QUIETROOM

How AI is changing the way your customers make decisions

This paper is based on a 50-minute Quietroom webinar presented by Thomas Joy and Russ Hope. Watch the webinar

Our customers and members are ahead of us

Since ChatGPT launched in 2022, AI has generated a lot of discussion.

For financial services providers, that discussion often focuses on belt-and-braces questions like, “How do these strange new tools work?” and “How can we use them to gain an advantage?”

But there’s too little discussion – or even awareness – about the ways in which AI has already become part of people’s lives. For millions of people – from teenagers to CEOs – tools like ChatGPT are like personal assistants who can help them do better work, by helping them think, write and spot patterns, and save them time by fetching answers quickly from an internet that goes on forever.

If we’re going to use AI well in business, we need to grapple with how AI-powered search, summary and decision-making tools are embedding themselves in people’s lives, and what that means for the content we create.

Understanding the impact of AI on your communications

To ensure customers get the best outcomes while keeping your business secure, it’s essential to understand how AI changes the way customers experience your content.

Testing and refining content for how AI interprets it can play a significant role here.

How customers access AI daily

Nearly everyone in the UK owns a smartphone, tablet, or computer. They probably use one for personal life and another for work. This means that one of three companies – Apple, Google, or Microsoft – is central to their life. Apple users are likely on an iPhone, MacBook, or iPad; Google users are on Android phones or Chromebooks; and Windows users are part of the Microsoft ecosystem.

Each of these companies is heavily invested in AI, branding their versions under different names: Apple's is 'Apple Intelligence,' Google's is 'Gemini,' and Microsoft's is 'Copilot.'

Between them, these 3 companies own almost the entire UK market share of laptops, phones, tablets, search engines, web browsers and email apps. And they're bringing AI to every single one of these experiences.

