowledge in society." The American economic review 35.4 (1945): 519-530.

Hsu, Shi-Ling. "A prediction market for climate outcomes." U. Colo. L. Rev. 83 (2011): 179.

IEA. “Net Zero by 2050: A Roadmap for the Global Energy Sector” (2021).

IPCC. “Climate Change 2021: The physical science basis”. Summary for policymakers (2021).

Keenan, Jesse M. "A climate intelligence arms race in financial markets." Science 365.6459 (2019): 1240-1243.

Kilgour, D. Marc, and Yigal Gerchak. "Elicitation of probabilities using competitive scoring rules." Decision Analysis 1.2 (2004): 108-113.

Lucas Jr, Gary M., and Felix Mormann. "Betting on climate policy: using prediction markets to address global warming." UC Davis L. Rev. 52 (2018): 1429.

McKinsey & Company. “Banking imperatives for managing climate risk”, Risk Practice, May (2020).

McLean, Bethany, and Joe Nocera. All the devils are here: unmasking the men who bankrupted the world. Penguin UK, 2010.

NGFS, “Progress report on bridging data gaps.” Network for Greening the Financial System Technical Document, May (2021).

Oreskes, Naomi, and Erik M. Conway. Merchants of doubt: How a handful of scientists obscured the truth on issues from tobacco smoke to global warming. Bloomsbury Publishing USA, 2011.

Othman, Abraham, and Tuomas Sandholm. "The Gates Hillman prediction market." Review of Economic Design 17.2 (2013): 95-128.

Overland, Indra, and Benjamin K. Sovacool. "The misallocation of climate research funding." Energy Research & Social Science 62 (2020): 101349.

Page, Lionel, and Robert T. Clemen. "Do prediction markets produce well-calibrated probability forecasts?" The Economic Journal 123.568 (2013): 491-513.

Page, Scott E. The difference: How the power of diversity creates better groups, firms, schools, and societies-new edition. Princeton University Press, 2008.

Pielke, Roger, and Justin Ritchie. "Systemic misuse of scenarios in climate research and assessment." Available at SSRN 3581777 (2020).

Roulston, Mark S. "Performance targets and the Brier score." Meteorological Applications: A journal of forecasting, practical applications, training techniques and modelling 14.2 (2007): 185-194.

Roulston, Mark S., and Leonard A. Smith. "Evaluating probabilistic forecasts using information theory." Monthly Weather Review 130.6 (2002): 1653-1660.

Roulston, Mark S., David J. Hand, and David W. Harding. "Establishing a Real-Money Prediction Market for Climate on Decadal Horizons." AGU Fall Meeting Abstracts. Vol. 2016. 2016.

Roulston, Mark S. "Development of a Climate Prediction Market." AGU Fall Meeting Abstracts. Vol. 2017. 2017.

Sandroni, Alvaro. "At least do no harm: The use of scarce data." American Economic Journal: Microeconomics 6.1 (2014): 1-3.

Schwalm, Christopher R., Spencer Glendon, and Philip B. Duffy. "RCP8. 5 tracks cumulative CO2 emissions." Proceedings of the National Academy of Sciences 117.33 (2020): 19656-19657.

Skreta, Vasiliki, and Laura Veldkamp. "Ratings shopping and asset complexity: A theory of ratings inflation." Journal of Monetary Economics 56.5 (2009): 678-695.

Sumner, Scott, and Aaron L. Jackson. "Using prediction markets to guide global warming policy." (2008).

TCFD, Task Force on Climate-related Financial Disclosures 2020 Status Report (2020).

U.S. Securities and Exchange Commission, The Enhancement and Standardization of Climate-Related Disclosures for Investors, 17 CFR 210, 229, 232, 239 & 249 (2022).

Vandenbergh, Michael P., Kaitlin Toner Raimi, and Jonathan M. Gilligan. "Energy and climate change: A climate prediction market." UCLA L. Rev. 61 (2013): 1962.

Vecchi, Gabriel A., Kyle L. Swanson, and Brian J. Soden. "Whither hurricane activity?" Science (2008): 687-689.

Witkowski, Jens, et al. "Incentive-compatible forecasting competitions." Management Science (2022).

Wolfers, Justin, and Eric Zitzewitz. Five open questions about prediction markets. No. w12060. National Bureau of Economic Research, 2006.

Zhang, Fuqing, Wei Li, and Michael E. Mann. "Scale-dependent regional climate predictability over North America inferred from CMIP3 and CMIP5 ensemble simulations." Advances in Atmospheric Sciences 33.8 (2016): 905-918.