

Comparative Perspectives on Deploying Computable Contracts

Dr. Megan Ma, Emeline Careme, Preston Carlson

Contracts, written in natural language, are the cornerstone of insurance processes. They are simultaneously the product, as well as the foundation on which the relationship between the policyholder and the insurer is defined. In the event of claims or a dispute, the terms and conditions of contracts can be a source of both parties' frustrations. However, investments in the modernization and digitization of the insurance contract have historically been scarce. This has resulted in a decentralization of contract data, which has raised both structural and technical barriers to the ability of insurers to fully automate underwriting and claims processes, and to create modular and highly customized insurance coverage. That is, until the advent of computable contracts.

Computable contracts – the representation of contract terms in a form executable by computers – are not merely a theoretical technical construct. Rather, they are ripe for adoption and finding implementation in contemporary commercial practices. In this paper, we consider questions of deployment, focusing on the experiences at the Stanford Center for Legal Informatics (CodeX) and at AXA, in developing and implementing computable contracts in real world contexts. In the case of CodeX, we explore how computable contracts can help Stanford students understand complex health insurance coverage. For AXA, we explore how the creation and application of such contracts have led to substantial improvements in the way insurance products are designed and operated. Furthermore, we reflect on how other frontier technologies, including large language models (LLM), are being integrated with computable contracts, and identify areas of future research. We then conclude with a few observations and next steps.

Computable Contracts (CompK) at CodeX

At CodeX, the computable contracts technology we develop is based on using logic programs to represent the terms and conditions of insurance policies. We call this technology Logic Program Insurance Contracts, or LPIC.

In deploying a computable contracts system that uses this technology, we aim to do the following. First, to demonstrate the full value of the LPIC approach by providing value to real consumers and by realizing the operational efficiencies afforded by the technology, such as:

1. Automatically determining whether real or hypothetical situations are covered by a portfolio of insurance policies.
2. The inherent explainability of the coverage determinations.
3. The ease of updating policies, since their computable representation more closely resembles their natural language form.
4. The ability to directly compare the coverage offered by a policy with (i) the coverage offered by other policies and (ii) the coverage needs of the consumer.

Second, we aim to better understand the complexities/difficulties associated with fielding a real system based on LPIC so that we can better support others who want to make use of the technology.

Computable Contracts at AXA

AXA has also been working on computable contracts and has done so since 2018. In collaboration with CodeX, research on computable contracts was initially conducted to address challenges across numerous stakeholders in the insurance industry. The most evident concern was a rising demand from consumers on increased transparency regarding their insurance policy coverage. An additional motivating factor was to capitalize on technological capabilities to maximize efficiency of internal processes.

Accordingly, AXA had developed a logical language to represent the legal wording of insurance contracts. Following research on business needs from various internal and external sources, AXA built a software suite that leverages this representation of the insurance bargain to automate or assist business operations, including contract underwriting, claims coverage check, etc. AXA Computable Contracts enables the transformation of written contracts (products, policies) into standardized digital models. These models are readable, understandable, and executable by machines. This comprehensive and unified modeling of the insurance policy can be leveraged consistently across IT systems and organizations, driving tangible improvements in areas of standardization, interoperability, simplification, automation, and transparency.

