October 22, 2024

ARTIFICIAL INTELLIGENCE/MACHINE LEARNING
HEALTH INSURANCE SURVEY

FILING GUIDANCE AND DATA DEFINITIONS

This document provides survey filing guidance and data definitions. It is organized by survey sections.

This survey aims to aid regulators in understanding what artificial intelligence/machine learning (AI/ML) techniques health insurers are using and where in their insurance operations they are using them. This survey does not aim to have insurers provide details of trade secret components of the various techniques they are implementing. This Health Insurance survey is restricted to interests in three operational areas: 1) pricing and underwriting; 2) marketing; and 3) risk management. Within each operational area, interest is restricted to a limited number of use cases. It is expected that leaders and data scientists in these operational areas will have significant involvement in responding to the survey items.

Any company licensed to write health insurance in one of the 16 participating states (Colorado, Connecticut, Illinois, Iowa, Louisiana, Maryland, Minnesota, Nebraska, North Dakota, Oklahoma, Oregon, Pennsylvania, Vermont, Virginia, West Virginia, and Wisconsin) and meeting one or more of the criteria below is required to complete the survey.

Criteria for Selecting Participating Insurance Companies:

1. Companies with more than $250,000,000 of earned premium on a countrywide basis which write business in one of the 16 participating states, and/or
2. Companies that represent a significant portion of the market share in one or more of the lines of business for one of the 16 states participating in the survey.

Survey Platform - Confidentiality and Security

The NAIC, pursuant to a separate agreement, will collect and compile the information on behalf of and under the direction of the Requesting States. The records will be maintained as confidential, and the compilation will be conducted in such a manner as to ensure the anonymity of the Company.

A public report of aggregated information will be issued, but the survey is being conducted pursuant to market conduct authority, and as such will be held confidential by the states and the NAIC.

The NAIC has contracted with Qualtrics, LLC, a secure third-party survey platform to collect the AI/ML Health Survey information. Qualtrics, LLC uses Transport Layer Security (TLS) encryption for all transmitted data and maintains the following security certifications: SOC 2 Type II Certification; ISO 27001, 27017, 27018 Certification; and FedRAMP Authorization. All survey responses will be retained in data centers located in the United States, and NAIC staff's access to the survey responses will be strictly limited to those granted administrative authority.

Definition of Artificial Intelligence/Machine Learning for This Survey – Applicable to All Sections

For purposes of this survey, AI is defined as models that can simulate learning in performing tasks. ML is a subset of algorithms that facilitate learning without being explicitly programmed to achieve a predetermined result. Models that are considered AI and built using ML include robotics, natural language processing, and sentiment analysis.

ARTIFICIAL INTELLIGENCE/MACHINE LEARNING
HEALTH INSURANCE SURVEY

FILING GUIDANCE AND DATA DEFINITIONS

Standard features of AI systems adopted for purposes of this survey include:

- Systems that adapt and adjust to new data and experiences without manual human intervention.
- Systems that arrive at results for which the outcomes and the stepwise approach toward the outcomes were not configured in advance by a human programmer.
- Systems that dynamically respond to conditions in the external environment without the specific nature of such responses being known in advance to the designers of the systems.
- Systems that use deep-learning neural network algorithms.
- Systems that learn under a supervised, semi-supervised, unsupervised, or reinforcement learning style.
- Systems that engage in automatic speech recognition, facial recognition, image recognition, text recognition, natural language processing, generation of customer-specific recommendations, automated customer communications (e.g., chatbots with non-preprogrammed prompts), or any other approach that does not require either preprogramming or human intervention in every instance of an action or decision.
- Systems that automatically generate adaptive responses based on interactions with a consumer or third party.
- Systems that determine which data elements to rely upon, in a non-preprogrammed fashion, among various possible alternatives.
- Generalized linear modeling (GLM) or generalized additive modeling (GAM) are considered ML. See the Appendix for an expanded list of ML techniques.