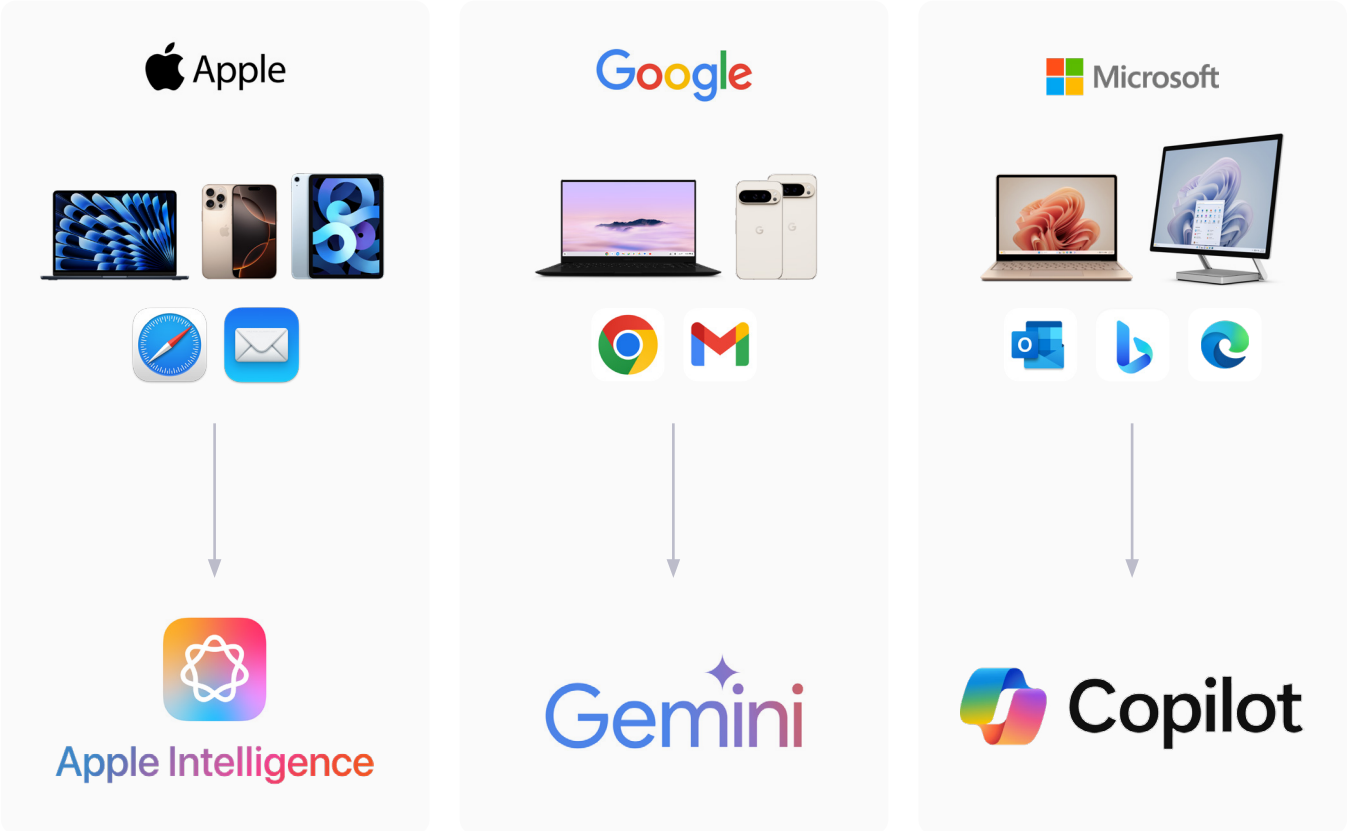

Three ways AI changes how people interact with content

At Quietroom, we've tested real life financial services content with the devices, apps, search engines and AI assistants that customers are using. We've looked at the research on how LLMs – and the AI assistants that sit on top of them – are developed. We want to understand how and why they change your content. We've brought in insights from research on the way that people make decisions. And we've used our own experience and expertise to explore how these tools are changing user behaviour.

From our research and testing, we've identified three major shifts in how people interact with information through AI:

1. Finding information
2. Consuming information
3. Engaging with information

1

How people find information with AI-powered search engines

Traditional search vs AI-powered search

Here's what a Google search results page used to look like.