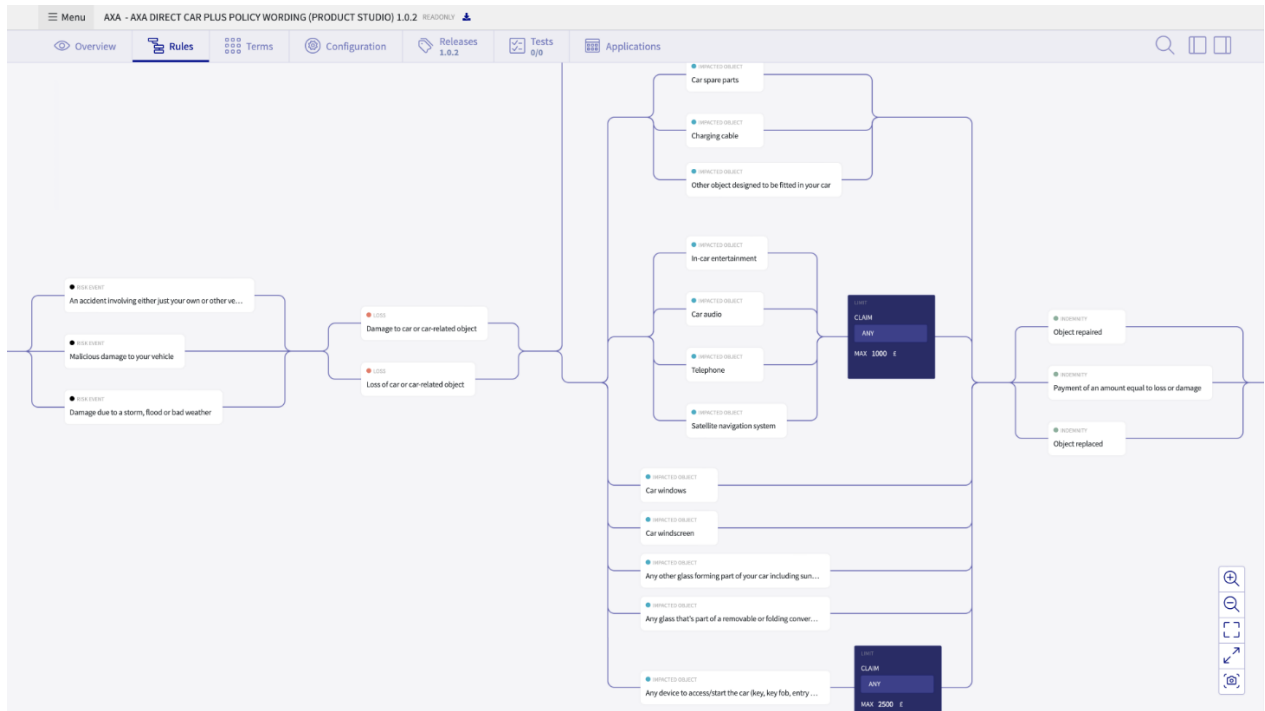


Illustration 1 – Motor product modeling via AXA’s graph editor interface

Three key products leverage AXA Computable Contracts: (1) Contract Studio; (2) Coverage Check (Instant Coverage Check and Claims Coverage Check); and (3) Contracts Analytics, each addressing a specific segment of the insurance value chain. Each product can be used independently or used in combination with the others to unlock additional value.

A. Contract Studio

The Contract Studio platform seeks to improve productivity of policy underwriters and to provide AXA sales teams with a tool to co-create in real-time insurance policies with clients and/or distribution partners. Previously, it had taken several days to design and validate contracts in line with customer needs and company policies. With AXA Computable Contracts, this can now be completed in a matter of hours. Moreover, Contract Studio reduces the insurers’ risk by translating core portions of the Underwriting guidelines into logical dependencies between clauses.

Contract Studio was successfully deployed in AXA Partners France (B2B2C partnerships subsidiary in France) and has enabled automating the generation of insurance notices for 15 different products. The solution has also been deployed at AXA Switzerland, for policy underwriting in Trade Credit Insurance and will likely be extended to other European subsidiaries in the coming years.

