Consultation: Publisher Conduct Requirement

Google's general search services

28 January 2026

The Competition and Markets Authority has excluded from this published version of the consultation document information which the CMA considers should be excluded having regard to the three considerations set out in section 244 of the Enterprise Act 2002 (specified information: considerations relevant to disclosure). The omissions are indicated by [✂]. Some numbers have been replaced by a range. These are shown in square brackets. Non-sensitive wording is also indicated in square brackets.

Contents

1. Introduction

1.1 This consultation sets out our proposals for a publisher conduct requirement (**Publisher CR**) which would provide publishers with informed and meaningful choice over how Google uses their content, crawled for general search, for the purpose of generative AI. That choice would be supported by improved transparency over how their content is used and engaged with, and measures to ensure effective attribution of content.

1.2 The draft of our proposed Publisher CR and related Interpretative Notes can be found in Section 3. For more information about the digital markets competition regime, Google's designation with SMS in general search services and the framework for considering conduct requirements, see the 'Introduction to the consultation' document published separately.

1.3 The document includes the following sections:

   (a) Section 2: Aim of our Publisher CR;

   (b) Section 3: Our proposed Publisher CR and Interpretative Notes;

   (c) Section 4: Effectiveness of our proposed Publisher CR;

   (d) Section 5: Provisional proportionality assessment for the Publisher CR; and

   (e) Section 6: Questions for consultation.

Issues we are seeking to address

Summary

1.4 Parties that make content available on the world wide web (**publishers**) make their web content discoverable to Google's Search crawler, enabling Google Search to present and link back to their websites in its general search results. Over the past two years, Google has introduced AI-generated responses to Google Search through features such as AI Overviews and AI Mode. It has also rolled out standalone generative AI services such as its Gemini AI assistant. These features and services were built using, and continue to rely on, content published on the web by individual creators and businesses and crawled by Google's Search crawler. Publishers are now faced with a decline in referrals back to their websites, and limited visibility as to how their content is being used in these novel systems.

1.5 The proposed CR aims to address three main issues arising from this trend:

(a) First, publishers currently do not have sufficient **choice** over how their content, gathered for search, is used by Google in its AI-generated responses. Given Google's SMS in general search services, publishers have no realistic option but to allow their content to be crawled. By not providing sufficient control over how this content is then used, Google can limit the ability of publishers to monetise their content, while accessing content for AI-generated responses in a way that its competitors cannot match.

(b) Second, publishers have limited **transparency** over how their content, gathered for search, is used by Google in AI-generated responses and how users engage with that content. This makes it harder for publishers to make informed decisions, including on whether to allow Google to use their content and for which purposes.

(c) Third, effective **attribution** of content in AI-generated responses is important for both consumers and publishers. For consumers, attribution can allow them to test the veracity of AI-generated content. For publishers, attribution can help ensure that consumers are aware of the sources of content, which in turn can allow them to sustain brand value and the creation of new material.

1.6 These issues matter because Google's general search services are a key gateway through which people access and navigate the world wide web, and businesses and content creators can reach consumers. As such, Google Search is important for virtually all content creators (as explained above we use the term 'publisher' broadly in this consultation to refer to all parties that make content available on the world wide web).[1]

1.7 In addition, general search services have become important to people as citizens, not least as a key route through which they access news. A well-functioning market would ensure that people can access a wide range of high-quality, accurate content which is attributed effectively, and that publishers are treated fairly where their content is used. In turn this can help support the long-term sustainability of content such as news, with wider benefits to society and democracy.

---

[1] Our analysis of a set of news and press publishers' data shows referrals from Google Search (including Google Discover) constitute 37% of online referrals, larger than all other third-party sources in combination. Source: CMA analysis of 13 responses to CMA's RFIs from: [✂].

The issues arising from Google's use of publisher content in generative AI

1.8 Our 2025 SMS Decision found that Google has SMS in the provision of general search services, due to its substantial and entrenched market power and position of strategic significance in that digital activity.[2] We explained in our SMS Decision that the Google products within scope of Google's general search services include its generative AI features such as AI Overviews and AI Mode; but that Google's Gemini AI assistant and Vertex AI product are not within scope of the designation.[3] We distinguish in the remainder of this document, where appropriate, between Google's:

(a) **search generative AI features**: generative AI-dependent features offered <u>within</u> Google's general search, such as AI Overviews and AI Mode; and

(b) **broader generative AI services**: generative AI-dependent products and services offering information retrieval capabilities <u>outside</u> of Google's general search, such as Gemini AI assistant and the Vertex AI API.

1.9 We refer to Google's search generative AI features and broader generative AI services together as Google's **generative AI services and features**.

