Catastrophe and Climate

CLIMATE RISK ASSESSMENT AND SCENARIO ANALYSIS

FEBRUARY | 2023

SOA Research INSTITUTE

SOCIETAL PURPOSE

Climate Risk Assessment and Scenario Analysis

A Case Study on Climate Change Impact on Home Prices Utilizing Bayesian Network Qualitative to Quantitative Analysis

AUTHORS Robert Lee, FCAS, MAAA

Chris Beck

Judith-Anne Brelih, ASA

Blake Fleisher

Ryan Huff

Leighton Hunley, MBA

SPONSOR Catastrophe and Climate Strategic Research Program Steering Committee

Table of Contents

Executive Summary

From physical risks like extreme wildfire events and hurricanes to transition risks like the stopping of internal combustion engine production (Els, 2021), climate-driven changes are becoming more and more intertwined with the business and financial risks managed by insurance companies and the actuarial profession. With the inherent uncertainty in both the changes in climate itself as well as the ways in which these changes impact organizations and society, actuaries need tools to evaluate all the moving pieces in a decision-useful way.

In this Society of Actuaries Research Institute (SOA) report, authored by consultants from Milliman, Inc. (Milliman), a case study is used to illustrate an iterative approach to building a model for scenario analysis. The case study develops a model of U.S. home prices, which takes the form of a Bayesian Network. The model starts as a simple and intuitive narrative, which is then incrementally supplemented with data and expert knowledge to produce a more refined assessment of the variable of interest. This report demonstrates how Bayesian analysis can be used to frame problems and perform sensitivity analysis around climate change and financial risks. We chose home price as an example, but the framework can easily be applied onto other topics of interest as well.

The final model in this case study incorporates components relating climate change variables and the housing market. It allows inputs of selected parameters from climate scenarios established by the Network of Central Banks and Supervisors for Greening the Financial System ("NGFS"). The dynamics between variables in the network are informed by data and expert knowledge, where distributions of and relationship between nodes are specified mathematically. The model then provides an integrated view of all of the components and allows researchers to 1) make inference on outcome variables, and 2) analyze the necessary conditions for the outcome variable to take a specific value.

The figure below shows a primary use case of the model: estimating the quarterly growth of Home Price Index ("HPI") in 2100 under different climate scenarios. In general, the Net Zero 2050 scenario produces the fewest extreme weather events, leading to lower home price growth. The opposite is also true: the Current Policies scenario is expected to result in the highest home price growth due to the increase in extreme weather events. Moreover, these extreme weather events generate high unpredictability in home prices in the Current Policies scenario, as indicated by the wide distribution of home price growth below.

Figure 1

DISTRIBUTION OF HPI QUARTERLY GROWTH UNDER DIFFERENT SCENARIOS

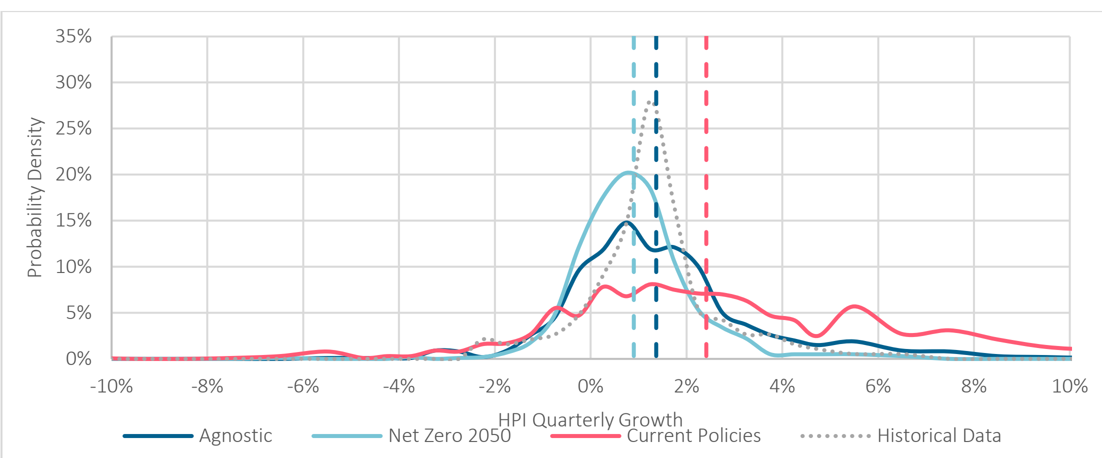

Introduction and Background

With the recent increase in frequency and severity of extreme weather events, organizations are becoming more aware of the various risks that climate change may bring. Traditionally, organizations focused their risk management efforts on physical risks—direct damages caused by extreme weather events like flooding, droughts, and storms. However, it is now apparent to many that climate-related risk extends far beyond damages to properties. For example, transition risks resulting from the transition to lower-carbon economy, as well as legal, compliance, and reputational risks can affect entities in ways other than direct physical damage to properties. In fact, even physical damages can cause economic impacts beyond just the affected properties. As markets acknowledge and begin to anticipate future risks, supply and demand can shift, causing valuation of assets to change.

While it is clear to many that there is an immediate need for climate risk assessments, it is not so clear how assessments can be done in a structured and quantitative manner. Moreover, actuaries are not only interested in inspecting the best- or worst-case scenarios, but perhaps more importantly in the distribution of the outcome as well as how that distribution changes with various moving pieces.

Scenario analysis is a core component in aiding the understanding of climate-related risks, as well as the assessments and disclosure of these risks. It helps organizations consider possible outcomes that may play out over the long term and helps drive the conversation about what may unfold in the future. The New York State Department of Financial Services ("NYDFS") recently issued guidance that New York-regulated domestic insurers are expected to use scenario analysis to inform business strategies and risk assessment and identification (Department of Financial Services, 2021). Scenario analysis can take on many forms; the Taskforce on Climate-related Financial Disclosures ("TCFD") writes in its report: "Scenario analysis can be qualitative, relying on descriptive, written narratives, or quantitative, relying on numeric data and models, or some combination of both (Task Force on Climate-related Financial Disclosures, 2017)."

There is a considerable learning curve for organizations to develop their risk management function in the context of climate change. Because of the requirement in upfront cost and knowledge, a qualitative assessment is often preferred, or simply necessary, when an organization first takes on the task of climate risk assessment. Scenario analyses of this form can be narratives or storylines that help management explore the potential range of climate change implications. As organizations gain better understanding and information on the topic, they should then be able to augment these narratives with data sets and quantitative models, and greater rigor and sophistication may be warranted. Moreover, it is often desirable to incorporate existing external scenarios and models, either developed by the entities' in-house modeling capabilities or by third-party providers. For example, as part of the NYDFS guidance, insurers are expected to shift from a qualitative approach to one rooted in quantitative assessments, using tools like geospatial and climate modeling.

Today, there is no uniform framework that is widely adopted across industries to perform scenario analysis related to climate change. Often, scenario analyses are sets of descriptive narratives asking "what-if" questions and seldom more sophisticated than varying input parameters to a financial model. This paper discusses the use of Bayesian networks as a flexible framework for organizations to represent the knowledge and information accumulated. Because a Bayesian network is a representation of a probability distribution, it contains within it the information required to answer any question that is applicable to the distribution. In addition, it can take into account the interdependency between parameters.

Because of this, there is value in a framework that allows a smooth transition from a qualitative assessment to a quantitative analysis as entities mature in their understanding of climate risk. As organizations gain more knowledge, they need to tune and refine the assessment efficiently. Moreover, the framework will need to incorporate information from multiple experts, models, and experiences in a coherent manner.

To integrate information from different sources, we can turn to Bayes' theorem. The Bayes' theorem states that the relationship between two dependent events A and B, given probabilities $P(A)$ and $P(B)$, can be described by the following relationship:

Equation 1
BAYES' THEOREM

$$P(B|A) = \frac{P(A|B)P(B)}{P(A)}$$

The simple yet elegant equation is the basis of many modern AI systems for probabilistic inference, one of which being the Bayesian Network. The Bayesian Network defines probabilistic independencies and dependencies among variables in the network. Because of its graphical representation, researchers can replicate the modular nature of the climate analysis. With a Bayesian Network, this paper uses a case study to outline a flexible framework that is suitable for organizations at any point in their maturity pathway, in which they can incrementally incorporate data and refine assumptions to produce assessments of increasing quality and granularity.

Methodology

BAYESIAN NETWORK

Represented as a network of interconnected nodes, Bayesian Networks are a way of defining complex phenomena and the associated dependency relationships between variables in that system. They can be thought of as an expert system that formalize the knowledge possessed by multiple experts. It provides an intuitive graphical framework to describe the relationships in a Directed Acyclic Graph ("DAG"). The dependence relationships are defined with probability distributions so that prediction and inference can be done on the variables in the network (Yang, 2019).

The DAG is a graphical representation of the Bayesian Network and outlines the dependency relationships. Figure 2 shows an example of a DAG of a portion of the case study model, with arrows indicating the hierarchical relationship between nodes. The DAG contains a set of nodes that correspond to the variables in the Hurricane Impact component of the model, which we determined to have an impact on home prices. Between the nodes are arcs, also called "edges", that define the dependency relationships between the variables/nodes. A direct relationship is indicated by an arc that connects two nodes from the parent node to the child node such as *Temperature* → *Hurricane Count*. An indirect relationship is indicated by a path of arcs between two nodes that contains intermediate nodes between, such as Temperature → *Hurricane Count* → *Hurricane Impact*. The nodes in a path must be distinct so as not to create cycles in the relationships. This property defines the acyclic nature of the DAG.

