Process automation in actuarial modelling: A PRT pricing case study

Adel Cherchali
Céline Francony
Claire Booth, FIA
Jessica Crowson, FIA
Mirakh Modasia, FIA
Maxence Colin
Russell Osman, FIA
Sam Burgess, FIA

Automation and workflow technologies have long enhanced actuarial processes, delivering measurable benefits but often requiring significant time and investment to implement—factors that have limited widespread adoption. Today, artificial intelligence (AI) represents a transformative accelerator: Modern AI applications can be deployed rapidly, at lower marginal cost and with markedly greater analytical precision. Early adopters are already reporting substantial efficiency gains and deeper insights, underscoring the urgency for organisations to embrace these advances or risk falling behind.

By implementing AI-driven automation, greater operational speed, efficiency, insights and control can be achieved, resulting in lower costs, higher quality and lower operational risk. Automating routine tasks also enables actuaries to dedicate more time to value-added activities, rather than spending their efforts on repetitive processes. The adoption of AI in the actuarial field is expanding rapidly, and there are many potential AI use cases for automation and workflow management. This paper offers case studies of AI applications to the pension risk transfer (PRT) pricing process, and how these lead to improved efficiency and outcomes. Speed and accuracy are particularly important in a competitive market such as the UK PRT market, where insurers are competing to promptly offer their best price whilst also needing to ensure their pricing is accurate to reflect the true risk and cost of the scheme. There may be additional time pressures on insurers who are pricing multiple schemes concurrently. Speeding up the basic operational steps of the pricing process also has the benefit of leaving more time to focus on areas of judgement—for example, judging the commerciality of the calculated premium and assessing how to manage the risks of unusual scheme features.

Whilst we focus on a specific PRT case study in this paper, the concepts and tools are very much applicable to a wide range of actuarial processes and product applications.

PRT pricing process

PRICING PROCESS OVERVIEW

The UK PRT market is large, with premium volumes of approximately £50 billion p.a., and highly competitive[1]. Insurers rely on efficient and accurate pricing processes to effectively and competitively price deals, usually having to issue multiple quotations simultaneously.

In Figure 1, we outline an example first-line actuarial workflow for the PRT premium generation process, in order to highlight the manual parts of the workflow that could benefit from greater automation. The figure is intended to be an example of the steps typically performed by the pricing team to calculate a quotation for a PRT deal. In practice, considerable collaboration will be required—drawing on expertise from risk, finance, investment, and both first- and second-line functions. The extent of this cross-functional involvement varies with the size and complexity of each transaction. Deals with unique or unusual features are also likely to require a modified approach.

FIGURE 1: TYPICAL STEPS IN AN INSURER'S PRT DEAL PREMIUM CALCULATION PROCESS

1. Data checks

Analyse and perform checks on data and benefit specification.

Raise queries.

2. Model setup

Input data and benefits into pricing model.

3. Assumptions

Input assumptions into pricing model: demographics, interest rate, inflation, etc.

4. Model run

Run pricing model (best estimate basis) to obtain cash flows.

5. Reinsurance

If using reinsurance, calculate reinsurer fee and fixed and floating leg cash flows.

6. Asset mix

Run asset optimisation tool to determine matching adjustment and asset cash flows.

7. Capital model

Input liability and asset cash flows to a capital model to calculate the balance sheet components—best estimate liability, SCR and risk margin.

8. Premium calculation

Calculate premium (allowing for expected expenses) that meets pricing targets (e.g., target internal rate of return or new business value).

9. Governance

Deal price presented to pricing committee for sign-off, with attaching risk opinion (if required).

1. The UK PRT market is explored in more detail in our paper 'Considerations for new entrants to the PRT market,' which can be found at https://www.milliman.com/en/insight/considerations-new-entrants-to-the-prt-market

AUTOMATING THE PREMIUM GENERATION

Insurers often have established models for running cash flows or calculating other balance sheet components such as the Solvency Capital Requirement (“SCR”) and risk margin. Use of proprietary software and/or open-source languages/tools means that insurers are typically further ahead in automating the later stages (stages 4–8) of the PRT deal pricing process.

Open-source languages, such as Julia, Python and R, offer myriad tools (libraries) for data science and automation. These tools can be used to automate data extraction and transformation, actuarial calculations, reporting, and documentation. The entire process can be chained together and managed as a workflow with these tools, and are most successful for processes which are repeated with minimal modification—that is, the processes which have standardised components.