1.10 Google obtains content through its main web crawler, Googlebot, for many purposes. This includes display of organic general search listings, search generative AI features, broader generative AI services, as well as the development of AI models powering these services. The main ways publishers can currently prevent Google from using content for generative AI purposes are through the following controls:[4]

(a) The 'noindex' control, which allows publishers to prevent their content entering Google's index,[5] meaning it won't appear or be linked to in either Google's organic listings or its search generative AI features.

---

[2] Strategic Market Status investigation into Google's general search services: Final Decision (SMS Decision), 10 October 2025.
[3] Strategic Market Status investigation into Google's general search services: Final Decision (SMS Decision), 10 October 2025, paragraphs 4.9-4.10 and footnote 205.
[4] Were Google to offer separate crawlers for these purposes, publishers could withhold consent for certain uses of their content at the point of crawling.
[5] Google currently offers publishers related functionality through Google Search Console to temporarily block a page from appearing in Google Search results, functioning in a similar way to 'noindex' (at least for general search services). This control can be used where a publisher wants to block Search Content from appearing in Google Search results faster than would be true were the publisher to update its publisher content controls to

(b) The 'nosnippet' control, which allows publishers to remain linked in Google's organic listings, whilst opting out of summaries of their content ('snippets') appearing as part of organic listings and in search generative AI features.

(c) The 'Google-Extended' control, which allows publishers to opt out content being used to train future generations of Gemini models powering Gemini Apps and Vertex AI API for Gemini, and for grounding in wider applications such as the Gemini AI assistant. Google states that use of this control will not be used as a ranking signal in Google Search.[6]

1.11 Based on the evidence we have seen to date, we consider that Google's existing controls do not provide publishers with sufficient **choice** over how their content, provided for general search (**Search Content**), is used by Google. In particular:

(a) Whilst publishers can use the nosnippet control to opt content out of being used as a direct input for AI Overviews and AI Mode,[7] use of this control is likely to affect their organic search ranking and reduce traffic because publishers lose the descriptive text that helps users decide whether to click onto the website. Evidence from a study produced by Google showed that removing snippets reduced traffic by nearly half.[8] Google internal data also shows [✂].[9] Given the role of Google Search as a key gateway for content discovery, such a reduction in referrals could have a substantial impact on publishers' ability to monetise their content. As such existing controls do not give publishers effective choice over the use of their content in search generative AI features.[10]

---

specify 'noindex' for example, which we understand would only take effect once Google had re-crawled their site. See Google, 'Removals and SafeSearch reports Tool: Temporarily block search results from your site, or manage SafeSearch filtering,' accessed by the CMA on 6 January 2026.

[6] The control is less granular than nosnippet and noindex, as it applies when the robots.txt file is read rather than being implemented through meta-tags which allow page or even sub-page level control. Source: Google, 'List of Google's common crawlers,' accessed by the CMA on 3 December 2025.

[7] Google said that 'content on web pages marked with nosnippet will not appear in AI Overviews' and that Search 'will not use content on web pages marked with nosnippet for creating AI Overviews or display alongside AI Overviews.' Google's response to the CMA's RFI.

[8] Google's research shows that reducing detail on publisher content shown (broadly in line with the display where publishers use nosnippet) reduces traffic to publishers by 45%. This research was published before gen AI features were included in Search. Source: archive.org capture of Google's blog titled 'Now is the time to fix the EU copyright directive', 7 February 2019, accessed by the CMA on 3 December 2025.

[9] Google's consolidated response to the CMA's RFI.

[10] For example, [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

(b) There is a lack of clarity over the scope of 'Google-Extended', which means that publishers are uncertain about how it operates in practice.[11] While Google has said that Google-Extended covers the use of publisher content for training Gemini models and grounding related applications, it is not clear how or whether it applies to other AI-generated content outside of search.[12] This ambiguity prevents publishers from making an informed choice as to how their content is used in broader generative AI services.

1.12 We have also seen evidence that Google provides insufficient **transparency** to enable publishers to exercise informed choice in relation to its controls.[13] First, there is a lack of transparency over the scope of Google's existing controls, including in relation to broader generative AI services, as noted above. Second, some publishers explained that their decision making in relation to these controls (as well as their content production and strategy more generally) is inhibited by the lack of visibility of user engagement with their content when displayed in search generative AI features.[14] We identified various specific limits to the information Google provides, as set out further in below.[15]

1.13 Finally, we have also seen evidence that prominent and accurate **attribution** of publisher content where it is used in search generative AI features can have a significant impact on publisher outcomes including click-through rates. Effective attribution is also important to enable users to verify the AI-generated responses that are provided, by clicking through to the underlying content if they wish.[16]

---

[11] For example, [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[12] Google's internal documents show that in updating the Google-Extended description [✂].