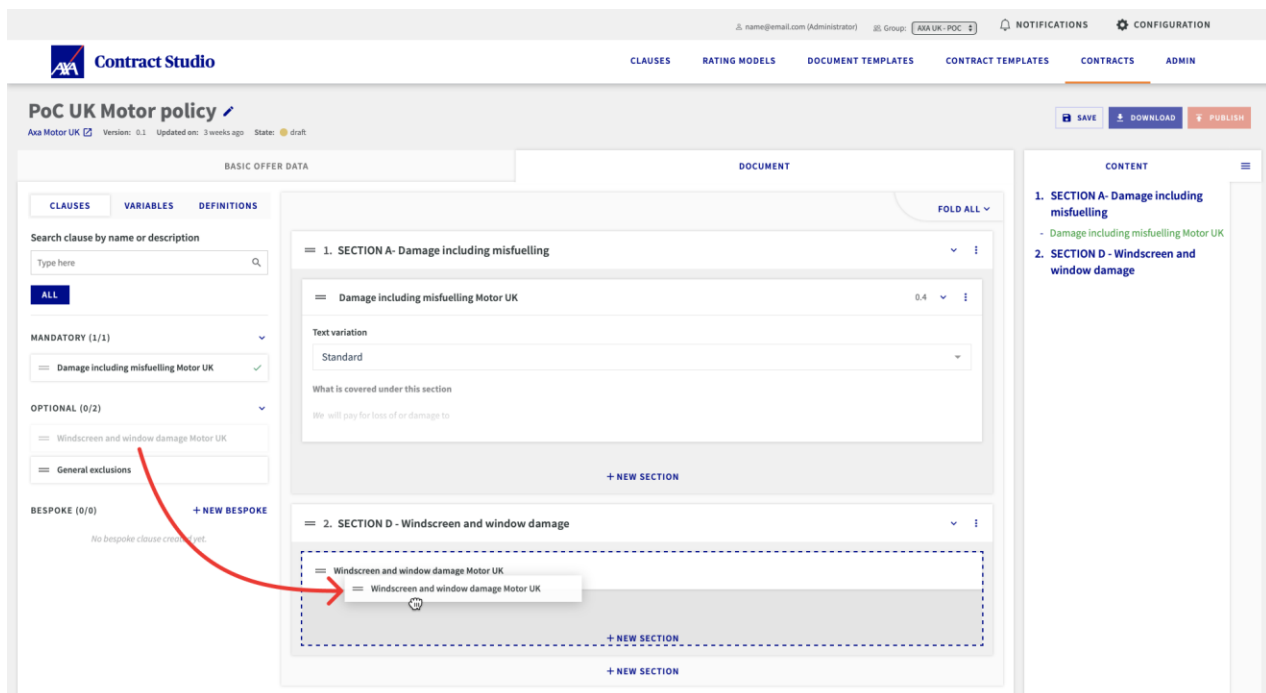


Illustration 2 – Contract Studio interface: drag & drop clauses

B. Coverage Check

Coverage Check leverages AXA Computable Contracts via two key functionalities. The first functionality is called Instant Coverage Check. It enables clients to assess instantly the extent of their policy coverage in light of a given risk. The second functionality, called Claims Coverage Check, is designed for claim management, automating the coverage check to provide an immediate calculation of the coverage. It enables improved efficiency of claim management, and improved certainty in the interpretation of coverage.

As of today, AXA Coverage Check products have been successfully deployed across three AXA subsidiaries: (1) AXA Partners; (2) AXA Switzerland; and (3) Direct Assurance. AXA Partners is now able to automatically execute Claims Coverage Check for two product categories:

1. Motor & Home assistance products (about 300 product variations) in five countries (Belgium, France, Germany, UK, Italy and Luxembourg), representing about 1.5M claims per year handled by 1000 employees in customer service centers
2. Travel products (+3000 active product variations) in three countries (Germany, UK, Spain), handled by 2500 collaborators for more than 300 retail partners

AXA Switzerland is leveraging Claims Coverage Check for First Notice of Loss and deductible validation in the Invoice check since 2022 representing about 200k claims per year. Direct Assurance has deployed the Instant Coverage Check for their Motor product since early 2023 powered by more than 2M calls to our API per year.