The requirement for standardised components cannot always be satisfied for some aspects of the PRT pricing process. The first three stages are very typically manual and can therefore be time-consuming. In practice, it has been difficult to automate tasks such as data cleansing, the extraction of information from documents and model setup.

Insurers receive various key files from the employee benefits consultant (EBC) managing the quotation process, including:

- A benefit specification—this may be in the format of either a PDF, MS Word or MS Excel, and describes the features (cohorts, benefit types, escalation types, etc.) of the defined benefit pension scheme.
- Pension scheme data—this is provided in various file formats (typically MS Excel) and involves seriatim member data fields, for example guaranteed minimum pension (GMP) and excess pension details, various key dates, and spouse details.

There is no consistent file format or layout for data received under each pricing exercise, making it difficult for pricing teams to create a uniform or automated process. Pricing teams spend a significant amount of time reading documents, manually extracting data to ascertain policy features and key assumptions from benefit specification documents. Then data must be reorganised into the required model format. Additionally, data commonly has deficiencies—for example, missing fields or incorrect dates—which take time to check, query, correct, or consider appropriate assumptions or replacements to use.

USE OF AI IN THE PRICING PROCESS

AI can therefore be used to automate the processes associated with these first three stages of the PRT pricing process. These processes include data cleansing (checking for errors, omissions, etc.), data manipulation (ascertaining what each field in the data provided means and translating data into the format needed for the insurer’s model) and unstructured data extraction (retrieving relevant details such as assumptions and scheme features from the benefit specification).

There are clear benefits to the use of AI in this context. Most notably, the benefits include increased process speed, reduced cost and reduced likelihood of operational risks crystallising (via missed data errors, for example). Reduced reliance on human knowledge of manual processes should reduce key person risk. In the context of a very competitive UK PRT market, efficiency in preparing quotes is key. If insurers utilising this technology can produce a quote in a shortened time frame, and hence with lower costs attached, then they will have a competitive advantage over insurers not utilising AI.

As mentioned previously, automation of these processes by the actuarial pricing teams allows team members to focus on value-added activities. This might include coming up with new solutions or features for PRT deals, which will overall improve the service provided to schemes.

Types of AI

The term AI refers to computer systems developed to perform tasks which require human intelligence, such as learning and perception. AI technologies therefore have the potential to transform a number of complex tasks currently conducted by humans. AI can be broadly categorised into:

- Narrow (‘weak’) AI
- General (‘strong’) AI
- Artificial superintelligence

Currently, AI is at the stage of being narrow AI, as it is designed to perform specific tasks or solve particular problems. Examples include image recognition, playing chess or translating languages. These types of AI can only do the tasks they were designed for and cannot adapt any further than their tasks. General AI and artificial superintelligence have not yet been developed but could be realised in the near future. General AI aims to match human adaptability and problem-solving abilities across diverse tasks, whilst artificial superintelligence, which is still theoretical, would surpass human intelligence and could have profound societal implications.

Narrow AI can be implemented using rule-based systems or machine learning. Rule-based AI refers to AI which operates through predetermined logic. Machine-learning methodologies include supervised learning (using labelled data), unsupervised learning (using unlabelled data) and reinforcement learning (optimising actions through trial and error).

Narrow AI can also be categorised into reactive machines and limited memory systems. Reactive machines (e.g., IBM’s Deep Blue chess computer) are AI systems that respond only to current inputs and do not store or use past experiences. Limited memory systems (e.g., self-driving cars) are AI systems that can use a small amount of past data to inform their decisions for a short period.

The narrow AI categories mentioned previously are high-level classifications, and within each of them there are a number of further subcategories.

A noteworthy field of narrow AI is natural language processing (NLP). This focuses on understanding and processing human language. Large language models (LLMs) represent a specialised subset of deep learning within NLP. Deep learning uses neural networks with many layers to process and understand text. Trained on vast datasets, these models utilise neural networks with many layers (transformers) to process, understand and generate human-like text with high fidelity. They are initially self-supervised (a subset of unsupervised learning), and are fine-tuned using supervised and reinforcement learning. For example, one of the most popular AI tools, ChatGPT, was developed using this technique.