Figure 2
EXAMPLE OF DIRECTED ACRYLIC GRAPH

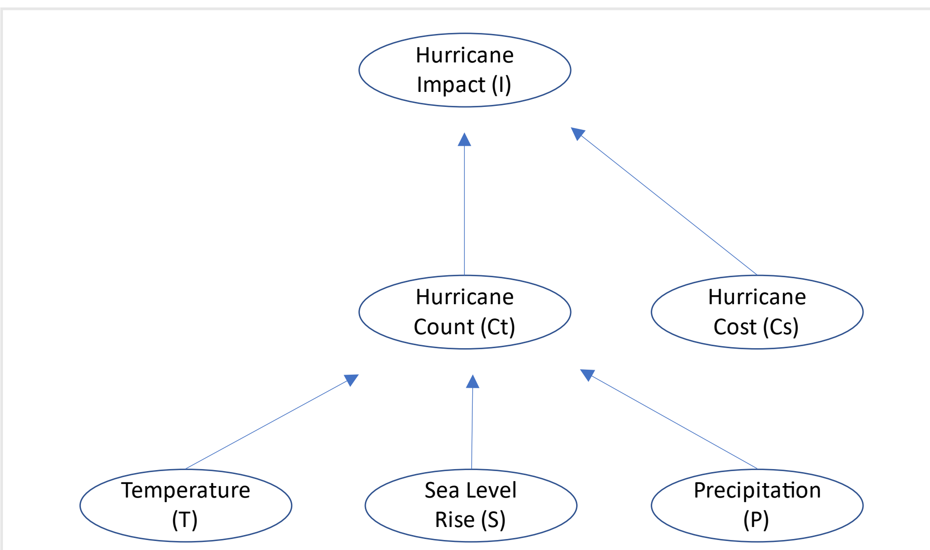

Beyond defining the relationships, the joint probability distribution for the network must be defined to perform meaningful quantitative analysis on variables. The global joint probability distribution for the whole network would have too many parameters to be completely defined, but for any set of nodes that have no arcs connecting them, one can assume independence and make the following probability statement:

$$P(T, S, P, Cs, Ct, I) = P(T)P(S)P(P)P(Cs)P(Ct|T, S, P)P(I|Cs, Ct)$$

The conditional independence assumption is fundamental to the Bayesian Network and simplifies all the calculations to a manageable level.

With conditional independence, one can then specify the relationship between nodes one at a time. Generally, relationships between nodes can be characterized in the form of node probability tables ("NPT"). The NPT specifies the nature of the node and its relationship with any parent nodes in the form of a distribution. If for example the child node has a Normal distribution, the mean of the Normal distribution could be calculated from the values of the parent nodes, which follow their own respective distributions. The variance of the Normal distribution can be either a fixed value or a value dependent on parent nodes, used to indicate the researcher's confidence in the relationship. By using the values of parent nodes to influence the values of the child nodes, the relationships of the network are quantified. In practice, these relationships can be defined by expert judgement or be informed by separate research or models.

Nodes can be either discrete or continuous depending on the variable being defined. Discrete nodes such as Boolean or Ranked nodes are limited in the number of states they can take. For example, a Boolean node can only be in one of two states, while a Ranked node has multiple states with a particular order such as low, medium, and high. Ranked nodes can be implemented using Truncated Normal ("TNormal") distributions, which have defined end points that can limit the range of values. The continuous values from the TNormal distribution can be discretized to specify the different states in an increasing order.

Continuous nodes can be specified using a full range of statistical distributions to define NPTs such as Uniform, TNormal, Beta, etc. as well as formulae involving real numbers. Historically, continuous nodes were manually discretized to make the calculations more manageable, but this always results in a loss of accuracy and requires significant effort. Recent developments in Bayesian networks have worked to solve this problem by Dynamic Discretization. Dynamic Discretization uses simulations to automatically determine an appropriate discrete form which approximates the true continuous values within a specified tolerance (see, for example, Zheng 2021).

Once the NPTs are defined for each of the nodes, the researcher is able to query the Bayesian Network to make either predictions or inferences. A prediction follows the paths in the DAG to forecast outcome variables by calculating the marginal probability for a node. Inference is when data is provided to child nodes, and distributions about the parent nodes are updated in light of the data. The process of updating probabilities in a Bayesian Network is called propagation. Propagation updates the marginal probabilities of the unobserved variables in the Bayesian Network.

In the context of scenario testing, observations or assumptions about one or more nodes can be provided, and the resulting distributions of other variables can be studied. This allows the researcher to ask if-then questions to the network. For example, given GDP, Carbon Price, Land Use Change, and Energy from Fossil Fuels, what is the distribution of Greenhouse Gas Concentration? On the other hand, given a specific Greenhouse Gas Concentration, what does that imply for GDP, Carbon Price, Land Use Change, and Energy from Fossil Fuels?

Below are some common terminologies in the context of Bayesian Networks:

- The *prior* is the assumed distribution of a node before evidence / data.
- The *posterior* is the updated assumption of a distribution of a node after evidence / data.

CRISALIS APPROACH

For this study, the CRisALIS[1] ("Complex Risk Analysis") approach was used to model climate risk. CRisALIS models nodes as a network of interconnected elements, representing a causal set of relationships between the drivers. For

[1] https://us.milliman.com/en/products/complexriskanalysis

this paper, CRisALIS is used to model the impact of climate risk on quarterly growth in home prices using the FHFA all-transaction home price index (“HPI”). The CRisALIS approach utilizes Bayesian Networks to construct networks and provides users with the ability to create nodes and specify relationships in the model. Using these models, we can perform analysis such as node sensitivity and influence metrics.

The modeling software provides a graphical interface for the user to specify the configuration of the network and nodes. Users can specify the distributions of nodes and use directed edges to specify causal relationships.

Figure 3 below shows an example of how users may specify the distribution of a child node, and Figure 4 provides a visualization of a sample network.

Figure 3
EXAMPLE OF NODE DEFINITION

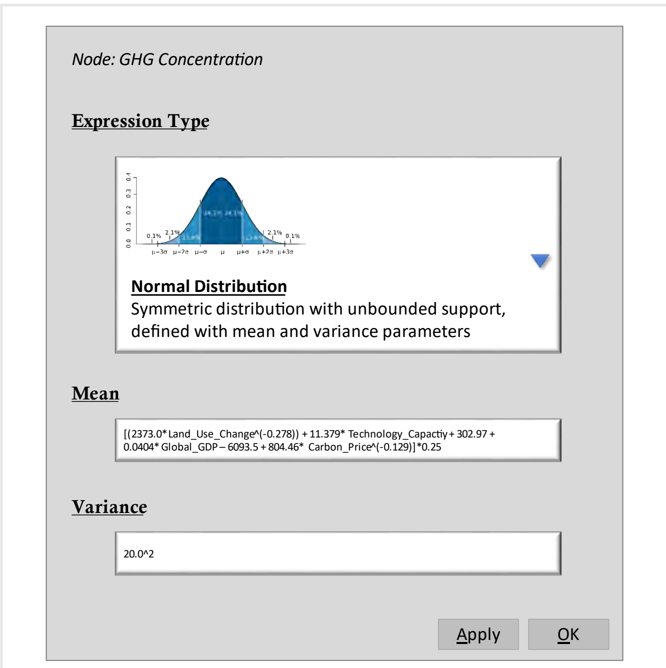

Figure 4
EXAMPLE OF NETWORK STRUCTURE

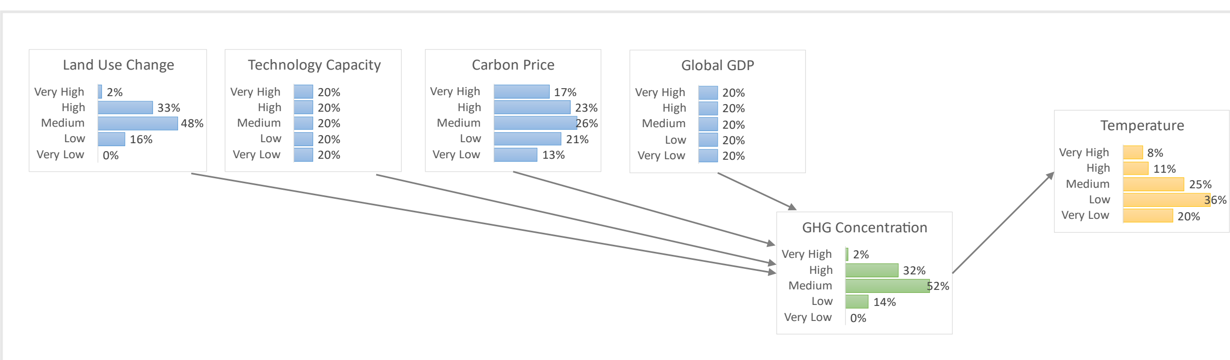

Please note that the use of specific software is not necessary for this framework—the framework is centered around Bayesian Networks, and the CRisALIS approach utilizes a software that implements the relevant algorithms and features around the probabilistic model. Other popular tools that implement Bayesian Network capabilities include the BUGS (Bayesian Inference Using Gibbs Sampling) project and PyMC3.

CLIMATE SCENARIOS

Climate change is driven by a myriad of human as well as nature factors. Because the future development of these factors is deeply uncertain, the resulting climate change is also difficult to predict. While we cannot predict the future, climate scenarios can be used to explore possible futures and the relationship between the assumptions they depend upon.

Climate change scenarios are projections of what can happen given plausible and internally consistent assumptions of factors that affect climate change. Each set of assumptions and the resulting projected climate is a pathway that provide insight on how the factors may evolve in the future.

The different components of the human factors that contribute to climate change are deeply intertwined. For example, the production of energy drives the economy, which both affects and is affected by population growth. Similarly, nature factors like atmospheric chemistry and ecosystems also have a complex relationship with each other and with the human factors. Integrated Assessment Models ("IAMs") integrate representations of multiple human and nature systems and provide a way to assess the highly complex and nonlinear relationships between the different components. Some of the most recognized IAMs include (Bishkek, 2018):

- Global Change Assessment Model ("GCAM")
- The Integrated model to Assess the Global Environment ("IMAGE")
- Model for Energy Supply Strategy Alternatives and their General Environmental Impact ("MESSAGE")
- Regionalized Model of Investments and Technological Development (REMIND)

Using one or more of these integrated models, researchers can explore plausible futures by means of "scenarios" or "storylines". Scenarios may range from different economic growth to attitudes towards the environment. Alternatively, scenarios can also restrict the outputs of the model such as a by setting a limit on emissions or warming.

The Network of Central Banks and Supervisors for Greening the Financial System ("NGFS"), the Intergovernmental Panel on Climate Change ("IPCC"), and the Internal Energy Agency ("IEA") are some of the most recognized public institutions that produce climate scenarios using IAMs. Each institution takes a different approach in generating scenarios in terms of granularity, assumptions considered, as well as the philosophy behind the construction of scenarios. Upon discussion with the POG, we elected to focus on the NGFS scenarios, which are used by regulators in many countries for climate stress testing for insurers.