For purposes of this survey, the following AI systems are excluded:

- Static "scorecards" that deterministically map consumer or other risk characteristics to treatments or decisions. (However, an AI/ML system may use the output of such static "scorecards" as input data for the AI/ML system to consider.)
- Systems with solely preprogrammed decision rules. (e.g., "If A, then B" applied invariably in all situations).
- Tables of point or factor assignments in risk classes.
- Deterministic "phone trees" that navigate consumers through prerecorded voice prompts.

Components of Survey

This Health Insurance survey is made up of the following sections:

- Survey Question Sequencing
- Respondent Information
- General Information on AI/ML Use

And, for each of the following product lines

- Individual Business - Comprehensive Major Medical
- Group Business - Comprehensive Major Medical – Single Employer – Small Employer
- Group Business – Comprehensive Major Medical – Single Employer - Other Employer
- Student (Individual Business – Student, and Group Business – Other Medical (Non-

Comprehensive) – Student combined),

the following sections:

a. Product Pricing and Plan Design
b. Claims Adjudication
c. Prior Authorization
d. Utilization/Severity/Quality Management
e. Risk Management
f. Risk Adjustment
g. Fraud Detection
h. Data Processing Issues
i. Marketing

And for all product lines in aggregate;

- Governance

The product line definitions come from the NAIC 2023 Annual Statement Instructions for completing the Accident and Health Policy Experience Exhibit:

Individual Business - Comprehensive Major Medical (Line A. 1.)
Group Business - Comprehensive Major Medical – Single Employer – Small Employer (Line B. 1.1.)
Group Business – Comprehensive Major Medical – Single Employer - Other Employer (Line B. 1.2.)
Individual Business – Student (Line A. 6.)
Group Business – Other Medical (Non-Comprehensive) – Student (Line B. 8.)

Comprehensive/Major Medical: Policies that provide fully insured indemnity, HMO, PPO, or Fee for Service coverage for hospital, medical, and surgical expenses. This category excludes Short-Term Medical Insurance, the Federal Employees Health Benefit Program and non-comprehensive coverage such as basic hospital only, medical only, hospital confinement indemnity, surgical, outpatient indemnity, specified disease, intensive care, and organ and tissue transplant coverage as well as any other coverage described in the other categories of this exhibit. Group business is further segmented under this category as follows (please note there is a separate category for Administrative Services Only/Administrative Services Contract business):

Single Employer: Group policies issued to one employer for the benefit of its employees. This would include affiliated companies that have common ownership.

Small Employer: Group policies issued to single employers that are subject to the definition of Small Employer business, when so defined, in the group's state of situs.

Other Employer: Group policies issued to single employers that are not defined as Small Employer business.

Individual Business: Health insurance where the policy is issued to an individual covering the individual and/or their dependents in the individual market. This includes conversions from group policies.

Student: Policies that cover students for both accident and health benefits while they are enrolled and attending school or college. These can be either individual policies or group policies sponsored by the school or college.

The remainder of this document will describe each of the above sections to help respondents understand the intent of the questions and expected responses.

It is important for the primary company contact for the survey to complete the Respondent Information and General Information sections before forwarding the survey link to those assigned to complete designated sections of the operational sections (Product Pricing and Plan Design, Claims Adjudication, Prior Authorization, Utilization/Severity/Quality Management, Risk Management, Risk Adjustment, Fraud Detection, Data Processing Issues, Marketing, and Governance) of the survey. The Respondent Information and General Information sections contain survey items that must be completed before the operational sections of the survey can be completed. A table of contents will appear after the questions in the General Information section are completed. Once respondents get to the table of contents, they may jump to any section of the survey that has been assigned to them. It is possible to navigate back to the Respondent Information and General Information sections from the operational sections of the survey. The "Previous Page" and "Next Page" buttons allow respondents to move freely between different sections of the survey as well. The table of contents is just a quicker way to get to a particular section of the survey. The table of contents discussion below goes into more detail regarding the navigational aspect of the table of contents.

Survey Block 2: Respondent Information

This section requests basic demographic information on the respondent, as depicted below.