Chain of thought (COT) refers to the intermediate reasoning steps—often expressed as a sequential flow of ideas or deductions—that connect an initial prompt or observation to a final conclusion or answer. A user can prompt an LLM for COT reasoning to help in complex tasks (e.g., commonsense reasoning or multistep decision making) to obtain improved performance and better transparency, and therefore better interpretability for the user.

Building on the language understanding and generation capabilities of LLMs, AI agents are autonomous systems that are able to perceive their environment, reason, and take actions to achieve desired outcomes. AI agents with self-verification layers can orchestrate and combine several tools or computational resources to perform more complex or large-scale operations. Unlike simple software, an AI agent can adapt, learn, and improve its performance over time.

In the following sections we consider some case studies for using NLP tools, primarily LLMs, for overcoming the current constraints in the early stages of the PRT pricing process. Typically, a combination of tools and/or training techniques would be used, with more human input added to refine and finalise the outcomes. Tools can be split into different functions where required—for example, checking data, organising into data types, assigning to columns—and different approaches may be applied to each of these purposes.

Case study: AI automation tools for PRT pricing

BASIS OF THE TOOL

For the first two case studies, we developed a Python-based automation tool for the input files (either the MS Excel data file containing the dummy pension scheme data in the case of case study #1 tool or the benefit specification document in the case of case study #2 tool). This solution leverages an LLM via an application programming interface (API), specifically the Azure OpenAI service. An API is a tool that allows different software applications to communicate with each other. In our case, it enables the code to send data/documentation to and receive responses from the Azure OpenAI service, which hosts the LLM.

For case study #3, we explore the further developments that would need to be made to the LLM functionality to enable the execution of more data-intensive tasks. In our PRT pricing use case, we discuss how we would operationalise the transformation of pension scheme data from one format to another.

CASE STUDY #1: DATA CHECKS

The data checks part of this case study aims to improve the accuracy of the data extract that is fed into the pricing model by identifying errors within it, and to reduce the volume of manual checking required at the data preparation stage. For example, an actuarial analyst manipulating and checking the pension scheme data in MS Excel will often spend several hours creating formulas to check for data inconsistencies, errors or missing data. Typically, they would create formulas for each type of error with a check per column in MS Excel. Whilst standard checking tools may exist, the fact that pension scheme data has its own unique formats means that manual adaption of these tools is necessary for each pension scheme considered. The checks themselves are subject to human error, such as manual input errors (keystrokes), incomplete checks or oversight. They also require a second review by another more senior actuary.

Ultimately, reducing the number of errors in the pricing data will result in more accurate modelling of the pension scheme cash flows and therefore a more accurate price. This helps to reduce the risk that the pension scheme is mispriced. Further to this, sophisticated tools could be used across different schemes, making the data-checking process far more efficient.

We created a dummy pension scheme dataset containing 500 policies to test whether our internally developed AI tool, which was based on using an LLM, could spot typical errors in pension scheme data. Figure 2 summarises the deliberate errors we included within the dummy data.

FIGURE 2: SUMMARY OF DATA ERRORS.

ERROR TYPE	ERROR DETAILS	
Missing Data	Missing date of birth Missing contribution amount Missing gender	Missing pension amount Missing postcode Missing spouse date of birth
Incorrect or Invalid Data	Invalid gender entry Invalid date of birth Invalid postcode	
Logical Inconsistencies	Pension commencement date prior to date of birth Pensioner with pension start date in the future Deferred pensioner with a pension start date in the past	
Unusual or Outlier Values	Negative pension amount Very large contribution	
Missing Required Information Based on Other Data	Pension starts before 1988, but no pre-/post-1988 GMP amount for a member	

The process to perform the data checks works as outlined in Figure 3.

