BUSINESS ROBOT.

# Agent-native insurance

Delegated buying, distribution and the changing role of digital brands

Academic-style conceptual white paper | Business Robot Research

UK and US motor insurance focus | 1 October 2026

---

**Abstract**

Insurance is beginning to move toward interfaces through which AI systems can retrieve information, compare offers and coordinate actions. This paper defines agent-native insurance as an operating and distribution model designed for bounded delegation, machine-readable transactions and accountable decision-making. It distinguishes that model from chatbots, conventional machine learning and unrestrained autonomy. Evidence reviewed through 1 October 2026 includes insurer deployments, an in-chat motor-insurance quotation service, regulator analysis, commerce specifications and academic security research.

The evidence supports a directional thesis: agent-compatible capabilities are becoming strategically relevant to insurers and intermediaries. It does not establish universal adoption, a timetable for full autonomy or the disappearance of human professionals. A hybrid model, in which agents prepare and coordinate transactions while consumers and regulated firms retain defined decision rights, is a better-supported near-term interpretation.

For digital brands, delegation can reduce the importance of URL recall while increasing the importance of trusted identity, readable product information and reliable transaction endpoints. A descriptive domain may retain human-facing campaign value even as delegation grows. No evidence reviewed establishes an automatic AI-ranking advantage or demand premium for any particular domain asset. The commercial case for such an asset should be tested rather than inferred from AI adoption.

**Keywords**

Agentic AI; motor insurance; delegated commerce; distribution; consumer protection; digital identity; domain names.

Status: conceptual synthesis, not peer-reviewed research. No original consumer study, integration test or commercial appraisal was conducted. Not legal advice or an investment recommendation.

---

## 1. Question, scope and method

The central question is not whether insurers will use AI. Many already do. It is whether buying, servicing and administering insurance will increasingly be organised around software agents acting within delegated authority. A second question follows: if agents mediate access to insurers, which parts of a brand's digital presence remain valuable?

### Operational definition

An AI agent is defined here as a software system that can select actions, use external tools and coordinate multiple steps toward a goal. "Agent-native" means that permissions, data structures, product interfaces, transaction states and oversight are designed for those interactions. It does not require every decision to be autonomous. An organisation can be agent-native while requiring approval for purchases, cover changes or claim payments.

Four levels should be separated: assistance, such as explanations; orchestration, such as retrieving and comparing information; bounded execution under explicit constraints; and adaptive autonomy, involving ongoing decisions without case-by-case approval. These are analytical categories, not a regulatory taxonomy. A chatbot may offer only assistance. An API-connected system may coordinate a transaction without independently choosing a policy.

### Evidence strategy

This is a purposive narrative review of fourteen public primary sources: regulator surveys and policy publications; company announcements; commerce specifications and platform documentation; and academic security research. Sources were read directly. Selection favoured

evidence that distinguishes deployed capability from aspiration. This is not a systematic literature review or a representative survey of insurers.

Regulator surveys support bounded population statements. Company announcements establish what a company reports, not independently verified effectiveness. Specifications establish interface designs, not market penetration. Academic benchmarks establish experimental failure modes, not production insurance incident rates. Predictions and proposed designs below are identified as analysis rather than observations.

#### Boundaries

The focus is personal motor insurance and distribution. Other insurance lines supply operational examples. European supervisory guidance is a comparator, not UK or US law. Domain examples are illustrative: this review does not assess the ownership, registration status, traffic, trademark clearance, site quality or financial performance of any specific domain.

---

## 2. What the evidence shows

#### AI adoption is not agent autonomy

The Bank of England and FCA's 2024 survey received 118 responses. Across responding financial-services firms, 75% reported using AI. Only 2% of reported AI use cases had fully autonomous decision-making. These are different denominators, and the survey covers financial services more broadly than motor insurance [1]. The figures cannot be restated as "75% of insurers are agent-native."

In the US, the NAIC reports that 88% of 193 responding auto insurers used, planned to use or planned to explore AI/ML. This broad measure includes intentions and exploration, and concerns AI/ML rather than specifically autonomous agents [2]. It demonstrates adoption and interest in underlying techniques, not universal agent-led distribution.

#### Operational deployments

Allianz reports that Project Nemo launched in Australia in July 2025 for low-value food-spoilage claims. Task-specific agents perform checks and prepare an audit summary, but a human makes the

final payment decision [3]. This is evidence of a bounded claims workflow in another line and geography, not autonomous UK or US motor-insurance purchasing.

AXA's 8 September 2026 announcement says its Global AI Hub's first version was delivered in July and was in use across five entities. Developing use cases include motor claims automation, customer email processing and enterprise knowledge management. AXA retains human oversight for business-critical decisions [4]. Deployment of a platform is not proof of production rollout of every developing use case.