Respondent Information
These are fill in the blank questions.

1. NAIC Company Code
2. Company Name
3. Contact Name
4. Contact Title
5. Contact Phone Number
6. Contact Email Address
7. Comments

Comments are optional but are encouraged if any of the respondent's responses need to be clarified.

Survey Block 3: General Information

This section attempts to understand to what extent the respondent uses AI/ML in nine key operational areas: 1) product pricing and plan design; 2) claims adjudication; 3) prior authorization; 4) utilization/severity/quality management; 5) risk management; 6) risk adjustment; 7) fraud detection; 8) data processing issues; and 9) sales and marketing. This survey primarily focuses on consumer impact models used in these operational areas. However, the respondent can include other operational areas in

the "Other" line in Question 3.

The respondent is referred to the definition of AI/ML above to respond to the questions in this section. There are six basic questions in this section of the survey. They are summarized as follows:

1. Does your company use (actively implementing AI/ML in insurance operations), plan to use (research has been completed and there are AI/ML models that exist that have not yet been implemented, but there are tentative or confirmed plan to implement), or is currently exploring the use (actively researching the benefits of AI/ML to improve, augment, or replace an insurance function) of AI/ML? Note: throughout the survey, instances of "your company" should be interpreted as "your company, affiliate of your company, or any contracted vendor performing services on your behalf." Provide the implementation status of AI/ML in each of the operational areas identified above. The status options are:
    a. N/A
    b. Already in Production
    c. < 1 Year
    d. 1–3 Years
    e. > 3 Years
2. List areas other than product pricing and plan design, claims adjudication, prior authorization, utilization/severity/quality management, risk management, risk adjustment, fraud detection, data processing issues, and marketing for which your company uses, plans to use, or is currently exploring the use of AI/ML.
3. Indicate whether your company has developed or plans to develop AI/ML systems internally, use vendors, or both internally with input from vendors.
4. Choose from the reasons given why your company does not use, plan to use, or explore AI/ML.

The respondent may further elaborate on any response in the "Additional Comments" column provided. The respondent is not required to add additional commentary.

The Table of Contents

Once respondents enter the operational section of the survey, they will see a navigational table of contents. The navigational table of contents allows respondents to jump to different parts of the survey. This feature is very useful for assigning sections of the survey to different subject matter experts to fill out. It is important for the primary contact of the company to complete the Respondent Information and General Information sections before sending the survey link to assignees to complete their sections. See Appendix B for the components of the table of contents.

Product Line-Specific Questions

For each of the four product lines listed in the Components of Survey Section above, a series of questions is asked concerning your company's use of AI/ML in each of the following operational areas:

1. Product Pricing and Plan Design
2. Claims Adjudication
3. Prior Authorization

4. Utilization/Severity/Quality Management
5. Risk Management
6. Risk Adjustment
7. Fraud Detection
8. Data Processing Issues
9. Sales & Marketing

For questions where a response indicates AI/ML is being used for a particular function, a follow-up question requests the respondent to identify the level of deployment of AI/ML for the use case. The levels are: research; proof of concept; prototype; or implemented in production. The respondent should select the highest level of deployment of AI. The deployment levels from the lowest level of use to the highest level of use are:

- Research: The investigation into and study of materials and sources to establish facts and reach new conclusions. This is the collection of information about a particular subject and is the lowest level of use.
- Proof of Concept (POC): A small exercise to test the design idea or assumption. The main purpose of developing a POC is to demonstrate the functionality and to verify a certain concept or theory can be achieved in development. It is testing the model for functional viability to be sure it runs and delivers a result.
- Prototype: Prototyping provides the opportunity to visualize how the product will function. It is a working interactive model of the end product that gives an idea of the design, navigation, and layout. Prototyping involves testing the model with actual data, in a limited, controlled environment. A prototype brings the POC idea to life.
- Implemented in Production: The model is being used in a live, production environment using real data. This is the highest level of use.