FIGURE 3: AUTOMATION PROCESS FOR DATA CHECKS

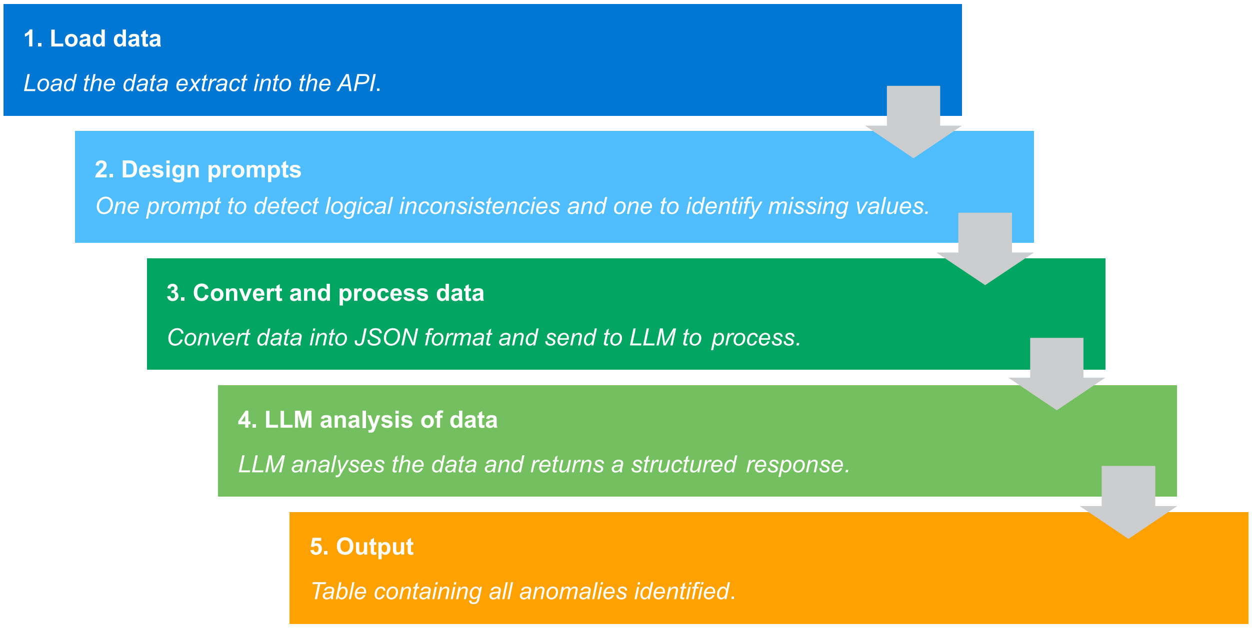

The prompts were split into several dimensions: missing data, format errors, timeline logic, biographical logic, financial checks and a holistic commonsense review. This helped the LLM systematically check all key aspects of the data quality, instead of only looking for obvious mistakes. For example, the prompts specified that certain key information—for example, member contributions—must always be present. The structured output file requested forces the model to be precise and clear, and to justify why something is an error, so the output is understandable and actionable. The prompts used the relevant pension administration terminology, gave examples of valid and invalid data and defined edge cases (unusual or extreme data values). For example, the prompts specified the correct format for a UK postcode and defined the required order for dates in the data extract to follow—e.g., a date of birth must be prior to service start date, which should be prior to the pension start date. The commonsense prompt asked the system to use its general intelligence to find anything else that looked wrong—for example, inconsistent genders, duplicate records or data fields that appear to have shifted to the wrong column.

The AI tool identified all the deliberate errors placed in the dummy data file. Zero false positives were obtained.

An example of the summary output obtained follows. This output in effect provides an audit trail of the LLM's work, conclusions and the reasoning for its conclusions, which supports human oversight and verification. The output includes the member's identifier number, a description of the error identified, the error type, the certainty level and an explanation as to why this is the error. The certainty level is a high/medium/low score assigned to each error and reflects the severity of the error. For example, a high confidence level reflects an error that is a mathematical impossibility or direct contradiction of a hard definition, such as a pensioner being born after they retired.

FIGURE 4: SUMMARY OF DATA CHECKS TOOL OUTPUT

DUMMY POLICY ID	ERROR DESCRIPTION	ERROR TYPE	CONFIDENCE	REASON
100027	Member contributions field is missing (null)	Missing values	High	Member contributions is a key column and cannot be missing. Record shows member contributions = null.
100380	Invalid postcode format 'ABC DEFG'	Common sense	Medium	The postcode does not conform to recognisable UK postcode structures, suggesting a data quality issue.
100439	Pensionable service start date before date of birth	Logical inconsistency—timeline	High	Date pensionable service commenced (circa 1945) is earlier than date of birth (circa 1950), which is chronologically impossible.