#### Direct distribution evidence

Insurify's August 2026 announcement reports real-time personalised car-insurance quotes inside ChatGPT. Depending on the carrier, shoppers proceed to the insurer's website or complete purchase through Insurify's agency. The announcement lists an initial eleven-state footprint and plans for expansion [5]. This demonstrates an agent-compatible journey with continuing intermediary participation, not universal in-chat policy binding.

The FCA's July 2026 Mills Review treats agent-led consumer journeys and changing control of customer interfaces as credible future shifts [6]. It supplies strategic analysis, not a guaranteed schedule. Together, these sources show a real transition in selected workflows and channels while leaving its breadth and speed unresolved.

---

## 3. A proposed agent-native transaction model

Motor insurance involves more than paying for an item. A transaction connects applicant and vehicle facts, cover requirements, price, disclosures, insurer acceptance and policy documents. The following is a proposed reference architecture, not an established industry standard.

#### A bounded sequence

1. Establish intent. Record drivers, vehicles, desired cover, start date, excess limits and affordability. Distinguish supplied facts from uncertain information. Specify what may be disclosed and what may be committed.

2. Discover providers. Identify eligible insurers or intermediaries and the scope of their panels. A partial panel should not be presented as the whole market.

3. Obtain and normalise offers. Preserve authoritative quote identifiers, timestamps, expiry, assumptions, total payable, payment schedules, excesses and exclusions. Explain differences without changing underlying terms.

4. Compare and confirm. Expose cover trade-offs and material declarations. Bounded delegation may allow some steps without new approval, but its limits must be explicit. "Find the cheapest" does not completely specify acceptable protection.

5. Execute and reconcile. The insurer or licensed intermediary validates the application and confirms cover. Payment success is not equivalent to cover being in force. Interrupted, expired or failed transactions require a clear final state, accessible documents and support.

#### Protocol opportunity and limitation

The Agentic Commerce Protocol provides an interaction model for programmatic checkout and payment coordination [7]. OpenAI's September 2025 Instant Checkout launch illustrates consumer-reviewed agent commerce in retail, not insurance [8]. Google's AP2 specification describes signed checkout and payment mandates and receipts, including human-present and autonomous flows [9]. These are relevant building blocks, not evidence that insurance-specific advice, quoting, binding or cancellation requirements are solved.

The design implication is layered: the agent interface sits above authoritative insurer systems and explicit permission controls. Probabilistic language generation should not become the source of truth for premium, eligibility, terms or effective cover. Protocol availability also does not establish an insurer's permission to participate or an agent's authority to commit the consumer.

---

## 4. Distribution economics and competition

#### Efficiency is plausible, not automatic

Delegated retrieval could reduce repeated data entry and make renewal comparison easier. Benefits depend on accurate risk facts, comparable cover, usable interfaces and willingness to delegate. No

conversion uplift, acquisition-cost reduction or aggregate savings estimate is offered: the reviewed evidence does not identify a causal market-wide effect.

Operational adoption can proceed without distribution transformation. An insurer may automate claims preparation while retaining its website, call centre and brokers. A distributor may offer conversational comparison while insurers retain conventional underwriting systems. These paths need not advance together [3-5].

### Disintermediation or new intermediation?

Agents could connect consumers directly to insurers, but could also strengthen existing comparison services. Insurify's reported journey retains an agency role [5]. Re-intermediation is therefore a plausible alternative to intermediary extinction: control may move toward the trusted consumer interface, data access and transaction integration.

The Mills Review identifies control of customer interfaces, concentration and vendor dependence as competition issues [6]. Agent-platform reach can be attractive while creating dependence on admission, ranking and commercial rules. An open protocol alone does not guarantee placement. The ACP site says each platform manages participation, and describes discovery mechanisms as a work in progress [7].

### Brand economics

An agent comparing price and cover may reduce advertising-led recall's role in the initial shortlist. Yet consumers may request a known insurer, reject an unfamiliar one or discover a service through advertising. These mechanisms can coexist. Brand trust differs from a memorable URL: a name can help people locate a provider without establishing authorisation or product quality.

In heavily delegated journeys, verified identity and reliable quote access may matter more than a catchy address. In mixed journeys involving radio, podcasts and consumer approval, verbal identity may still matter. These are competing hypotheses, not proven price premiums. Insurers should treat agent compatibility as a distribution and governance capability. Domain buyers should distinguish a name from traffic, reputation, integrations and customer relationships.

---

## 5. Governance and reliability

### Continuing responsibility

The FCA's published AI approach points to existing frameworks, including Consumer Duty and senior-manager accountability. Firms must provide products suited to target customers, fair value, appropriate communications and support [10]. A conversational interface changes delivery, not applicable duties. UK advice, distribution and data-protection questions require activity-level legal analysis.