At the end of each product line section, two tables identical to those shown below will appear:

ARTIFICIAL INTELLIGENCE/MACHINE LEARNING
HEALTH INSURANCE SURVEY

FILING GUIDANCE AND DATA DEFINITIONS

Functions	Machine Learning Techniques												
	DL	ENS	NN	REG	RS	RGS	BAY	DT	DR	IB	CLU	COX	Other
Product Pricing and Plan Design	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	
Claims Adjudication	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	
Prior Authorization	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	
Utilization/Severity/Quality Management													
Fraud Detection	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	
Risk Management													
Risk Adjustment	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	
Data Processing Issues	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	☐	

Functions	List the Names of AI\ML Model in Use for Each Functional Areas
Product Pricing and Plan Design	
Claims Adjudication	
Prior Authorization	
Utilization/Severity/Quality Management	
Fraud Detection	
Risk Management	
Risk Adjustment	
Data Processing Issues	

You will be asked to indicate by checking the boxes in the first table for the machine-learning techniques that are used in each of the listed operational areas for the product line. If any machine-learning techniques are used, the second table asks for the model name(s) in use for each operational area.

Definitions of the machine learning techniques in the first table can be found in Appendix A: Machine Learning Taxonomy

Definitions by Operational Area

Product Pricing and Plan Design

- **Provider:** A doctor, hospital, health care professional, or health care facility.
- **Affiliated Provider:** A health care provider or facility that is paid by a health plan to give service to plan members.

Claims Adjudication

- **Step Therapy:** A process by which insurers (public or private) require patients to take one or more

alternative medications before they can access the medicine prescribed by their provider, and refusing to pay for a higher-cost therapy until it is shown that a lower-cost therapy is not effective (commonly known as "step therapy protocols" or "fail-first policies")

Risk Management

- **Wearable Devices**: Wearable devices refer to smart electronic devices with sensors that collect and deliver biometric information. The technology includes devices that are worn on the wrist and other forms such as jewelry, glasses, clothing, shoes, and implanted devices. The main category of wearables in the market are fitness trackers and smartwatches, which gather metrics associated with physical activity: step count, activity minutes, distance traveled, floors climbed, calories burned, heart rate, and sleep patterns.
- **Wellness Initiatives**: A program intended to improve and promote health and fitness that's usually offered through the workplace, although insurance plans can offer them directly to their enrollees. The program allows employers or plans to offer premium discounts, cash rewards, gym memberships, and other incentives to participate. Some examples of wellness programs include smoking cessation, diabetes management programs, weight loss programs, and preventative health screenings.
- **Discount Medical Programs**: Programs that provide a discount on medical services or prescription drugs.
- **Technology to Detect Smoking:** Wearable sensor technology used to detect and help people quit smoking.
- **Disease Detection:** AI and ML programs designed to diagnose disease using training data— such as the patient's history, lab results, scans, symptoms, and images of confirmed and susceptible cases— or real-time facial recognition images.

Risk Adjustment

- **Risk Adjustment:** The way that payments to health plans are changed to take into account a person's health status.

Data Processing Issues

- **Accuracy:** A measure of the extent to which data represents the true value of the attribute it is intended to measure.
- **Validity:** The correctness and appropriateness of data for a specific use.

Sales & Marketing

- **Targeted Online Advertising**: Determination of which individuals on the internet should receive or see advertisements from the insurer.
- **Identification of Recipients of Mail or Phone Advertising**: Determination of which individuals would be desirable recipients of an insurer's advertisements via the telephone or physical mail.
- **Provision of Offers to Existing Customers**: Determination of which customers should be notified of new insurance products, discounts, options to be written in a different book of business, or any other benefit or favorable treatment that the insurer seeks to extend.

- **Identification of Potential Customer Groups**: Determination regarding which consumer sub-populations could likely become additional customers of the insurer and/or benefit from the insurer's products and services.
- **Demand Modeling**: Identification of consumers' needs for and interest in specific types of insurance and insurance products that the insurer is offering or whose development or sale the insurer may be considering or exploring.
- **Direct Online Sales**: Selling insurance policies to consumers through a direct internet-based channel in a manner that does not rely solely on preprogrammed decision rules.