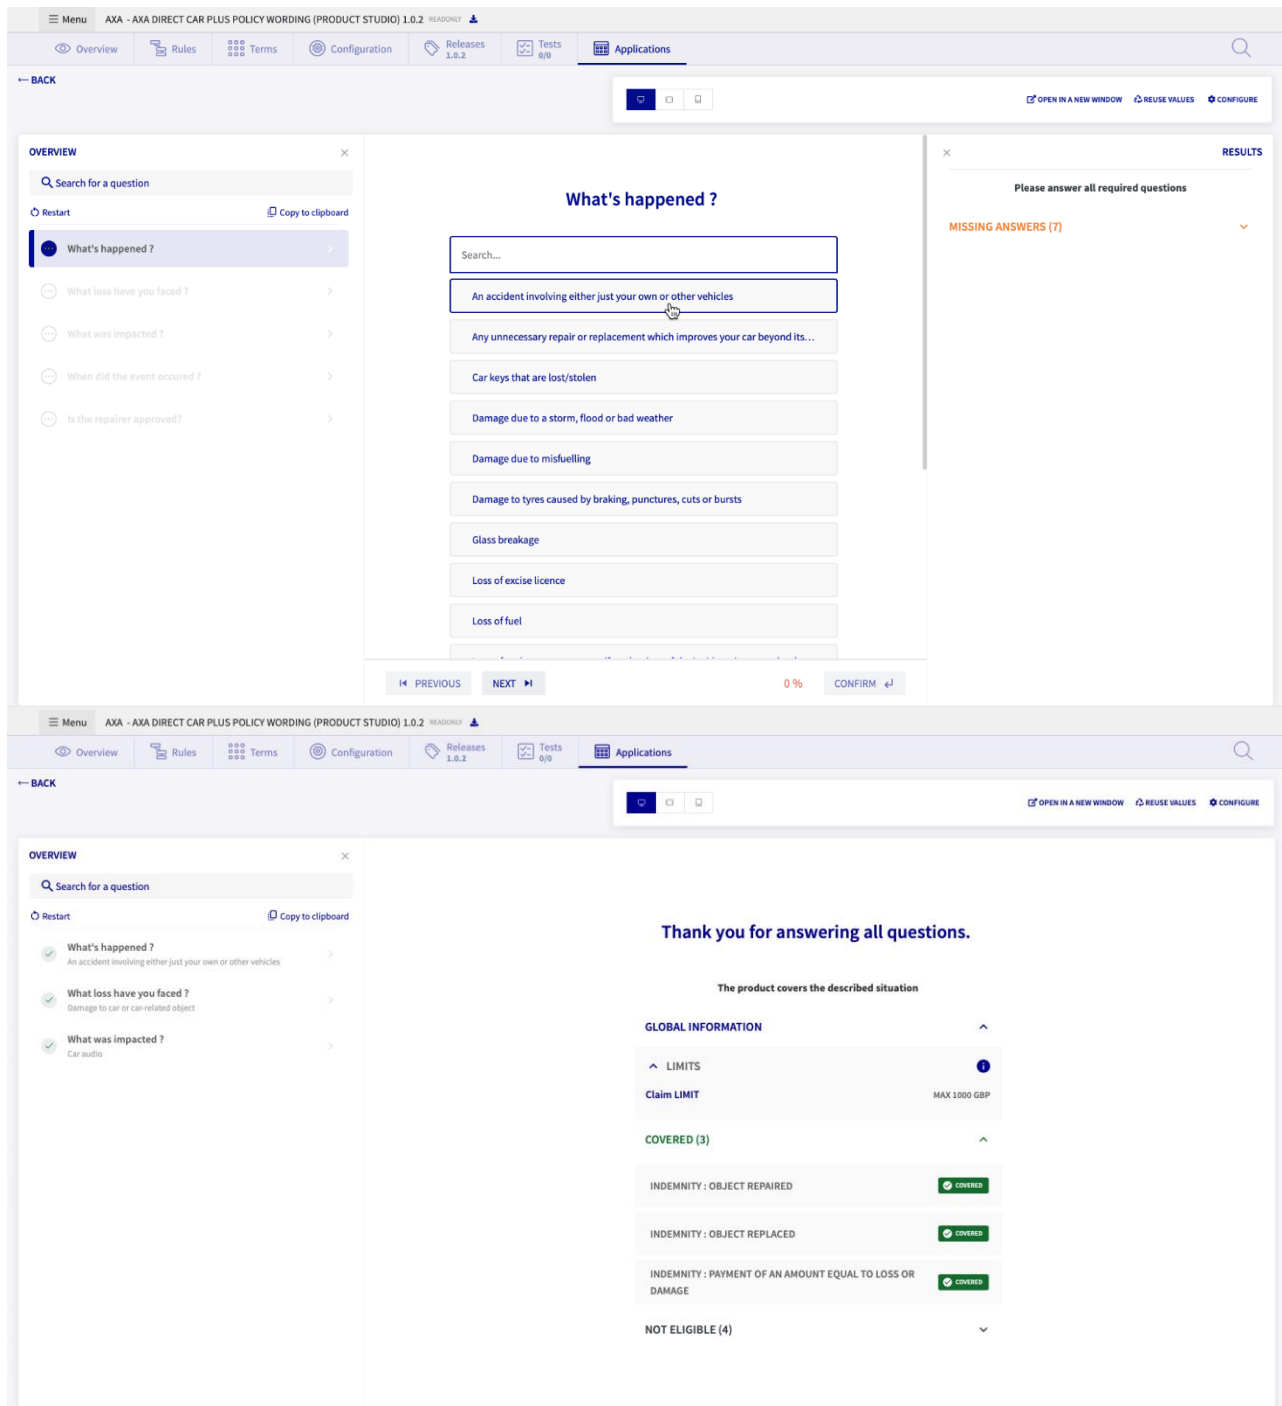


Illustration 3 – Claims coverage check: dynamic questionnaire form and coverage result

C. Contract Analytics

Finally, the Contract Analytics tool is capable of portfolio risk monitoring, provides market insights, and benchmarks insurance products. This automated review is made possible by the data generated by the Computable Contracts approach. Contract Analytics provides benefits such as [improved risk management](#) (e.g. optimized management of reserves) and improved product performance. Experimentations are conducted within AXA to extract a

simplified Computable Contract from commercial lines contracts, both to improve the reporting of their risk exposure, and to identify problematic policies, e.g. policies with unlimited coverage.