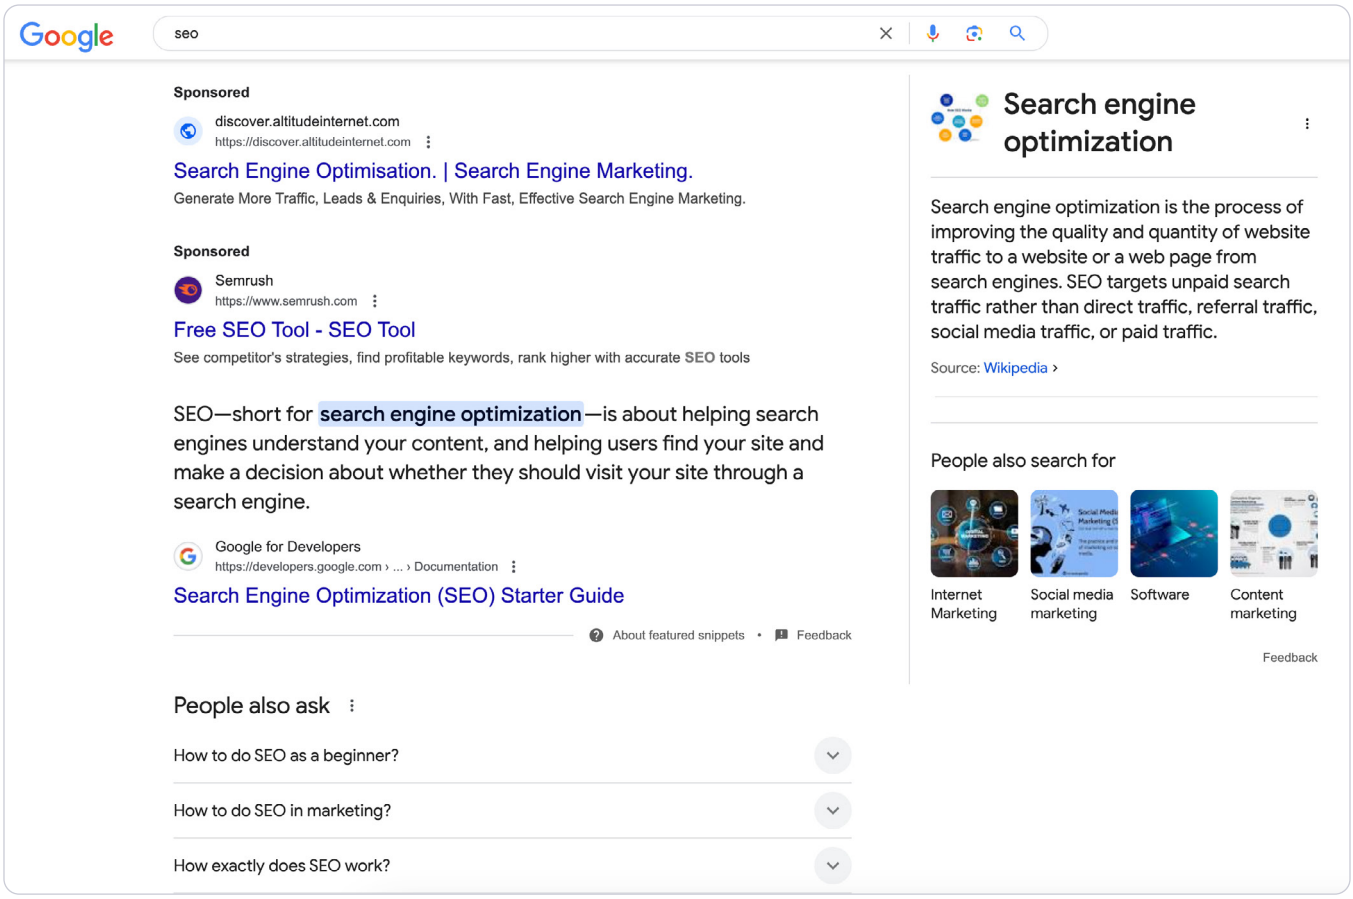

When we search, we're often looking for answers. These can range from simple questions like "What's my bank's phone number?" to complex questions like "How should I invest my money?"

In traditional Google search, you enter your query, and Google presents a list of links. The search engine expects you to click through to these links to find the answer.

AI-powered search changes this

Google is starting to generate AI answers to search queries. Users do not opt-in to these answers – they're appearing automatically for the questions that Google decides they're appropriate for.

So instead of sending a user to your website, Google reads your website instead. It picks out the parts that are relevant to a user's question, then summarises, condenses and rewrites the content in the search window. That means a user sees an answer that attributes you as the source, but which does not contain the exact content that you wrote and signed off.

Users often take these answers at face value. They do not place much weight on the disclaimer that says 'Generative AI is experimental', and they do not thumb through the list of sources to check the answer is right. If AI-generated answer intuitively seems correct then users will stop searching any further because they believe their question has been answered – even if the answer that Google gave them is wrong.

These answers can go wrong in lots of ways.

They can omit important details, like a letter from a postal address which could mean a user posts a form to the wrong address. Even minor errors like this matter, especially if an answer mistakenly pulls details from a misleading source or scam website designed to look like yours.

They can use sketchy sources – like being based on information from companies who say they're associated with you. We found that one advice company had smartly engineered their web pages so that, when a user searches about transferring their defined benefit pension, the adviser would appear as a source of help and information.

And AI can also simply get it wrong. A great deal of online financial services content is hard to read, contradictory, and out of date. When we looked at examples of incorrect answers in more detail, we found that a human wouldn't have had a chance at understanding the content – let alone an early version of an AI-powered search engine.

Of course, AI generated answers don't just change how users consume content about you, they change whether users even find out about you in the first place. In the past, a user who searched for a pension or insurance provider might scroll through a long list of search results and click in to each company until they found the company they want to do business with. In the future, it's the companies that appear in these AI answers that customers will be more likely to do business with – because it's who the customer will see first.