Governance

The purpose of the questions related to AI/ML governance is to obtain a better understanding of a company's awareness of specific risk areas tied to the National Association of Association's (NAIC's) Artificial Intelligence Principles. The questions in this section apply to the nine operational areas: 1) product pricing and plan design; 2) claims adjudication; 3) prior authorization; 4) utilization/severity/quality management; 5) risk management; 6) risk adjustment; 7) fraud detection; 8) data processing issues; and 9) marketing.

The NAIC recommends that insurance companies and all persons or entities facilitating the business of insurance that play an active role in the AI system life cycle—including third parties such as rating, data providers, and advisory organizations (hereafter referred to as "AI actors") —promote, consider, monitor, and uphold the NAIC AI Principles according to their respective roles.

The principles are intended to establish consistent high-level guiding principles for AI actors. These principles are guidance and do not carry the weight of law or impose any legal liability. This guidance can serve to inform and establish general expectations for AI actors and systems emphasizing the importance of: accountability; compliance; transparency; and safe, secure, fair, and robust outputs.

In addition, the survey seeks information to understand if guidelines or best practices are documented. Specifically, if the company is involved in using AI models, does it have a documented process in place to ensure that AI actors are:

> **Fair and Ethical**
> AI actors should respect the rule of law throughout the AI life cycle. This includes, but is not limited to, insurance laws and regulations, such as those relating to trade practices, unfair discrimination, access to insurance, underwriting, privacy, consumer protection and eligibility practices, ratemaking standards, advertising decisions, claims practices, and solvency.
>
> Consistent with the risk-based foundation of insurance, AI actors should proactively engage in responsible stewardship of trustworthy AI in pursuit of beneficial outcomes for consumers and to avoid proxy discrimination against protected classes. AI systems should not be designed to harm or deceive people. They should be implemented in a manner that avoids harmful or unintended consequences and corrects and remediates for such consequences when they occur.

Accountable
AI actors should be accountable for ensuring that AI systems comply with these principles consistent with the actors' roles, within the appropriate context and evolving technologies. Any AI system should comply with legal requirements governing its use of data and algorithms during its phase of the insurance life cycle. Data supporting the final outcome of an AI application should be retained and be able to be produced in accordance with applicable insurance laws and regulations in each jurisdiction. AI actors should be responsible for any AI system's creation, implementation, and impacts, even if the impacts are unintended. AI actors should implement mechanisms and safeguards consistent with the degree and nature of the risks posed by AI to ensure all applicable laws and regulations are followed, including ongoing (human or otherwise) monitoring and, when appropriate, human intervention.

Compliant
AI actors must have the knowledge and resources to comply with all applicable insurance laws and regulations. AI actors must recognize that insurance is primarily regulated by the individual states and territories of the U.S., as well as by the federal government, and that AI systems must comply with the insurance laws and regulations within each jurisdiction. Compliance is required whether the violation is intentional or unintentional. Compliance with legal requirements is an ongoing process. Thus, any deployed AI system must be consistent with applicable laws and safeguards against outcomes that are either unfairly discriminatory or otherwise violate legal standards, including privacy and data security laws and regulations.

Transparent
To improve the public's confidence in AI, AI actors should commit to transparency and responsible disclosures regarding AI systems to relevant stakeholders. AI actors must be able to protect the confidentiality of proprietary algorithms, provided adherence to individual state law and regulations in all states where AI is deployed can be demonstrated. These proactive disclosures include revealing the data being used, the purpose of the data in the AI system and the consequences for all stakeholders.

Consistent with applicable laws and regulations, stakeholders (which includes regulators and consumers) should have a way to inquire about, review, and seek recourse for AI-driven insurance decisions. This information should be easy-to-understand and describe the factors that lead to the prediction, recommendation, or decision. This information may be presented differently and should be appropriate for applicable stakeholders.