As a high-level summary, three main types of variables are available in NGFS climate scenarios:

- **Physical risk variables** such as extreme weather events like flooding and extreme temperature changes, and gradual changes in climate like sea-level rise and changes in agricultural yields
- **Transition risk variables** such as governmental policies like carbon tax and subsidies, and technological trends like renewal energy electrification of motor vehicles.
- **Macroeconomic risk variables** such as GDP growth, unemployment, interest rates, and inflation.

Case Study Problem

This paper uses a case study problem to illustrate how the framework can be implemented in practice. Below are the background and description of the problem:

Impacts of extreme climate events may extend far beyond losses of lives and the direct damage caused to properties at the locations that they occur. With the increased frequency and severity of "unprecedented" events like the wildfires in California and sandstorms in China, there is increased awareness of climate change and the risk that it brings. It is potentially only a matter of time when climate risk will become a standard factor used to evaluate real estate investments just like school district and property size for residential properties. In fact, the impact of the perception of climate risk can already be seen in some areas:

- Researchers found that in Florida, there is evidence of price decline in census tracts that are more exposed to sea level rise after 2016. By 2020, prices in these markets were 5-10 percent below trend (Keys and Mulder, 2020).
- A study found that home buyers and sellers tend to revise their perceptions of fire risk after a major fire, such that house prices drop by up to 15% in neighboring unburned towns (Loomis, 2004).

However, the impact of climate risk on housing prices is not necessarily one sided. In Napa County, California, where the LNU Lightning Complex Fire burned more than 363,000 acres in 2020, home prices have seen a more than 30% increase because of higher buyer demand and low housing supply (Trapasso, 2020). Separately, those that are displaced need to find an alternative location to settle down, driving house prices up in areas where there is less perceived climate risk.

With uncertainties in the occurrences of future extreme climate events as well as in the complex dynamics of how these events affect home prices, climate change then becomes an asset risk for organizations that have large real estate portfolios. Research on this topic is limited and there is no widely accepted model to relate climate change impacts and home prices. The rest of this paper discusses how one may address this task.

Model Building Process

MAIN GOAL OF THE MODEL

We will demonstrate how an organization with no prior knowledge and research, gather, and eventually synthesize information to create a coherent view of the topic. This discussion pertains specifically to this case study, but the approach is applicable to any problem.

The purpose of scenario analysis is to understand, and subsequently plan for the uncertain future. While there is plenty of existing research and established information on climate change, no one research paper encompasses every single aspect of climate change. In addition, entities need to have a way to incorporate new knowledge as data becomes available and as they learn new information.

With that in mind, the goal of this framework is to create a model—a "subject matter expert" that consumes and reflects research, expert knowledge, and data in a systematic manner. The overall research process consists of the following 4 steps:

1. Define the variable of interest
2. Literature review
3. Create a network of nodes and relationships
4. Assess network feedback and refine nodes and relationships

DEFINING THE VARIABLE OF INTEREST

As with typical research processes, this framework begins with an end goal in mind. What is different is that this research process does not start with a hypothesis but instead with a variable. The initial phases of climate analysis are often "fuzzy", where stakeholders may not yet have a good idea of what questions need to be answered. The value of scenario analysis is that it can answer not just one but many questions in the form of scenario inputs.

Once an initial variable of interest is selected, it is important to interview stakeholders to understand and address any concerns. For example, having decided that the topic of interest is the impact of physical climate risk on the housing market, we consulted the Project Oversight Group ("POG") on the initial selection of home price as the variable of interest. Feedback on the selection included:

- Home prices do not fully reflect general housing costs, which includes rent
- Home prices do not reflect changes in the economy as quickly as rent prices
- Home prices evolve over time

In our case, the all-transactions home price index ("HPI") was selected as we believe it is a suitable barometer for the housing market at large. Other housing costs like rent prices tend to be correlated (possibly lagged, but correlated nonetheless) with home prices, so the model could be extended to analyze rent prices. To address the issue of changing home price index over time, the variable was transformed into a quarterly growth rate, so the variable is comparable over time.

LITERATURE REVIEW

PURPOSE OF LITERATURE REVIEW

Once the variable of interest is defined, the next step is to determine what other variables should be included in the model. This step serves as the exploratory phase of the model building process, where researchers conduct an initial review of information available on the topic of interest.

In typical modeling exercises, it is often necessary to have a dataset that contains all variables in the model. The use of Bayesian Network in this framework, however, allows ingestion of data in a different format. As will be discussed later, relationships between variables can be set up one at a time, so researchers are free to conduct this research in a "piecemeal" manner and are not limited by the availability of a dataset that encompasses all variables.

One of the purposes of this exploratory phase is to amass information and knowledge on the topic. The most effective way to do this is to learn from existing literature. We conducted an extensive research process whereby we identified many scholarly papers that referenced the impact of climate change on home prices. Specifically, we focused on literature that provided a causal effect of climate risk on home prices such that these relationships can be implemented in our model framework. Through our research the perils of climate risk with the most literature coverage were flood, wildfire, and hurricane. We focused on these three perils in our model and aggregated the causal relationships to determine overlap across papers and ultimately selected the relationships that were most transparent to the end user of the model.

THE TREE BUILDING METHOD

Climate change is a very broad topic and there is a large amount of information available, so it is important to have a systematic way of scanning and consolidating research results. One way to effectively survey the landscape is the Tree Building Method. The "tree" discussed here refers to a tree data structure whose nodes represent variables or concepts in existing research. The building of this tree consists of the following steps:

1. Start with the variable of interest as the root node
2. Identify existing literatures that mention the variable. Add any relevant concepts or variables in these literature as child nodes to the root node.
3. Repeat step (2) for each newly added child node. Repeat until the depth of the tree is satisfactory.

Below is an illustration of how this may work in practice, using the case study as an example.

First Level

The first source discusses how the insurance premium and the belief in climate related risks ("Climate Beliefs" below) can affect home prices. Figure 5 below shows how we begin to build the first level of the tree with this information:

Figure 5
FIRST LEVEL OF TREE

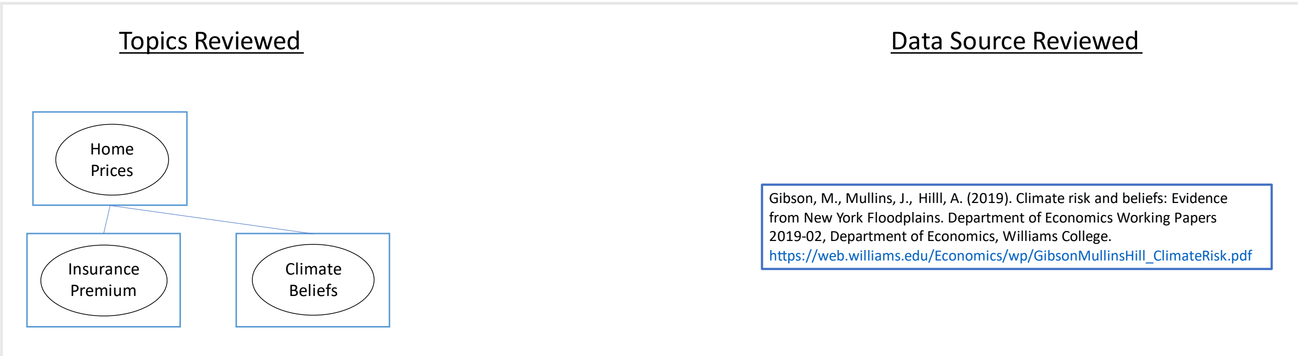

Second Level

Then, we drill down further into the first child node (Insurance Premium) in the context of climate change. A prior Society of Actuaries Research Institute study provides estimates of the uninsured flood exposure in the contiguous United States and the relationship between increased severity of extreme flooding events due to sea level rise. Two selected variables are added as child nodes. It is up to the researcher what and how many variables are within the scope of the modeling process and should be added to the tree. Reliance on existing research, expert knowledge, and deliberation on the problem scope allows us to filter out variables extraneous to the framework. For example, the study also discusses National Flood Insurance Plan (NFIP) take up rate, which we elected to not give further consideration—while it is relevant to the topic of climate change, we did not think it was central to our investigation of how climate change can affect home prices. We add this information to the existing tree, creating the second level like illustrated in Figure 6:

Figure 6
SECOND LEVEL OF TREE

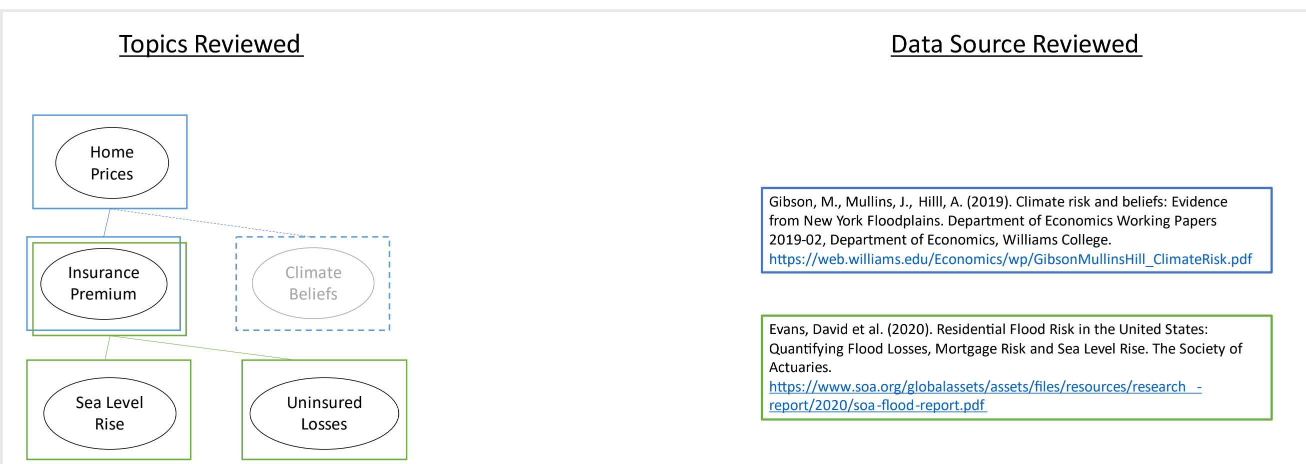

Ultimate Level

We continued this process until we reached a level which we considered a good stopping point for this branch. What is considered a good stopping point is entirely up to the researchers. While the tree can include very granular variables and causal relationships in theory, it is helpful to avoid over-complicating the framework with overly detailed information that is not helpful for addressing the problem. Since the ultimate goal of the model is to conduct scenario analyses, it is helpful to ask, “how are the scenarios defined?” Once the tree reaches the variables that are used to define the scenarios, there is little benefit to going beyond that. In our case study problem, we wanted to look at scenarios defined by NGFS variables. For example, we would like to include global GDP as a scenario parameter, but the “how” and “why” of GDP were outside the scope of our study. Figure 7 illustrates the tree reaching the ultimate level.