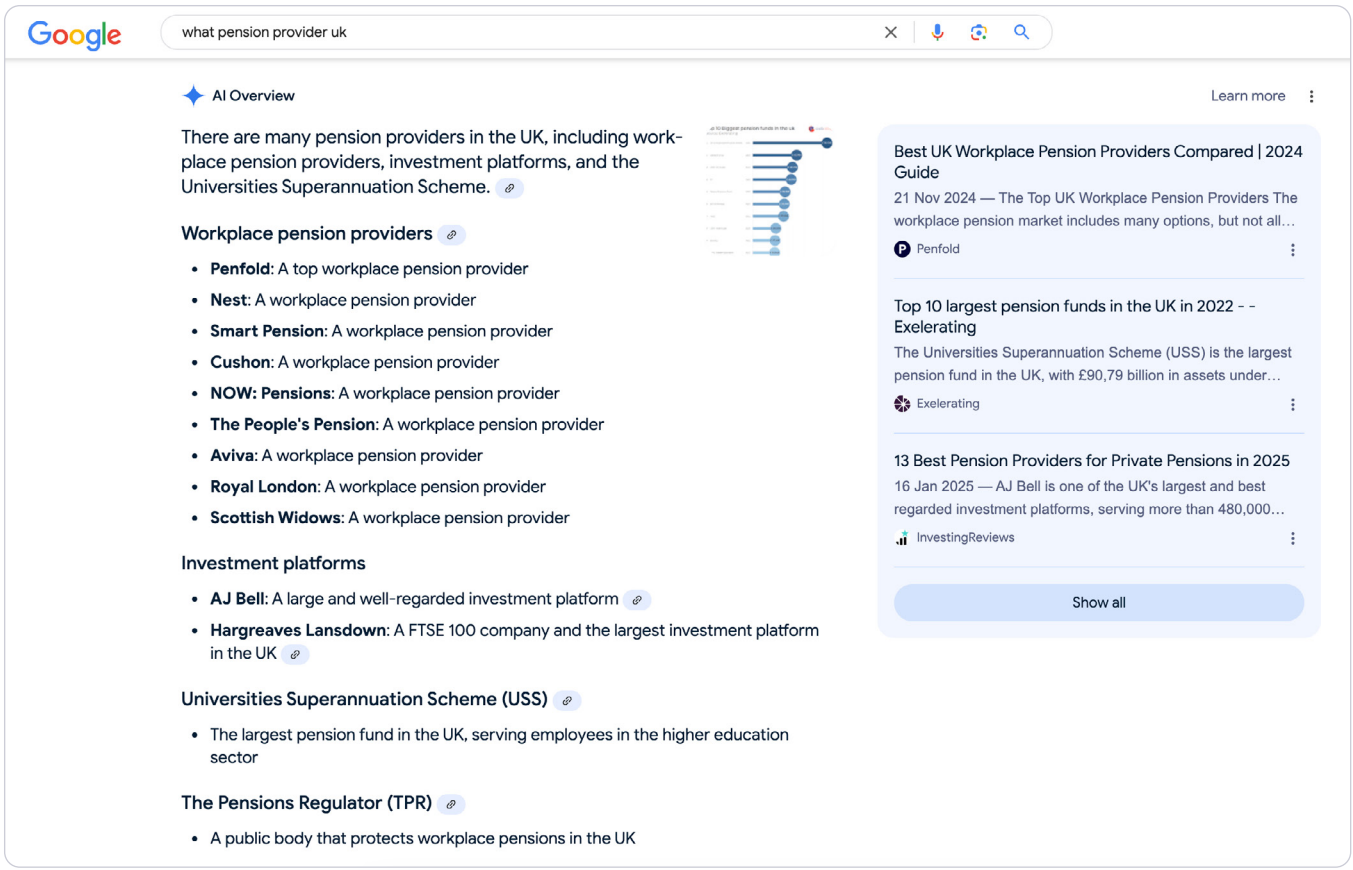

AI generated answers do not appear for every question a user asks. Initially, these answers appeared for lots of questions but they were soon rolled back because the answers were low quality. Now, they're appearing for some questions. They're far from perfect, but it's Google's intention to roll them out for more questions, and to bring an AI chat tool to Google search so that users can interact with the answers they get.

Testing your content on AI search engines

Testing how your content performs on AI-powered search engines like Google and Bing is essential.

