

To: Competition and Markets Authority (“CMA”)
Re: Feedback on Publisher Conduct Requirements on Google’s general search services
Date: March 2026

We agree with the CMA’s conclusion that Google’s generative search features pose risks to publishers and to consumers. However we believe the risks are greater than those identified, and that the proposed Conduct Requirements (CRs) do not address either the risks outlined by the CMA, or the greater risks we see.

We have broken out our response into two sections:

1. The greater risks.
2. Detailed responses to the proposed CRs.

Section 1: Risks

1. Google’s generative AI search features are a risk to the UK economy and remove citizens’ choices

Google has created a great business providing the infrastructure that connects businesses that create information to those that need it. By being the highway rather than the destination - they charge a toll on every search journey, and this model has created one of the most successful companies in history.

That model is over. Google is now rebuilding themselves as the destination, not the highway.

Google is doing two damaging things by taking the information from all websites and providing the answer, rather than a search:

- Without traffic, there is a lower ROI in investing in new knowledge and information. If traffic to creators approaches zero, the ROI on new knowledge creation approaches zero. This directly damages all information-creation industries, and removes quality information from all industries that need it.
- Consumer choice is replaced by Google’s choice. In a search engine, people choose their sources. In a generative answer, such as an AI Overview, Google chooses the sources, and presents their interpretation of that information as fact¹.

¹ The EBU and BBC study, [News Integrity in AI Assistants](#), addresses the risks in detail. 45% of all generative answers have significant issues, and 81% of answers have issues of some form. The most significant issues are in sourcing, and in accuracy. So - when there are quality answers in the source material, answer engines fail to use those sources, misrepresent the sources they are using, and inaccurately quote or represent facts. Gemini was found to be the worst - for example having significant accuracy issues 20% of the time, or having issues with sourcing 72% of the time. The generative answers are worse than the quality of the information in the sources - so these generative features are worse for consumers than the underlying journalism. This also presents a risk to trust in the sources - the study found that users blame the source for any mistake 42% of the time, when the error was in the answer engine. Though this study is for news, the conclusions likely hold for all consumers seeking all information - resulting in a risk to all businesses that answer-engines disintermediate.

Everyone who creates information or attracts customers for that information is impacted by choices Google makes. This affects the entirety of the UK economy. Any impact that publishers see is an indication of the size of impact that can apply to the overall economy, and the costs and benefits of regulation should account for the larger impact. The CMA has rightly identified that 'publishers' should be read as 'anyone with a website'. But the cost/benefit calculations only account for the narrow definition of publishers - for example in section 5.34. When balancing the cost of remedies with the benefits - the calculations should account for the larger UK information economy.

Google's current trajectory is a direct risk to both the UK economy, and to an informed democratic electorate.

2. The risk is here today

We are one of the world's largest publishers. We see changes to search user experiences that extend beyond news, and our findings are complementary to data from news publishers.

Over the last 2 years, it has been widely reported and we can independently corroborate the loss of Google search-sourced traffic in the range of 50% to as much as 90%² (citation [study](#)). Over that same time, we see an increase in AI Overviews from covering 0% of relevant searches, to covering over 65%+ of relevant searches. When an AI Overview is shown to a user, clicks off of Google decline by an [average of 61%](#).

We do not expect the impact to end here - and are preparing for a world where Google replaces search with answers on all search queries. At that point, Google's AI Overviews become the default facts - removing user choice. This also directly benefits Google's business as the only way to consistently attract a search user when AI Overviews are present, is to pay Google for that user through Google Ads.

This trend changes the investment calculation for creating new content. As this trend continues, the ROI on investing in quality content approaches zero. This is a direct threat to an AI economic growth agenda, because high quality reference content is exactly what is needed to ground AI in real time, stop hallucinations, and result in reliable AI. By removing the ROI for knowledge creation, Google has guaranteed there will be less investment in the information needed for an AI growth economy.

The UK is one of the most advanced knowledge economies in the world. AI adoption should be a growth driver that connects businesses that need information to those that create it - funding new knowledge creation in the same way AI drives chips and energy investment. Google's disintermediation of the information supply breaks the information supply chain and puts the AI growth agenda at risk.

3. The worst case scenario is that Google Search defaults to AI Mode

² We track our search rankings in Google search. There was no material change in search rankings over this time period. Additionally - every other source of traffic grew. This traffic pattern is isolated to Google and correlated only to Google's product changes to search. This is not a change in user behaviour.