Figure 7

ULTIMATE LEVEL OF TREE

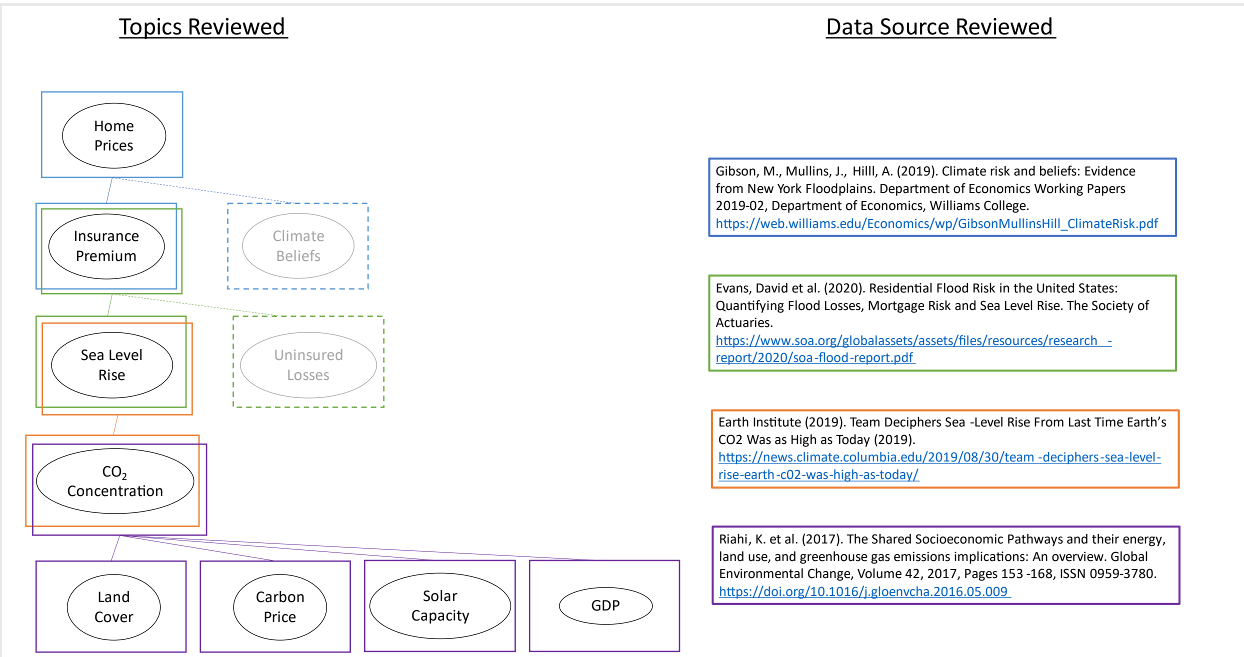

Other Branches

After reaching the desired depth, we can begin building the other branches of the tree. This process continues until the researchers believe they have achieved a satisfactory coverage of information on the topic. Figure 8 shows what a tree with multiple branches may look like:

Figure 8
OTHER BRANCHES OF TREE

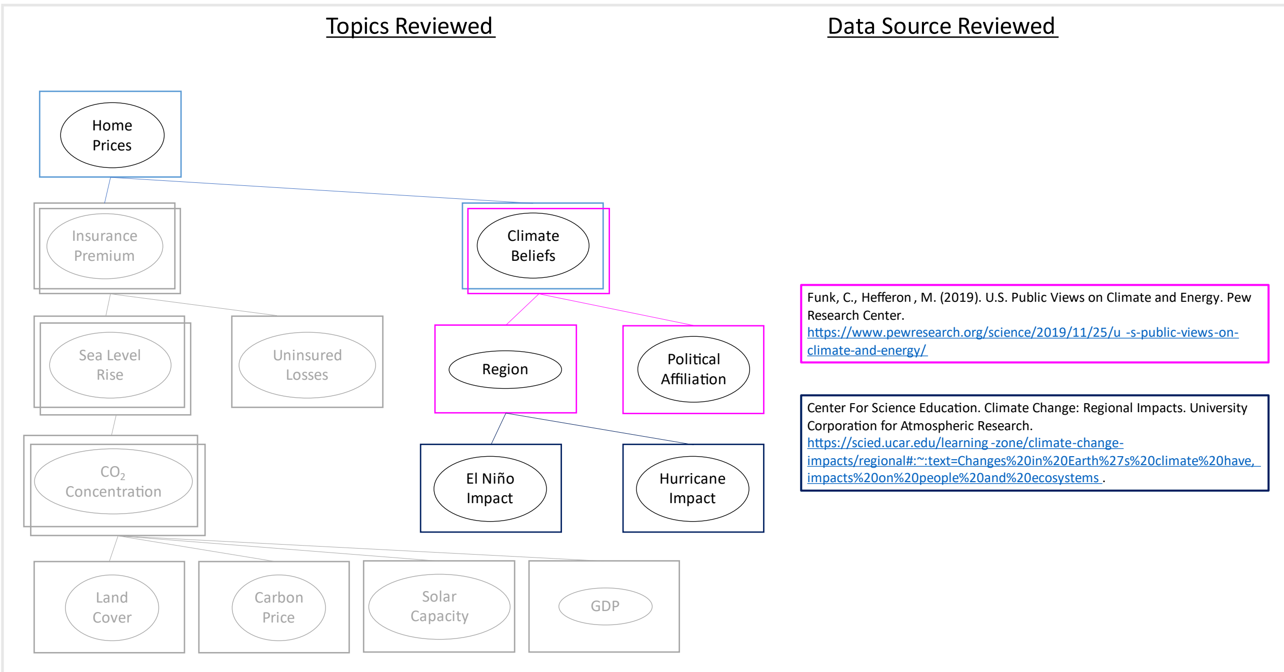

In general, the tree needs not be directed—there does not need to be a hierarchical relationship between child and parent nodes. For example, in the article *Climate Change: Regional Impacts* by the Center of Science Education, "El Niño Impact" does not drive "Region", nor does "Region" drive "El Niño Impact". The two are simply related.

At this point, there is no need to be overly parsimonious on what variables to include—this is still the exploratory phase, and it is often easier to cast a wide net on potentially relevant variables than to miss out on those that appear tangential at first glance, which may turn out to lead to a whole branch of other important information.

The Tree Building Method is an effective way to consolidate and summarize research on multiple pieces of information in an organized manner. It is especially applicable for building Bayesian Networks because the tree itself resembles the structure of the network that will be built.

In this illustration we discussed research results from existing literature, but nothing prevents us from incorporating expertise from subject matter experts or even datasets, which makes this exploratory research process very flexible and efficient.

CREATE A NETWORK OF NODES AND RELATIONSHIPS

With the tree built and potential variables listed, the next step is to create the network that will eventually become the model. This consists of two steps: defining the structure and defining the nodes.

DEFINING THE STRUCTURE

The first step of creating the model is to define the network structure, which involves placing each variable in the tree from the previous step into a directed and acyclic graph. As described in the methodology section, Bayesian Networks are DAGs, where 1) a hierarchy exists between any two nodes, and 2) following the directions of the edges connecting nodes will never form a closed loop.

Note that at this point, variables, or nodes in the network, are largely still concepts: all that is needed is a qualitative description of what the variables represent, and datasets are not yet needed. In the next step, *Defining the Nodes,* assumptions or information of each variable will be used to give a precise definition of each variable.

Unlike the Tree Building Method where researchers are encouraged to cast a wide net on potentially relevant variables, this step requires more careful judgement on what variables make it into the model. With multiple researchers, this process often involves multiple rounds of discussions. One exercise to streamline these discussions is the Delphi Method. The Delphi Method is a structured communication technique, in which a panel of participants express their opinions in a guided manner. The method can roughly be described with the following cycle:

Figure 9
ILLUSTRATION OF THE DELPHI METHOD

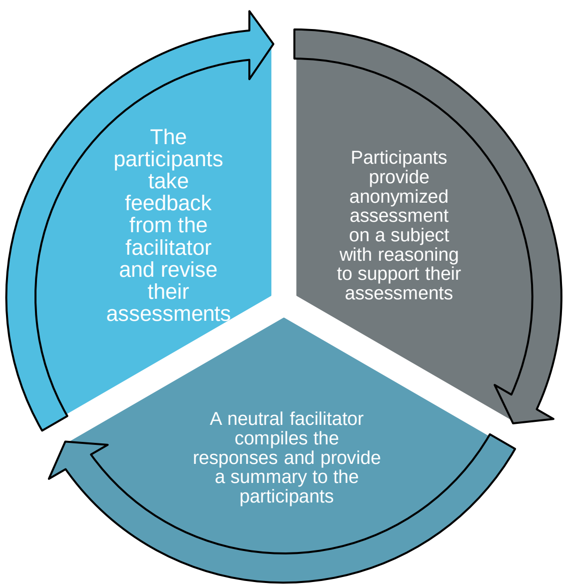

In this use case, the participants are the researchers, and their assessments can be variables to include in the model or the proposed model structure. The conventional Delphi Method calls for all assessments to be anonymous to prevent personal biases from the process, however a modified approach can be taken if that is not a concern. To further simplify the process, one of the researchers can also take the role of the facilitator. Ultimately, the key to this process is to conduct multiple rounds of discussion with feedback between each, where participants are encouraged to comment on the responses of others and revise their opinions in real time.

For this case study, we conducted multiple meetings including internal meetings within the research team as well as external meetings with the POG. Because of separable nature of networks, not all researchers have to be at all meetings. In our case, subject matter experts were assigned to various components of the model development efforts.

DEFINING THE NODES

The second step of creating the model is to define the nodes and the relationships between them. In a Bayesian Network, each node is defined by a distribution, and the existence of an edge between two nodes is an indication that the child node depends on the value of the parent node—to be more precise, child nodes are specified as conditional distributions or arithmetic expressions that depend on the parent node's value.