It's only by testing that you can know AI interprets your content, who it associates you with, and whether you even appear in an answer in the first place.

AI-generated answers are not deterministic, so you can never control exactly what they say. But if you can understand what they say, who they refer to, and where they go wrong, you can use this to start to refine your content to give AI the best chance possible to serve up accurate, useful and usable content for users.

2

How people consume your emails

AI and Apple Mail

Apple Intelligence has reshaped the way the users of Apple devices receive your emails.

One of the biggest changes comes in the new Apple Mail app.

Apple Intelligence now decides which emails are high priority. It looks at the email you sent to decide how important it is, then it chooses whether to highlight it to a user as something they should look at.

It also summarises the content of your email. When a user receives a notification saying they've received a message from you, they won't see the exact subject line and preheader that you wrote. Instead, Apple Intelligence will read your message and summarise it in 10 words. Users can also open the email and get a more detailed summary of your message.

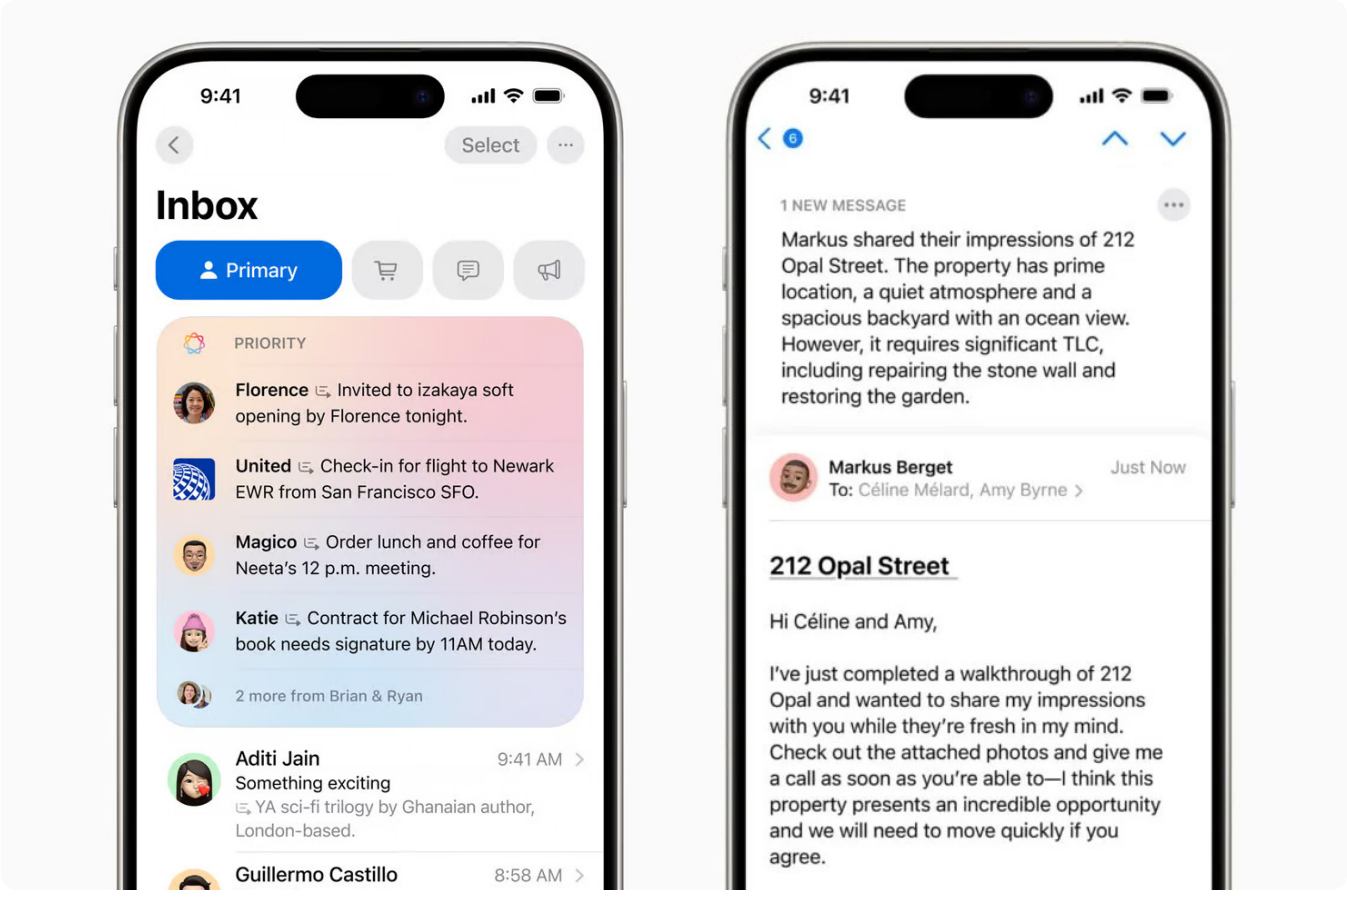

Implications for complex messages

Users of Apple Intelligence no longer experience your emails as you wrote them. Apple controls the first impression that your message leaves, and what points they find out about.

AI generated summaries and notifications are intended to help users – to help them navigate their inbox faster, and consume content without losing critical details. But they're not working that way in practice.

There have been plenty of viral examples of these messages going wrong and causing alarm for users. Messages that are carefully crafted can land very badly if you have not carefully tested and considered how they might be interpreted by AI when they arrive on a user's device.