Human oversight is required to ensure that AI tools work as expected and are not producing erroneous results. This is particularly important given that AI is known to have issues with lack of explainability and transparency, so human oversight has a key role in ensuring that an AI tool produces results that are traceable and understandable. In this case study, there is clearly the need for human input to set up the tool and check the tool's results. In particular, there would be some initial work to set up and refine the prompts, specifying which values are valid, invalid or too extreme to be plausible. However, to a large extent, this logic is likely to already exist in data checks in existing models. Additionally, there would need to be some high-level sense checks that the AI output is correct, for example:

- Further investigation into values that are missing or invalid
- Spot checks that a reasonable number of errors have been identified

Considering further development opportunities, an extension of the tool would be for each issue or group of similar issues to draft queries that can be sent to the EBC. There could also be an extension of the tool to cross-check with any existing query logs or the benefit specification to see if errors identified have already been covered or if there is additional information that means the potential error is actually correct data.

CASE STUDY #2: DATA EXTRACTION FROM PENSION SCHEME BENEFIT SPECIFICATION

We then tested the AI tool to see if we could output the key features that an actuarial modeller would need to know to model the scheme in our pricing model. LLMs can be used to understand unstructured text and extract information into more structured formats. Specifically, we wanted to output:

- Key features of the pension scheme required to allow us to set up the pricing model—for example, we extracted details such as the timing of payments, death benefits (pre- and post-retirement), spouse benefits and definitions and guarantee period.
- Assumption tables in the correct format for inputting into our pricing model—for this case study, we extracted an assumption table covering details of the escalation and revaluation rates, as well as associated caps and floors for the benefits within the scheme, including reference months, increase dates, etc.

We tested the AI on three benefit specifications to ensure that the prompts functioned adequately when faced with different formats of documents and terminology.

To output the key scheme features, we created a standard template covering the key modelling details required. This created a standardised template which allows schemes to easily be compared, and the key features to be quickly and easily input into the model. We found that the tool accurately output all the required details, working well for each of the three documents analysed. With all three benefit specification examples, there were minor elements that could have been elaborated on; some manual checking was therefore required to corroborate that the output was correct and that important information was not being missed. That being said, the tool removed the need for an initial in-depth review of the documents and manual extraction of common inputs, and instead only required the second layer of review by a more experienced person.

To output the assumption tables for the model, we also found that this could be achieved quickly and accurately. We asked the tool to extract the following details for each relevant benefit tranche:

- Index type—CPI/RPI/fixed rate
- The fixed rate—e.g., 3%
- The floor and cap for RPI and CPI index types, where relevant
- The application—either escalation or revaluation
- The pension increase month
- The reference month for increases
- The time lag in months between the escalation reference month and the application of the escalation rate

The information was output into MS Excel in the assumption table format required by Milliman's Integrate model[2], meaning that the MS Excel file could be immediately uploaded into the Integrate cash-flow model used for pricing. This automates the production of the assumption tables used in the pricing model, taking the information directly from the benefit specification. The same approach could obviously be applied for other cash-flow modelling systems.

The tool produced the output quickly, taking under a minute (per document) to run the prompts and produce the output tables. Typically, it would take an actuary several hours to go through such a document manually in enough detail to have comprehensively compiled a list of all relevant scheme features.

The tool also outputs an AI-generated description of the escalation type, as well as a description of the benefits and a reference to the document that the information has been extracted from, for each assumption. All this information can be utilised to support human checking of the output.

Whilst we have output the assumption table for escalation and revaluation rates, other assumption tables could also be extracted to provide the full suite of assumption tables required for pricing. This, in conjunction with the output on the key features of the scheme as described previously, should give the actuarial user the full suite of information required to set up the model.

An area of future development for the tool will be to consider the extent to which it can pick out non-standard features within the benefit specification that may require bespoke modelling or an out-of-model adjustment. For example, the scheme rules may contain complex underpinnings within the indexation, which the tool would output. Further testing on schemes that have these more complex features would determine the extent to which this is possible, and therefore reduce the time required for the second layer of skilled human review. Indeed, it may be possible for the tool to suggest how a non-standard feature could be modelled. Additionally, if the tool identifies that there is missing information required in order to model the cash flows accurately, or ambiguity in the information provided which has a direct modelling impact, then the tool could be extended to draft a query to send to the EBC.

The tool could also be extended to analyse both the scheme data and the benefit specification together, to look for anomalies, information gaps and, again, draft queries for consideration by the actuarial modeller. For example, there may be a field within the scheme data that the tool deems relevant to the cash flows, but there is no information in the benefit specification on how it impacts the benefits paid.