Nodes are like statistical variables—they can either be discrete or continuous, and discrete nodes can further be classified as numeric, ordinal, or nominal. Depending on the type of the node, the researcher may have the option to use a parametric distribution to define the node, or manually specify a non-parametric probability distribution of the variable.

Unlike supervised learning techniques like linear models or random forest where models are fit on one single dataset, the Bayesian Network approach is significantly more fragmented. Users have complete freedom to specify (unconditional or conditional) distributions of nodes however they seem fit. This is especially helpful for climate analysis, where there may be a mixture of variables with readily available data and concepts that are not precisely defined.

For example, in this case study, we wished to represent "Technology Capacity" as a concept of the society's capability to efficiently generate green energy but did not know of any good data sources or even a way to precisely define it in the initial phases of the modeling process. We set up a "ranked node"—an ordinal node—and specified a weak prior by assigning equal probability for each outcome for this variable. On the other hand, there was ample data for the node representing interest rate, so we were able to fit a distribution to that data and give the node a more precise prior distribution.

As entities gain more experience and expertise on the topic or as data becomes available, researchers may revise the node distribution or even the node type. In the example of *Technology Capacity*, we opted to use data from NGFS scenarios, suggested by one of the POG members, to assume a distribution of percentage of electricity production capacity that depend on fossil fuels as compared to green energy. The node was renamed to "*Energy from Fossil Fuels*".

Because of the "fragmented" nature of Bayesian Networks, only the child nodes of the revised node need to be updated, and nothing further upstream or downstream. Equipped with this flexibility, we found it helpful to not get fixated on finding the "perfect" distributions. Setting up the model with reasonable node distributions and iterating quickly can often yield fruitful insights, as will be discussed in the next section.

The ultimate parent nodes (nodes with no parent nodes) do not depend on other nodes, so specifying their distributions is fairly straightforward. On the other hand, nodes with one or more parents need to be defined as conditional distributions. Depending on the type of parent and child nodes, different types of conditional distributions are needed.

Table 1 illustrates the format of the conditional distribution for combinations of parent and child node types, and Table 2 through Table 5 give an illustration of each case. Please note that these values are for illustrative purposes only.

Table 1
NODE SPECIFICATION BY PARENT AND CHILD NODE TYPES

		CHILD NODE	
		CATEGORICAL	CONTINUOUS
PARENT NODE	CATEGORICAL	Joint probability table (Table 2)	One child node distribution for each possible value of parent node (Table 3)
	CONTINUOUS	"If/Then" rules or binning (Table 4)	Parameter of child node distribution as a function of parent node value (Table 5)

Table 2
EXAMPLE OF NPT WITH CATEGORICAL PARENT NODE AND CATEGORICAL CHILD NODE

		CATEGORICAL CHILD NODE				
CATEGORICAL PARENT NODE	IF CO2 CONCENTRATION IS…	THEN TEMPERATURE ANOMALY WILL BE…				
		VERY LOW	LOW	MEDIUM	HIGH	VERY HIGH
		with probability				
	LOW	15%	60%	15%	6%	4%
	MEDIUM	10%	10%	50%	20%	10%
	HIGH	5%	5%	10%	60%	20%

Table 3
EXAMPLE OF NPT WITH CATEGORICAL PARENT NODE AND CONTINUOUS CHILD NODE

		CONTINUOUS CHILD NODE
CATEGORICAL PARENT NODE	IF CO2 CONCENTRATION IS…	THEN TEMPERATURE ANOMALY (°C) WILL FOLLOW A DISTRIBUTION OF
	LOW	Normal(0.5, 0.5)
	MEDIUM	Normal(1.5, 0.75)
	HIGH	Normal(3.5, 1.0)

Table 4
EXAMPLE OF NPT WITH CONTINUOUS PARENT NODE AND CATEGORICAL CHILD NODE

		CATEGORICAL CHILD NODE				
CONTINUOUS PARENT NODE	IF CO2 CONCENTRATION (PPM) IS…	THEN TEMPERATURE ANOMALY WILL BE…				
		VERY LOW	LOW	MEDIUM	HIGH	VERY HIGH
		if…				
	x	$x < 400$	$400 \leq x < 450$	$450 \leq x < 500$	$500 \leq x < 600$	$600 \leq x$

Table 5

EXAMPLE OF NPT WITH CONTINUOUS PARENT NODE AND CONTINUOUS CHILD NODE

CONTINUOUS PARENT NODE		**CONTINUOUS CHILD NODE**
	IF CO2 CONCENTRATION (PPM) IS…	**THEN TEMPERATURE ANOMALY (°C) WILL FOLLOW A DISTRIBUTION OF**
	X	Normal(0.0143 * x - 4.766, 0.215)

In cases where a child node has both categorical and continuous parent nodes, the definition of the child node will be a combination of Table 3 and Table 4 above, where each distribution in Table 3 will be modified to depend on the continuous parent variable (i.e., include the parent node's value in one or more of its parameters).

While categorical nodes are easy to set up, especially when there is no reliable information with which to fit a distribution, it should be noted that size of the joint probability table of a child node increases exponentially with the number of categorical parent nodes it has. Specifically, the size of the joint probability table of the child node equals the product of the number of categories from both the parent and child nodes. For example, a 5-category child with a 5-category parent requires a joint probability table of size 25 (= 5*5), while a 5-category child with two 5-category parents requires a joint probability table of size 125 (=5*5*5).

FEEDBACK AND REFINE NODES AND RELATIONSHIPS

After the network structure and node distributions are defined, researchers can query the model. Since the model is built to represent the information gathered from the research and data in prior steps, querying the model is comparable to learning from an expert that possesses all that knowledge.

The rest of this report discusses different ways that researchers can query the model. Based on the learnings from the answers to those queries, researchers can then update the network structure, node definitions, or both.

The value of this approach is that the updates can be done in a piecemeal manner. For example, if a new GCAM model becomes available indicating another view of the relationship between *GDP Growth* and *Greenhouse Gas (GHG) Concentration*, the researcher only needs to update that portion of the model, and the remaining parts are unaffected. Similarly, if the researcher wishes to update the assumed distribution (conditional or unconditional) of a node, there is no need to revisit all the other nodes in the model. As alluded to in an earlier section, this "fragmented" approach is especially applicable for climate analysis, where knowledge can be scattered, and the science is constantly changing.

CONSIDERATIONS

It should be noted that the process described above has the ultimate goal of representing the researchers' expertise in a model. To the extent that the researchers are not able to codify knowledge or insights into the model, the model will not be able to reflect those pieces. For example, in the tree building method, there is a risk that the researcher does not identify or capture key relationships, and the resulting model will not reflect said relationships. Similarly, if node probability distributions are mis-specified, the model estimates will be biased.