Google is on a path to fully replacing searches with answers:

- The AI Overview trend is on track to cover 100% of relevant searches within 2026.
- Google has rolled out AI Summaries in Discover.
- Google interrupts consumer searches to drive them to AI Mode.
- AI Mode takes the AI Overview one step further and does away with any search response page replacing it with a fully written answer. While those answers do have cited sources, clicks to cited sources are infrequent and we estimate occur [less than 1% of the time](#).
- Chrome is agentically enabled for Gemini Pro users, replacing user search paths through the browser.

Google can - at any point - fully replace search by defaulting to AI Mode. Every indication is that this is the direction Google is pursuing³. This would fully break the information supply chain, remove user choice on sources, and replace searches with Google's default answers from Google's chosen sources.

If this were to happen, Google becomes the destination and effectively privatizes the Internet as we know it. Google owns the user, the answer, and the monetization. This could play out across not only the information economy but also across commerce goods and services, health, travel and all other use cases formerly reserved for search based discovery on the Internet.

This would be disastrous to the UK knowledge economy, sovereignty over the information supply, and citizen choice. Nothing in the proposed Publisher Conduct Requirements guards against this future.

4. Blocking crawlers generates leverage - Google's SMS in search blocks leverage

Blocking crawlers from accessing content for AI use creates leverage. We cannot block Google's crawler as they have tied AI use to search. No other company poses the same risk, or has the same access. The only way to address the risk is to require Google to split crawling for AI use, from crawling for search discovery.

Blocking creates leverage and drives deals. In the summer of 2025 Cloudflare partnered with a large number of publishers to block AI crawlers from accessing content. Many deals have been signed since blocking became a real option. These deals are direct results of the leverage from blocking crawlers. No one has that leverage with Google as long as they tie access to search indexing with crawling for AI use.

5. Google's crawler exception is stopping a market from emerging

Grounding as a service is a growing market. Multiple companies are now selling grounding or AI search as a service. The pricing of these services [ranges from \\$1/1000 searches](#) to

³ Google is public about this ambition. [From September 2025](#) 'Logan Kilpatrick, lead product manager for Google, said [...] that Google's AI Mode will be the "default" search experience for Google Search "soon."'

[\\$30/1000 searches](#). These services operate by crawling without permission, and not providing any economics to the sources. In negotiations with web-crawler companies, they ask ‘why should we pay, when Google gets it for free?’

This market is valuable today, and has the potential to be even more meaningful in the future. AI systems regularly use 10-50 articles to answer a question where an individual would previously have read 1-3. As more people turn to AI systems to answer questions, there is a 10x larger search market being built. If that were to provide \$1-30/1000 visits to the IP rights holders, that creates a large market that would justify increased investment by knowledge creators and rights holders. If the economics are held by the crawler and not passed through, that’s a multi-billion dollar global market built on information the crawler does not own. A premium grounding industry would support a growth market for UK creators, and would result in higher quality AI outputs.

Google’s behavior is directly cited by companies as the reason they cannot pay for content. Google is a blocker to a functioning AI information supply chain that would power investment, innovation, and higher quality AI products.

6. The fine-tuning exception ties search and AI, defeating the aim of the proposed CRs.

The proposed CRs do not provide the same control or leverage as separating the crawlers.

In the current proposals:

- Google-extended controls allow publishers to opt content out of training Gemini models, and any grounding of non-search products.
 - Google claims this will not impact search results, while also acknowledging that the Gemini models *are* used in search ranking.
- Google retains the ability to use all content indexed in search to fine-tune the model.
 - Google argues that without being able to fine-tune the model on content, ranking would be harmed.
 - Google also argues that fine-tuning is required to minimise hallucination. Therefore fine-tuning does power generative search features, with no opt-out.
- New controls would be introduced to allow publishers to opt-out of grounding for AI search experience.
 - Google notes that they are unable to control for any second-order impact on search ranking if a publisher is not in AI products such as AI Overviews or AI Mode. This statement is clearly false - as Google has 25+ years of expertise in controlling for second-order impacts on search ranking.

The fine-tuning exception ties search performance to use in AI products - the exact problem the CRs are attempting to solve. This would be resolved by splitting the crawlers.

From Google’s responses regarding hallucination, it is clear that fine-tuning on the search index will result in publisher content appearing in generative AI features, negating any grounding opt-out. Hallucinations result when a model generates inaccurate information.

Reducing hallucination means answers are more accurate⁴. Therefore the fine-tuning must reproduce publisher information in answers. That means the current opt-out from grounding, with no opt-out from fine-tuning, will still result in publisher data being directly used in generative answers. If the fine-tuning exception is left in, there is no functional opt-out for publishers.