The NAIC says insurers remain responsible for compliance when AI makes or supports decisions. Its Model Bulletin was adopted in December 2023 [2]. That is not a single federal insurance law or evidence of identical adoption in every state. US implementation requires jurisdiction-specific review.

EIOPA's August 2025 opinion clarifies existing governance principles including record-keeping, fairness, cybersecurity, explainability and human oversight [11]. It is an EU comparator. Its discussion of high-risk life and health pricing under the AI Act must not be transposed into a claim that all motor-insurance agents have the same classification.

### Security and factual integrity

INJECAGENT and AgentDojo demonstrate that malicious instructions embedded in external material can redirect tool-using language-model agents [12,13]. Their results are not insurance production loss rates. They nevertheless justify treating retrieved pages and third-party material as potentially hostile inputs, not authority to change consumer instructions.

Proposed controls include validated quote schemas; separation of source data from executable instructions; minimum necessary permissions; independently checked transaction parameters; explicit commitment gates; transaction identifiers that prevent unintended repeats; and records of user authority. These are design recommendations, not a guarantee that prompt injection can be eliminated.

### Meaningful oversight

A nominal human reviewer is insufficient if the reviewer cannot understand or challenge the recommendation. Review should expose cover differences, uncertain declarations and authoritative

sources. Customers need corrections, complaints support and a non-delegated route. Model and supplier changes require renewed outcome testing.

The accountability chain should distinguish consumer agent, distributor, insurer and payment provider. It must remain possible to recover who supplied a fact, selected an option, accepted risk and confirmed cover. Payment mandates [9] can evidence spending authority, but do not alone establish informed insurance consent or appropriate cover.

---

## 6. Descriptive domains under delegated buying

### A conditional commercial hypothesis

A descriptive insurance domain may combine an urgency cue with a product description. That linguistic property alone does not demonstrate recall, conversion or value. The positive case is strongest in a mixed journey: someone hears an advertisement, remembers or speaks the name, then asks an assistant to find the associated provider. A category cue may clarify the intended service. Whether an urgency cue is heard as a word or as spelled letters, and whether listeners reproduce the address correctly, are empirical questions.

### The counter-case

If someone asks "compare suitable car insurance for me," the agent need not know or type any particular address. It may consult integrated distributors, indexed content or provider registries. A descriptive domain creates neither insurer authorisation nor an API, underwriting capacity, reputation or platform inclusion. A credible endpoint can exist under many names.

Google's guidance says normal SEO foundations remain relevant to AI Overviews and AI Mode, without additional technical eligibility requirements. It recommends accessible text and structured data that matches visible content [14]. This supports credible information, not a promise that a keyword-rich domain gets preferential AI ranking. The guidance concerns Google Search, not all agents.

### Identity is more than an address

A useful agent-facing presence connects its domain to a verifiable business, actual provider or intermediary, applicable authorisation, product scope and reliable service interfaces. A domain names that presence; it does not authenticate the economic counterparty or policy claims by itself. Machine-readable endpoints may support integration, while platform onboarding and insurance permissions remain separate [7,9].

For an established insurer, another domain may also create migration, consistency or confusion costs. A campaign redirect has different economics from a new brand. Without buyer-specific plans, comparable transactions and incremental performance data, a defensible commercial assessment is not possible.

The supported claim is conditional: a descriptive domain might be useful for a human-facing campaign or as a named entry point to an agent-compatible service. The unsupported claim is that agent-native insurance necessarily makes any such domain valuable, guarantees discovery or obliges insurers to adopt it. AI adoption alone supports neither a price nor an expression of market interest.

---

## 7. Scenarios, tests and conclusion

### Alternative futures, not probability forecasts

Assisted continuity: agents help people while websites, brokers and human approval remain central. Hybrid delegation: agents obtain quotes and prepare transactions, with bounded approvals and escalation. Platform-led delegation: persistent agents control much of discovery and execution; manual URL recall becomes less central, while provider identity, integration and platform access gain weight.

These scenarios have no assigned probabilities or deadlines. Current evidence demonstrates elements of assistance and hybrid delegation [3-5]. Useful future indicators include actual agent-channel policy bindings, repeat consumer use, complaints, published insurer interfaces and the scope of delegated authority.

### Falsifiable propositions

P1: agent-assisted comparison reduces effort without worsening cover suitability. Compare time, completion, omitted risk facts, comprehension and subsequent corrections against conventional journeys. Fewer clicks alone do not establish better outcomes.

P2: agent-channel availability improves quote-shortlist inclusion. Test matched providers and eligibility conditions. Record panel restrictions, failures, platform and model versions. Search mentions are not purchasable quotes.