Secure, Safe, and Robust
AI systems should be robust, secure, and safe throughout the entire life cycle so that in conditions of normal or reasonably foreseeable use, or adverse conditions, they can function in compliance with applicable laws and regulations. To this end, AI actors should ensure a reasonable level of traceability in relation to datasets, processes, and decisions made during the AI system life cycle. AI actors should enable analysis of the AI system's outcomes, responses, and other insurance-related inquiries, as appropriate

in keeping with applicable industry best practices and legal requirements.

AI actors should, based on their roles, the situational context, and their ability to act, apply a systematic risk management approach to each phase of the AI system life cycle continuously to address risks related to AI systems, including privacy, digital security, and unfair discrimination as defined by applicable laws and regulations.

It is understood that governance models vary in terms of components and terms used to describe these risk areas. However, there is a common thread across most governance models, and this language was specifically used in this survey as it ties directly to the NAIC's AI Principles. Where there may be concerns about overlap, the intention is for this additional information to clarify the unique intent of each. The company should reply to each component as precisely as possible.

The governance questions are designed to gain an understanding of the components of governance frameworks insurers have in place to protect consumers from harms posed by AI/ML. The numbering of the questions reflected below may differ from the numbering reflected in Qualtrics.

1. **If you are using data, scores and/or AI/ML models aggregated or developed by a third-party, do those contracts include any conditions that would limit disclosure or otherwise limit transparency to regulators.** The expected response to this question is "Yes" or "No" as provided by the radio buttons.

2. **Follow-Up Question:** If model governance is in place, please check (Yes/No) in the following areas are documented in your governance program. Refer to the NAIC AI Principles for a description of each of the five NAIC AI principles reflected below and the Survey Guidance document more information regarding completing this question. PLEASE NOTE: It is expected the governance levels and maturity to be commensurate with the maturity level and risk associated with the program. This follow up question seeks to understand whether the respondent has governance frameworks in place to address the five fairness and ethics considerations defined in the NAIC AI Principles for each operational area.

3. **Have you adopted practices with respect to Accountability for Data Algorithms' Compliance with Laws for each operational area below?** This question seeks to understand whether the respondent has practices in place that ensure accountability for data algorithms' compliance with laws. The respondent is expected to answer "Yes" to indicate they have practices in place that ensure accountability for data algorithms' compliance with laws or "No" if they do not have any practices in place using the radio buttons. If the respondent has practices in place, the respondent is expected to briefly describe those practices in the space provided.

4. **Have you adopted practices with respect to Accountability for Data Algorithms' Intended Impacts for each operational area below?** This question seeks to understand whether the respondent has practices in place that ensure accountability for data algorithms' intended impacts. The respondent is expected to answer "Yes" to indicate they have practices in place that ensure accountability for data algorithms' intended impacts or "No" if they do not have any practices in place using the radio buttons. If the respondent has practices in place, the respondent is expected to briefly describe those practices in the space provided.

5. **Have you adopted practices with respect to Accountability for Data Algorithms' Unintended Impacts for each operational area below.** This question seeks to understand whether the respondent has practices in place that ensure accountability for data algorithms' unintended impacts. The respondent is expected to answer "Yes" to indicate they have practices in place that ensure accountability for data algorithms' unintended impacts or "No" if they do not have any practices in place using the radio buttons. If the respondent has practices in place, the respondent is expected to briefly describe those practices in the space provided.

6. **Have you adopted practices with respect to Accountability for Appropriate Resources and Knowledge Involved to Ensure Compliance with Laws Including those Related to Unfair Discrimination for each operational area below.** This question seeks to understand whether the respondent has practices in place that ensure appropriate resources and knowledge to ensure compliance with laws, including those related to unfair discrimination. The respondent is expected to answer "Yes" to indicate they have practices in place that provide appropriate resources and knowledge to ensure compliance with laws, including those related to unfair discrimination, or "No" if they do not have any practices in place using the radio buttons. If the respondent has practices in place, the respondent is expected to briefly describe those practices in the space provided.