All three solutions have different impacts: (1) on the operating margin, such as improving efficiency of sales and claims agents, (2) on the technical margin, such as reducing leakage and improving risk management, and (3) on net new business, such as improving time to market, conversion, better customer satisfaction and better performance of products.

Considerations for Initial Deployment of CodeX CompK

With CompK at CodeX, we ultimately aim to deploy a comprehensive insurance portfolio dashboard based on LPIC technology. Such a system can “provide its user with an overview of [their] current policies and needs, with an analysis of overlap and gaps” ([Genesereth 2022](#)). However, deploying a fully-featured insurance dashboard requires encoding many different insurance policies as logic programs (potentially hundreds or thousands), so that users can represent their insurance portfolios in full. Currently, our LPIC approach does not scale out-of-the-box to encoding hundreds of policies without significant human labor. So, we chose to initially deploy a smaller-scale system.

The initial system is focused on the Stanford student body as its user base and is built upon the encoding of a single insurance policy: the Cardinal Care student health insurance policy offered by Aetna Student Health. We decided to focus on Cardinal Care for this initial system for a few core reasons.

First, health insurance is a very complex domain. Health insurance is a good stress-test of our technology because (i) the conditions for coverage are complex, (ii) the policies are large, and (iii) developing a useful health insurance policy LPIC requires encoding many documents and guidelines that are external to the core policy documents. And since the complexity of the policies leads to a lot of consumer confusion, we can provide a lot of value by helping clarify the coverage offered by such policies.

Second, there are many similar policies offered by the same company (Aetna Student Health), so our work encoding Cardinal Care would generalize well to the encoding of many other policies. **Third,** the Stanford student body is quite homogeneous in terms of their health insurance policies, since students are required to enroll in Cardinal Care unless they have sufficient other insurance to waive enrollment. This allows us to serve many users by encoding a single policy. **Finally,** as CodeX is situated on Stanford campus, we have ready access to our user base, which allows us to easily get high-fidelity feedback on the system and on user needs.

COVID-19 Vaccine Needs More Info

Constraints

^ Date of Birth

^ Are you immunocompromised?

^ Vaccination History
Add an entry for each Covid vaccine you've received. If you received it before September 11, 2023, indicate its type. If it was the most recent formulation (i.e. received since September 11, 2023), indicate that instead.

^ What brand of vaccine did you receive?

^ When was the vaccination performed?

^ Where was the vaccine administered?

COVID-19 Vaccine Covered

Constraints

^ Date of Birth

^ Are you immunocompromised?

^ Vaccination History
Add an entry for each Covid vaccine you've received. If you received it before September 11, 2023, indicate its type. If it was the most recent formulation (i.e. received since September 11, 2023), indicate that instead.

^ What brand of vaccine did you receive?

^ When was the vaccination performed?

Illustration 4 – COVID-19 Vaccine Coverage Check (CodeX CompK)

This smaller-scale system has allowed us to demonstrate most of the value of our approach (apart from the ability to compare coverage between policies), while enabling us to iterate quickly on our encodings as we (i) have been faced with the software development constraints associated with deploying a consumer-facing system, (ii) discover new best practices for encoding policies, and (iii) come to better understand the line-of-business and specific product we're encoding. And the lessons we've learned in developing the Cardinal Care encoding are informing our research into how to encode policies at the scale necessary to enable a fully-featured insurance portfolio dashboard.

Challenges in Deployment of AXA Computable Contracts

One of the primary challenges of deployment is identifying the right vocabulary to articulate the long-term impact of computable contracts. That is, computable contracts represent a fundamental transformation of core insurance processes. This form of [change management](#) targets the core value chain and, therefore, requires strong commitment from multiple stakeholders: the business teams of the product, the legal team, underwriting, as well as claim management.

Achieving success in change management in our context is particularly difficult, considering the complexity of communicating a concept of technical novelty such as computable contracts. This concept is very different from traditional IT configurations, which occasionally made the understanding and value of implementation difficult to grasp. However, it is paramount to engage stakeholders in knowledge sharing in order to understand the underlying benefits of representing the terms and conditions of an insurance contract in logical form, agnostic from any IT system.

Managing Expectations with Deploying CodeX CompK

CodeX faced similar challenges in inter-institutional knowledge sharing and engagement of stakeholder buy-in. As healthcare decisions are incredibly important and consequential, we interfaced with numerous non-user stakeholders in advance of deployment and are continuing to do so. Particularly, we have engaged with two non-user groups: Vaden and the VSHAC.