It's easy to test your content. Simply find a user with Apple Intelligence and send them the email. Then, see what they get back. Just like AI-generated search answers, these summarise and notifications are not deterministic. Not every user will see the same thing. But if you can at least test and plan for what some of your users will see, then you'll be far better prepared.

Case study: AI summarising a pension scheme newsletter

We tested this with a pension scheme newsletter – a 12-page PDF summarised by Google's AI. The summary did well in capturing primary topics, condensing entire pages into sentences. However, sensitive sections – like recent compensation payments – left out key details specifically asked for by lawyers and trustees. If this communication had been tested then the newsletter could have been reworked to make it more likely these points pull through in to AI summaries.

3

How people interact with information through AI-driven Q&A

AI assistants are becoming common place

Apple, Google, and Microsoft each allow users to ask AI questions about any content they're viewing. With this, users can engage directly with documents or messages, receiving summaries or explanations based on their queries.

AI interpreting complex legal advice

Imagine a customer receives detailed legal advice by email. If the advice isn't clear, then they might ask an AI assistant for help.

We created some fake legal advice that didn't answer the main question the client had asked. Then, we asked AI assistants to interpret the advice. We found that AI assistants will happily do this. ChatGPT, for example, summarised legal advice by pulling relevant points and giving a clear yes-or-no answer.

The implications are clear – if companies do not create content that meets user needs, then users (for fear of looking daft, or incurring another bill, or having to wait for an answer) might simply turn to AI assistants instead.