7. **Have you adopted practices with respect to Ensuring Transparency With Appropriate Disclosures Including Notice to Consumers Specific to Data Being Used and Methods for Appeal and Recourse Related to Inaccurate Data for each operational area below.** This question seeks to understand whether the respondent has practices in place that ensure transparency with appropriate disclosures, including notice to consumers specific to data being used and methods for appeal and recourse related to inaccurate data. The respondent is expected to answer "Yes" to indicate they have practices in place that ensure transparency with appropriate disclosures, including notice to consumers specific to data being used and methods for appeal and recourse related to inaccurate data, or "No" if they do not have any practices in place using the radio buttons. If the respondent has practices in place, the respondent is expected to briefly describe those practices in the space provided.

8. **Have you adopted practices with respect to AI Systems are Secure, Safe and Robust including Decision Traceability and Security and Privacy Risk Protections for each operational area below.** This question seeks to understand whether the respondent has practices in place that ensure AI systems are secure, safe, and robust, including decision traceability and security and privacy risk protections. The respondent is expected to answer "Yes" to indicate they have practices in place that ensure AI systems are secure, safe, and robust, including decision traceability and security and privacy risk protections, or "No" if they do not have any practices in place using the radio buttons. If the respondent has practices in place, the respondent is expected to briefly describe those practices in the space provided.

9. **Do you follow some other existing standards or guidance in regard to governance framework?** This question seeks to understand whether the respondent follows existing guidance that was provided by a third party or a hybrid of an internal governance structure and a third party. The expected responses are "Yes" or "No" as indicated by the radio buttons.

10. **This is a follow-up question to number 9 above: If yes, were the existing standards or guidance developed internally, provided by a third-party, or a hybrid of internally developed and third-party components?** The respondent is asked to select the appropriate radio button to respond to this question.

11. **This is a follow-up question to number 9 above: If any standards or guidance are provided by a third party, please provide the following:**
    a. **Name of the Third-Party**
    b. **Components of the Governance Framework provided by the Third-Party.**

12. **What is the process, if any, that your company has for providing consumers with the data elements used in consumer impact models?** This question seeks to understand the processes the respondent has in place to provide consumers with the data elements used in AL/ML models that affect them in the operational areas that are the focus of this survey.

13. **What do you do to ensure compliance with regulatory frameworks, such as Model 880 and the Unfair Trade Practices Act, when using non-FCRA data?** This question seeks to understand how the respondent ensures compliance with regulatory frameworks, such as the *Unfair Trade Practices Act* (#880), when using non-Fair Credit Reporting Act (FCRA) data. The respondent is expected to supply the information in the spaces provided.

14. **Are the consumers made aware of what non-FCRA data is collected, when it is used, and how it is used?** This question seeks to understand whether the respondent makes consumers aware of the non-FCRA data collected on them and used in AI/ML models. There are follow-up questions that asks the company to explain the processes in each of the three operational areas.

15. **If you are using data not included in FCRA, do consumers have an opportunity to correct it?** This question seeks to understand whether the respondent allows consumers to dispute and correct data used by the company. Expected responses are “Yes” or “No.” Respondents can indicate their answers using radio buttons.

16. **How does human intelligence influence the decision making based on the AI/ML results? How much human intervention is involved?** This question seeks to understand the extent to which a human is involved in algorithmic decision-making. Does the AI/ML inform decision-making (i.e., the human is in the loop)? Does the AI/ML model replace decision-making (i.e., the human is out of the loop)? Or is it a combination of the two (i.e., the human is on the loop)? The follow-up question seeks to understand how much intervention is involved. The respondent is expected to provide a response in the space provided.

17. **There is a follow-up question to 18 above: How are results monitored?** This question seeks to understand how AI/ML model results are monitored. For example, one type of monitoring could involve the metrics collected to ensure the model is working as designed. The respondent could also monitor the deployment of the AI/ML model to ensure end users are deploying the model as intended. The respondent is expected to provide a response in the space provided.