Vaden is the campus student health care provider and is responsible for communicating with students regarding Cardinal Care coverage and on-campus healthcare. In discussing deployment with them, we observed that integrating computable contracts into existing systems would be unduly burdensome for their team, given their difficult position as the institutional intermediary, accountable for the communication of insurance policies that they do not themselves administer. As well, they expressed concern about bearing responsibility in the event that these systems could inadvertently misinform Stanford students about their coverage. Therefore, we have decided to deploy a standalone system, taking significant efforts to align our messaging around Cardinal Care with theirs.

The VSHAC is the Vaden Student Health Advisory Committee, and acts as the students' voice regarding their healthcare and health insurance needs. The VSHAC has provided an invaluable perspective for our user research by helping us learn what portions of Cardinal Care coverage are most relevant to the student body and by showing us the expectations that students are likely to have of our LPIC system.

In collaboration with the VSHAC, we are now testing the system with subsets of the student body. It is currently being presented as a research system while we learn (1) the limits of our encoding and (2) how to improve the service to meet the needs of our user base. The initial system has encodings of the general conditions and exclusions of Cardinal Care and full encodings of particularly complex benefits, including Flu and Covid vaccine coverage and treatment of infertility. Early testing has shown us the importance of managing the expectations of users regarding the capabilities and intended use cases of the LPIC system.

Integrations with Frontier Technology at AXA: Large Language Models (LLMs)

Large Language Models (LLMs) have demonstrated impressive capabilities of interacting in natural language with consumers in a highly accessible manner. Leveraging the relative strengths of LLMs to process unstructured data, AXA has recently piloted a number of initiatives that integrate LLMs with computable contracts. These include:

- Automating or assisting computable contracts generation from existing Terms and Conditions by leveraging the capabilities of LLM to extract concepts as well as logical rules from natural language
- Extracting automatically, risk exposure information from policy documents, as an input for risk analytics at policy and portfolio level
- Assisting the comparison of two insurance policies among different dimensions, to identify their respective strengths and weaknesses, and capture new customers from competitors
- Improving the claims experience by leveraging LLMs to extract unstructured data from voice recording, proof documents or form text.

AXA is actively researching questions of technological interoperability, particularly deepening how LLMs can play a role in scaling and enabling new user experiences.

The Road Ahead: Scaling CodeX CompK Pilot for Deployment

The deployed LPIC system for Stanford students is a first step before scaling to a full insurance portfolio dashboard. We have many possible next steps in the next phase of deployment.

- a. We could scale to serve other universities.** As mentioned, Aetna Student Health offers many similar policies. Accordingly, our work encoding Cardinal Care could easily be adapted to encoding other university student healthcare policies. Moreover, our experience fielding a system for student health insurance could better situate us to encode student health insurance policies in general.
- b. We could serve the broader Stanford community.** University staff and faculty are also relatively homogeneous in their health insurance, since Stanford offers a handful of employer-sponsored health insurance plans. And as we have ready access to the Stanford community as a user base, we could easily get a sense of which other (i.e. non-health) insurance policies are of interest to them.
- c. We could scale directly to a general insurance portfolio dashboard.** If we are soon able to scale the encoding of insurance policies, we could start serving users

far beyond the university environment. (Notably, LLMs may present an opportunity for scaling the encoding of insurance policies. In our [initial exploratory efforts](#), we found that for a simplified hospital cash policy, the o1-preview model from OpenAI was able to generate and use LPIC encodings to get correct answers to coverage queries 83% of the time, on average.)

Concluding Remarks

Computable contracts represent a pivotal advancement in insurance technology, converting static policy documents in complex verbiage into dynamic, precise logic that defines coverage terms and claims handling with a high degree of accuracy and specificity. Implementations of these solutions, showcased in this paper, demonstrate tangible benefits through real-time policy verification.

By encoding policy terms as executable logic, every coverage determination and claims decision can be traced and validated. This eliminates the ambiguity that has historically plagued insurance contracts and provides both carriers and policyholders with clear, auditable proof of how specific situations are covered.

Early implementations, as illustrated by AXA and CodeX, have shown dramatic improvements in accuracy and consistency. In effect, computable contracts are not a theoretical construct and are actively transforming insurance operations today. We consider this technology as not just an incremental improvement, but as a fundamental redesign of how insurance products may be structured and serviced to the broader public.