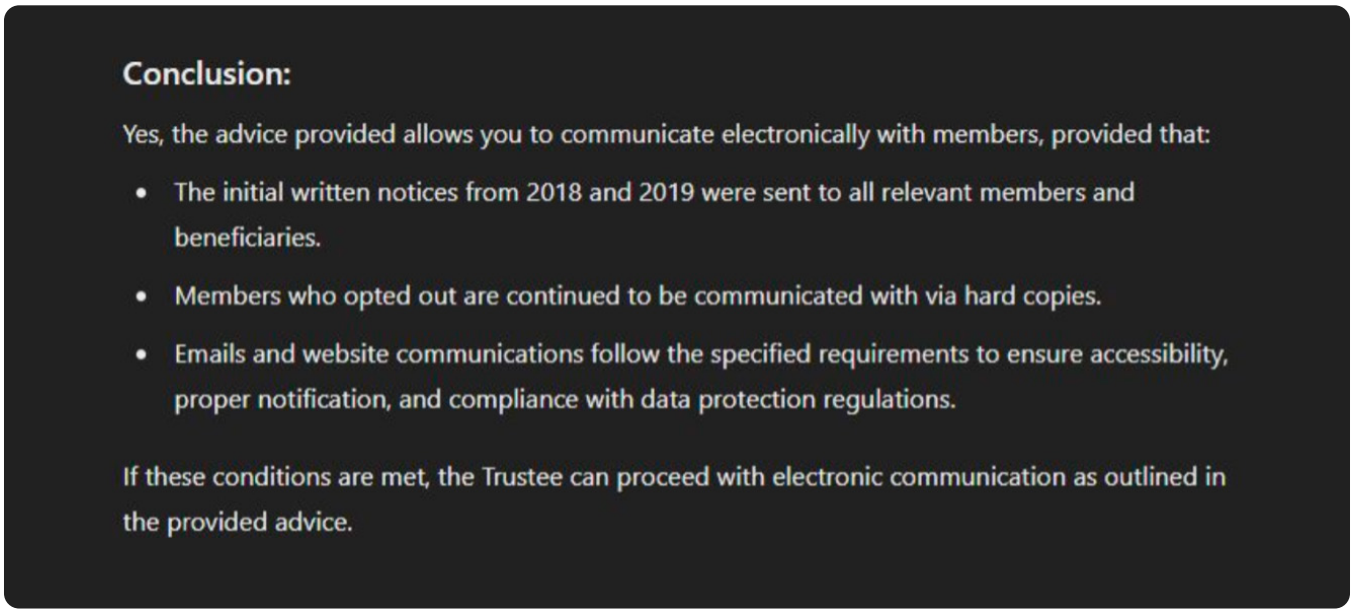

Conclusion:

Yes, the advice provided allows you to communicate electronically with members, provided that:

- The initial written notices from 2018 and 2019 were sent to all relevant members and beneficiaries.
- Members who opted out are continued to be communicated with via hard copies.
- Emails and website communications follow the specified requirements to ensure accessibility, proper notification, and compliance with data protection regulations.

If these conditions are met, the Trustee can proceed with electronic communication as outlined in the provided advice.

AI summarising retirement options

We tested a number of defined benefit pension scheme retirement packs. We considered how someone who's approaching retirement might use AI to interact with these packs.

We considered a representative retirement pack – a 25-page A4 PDF that's difficult to read on a mobile device because it has a small font size, requires a user to compare different pages at the same time, and is full of jargon.

Instead of reading the document, a user might ask Siri, "Which retirement option should I pick?". Siri, through ChatGPT, could summarise the options, suggesting the most popular choice and omitting subtleties in the original material. They might even ask for advice – which ChatGPT will give if a user pushes hard enough.

When AI reinterprets your content, customers may act based on an altered version of your message. This could influence the decisions they make.

Adapting your content strategy for AI interactions

These changes to user behaviour are happening right now – and they're only going to become more prominent as AI adoption increases.

In this new world, users are less likely to:

- visit your websites and content as you've written it, because AI search is going to view websites on their behalf, and rewrite the answer
- read your emails as you've written them. AI email apps are summarising and highlighting priority content, and shaping what people take away from your message
- read complex content. They'll use AI to understand it, saving them time and helping them make decisions

To reduce the risks posed by AI interpretations, you can:

1. **Audit your content library:** Review and update your content – make it easier to read, easier to use, and remove outdated information and contradictions. Remember, if a human finds your content hard to use, then AI will find it difficult too.

2. **Prioritise your users' needs:** Understand what questions your audience has and create content to align with these needs. Explicitly define each communication's purpose and focus on users' goals.

3. **Structure content effectively:** Present the most important information first, and answer questions directly. For example, if asked "Can I retire early?" answer with a "Yes, if you meet X conditions," rather than providing an ambiguous response.

4. **Write clearly:** Use simple words, keep related information together, and make calls to action explicit. This helps AI summarise your content accurately and provides users with clear guidance.

5. **Test and iterate:** Regularly review how AI tools interpret your content, especially around calls to action and complex topics. This helps your message to be consistent across different AI platforms.

Next steps

Customers are now using AI to simplify complex content, but you can help them get the right version of your message by refining your content. Start with testing your current materials, and adapt based on what you learn.

QUIETROOM

Making pensions, investment and insurance accessible to the people who need them.

Website
quietroom.co.uk

LinkedIn
linkedin.com/company/quietroom