2. Integrate is Milliman's life insurance actuarial software solution: https://integrate.milliman.com/en

CASE STUDY #3: DATA FORMATTING

Another application of AI within the PRT pricing process is data reformatting. As the data provided for different pensions schemes will inevitably come in various formats, pricing teams are required to reformat this into a format compatible with their own pricing model. If attempting to quickly onboard and price a number of schemes, this can become highly time-consuming.

For example, the pricing model may require the data field ‘Age at Valuation Date,’ but the pension scheme data only contains the data field ‘Date of Birth.’ We could then use an LLM to translate dates of birth to ages at valuation date by setting up a specific LLM prompt. For example, the prompt could specify something along the lines of *Valuation Date = 31/12/2024, DOB = Date of Birth. Please can you give the age of each pensioner in the column “Age”.*

However, when using LLMs for data-related tasks, particularly where there are large volumes of data, it is important to be aware of their limited computational performance. LLMs are excellent at understanding context, interpreting instructions and manipulating textual data, but they are not optimised for heavy numerical calculations or complex data processing at scale. Instead, calculations typically work best for small datasets or individual cases.

To address these limitations, there are a number of alternative tools that can be utilised. In the context of large-scale data transformation, tools such as AI agents and COT can be utilised. In this section, we explore how these tools could be leveraged to automate data formatting. The development of these tools will be considered an area of further research.

As described earlier in this paper, an AI agent is a computer program capable of perceiving its environment, analysing information, making decisions and acting autonomously to achieve a specific goal. It can take various forms: virtual assistant, chatbot, PowerPoint slide creator or Excel file editor/enhancer. What differentiates an agent from a simple LLM is the ability to act on the environment and perform actions. An LLM simply produces textual responses and is not ‘connected’ to the system. An AI agent is an LLM equipped with memory, a goal, and the ability to perceive and act on its environment. In the context of our example of translating dates of birth to ages at valuation date, the AI agent will create a Python code to execute these calculations. Moreover, the AI agent could be used to create code for all columns that require mapping and translation from one format to another.

COT is a reasoning process that breaks down complex tasks into smaller logical steps, ensuring that each instruction is systematically followed. It is beneficial to use with LLM, as it helps with making decisions step-by-step. Figure 5 illustrates the difference between prompt chaining and COT.

FIGURE 5: DIFFERENCE BETWEEN COT AND PROMPT CHAINING

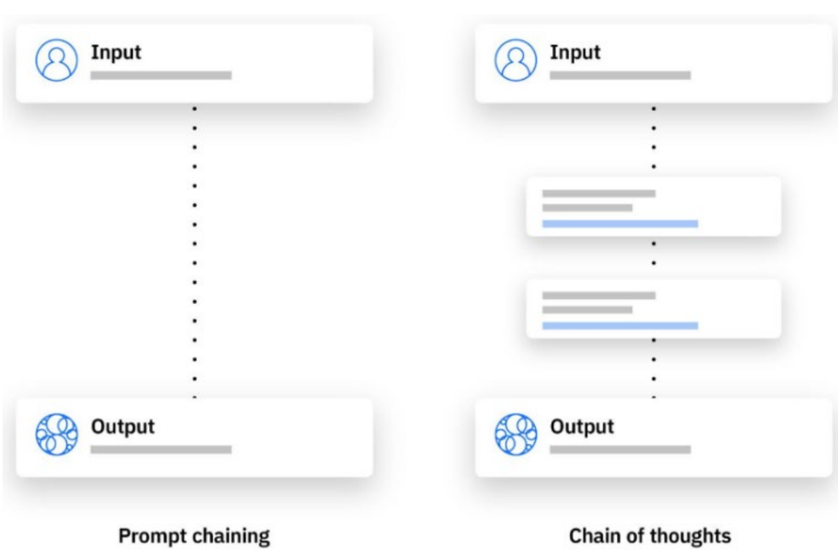

In particular, with respect to our use case, it can be used to generate intermediate format data files whereby certain columns of data are mapped or translated into the required format. It could also be used to adapt the script after detection of errors within the data, as in case study #2.

To address the challenge of data conversion within the PRT pricing process, we propose the implementation of an AI-based tool, as outlined below.

Insurers typically possess pre-existing data processing scripts (e.g., written in R or SQL) designed to generate required data files from pension scheme data. In instances where such scripts do not exist, an LLM can be employed to generate these scripts based on the desired output format. The proposed process combines the capabilities of LLMs with COT reasoning or an AI agent to facilitate data conversion. The workflow consists of the steps outlined in Figure 6.