Google argues that fine-tuning is necessary to ensure the search results are high quality, and opting out of fine-tuning risks harming search rankings for a publisher. This repeats the problem of tying search rankings and generative experiences that these CRs were created to solve.

Google is making their engineering choice a publisher problem⁵. If they can enable opt-out of Google-extended without impacting search results, then they can allow an opt-out of fine-tuning without hurting search results. If Google cannot separate model training and generative results from search rankings, then they need to have two crawlers.

In the notes on fine-tuning (4.10(d)) "If Google were to try to maintain the hard distinction between generative and ranking models, it would likely need to develop and maintain duplicative models to achieve the same quality improvements". In our view, we agree. If Google cannot separate generative and ranking models, then the models - and the crawlers - need to be split.

7. The risks, and proposed solutions

- **Risk:** The risks and benefits apply to more industries than publishing - this impacts everyone who relies on search for their business
 - **Recommendation:** update the cost and benefit analysis for the wider UK information economy, rather than limiting analysis to digital publishing.
- **Risk:** Google is pursuing a path that replaces searches with answers
 - This removes user choice and makes Google the arbiter of sources and answers.
 - This breaks the information economy by disintermediating information creators and those that need information.
 - **Proposal:** preserve user choice by requiring a commitment that Google will not switch to AI Mode as the default - or any other experience that would fully replace searches with answers.
- **Risk:** Google is tying fine-tuning for AI experiences to search rankings, repeating the tying that the CRs are aimed at addressing.
 - **Proposal:** require that Google split the crawlers, and run different instances of their models for search ranking and for generative answer experiences.

Section 2: Detailed Responses to the Conduct Requirements

⁴ AI research labs compete to be the best at minimising hallucinations. With a fine-tuning exception, Google would have a unique advantage that others are paying for.

⁵ We have partnerships with a number of frontier LLMs. Training and grounding are clearly defined in each case - and highly valued. No-one other than Google has argued for special status for fine-tuning.

In this section we address the specific questions raised by the CMA about the proposed conduct requirements.

6.2 (a) What are the benefits and risks of Google providing separate controls over training and grounding outside of general search?

In all our partnerships with AI companies, training and grounding are treated separately in both rights and economics. We do not see any reason for Google to be an exception to what is now an established separation.

Training rights are tied to the model, and once content is provided, it is hard to untrain, or control how that information is used or attributed going forward.

Grounding is an ongoing, real-time process - with straightforward attribution, direct use of full content at the time of the consumer need, and rights and controls on verbatim use of that content in displayed answers. As we have discussed above, there are emerging grounding markets that sell this as a service separate to the model itself. [Google sells grounding as a service at \\$14/1000 searches](#). They are taking the web for free, and selling it for a premium.

The rights and economics of grounding and training are different and should not be tied in a single control.

This distinction is not unique to any single deal. Emerging global licensing standards - [e.g., RSL](#) - treat grounding and training as separate uses and rights. Controls should reflect the granularity of web-scale industry standards.

6.2 (b) What are the benefits and risks of Google providing page level controls outside of general search (i.e., Google extended)

Publishers generally have different rights to different pieces of content across their properties. For example: fully-owned photography from Getty Images carries different restrictions from photos only licensed to Getty, and a licensed article from the Associated Press has different restrictions to a piece written by an in-house writer. Google's blanket opt-in/opt-out at the site level ignores these distinctions, and makes it impossible to ensure that the content that is used is within the rights to be used.

In deals with Google's competitors we handle this by providing a licensable version of our content directly to our partners. We remove any articles - or elements within articles - for which we cannot provide the appropriate rights. This is uploaded to an S3 bucket. This approach is clean, simple, and requires no crawling. This is a modern way of solving these questions - rather than relying on re-purposing a public search web crawler.

We use per-article controls for our AI partnerships, and apply controls by page, and by element on the page. Google should not be granted an exception to these practical necessities.

6.3 (a) What are the benefits and risks of Google providing performance and engagement information on a 'per-feature' basis within general search?

Google's competitors are already providing reporting at the individual feature level. For example, we get feature level reporting from one partner, at the level of the individual URL cited with performance metrics that reflect the specific feature in which the content was used, such as inline citations, sidebar citations, TLDR responses, and search results. These reports quantify how often content was shown, how often it was clicked, and the CTR - by each generative feature.