Final Model

STRUCTURE OF THE MODEL

Figure 10

FINAL MODEL STRUCTURE

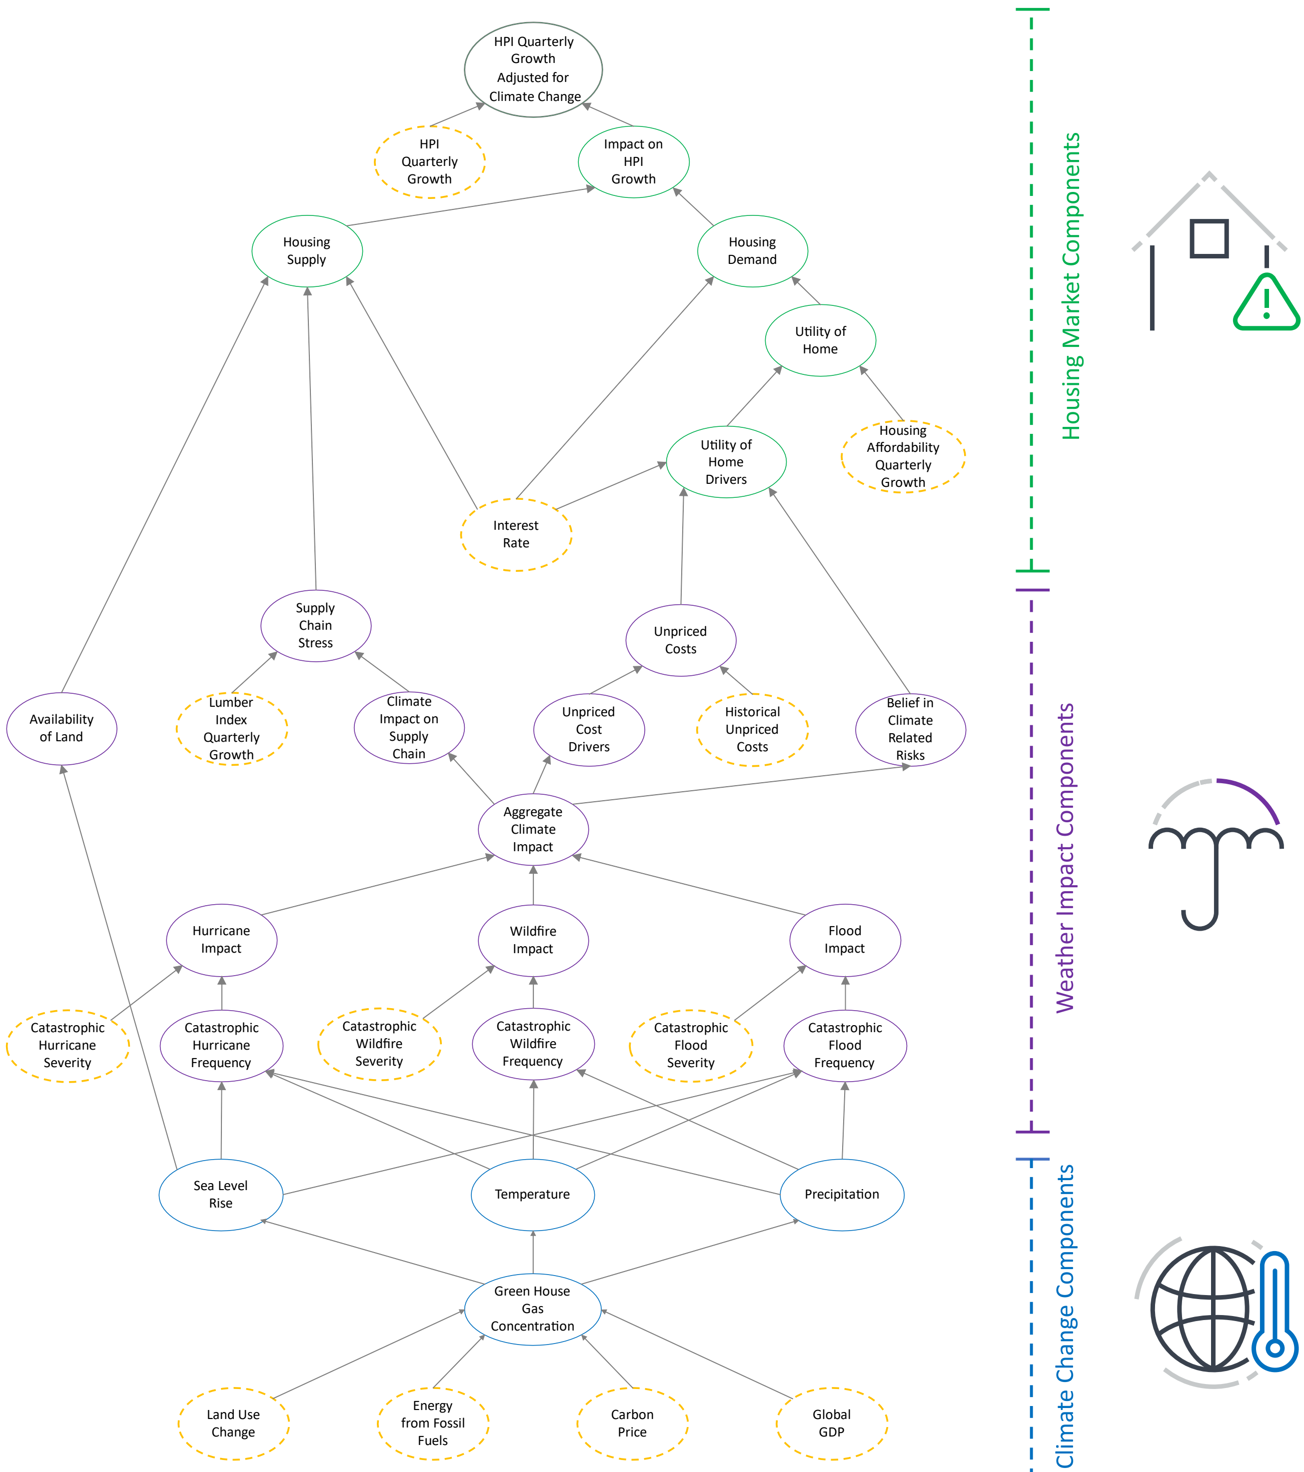

The structure of the final model is shown in Figure 10 above. Nodes with no parents are highlighted with orange dashed borders, and the arrows between nodes indicate the parent → child relationship.

The final model consists of three main components: a climate change component, a weather component, and a housing market component.

The structure and node definitions of the climate/weather components are primarily informed by 1) the NGFS scenario database, and 2) historical extreme weather event data. Four selected variables from the NGFS scenarios constitute the four main inputs for scenario analysis, which drive the catastrophic weather event frequency nodes. Severities are assumed to be independent of climate and distributions are fitted using historical data from the National Oceanic and Atmospheric Administration (NOAA). The frequency and severity of each peril are then combined, and aggregated, to form the Aggregate Climate Impact node.

The housing market component consists of supply and demand variables adjusted by the *Aggregate Climate Impact* node. For housing supply, we are interested in capturing the effect of extreme weather events as it pertains to the general availability of land (as a result of sea level rise) and the stress on supply chain in building materials. For housing demand, we wanted to capture potential buyers' reaction to extreme weather events, which considers unpriced costs and belief in climate related risks. Depending on whether a region is affected or unaffected by extreme climate events, the effect of Belief in Climate Related Risks can either have a positive or negative impact on Utility of Home. For example, in catastrophe-prone areas, the belief in climate related risks will decrease the utility of owning a home; on the other hand, belief in climate related risk will likely increase one's utility of owning a home in a relatively catastrophe-free area nearby. In this model, we opted to assume a positive relationship between Belief in Climate Related Risks and Utility of Home. The model can be modified to consider the opposite case or be customized to provide a more granular view in regions by adding peril-specific Belief in Climate Related Risks nodes.

Interest rate is included in both the supply and demand side because we believe it is an important variable to consider. *Housing Supply* and *Housing Demand* are then combined to modify the historical distribution of *HPI Quarterly Growth*, where the final node, *HPI Quarterly Growth Adjusted for Climate Change* has a positive correlation with *Housing Demand* and a negative correlation with *Housing Supply.*

We noted that other variables such as population growth and age trends also have an impact on home price growth. These variables were intentionally left out because the goal of this model is to evaluate the impact of climate change on home price growth, with the average home price growth reflective of historical data. In other words, by excluding these variables, we are able to focus on the direct impact of climate change.

TIME HORIZON

The purpose of scenario analysis is to answer the question: what will happen to *y* in the future given observations *x* now? Time horizon is especially relevant in the context of climate scenario analysis, where variables of interest are typically far into the future. To address this in the model, the climate/weather component reflects the 2100 Greenhouse Gas Concentration as a result of observed climate variables in 2030. In other words, the model takes input variables Land Use Change, Energy from Fossil Fuels, Carbon Price, and Global GDP in 2030, and estimates the resulting Greenhouse Gas Concentration in 2100. Since the trajectories of these variables are already projected by the NGFS scenarios (each NGFS scenario has a 2100 GHG concentration and corresponding 2030 input variables), we are able to extract the relationship and place it in the model. In cases where the full projection is not available, one can use a smaller time step (e.g., parent nodes at time t and child node at time t+1) and simulate the model iteratively to investigate scenarios of any future time.

On the other hand, the housing market component of the model is assumed to be time invariant—that is, the relationship between nodes within the housing market component is the same regardless of time period. Within the scope of this analysis, we believe that is a reasonable and accurate approach.

Combining these components, the model can then answer the question: What will happen to HPI Quarterly Growth in 2100, given observations of Land Use Change, Energy from Fossil Fuels, Carbon Price, and Global GDP in 2030?

QUERYING THE MODEL

There are a few ways the model can be used to make predictions and inferences. The following are examples of each in the context of climate analysis:

PREDICTION

Prediction in this context is the updating of distribution of the variable of interest given information, also called "evidence," for the parent nodes. Researchers can specify values for nodes in the network to observe how it affects downstream variables. This allows for a transparent and internally consistent way of conducting scenario analysis: the assumptions are clearly laid out in the node definitions, and changes in inputs only affect outputs to the extent the network specifies (compared to querying an expert, which may involve subjective judgement). In addition, researchers are able to infer a distribution of variables instead of just a point estimate, which is necessary for deriving metrics like Value-at-Risk ("VaR") and expected utility.

The focus needs not only be on the variable of interest, however. Nothing in this process prevents us from learning more about the distribution of intermediate variables. This is valuable because knowledge of different variables (e.g., A→B, B→C, C→D) may be scattered across different experts or datasets, and the interconnectedness of the network allows us to infer the relationship between any of those variables (e.g., A→C or A→D).

Alternatively, researchers can also assume a prior distribution for the ultimate parent nodes instead of setting them to a value. This is a more general case of the scenario analysis use case described above, where the "prior" distribution is a degenerate one with $P(parent\ node = evidence) = 1$.

Recall that the network is specified with conditional distributions; one way to get the unconditional distribution from the conditional distributions is to sample the model following the steps below:

1. Sample the ultimate parent nodes
2. Use values from step (1) to sample child nodes of those from step (1)
3. Repeat (2) until we reach the ultimate child node (i.e., the variable of interest)
4. Repeat steps (1) to (3) until there are enough samples to approximate a distribution of the variable of interest

In the final model, we examined two NGFS scenarios: the Net Zero 2050 scenario and the Current Policies scenario. Each NGFS scenario explores a different set of assumptions for how climate policy, emissions, and temperatures evolve. The Net Zero 2050 scenario limits global warming to 1.5°C through stringent climate policies and innovation, reaching global net zero CO_2 emissions around 2050. On the other hand, the Current Policies scenario assumes that only currently implemented policies are preserved, leading to high physical risks. To incorporate these scenarios into the model, we set the values of the ultimate parent nodes to the values in the NGFS scenario dataset.

Separately, we included an "agnostic scenario" using prior distributions for the ultimate parent nodes. Weak priors are selected for the climate-related ultimate parent nodes, where they are assumed to follow uniform distributions with reasonable upper and lower bounds. This represents a view where there is no strong inclination to believe certain values will happen. For example, in the case of Carbon Prices, the prior of $Uniform(\$5.0, \$300.0)$ represents a view in which we believed it is equally likely for Carbon Prices to fall anywhere in the range between $5 and $300.

Table 6 shows the selected settings for *Land Use Change*, *Energy from Fossil Fuels, Carbon Prices, and Global GDP*. After specifying the ultimate parent nodes, we sampled the network to produce distributions for the intermediate and ultimate variables. Using the results of these simulations, we then estimated the distribution of the resulting HPI Quarterly Growth under each scenario, shown in Figure 11 below. The mean of the distribution of each scenario are shown as vertical dotted lines, and the distribution of historical HPI quarterly growth is included as a grey dashed line for comparison.

Table 6
SETTINGS FOR NGFS SCENARIOS AND AGNOSTIC VIEW

VARIABLE	SCENARIO		
	AGNOSTIC	NET ZERO 2050	CURRENT POLICIES
LAND USE CHANGE	Uniform (100 Mt CO_2/year[2], 700 Mt CO_2/year)	491 Mt CO_2/year	108 Mt CO_2/year
ENERGY FROM FOSSIL FUELS	Uniform (5.0%, 30.0%)	9.6%	29.2%
CARBON PRICES	Uniform ($5.0, $300.0)	$115	$6
GLOBAL GDP	Uniform ($159T, $165T)	$162T	$164T

Figure 11
DISTRIBUTION OF HPI QUARTERLY GROWTH UNDER DIFFERENT SCENARIOS

Comparing the Agnostic view, the Net Zero 2050 scenario, and the Current Policies scenario, we observe that the Net Zero 2050 scenario produces the lowest mean, followed by the Agnostic view and then the Current Policies scenario. This is expected given the narrative behind the scenarios: the Net Zero 2050 scenario leads to the least impact from extreme weather events, and therefore there is less stress on housing supply and demand. In Current Policies, the opposite is true, so we expect a steeper increase in home prices. The Agnostic view, being a mix of the two (and other) scenarios, falls in the middle.