FIGURE 6: WORKFLOW DATA STEPS FOR DATA CONVERSION

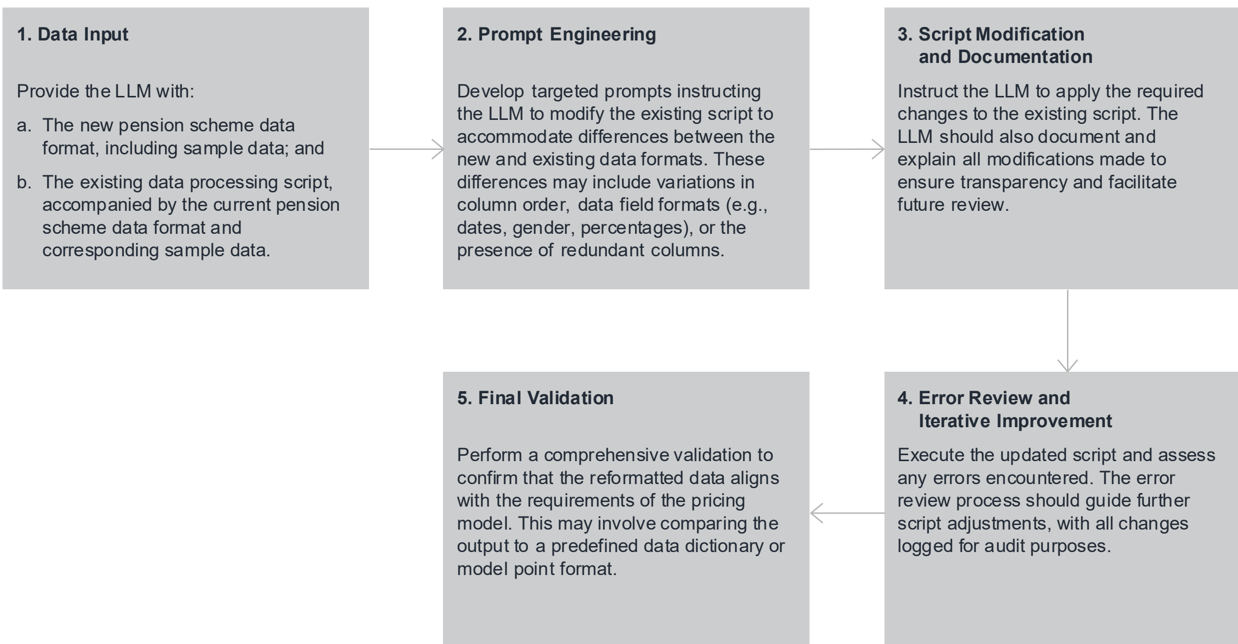

Upon completion of these steps, the user may utilise the case study #1 tool to verify the absence of data errors in the final output file. It is anticipated that steps 4 and 5 will leverage the combined strengths of LLMs and AI agents or COT reasoning to address complex, large-scale, and logical tasks inherent in the data conversion process.

Risks and limitations

As with all progressive changes, automation and use of AI presents several risks. When developing a tool, there are a number of considerations to assess to ensure that benefits of a tool are realised in a secure, compliant and methodical way. Key areas that may need to be considered are:

1. Data security—pricing PRT deals involves using scheme data that contains personally identifiable information (PII) (as defined under the General Data Protection Regulation (GDPR)), including date of birth, gender, postcodes and other identifying details. It is imperative that the tool is built such that this information is handled securely and complies with the relevant regulatory standards. There is also a need to consider how the approach and sophistication of malicious agents could evolve as AI applications expand. To address these issues, the tools may be “air-gapped,” meaning that the network the tool runs on is isolated from other connections.
2. Effectiveness—the effectiveness of the AI automation tool is closely tied to the quality and relevance of its training data. A tool is most reliable when the processing data is similar to the data that it was trained on. Given the complexity of PRT schemes and the large volumes of data with intricate or unique policy features, ensuring you have comprehensive and representative training data may be challenging. That being said, established providers should have access to a large volume of pension scheme data from previous deals which they can use to train the AI tool, if data storage and processing agreements allow.
3. Initial investment—key costs include training and developing staff to ensure that they know how to use AI tools and understand the dynamics. In terms of our case studies, actuaries may need to develop prompt-writing skills, and/or the ability to interpret new coding languages or output. However, it’s worth noting that the types of tools proposed in this paper are likely to be relatively low cost.
4. Potential for bias—there is also a risk that an AI model will produce biased or discriminatory outputs, particularly if it is trained on a narrow set of data. Whilst this is perhaps not as relevant in our PRT pricing use case, for wider uses such as underwriting processes, this could cause significant issues in terms of customer harm.
5. Compliance—as AI is an emerging technology, the regulatory landscape is still evolving and future requirements remain uncertain. Specific UK regulations on AI to date have been limited with only existing regulations in related areas applying (for example, GDPR). AI tools may be developed now and subsequently become non-compliant with regulations introduced in the future. There is therefore a need for firms to thoroughly document AI tools, and ensure that key risks such as security of personal data, model explainability and the risk of bias are adequately addressed. It is worth noting that in the UK the Technical Actuarial Standard (TAS) 100 (published by the Financial Reporting Council) apply to technical actuarial work that uses AI and machine-learning techniques[3] if they meet the TAS definition of a model. TAS 100 includes requirements that actuaries should ensure that they understand the models they use, that appropriate model governance is in place, that any material model bias is assessed, and that the implications of model limitations are assessed.

3. TAS 100 Model Guidance, section 2.14, https://media.frc.org.uk/documents/Technical_Actuarial_Guidance_Models_October_2024.pdf

Application to other actuarial processes

The case study provided in this paper covers PRT pricing but, in reality, AI automation tools can be used for a range of other actuarial modelling processes. This is especially the case where processes require digestion of and extraction from unstructured data, as is the case with the initial stages of the PRT pricing process. Other applications could include:

- Experience analysis
- Reserving
- Stress and scenario testing
- Financial reporting—automating of regulatory reports
- Document processing—extracting information from policy documents and contracts
- Analysing customer behaviour

Conclusions

This paper has demonstrated some examples where AI could automate and enhance key processes within actuarial modelling, focusing specifically on the PRT pricing process. By leveraging AI—particularly LLMs and related automation tools—insurers can achieve substantial improvements in operational speed, accuracy and cost-effectiveness. Our case studies showed that AI-driven tools can successfully automate data checks, extract key information from unstructured benefit specification documents and reformat data for model input. These advancements not only reduce manual effort and operational risk, but also enable actuaries to dedicate more time to higher-value analytical and strategic activities. Importantly, the approaches outlined are broadly applicable and can be extended to other actuarial processes that involve handling complex or unstructured data.

While the results are promising, there remain important areas for further development and consideration. Future work should focus on expanding the capabilities of AI tools to handle more complex and non-standard scheme features, and further integrating AI agents for large-scale data transformation. Additionally, ongoing human oversight and prompt refinement are essential to maintain accuracy and reliability. As the actuarial profession continues to embrace AI, ongoing research and practical implementation will be key to realising the full benefits of automation—ultimately contributing to a more resilient and efficient insurance market

Solutions for a world at risk™

Milliman leverages deep expertise, actuarial rigor, and advanced technology to develop solutions for a world at risk. We help clients in the public and private sectors navigate urgent, complex challenges—from extreme weather and market volatility to financial insecurity and rising health costs—so they can meet their business, financial, and social objectives. Our solutions encompass insurance, financial services, healthcare, life sciences, and employee benefits. Founded in 1947, Milliman is an independent firm with offices in major cities around the globe.

milliman.com

CONTACT

Claire Booth
claire.booth@milliman.com

Sam Burgess
samuel.burgess@milliman.com

Adel Cherchali
adel.cherchali@milliman.com

Maxence Colin
maxence.colin@milliman.com

Jessica Crowson
jessica.crowson@milliman.com

Céline Francony
celine.francony@milliman.com

Mirakh Modasia
mirakh.modasia@milliman.com

Russell Osman
russell.osman@milliman.com

© 2026 Milliman, Inc. All Rights Reserved. The materials in this document represent the opinion of the authors and are not representative of the views of Milliman, Inc. Milliman does not certify the information, nor does it guarantee the accuracy and completeness of such information. Use of such information is voluntary and should not be relied upon unless an independent review of its accuracy and completeness has been performed. Materials may not be reproduced without the express consent of Milliman.