We have had access to feature level data in Google search before - such as whether content appears in a Featured Snippet or a FAQ Block, or shows Rich Snippets. This level of visibility allowed us to invest into content and meta-data for those features. Without that visibility, the benefits, if any, of such investments are unknown.

More importantly - without feature level visibility, we cannot make decisions about opting in or out of Google's generative experiences, which removes transparency and therefore any informed publisher choice. Currently, 'search generative features' cover at least:

- AI Overviews in Search
- AI Mode
- AI Summaries in Discover

These are completely different user experiences - with totally different value calculations to a publisher. Without feature level reporting, Google is once again tying different experiences together and maintaining an information asymmetry. This only serves to make it harder for publishers to make any decision, which only serves Google.

As Google rolls out new generative search features, they will come with a cost to publishers. Without feature level reporting, that cost will be obscured within the whole, making it impossible to make any informed decision. Feature level reporting is necessary for the Conduct Requirements to be actionable for a publisher.

6.3 (b) Views on effective ways for Google to provide publishers with information to enable them to understand the quality of clicks from search generative features.

It is common practice to provide data parameters when a user moves from one site to another that allow publishers to understand where a session came from. These are often in the 'source' or 'medium' parameters, or in a 'utm parameter'. Google is expert in providing this level of detail, it is a key element of Google Analytics, an industry standard for understanding the quality of traffic from different sources. There are no such fields for traffic from different Google products, therefore we cannot segment and analyse these different traffic sources by feature. The lack of this data is a choice.

We fully reject Google's argument that generative search features send fewer but higher quality sessions.

As discussed previously, the emergence of AI Overviews has correlated with a drop by half in Google search sessions. We cannot break out AI Overview-sourced sessions from classic search-sourced sessions. But we can say with confidence that we have not seen a doubling of session duration, a doubling of ad clickthrough rate, or a doubling of e-commerce

conversions over this time period. Any change in the value of a session from AI overviews is too small to be visible, and certainly too small to compensate for the traffic loss from these changes. Recent studies support this, finding that LLM referrals convert [either no better than](#) or [possibly worse than](#) classic search referrals.

Even if we were able to identify AI Overview sessions as distinct from classic search, and we did see differences in behavior between the two, that would still not support Google's claim for three reasons:

First, it is clearly the case that lower engagement traffic no longer makes it to our sites. If someone simply needs an answer, they can now get that answer on the Google search page, rather than on our sites. We create the information in both cases, but Google gets the value and we lose the value. By removing these lower engagement (but still valuable) visitors from our sites, the remaining average mix of sessions is more engaged. This is not an increase, it's just a removal of the lower end of the distribution. If I have a pocket full of £1 coins, and £10 notes, and someone steals my £1 coins, I don't say 'thank you'.

Second, there's a selection bias that comes from barriers to action. If you place a high barrier in front of someone to do something, those that do make it over the barrier are usually very persistent. Therefore, there will be a correlation between Google making attribution links harder to find, and an increase in downstream engagement from those that overcome the barrier. This is not growth.

Third, any value they claim they are creating, comes from a user reading the content in the summary, which comes from publishers and was not created by Google. When a user comes to us from search, they have a question. They are a relatively low-value user. We help them answer their questions, and then they are ready to take action. Content is a funnel that increases the value of a user by educating, helping, and inspiring them. Google's claim that AI Overviews increase the value of a click is just a restatement of this fact. They have discovered the value of publishers, and they are claiming this value as their own - even though they are not a publisher, and don't create anything.

Google is taking publisher content, and behaving like they are the first to discover content has value. This speaks to a breathtaking level of hubris from a company that does not create information. We refuse to be paid with our own work.

6.4 (c) Views on the extent to which the proposed conduct requirement can be expected to result in the identified consumer benefits, including ensuring that users are able to assess and trust content they read on the web.

As we have mentioned above, we agree with the CMA's identification of the problem. However, we believe that the potential consumer harms are greater than the proposed solutions would suggest, and the current CRs do not go far enough.

However, If they are strengthened, there is an opportunity to:

- Ensure the UK information supply chain is robust - enabling increased investment in new knowledge, in turn powering better AI products that accelerate the AI economy.

- Protect against a future in which Google no longer provides search services - but instead creates the answer, and chooses the source of that answer, for all UK consumers for all questions.
- Require that generative answers are well sourced and accurate, and accurately reflect trustworthy sources.

As long as Google can tie generative features to their existing search engine - either directly or via fine-tuning - there is no way to ensure that they will not become the provider of all answers to all questions, breaking the UK information economy, and removing consumer choice.