[2] Metric ton CO_2/year

One interesting observation is that the widths of the distributions also follow the same order. It is intuitive for the Agnostic view to have a wider distribution than Net Zero 2050 because of the increased uncertainty in the four parent variables, but the same conclusion did not apply for Current Policies. This is because of a separate factor that also plays a role: in the Agnostic and Net Zero 2050 scenarios, there are fewer extreme weather events, and therefore less associated uncertainty from the impact of the events. In contrast, under the Current Policies scenario, there is a higher expected number of catastrophic floods, wildfires, and hurricanes, and the uncertainties in the effect of these events are amplified. For all three scenarios, the distributions are wider compared to the distribution of historical HPI Quarterly Growth. This is a display of how the network reflects the volatility that climate risk brings—there is uncertainty in both the events as well as the effect of those events. As a result, uncertainties are amplified downstream. The opposite is also true (fewer extreme events → less uncertainty propagated downstream), as evidenced by the narrower distribution in the Net Zero 2050 scenario.

Table 7 below shows the descriptive statistics of HPI Growth under different scenarios.

Table 7. Descriptive Statistics of HPI Quarterly Growth

	Agnostic	**Net Zero 2050**	**Current Policies**	**Historical**
Mean	1.37%	0.90%	2.41%	1.10%
Median	1.13%	0.81%	1.99%	1.04%
90th Percentile	3.62%	2.27%	6.82%	2.27%
Standard Deviation	2.08%	1.30%	3.44%	1.11%

INFERENCE

What was described as "prediction" above is the learning of resulting node distributions under different scenarios, querying the network in a "forward" manner, where evidence is specified in parent nodes and effects are propagated to child nodes. Alternatively, researchers can also query the network in a "backward" manner, where evidence can be inputted into child nodes, and the effect on parent nodes can be observed. We will refer to this as "inference" in this discussion.

For example, one may want to understand the distribution of Interest Rate in a situation in which *HPI Quarterly Growth Adjusted for Climate Change* takes a particular value. We can set the value *of HPI Quarterly Growth Adjusted for Climate Change* to a value and observe what that implies for the distribution of Interest Rate. Because of Bayes' theorem, even though the distribution of *HPI Quarterly Growth Adjusted for Climate Change* is set up to depend (indirectly) on *Interest Rate*, the opposite is also true, and we can infer the conditional distribution of *Interest Rate* given evidence on its downstream nodes.

Figure 12 shows a comparison of the implied distributions of Interest Rate in two scenarios where *HPI Quarterly Growth Adjusted for Climate Change* is +/-10%. As we can see, higher HPI growth tends to occur when interest rates are higher.

Figure 12
DISTRIBUTION OF INTEREST RATE FOR HPI QUARTERLY GROWTH ADJUSTED FOR CLIMATE CHANGE = +/-10%

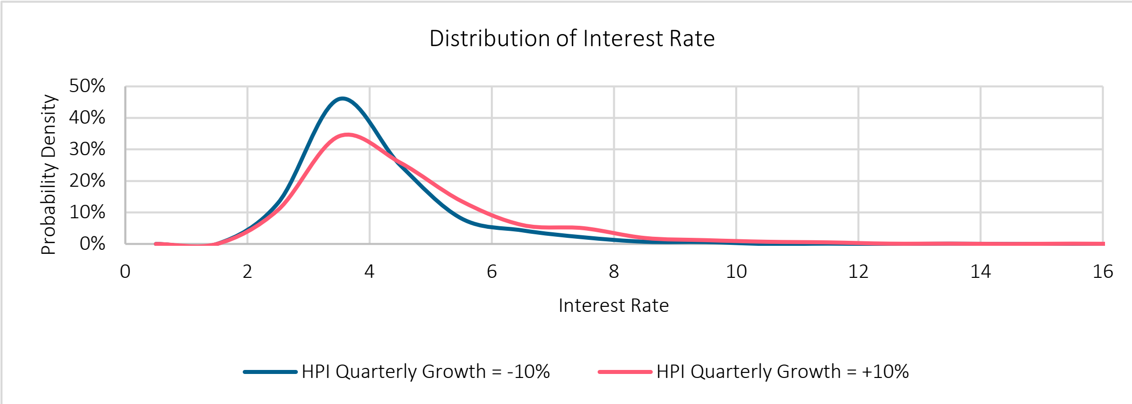

SENSITIVITY ANALYSIS

One question often asked in the context of climate scenario analysis is the sensitivity of results to changes in related observable variables. There are many types of sensitivity analysis, with varying level of complexity. One example of sensitivity analysis involves assessing the sensitivity of target node to perturbations in the source nodes. Effectively, sensitivity values are partial derivatives of output probabilities with respect to the source nodes being varied.

Sensitivity analysis can be used to drive decisions to focus on those parameters having the most impact on the output variable of interest. Moreover, it is an important tool to assess if the model results are reasonable. A Tornado diagram, also called a Tornado Chart, is a two-way bar chart that is a convenient way to visualize the sensitivity. Variables are listed vertically and ordered by descending importance from top to bottom, where importance is measured by the potential influence it can have on the output node.

Figure 13 shows an example of a Tornado diagram of selected variables. The diagram shows that the order of importance of hurricane, flood, and wildfires follow the typical magnitude of economic damage that they typically cause. On the other hand, the diagram suggests that the model is more sensitive to Availability of Land than it is to the catastrophes—this can be explained by Availability of Land having fewer degrees of separation from the final node, and therefore its relationship with the outcome variable is much more direct than the others.

Figure 13
TORNADO GRAPH: VISUALIZATION OF SENSITIVTY ANALYSIS OF SELECTED VARIABLES

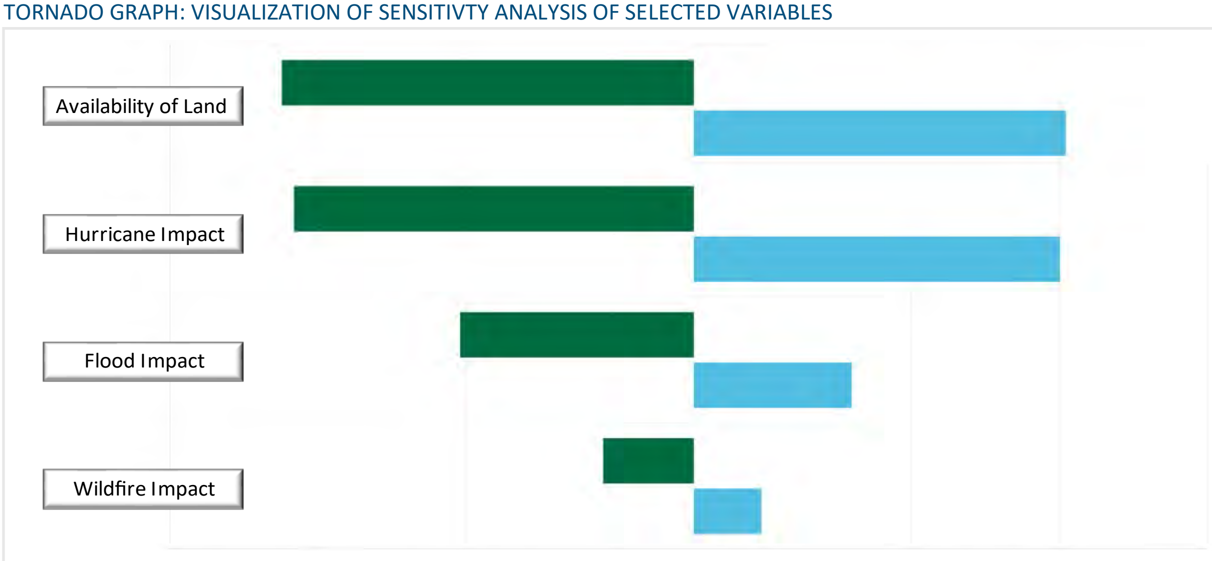

OTHER MODEL USE CASES

Similar to obtaining the distribution of other variables given evidence of one or more variables, researchers can also use the concept of Most Probable Explanation (MPE). The MPE is the set of most probable configuration of values for the nodes without an observation and is similar to the more general concept of maximum a-posteriori probability (MAP). The MPE is a concise way of summarizing plausible cases of other variables given some observations, without having to investigate the joint probability distribution of all those variables. The search for MPE is an NP-hard problem and intractable in general, but there are algorithms to find an explanation that is "good enough". An in-depth discussion of MPE is outside the scope of this paper, and we did not utilize the concept of MPE in this case study. Interested readers are referred to the paper *Most Probable Explanations In Bayesian Networks: Complexity And Tractability* by Johan Kwisthout (Kwisthout, 2011).

Related to the concept of MPE is the search for the Maximum Expected Utility (MEU) decision. Utility functions are seldom linear with respect to outcomes, so evaluating the utility function of the outcome variable at an expected value is not equal to the expected utility function, given a scenario. The MEU can be described as:

$$(d_1^*, \dots, d_k^*) = \ arg \max_{(d_1,\dots,d_k)} \int P(x, evidence) \cdot Utility(x)$$

Where $(d_1, \dots, d_k)$ are nodes related to decision variables, and $Utility(\vec{x})$ is the utility function of the outcome variable x.

Conclusion

In this report, we use a case study to demonstrate how Bayesian Networks can be used to conduct scenario analysis in the context of climate change. We describe a step-by-step process that entities can follow to build their climate risk assessment function. The proposed framework is flexible and can be adapted for both qualitative and narrative-based, as well as quantitative and data-driven analyses.

The case study provides a practical example of how Bayesian analysis can be used to frame problems and formalize concepts in a systematic manner. Specifically, we combine information from established climate scenarios, historic catastrophe loss information, as well as housing market datasets to create one integrated model. Starting with a simple and intuitive narrative, we use a network of interconnected nodes to relate different concepts. We then refine the network and nodes with data and expert knowledge to produce an improved version of the model.

The final model in this case study encompasses all the research and assumptions we consider within the scope of the problem. The model accepts inputs of selected parameters from NGFS climate scenarios and produces a distribution of future home price growth given the specified climate scenario parameters. Using this, we are able to make predictions and inferences of the distribution of future quarterly home price growth under different climate scenarios. This specific case study uses home prices as an example, but the framework can be readily extended for other analysis around climate change and financial risk analyses.