P3: a descriptive domain improves human-to-agent handoff. Randomise comparable audio advertisements using a descriptive category domain, an established brand address and a neutral campaign name. Hold offer and placement constant. Measure unaided recall, voice transcription, correct provider identification, qualified visits and completed quotes. Accept null or negative results.

P4: deeper delegation reduces URL recall's contribution relative to trusted identity and interface access. Compare behaviour across delegation levels while controlling for brand familiarity, price and cover. Correlation between AI adoption and domain demand would not prove causation.

### Conclusion

Invest first in accurate product data, controlled interfaces, permissions and outcome testing. Evaluate a domain against a defined use case and measured incremental contribution, net of deployment costs. Selected insurers and distributors are already moving toward agent-compatible workflows. "All insurers will go agent-native" remains an unproven universal forecast. The stronger thesis is that competitive insurance businesses may increasingly need to serve both people and their authorised agents. That changes distribution requirements; it does not settle the value of a name.

## References

All sources accessed and read on 1 October 2026. Company announcements are not independent audits. Mutable documentation describes the version read, not a guaranteed future standard.

[1] Bank of England and FCA (2024). Artificial intelligence in UK financial services - 2024. 21 November 2024. Regulator survey: adoption, autonomy and respondent counts. https://www.bankofengland.co.uk/report/2024/artificial-intelligence-in-uk-financial-services-2024

[2] NAIC (2026). Insurance Topics: Artificial Intelligence. Page labels last update 3 April 2026. Official overview: survey denominators, compliance and Model Bulletin. https://content.naic.org/insurance-topics/artificial-intelligence

[3] Allianz (2025). When the storm clears, so should the claim queue. 3 November 2025. Company account: Project Nemo and human payout review. https://www.allianz.com/en/mediacenter/news/articles/251103-when-the-storm-clears-so-should-the-claim-queue.html

[4] AXA (2026). AXA and Publicis Sapient collaborate to deploy Global AI Hub at scale. 8 September 2026. Company announcement: platform deployment and developing use cases. https://www.axa.com/en/press/press-releases/axa-and-publicis-sapient-collaborate-to-deploy-global-ai-hub-at-scale

[5] Insurify (2026). Insurify Expands ChatGPT Plugin with Real-Time Personalized Quotes and In-Chat Shopping. Updated 18 August 2026. Company announcement: quotes, carrier-dependent purchase route and state coverage. https://insurify.com/press/news/insurify-expands-chatgpt-plugin/

[6] Mills, S. / FCA (2026). The Mills Review: AI and the future of retail financial services. July 2026. Regulator strategic review: consumer journeys and competition. https://www.fca.org.uk/publication/corporate/the-mills-review.pdf

[7] Agentic Commerce Protocol (undated). Agentic Commerce Protocol. Official specification site: checkout model, participation and evolving discovery. https://www.agenticcommerce.dev/

[8] OpenAI (2025). Buy it in ChatGPT: Instant Checkout and the Agentic Commerce Protocol. 29 September 2025. Historical retail launch, not insurance availability. https://openai.com/index/buy-it-in-chatgpt/

[9] Google Agentic Commerce / AP2 contributors (undated). Agentic Payment Protocol (v0.2), specification. Mutable repository document: mandates, receipts and delegated payment flows. https://github.com/google-agentic-commerce/ap2/blob/main/docs/ap2/specification.md

[10] FCA (undated). AI and the FCA: our approach. Official policy page: existing frameworks, Consumer Duty and accountability. https://www.fca.org.uk/firms/innovation/ai-approach

[11] EIOPA (2025). EIOPA publishes Opinion on AI governance and risk management. 6 August 2025. EU supervisory principles, not UK or US law.
https://www.eiopa.europa.eu/eiopa-publishes-opinion-ai-governance-and-risk-management-2025-08-06_en

[12] Zhan, Q., Liang, Z., Ying, Z. and Kang, D. (2024). INJECAGENT: Benchmarking Indirect Prompt Injections in Tool-Integrated Large Language Model Agents. Findings of ACL 2024. Academic security benchmark.
https://aclanthology.org/2024.findings-acl.624.pdf

[13] Debenedetti, E., Zhang, J., Balunovic, M., Beurer-Kellner, L., Fischer, M. and Tramer, F. (2024). AgentDojo: A Dynamic Environment to Evaluate Prompt Injection Attacks and Defenses for LLM Agents. arXiv:2406.13352. Academic evaluation framework.
https://arxiv.org/abs/2406.13352

[14] Google Search Central (undated). AI features and your website. Official documentation: Google AI-search eligibility and SEO. Not a universal agent-ranking specification.
https://developers.google.com/search/docs/appearance/ai-features