18. **Do you have a process for applicants for health insurance to contest an adverse underwriting**

decision? This question seeks to understand whether the respondent has a process in place for consumers to contest an adverse underwriting decision. This question only applies to pricing and underwriting. The expected response to this question is "Yes" or "No."

19. **Do you keep a log of the number of contested underwriting decisions?** This question seeks to understand whether the respondent keeps a log of contested underwriting decisions. The expected response to this question is "Yes" or "No."

Appendix A: Machine Learning Taxonomy

The machine learning (ML) taxonomy provides a listing of ML algorithms that fall under each category that makes up an ML taxonomy. The list is representative but not exhaustive. This is one of many existing ML taxonomies. Some include more or fewer techniques than depicted below, more out of convenience than structure.

Machine Learning Taxonomy Techniques

Generally, the most common techniques are reflected in each category below. If the technique you are using is not represented by one of the categories below, select the other category and enter the ML type in the indicated column. You may enter more than one ML type separated by commas. Each category below is represented in the spreadsheet by the abbreviations in parentheses.

1. Deep Learning (DL)
    a. Deep Boltzmann Machine (DBM)
    b. Deep Belief Network (DBN)
    c. Convolutional Neural Network (CNN)
    d. Stacked Auto-Encoder
2. Ensemble (ENS)
    a. Random Forest
    b. Gradient Boosting Machine (GBM)
    c. Bootstrapped Aggregation (Bagging)
    d. AdaBoost
    e. Stacked Generalization (Blending)
    f. Gradient Boosted Regression Trees
3. Neural Networks (NN)
    a. Radial Basis Function Network (RBFN)
    b. Perceptron
    c. Back-Propagation
    d. Hopfield Network
4. Regularization (REG)
    a. Ridge Regression
    b. Least Absolute Shrinkage and Selection Operator (LASSO)
    c. Elastic Net
    d. Least Angle Regression (LARS)
5. Rule System (RS)
    a. Cubist
    b. One Rule (OneR)
    c. Zero Rule (ZeroR)
    d. Repeated Incremental Pruning to Produce Error Reduction (RIPPER)
6. Regression (RGS)
    a. Linear Regression
    b. Ordinary Least Squares Regression (OLSR)
    c. Stepwise Regression
    d. Multivariate Adaptive Regression Splines (MARS)
    e. Locally Estimated Scatterplot Smoothing (LOESS)

   f. Logistic Regression
7. Bayesian (BAY)
   a. Naïve Bayes
   b. Averaged One-Dependence Estimators (AODE)
   c. Bayesian Belief Network (BBN)
   d. Gaussian Naïve Bayes
   e. Multinomial Naïve Bayes
   f. Bayesian Network (BN)
8. Decision Tree (DT)
   a. Classification and Regression Tree (CART)
   b. Iterative Dichotomiser (ID3)
   c. C4.5
   d. C5.0
   e. Chi-Square Automatic Interaction Detection (CHAID)
   f. Decision Stump
   g. Conditional Decision Trees
   h. M5
9. Dimensionality Reduction (DR)
   a. Principal Component Analysis (PCA)
   b. Partial Least Squares Regression (PLSR)
   c. Sammon Mapping
   d. Multidimensional Scaling (MDS)
   e. Project Pursuit
   f. Principal Component Regression (PCR)
   g. Partial Least Squares Discriminant Analysis
   h. Mixture Discriminant Analysis (MDA)
   i. Quadratic Discriminant Analysis (QDA)
   j. Regularized Discriminant Analysis (RDA)
   k. Flexible Discriminate Analysis (FDA)
   l. Linear Discriminant Analysis (LDA)
10. Instance-Based (IB)
   a. k-Nearest Neighbor (KNN)
   b. Learning Vector Quantization (LVQ)
   c. Self-Organizing Map (SOM)
   d. Locally Weighted Learning (LWL)
11. Clustering (CLU)
   a. k-Means
   b. k-Medians
   c. Expectation Maximization
   d. Hierarchical Clustering
12. Cox Regression (COX)