Acknowledgments

The researchers' deepest gratitude goes to those without whose efforts this project could not have come to fruition: the Project Oversight Group and others for their diligent work overseeing this research and reviewing and editing this report for accuracy and relevance.

PROJECT OVERSIGHT GROUP MEMBERS

Timothy Cheng

Sam Gutterman

Aadit Sheth

David Stoddard

Remi Villeneuve

Betty-Jo Walke

Nathan Worrell

AT THE SOCIETY OF ACTUARIES

Rob Montgomery, ASA, MAAA, FLMI, Consultant – Research Project Manager

Appendix A: Data Sources

Node	Data Source
Land Use Change	NGFS | REMIND-MAgPIE 3.0-4.4
Energy from Fossil Fuels	NGFS | REMIND-MAgPIE 3.0-4.4
Carbon Price	NGFS | REMIND-MAgPIE 3.0-4.4
Global GDP	NGFS | REMIND-MAgPIE 3.0-4.4
Greenhouse Gas Concentration	NGFS | REMIND-MAgPIE 3.0-4.4
Sea Level Rise	NOAA Global Sea Level Time Series (annual signals removed)
Temperature	NOAA NCEI Contiguous U.S. Average Temperature
Precipitation	NOAA NCEI Contiguous U.S. Precipitation
Catastrophic Hurricane Severity	NOAA NCEI Billion-Dollar Weather and Climate Disasters (CPI-Adjusted)
Catastrophic Hurricane Frequency	NOAA NCEI Billion-Dollar Weather and Climate Disasters (CPI-Adjusted)
Catastrophic Wildfire Severity	NOAA NCEI Billion-Dollar Weather and Climate Disasters (CPI-Adjusted)
Catastrophic Wildfire Frequency	NOAA NCEI Billion-Dollar Weather and Climate Disasters (CPI-Adjusted)
Catastrophic Flood Severity	NOAA NCEI Billion-Dollar Weather and Climate Disasters (CPI-Adjusted)
Catastrophic Flood Frequency	NOAA NCEI Billion-Dollar Weather and Climate Disasters (CPI-Adjusted)
Lumber Index Quarterly Growth	Baseline Scenario (January 2023): PPI: Lumber and wood products - Lumber, (Index 1982=100, NSA) **Mnemonic: FXPPILU1.IUSA**
Historical Unpriced Costs	Unpriced Flood Damages and Flood Insurance Premiums
Interest Rate	Baseline Scenario (January 2023): Interest Rates: 30-Year Fixed Rate Mortgage Commitment Rate - National, (%, NSA) **Mnemonic: FRFHLMCFM.IUSA**
Housing Affordability Quarterly Growth	Baseline Scenario (January 2023): Fixed rate Housing Affordability Index, (Index, SA) **Mnemonic: FHXAFF.IUSA**
HPI Quarterly Growth	Baseline Scenario (January 2023): FHFA All Transactions Home Price Index, (Index 1980Q1=100, SA) **Mnemonic: FHOFHOPIQ.IUSA**

References

Bishkek, K. 2018. Overview of the Global Change Assessment Model (GCAM). Joint Global Change Research Institute. Retrieved from https://unece.org/fileadmin/DAM/energy/se/pdfs/CSE/PATHWAYS/2019/ws_Consult_14_15.May.2019/supp_doc/PNNL-GCAM_model.PDF. Accessed January 4, 2023

Center For Science Education. Climate Change: Regional Impacts. University Corporation for Atmospheric Research. https://scied.ucar.edu/learning-zone/climate-change-impacts/regional

Department of Financial Services. 2021. *"Acting Superintendent Adrienne A. Harris Announces DFS Issues Final Guidance To New York Domestic Insurers on Managing The Financial Risks from Climate Change"*. New York State. https://www.dfs.ny.gov/reports_and_publications/press_releases/pr202111151. Accessed January 10, 2023.

Earth Institute (2019). Team Deciphers Sea-Level Rise From Last Time Earth's CO2 Was as High as Today (2019). https://news.climate.columbia.edu/2019/08/30/team-deciphers-sea-level-rise-earth-c02-was-high-as-today/

Els, Peter. 2021. *"This Is Why Leading Car Companies Will Be The First To Stop Producing Internal Combustion Engines"*. Hotcars.com. https://www.hotcars.com/this-is-why-leading-car-companies-will-be-the-first-to-stop-producing-internal-combustion-engines/. Accessed November 19, 2022.

Evans, David et al. (2020). Residential Flood Risk in the United States: Quantifying Flood Losses, Mortgage Risk and Sea Level Rise. The Society of Actuaries. https://www.soa.org/globalassets/assets/files/resources/research-report/2020/soa-flood-report.pdf

Funk, C., Hefferon, M. (2019). U.S. Public Views on Climate and Energy. Pew Research Center. https://www.pewresearch.org/science/2019/11/25/u-s-public-views-on-climate-and-energy/

Gibson, M., Mullins, J., Hilll, A. (2019). Climate risk and beliefs: Evidence from New York Floodplains. Department of Economics Working Papers 2019-02, Department of Economics, Williams College. https://web.williams.edu/Economics/wp/GibsonMullinsHill_ClimateRisk.pdf

Keys, Benjamin J. and Mulder, Philip. 2020. *"Neglected No More: Housing Markets, Mortgage Lending, and Sea Level Rise"*. National Bureau of Economic Research. https://www.nber.org/papers/w27930.

Kwisthout, J. 2011. "*Most Probable Explanations In Bayesian Networks: Complexity And Tractability"*. Radbound University Nijmegen Institute for Computing and Information Sciences. https://www.socsci.ru.nl/johank/MPE_complexity_pp.pdf

Loomis, John. 2004. *"Do nearby forest fires cause a reduction in residential property values?"* Journal of Forest Economics. https://www.sciencedirect.com/science/article/abs/pii/S1104689904000340.

Riahi, K. et al. (2017). The Shared Socioeconomic Pathways and their energy, land use, and greenhouse gas emissions implications: An overview. Global Environmental Change, Volume 42, 2017, Pages 153-168, ISSN 0959-3780. https://doi.org/10.1016/j.gloenvcha.2016.05.009

Task Force On Climate-related Financial Disclosures. 2017. *"Final Report – Recommendations of the Task Force on Climate-related Financial Disclosures"*. https://assets.bbhub.io/company/sites/60/2020/10/FINAL-2017-TCFD-Report-11052018.pdf.

Trapasso, Clare. 2020. *"Why California's Devastating Wildfires Will Push Home Prices Even Higher."* Realtor.com. https://www.realtor.com/news/trends/why-californias-devastating-wildfires-will-push-home-prices-even-higher/. Accessed on November 28, 2022

Yang, X.-S. 2019. Mathematical foundations. Introduction to Algorithms for Data Mining and Machine Learning, 19–43. https://doi.org/10.1016/b978-0-12-817216-2.00009-0

Zheng, Hanwen. 2021. NAT-Modeled Dynamic Discretization for Inference with Sum of Continuous Variables. University of Guelph. https://atrium.lib.uoguelph.ca/xmlui/bitstream/handle/10214/26524/Zheng_Hanwen_202110_MSc.pdf?sequence=1.

Limitations

In performing this analysis, we relied on data and other information provided by others. We have not audited or verified this data and information. If the underlying data or information is inaccurate or incomplete, the results of our analysis may likewise be inaccurate or incomplete. In that event, the results of our analysis may not be suitable for the intended purpose.

We performed a limited review of the data used directly in our analysis for reasonableness and consistency. We did not find material defects in the data. If there are material defects in the data, it is possible that they would be uncovered by a detailed, systematic review and comparison of the data to search for data values that are questionable or relationships that are materially inconsistent. Such a detailed review was beyond the scope of our assignment.

Any reader of these workpapers must possess a certain level of expertise in areas relevant to this analysis to appreciate the significance of the assumptions and the impact of these assumptions on the illustrated results. The reader should be advised by, among other experts, actuaries or other professionals competent in the area of actuarial projections of the type in these workpapers, so as to properly interpret the projection results.

Feedback

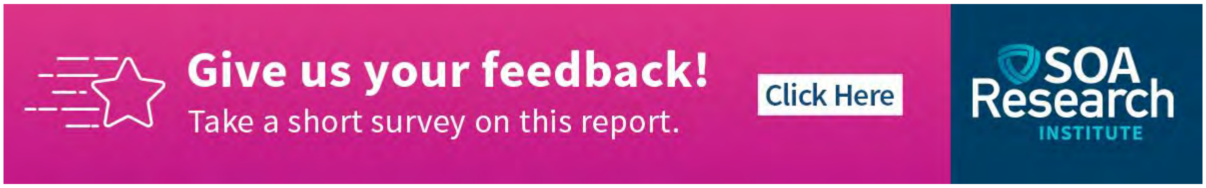
Give us your feedback!
Take a short survey on this report.
Click Here
SOA Research
INSTITUTE

About The Society of Actuaries Research Institute

Serving as the research arm of the Society of Actuaries (SOA), the SOA Research Institute provides objective, data-driven research bringing together tried and true practices and future-focused approaches to address societal challenges and your business needs. The Institute provides trusted knowledge, extensive experience and new technologies to help effectively identify, predict and manage risks.

Representing the thousands of actuaries who help conduct critical research, the SOA Research Institute provides clarity and solutions on risks and societal challenges. The Institute connects actuaries, academics, employers, the insurance industry, regulators, research partners, foundations and research institutions, sponsors and non-governmental organizations, building an effective network which provides support, knowledge and expertise regarding the management of risk to benefit the industry and the public.

Managed by experienced actuaries and research experts from a broad range of industries, the SOA Research Institute creates, funds, develops and distributes research to elevate actuaries as leaders in measuring and managing risk. These efforts include studies, essay collections, webcasts, research papers, survey reports, and original research on topics impacting society.

Harnessing its peer-reviewed research, leading-edge technologies, new data tools and innovative practices, the Institute seeks to understand the underlying causes of risk and the possible outcomes. The Institute develops objective research spanning a variety of topics with its strategic research programs: aging and retirement; actuarial innovation and technology; mortality and longevity; diversity, equity and inclusion; health care cost trends; and catastrophe and climate risk. The Institute has a large volume of topical research available, including an expanding collection of international and market-specific research, experience studies, models and timely research.

Society of Actuaries Research Institute
475 N. Martingale Road, Suite 600
Schaumburg, Illinois 60173
www.SOA.org