[13] Several publishers suggested they lack clarity on the impact of making any choice in relation to crawler access. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[14] [✂] publishers mentioned benefits to content production and strategy from receiving more information on how Google makes use of publisher content. For example [✂] internal document [✂] and [✂] response to the CMA's RFI.

[15] For example, Google does not tell publishers how often their content appears in AI Overviews. Scraping, a method which could allow publishers to verify this themselves, is not compliant with Google's terms of service. See Google Terms of Service – Privacy & Terms – Google. Google confirmed that scraping is not compliant with its terms of service, for reasons including that it increases server load. Google's submission to the CMA.

[16] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

1.14   We recognise that Google has made, and is continuing to make, concrete steps to improve attribution[17, 18] and has seen materially higher model accuracy in at least some settings.[19] We also recognise that the approach to attribution needs to balance user and publisher interests, and that there are inherent technical difficulties in ensuring uniform accuracy.[20] However, publishers' inability to withdraw content contributes to insufficient incentives for Google to assure its attribution of publishers in search generative AI features remains accurate and sufficiently prominent.[21] Further, given its SMS status in general search services, there is some risk that publisher choice and the need to retain users provides too little pressure on Google to ensure sufficient attribution in general search services in future.

1.15   Taken together, these issues mean that publishers are unable to exercise sufficient choice over how their Search Content is used by Google across its generative AI services and features. These issues arise in a context in which Google's use of Search Content in generative AI services and features is having a significant impact on traffic to publishers' websites, and therefore their ability to monetise their content. Evidence from an internal Google experiment indicates that, for 'in slice' data (ie queries which would have triggered an AI Overview), clicks to publisher sites decrease by approximately [✂]% when an AI Overview is shown on a search results page, relative to if the AI Overview was not shown.[22] As set out further in our proportionality assessment, we consider that this inability to exercise sufficient choice can

---

[17] See for example: Google, 'AI Overviews: About last week', 30 May 2024, accessed by the CMA on 3 December 2025. Google, 'Supporting the web with new features and partnerships', 10 December 2025, accessed by the CMA on 11 December 2025.

[18] A Google internal document on corroboration in AI Overviews and AI Mode indicates that [✂].

[19] Google internal documents from July 2024 report that factuality levels for queries related to [✂] provided in response to the CMA's RFI.

[20] The probabilistic nature of the models means that it would be likely impossible for anybody to develop models which never return inaccurate responses. Google also told us that it continues to experiment with prominence, but that it must make trade-offs between the needs of users and publishers. See Google's response to the CMA's RFI.

[21] For example, one publisher told us that being able separately to identify and opt out of grounding and training of search generative AI features would put pressure on Google to improve attribution, referral of traffic and the factuality of results in its AI interfaces, to encourage publishers to opt in. See [✂] response to the CMA's RFI.

[22] Evidence from an internal Google experiment indicates that organic clicks decrease by approximately [✂] % for the slice of traffic where an AI Overview is shown on a search results page. Google's experiment is a randomised control trial [✂]. This experiment does not take into account [✂]. The overall fall in clicks is [✂], this would in part appear to be because AI Overviews only show on a minority of queries. Google's internal document. Google's response to the CMA's RFI.

lead to the following adverse impacts on web publishers and end-consumers of web content:[23]

(a) Publishers cannot exercise sufficient choice in relation to Google's controls, including opting out where AI-generated responses do not benefit the publisher;[24]

    (i) Where publishers' Search Content is used in search generative AI features, this can affect consumers' engagement and the extent to which they click through to the underlying website – and so affect publishers' ability to monetise their content;

    (ii) Where publishers' Search Content is used in broader generative AI services, this can also affect publishers' ability to monetise their content and allow Google to benefit from its position in search when competing in markets for generative AI-dependent products and services outside search;

(b) Publishers are constrained in their ability to optimise their content or monetisation strategies effectively, for example in how they allocate editorial resources to produce content which best supports click-through rates;[25]

(c) Publishers' bargaining positions are likely undermined because Google is using their Search Content in new ways without providing additional information or controls on how it does so, inhibiting publishers' ability to maintain and improve their compensation for such content; and

(d) Google's incentives to maintain and improve the factuality and prominence of attribution of publishers in search generative AI features

---

[23] For example, several publishers suggested ineffective choice in controls inhibits commercial or financial benefits they receive from providing content to Google. We expect this would have at least some knock-on effect for end-consumers. See [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[24] When AI Overviews are shown, Google's data shows [✂] referrals are from organic links under the AI Overview, rather than from links within the AI Overview itself. Publishers told us that being able to opt their content out of AI Overviews would help support their ability to invest in producing content, as set out in paragraph 5.31 below. Source: CMA Analysis of data submitted in [✂] Google's response to the CMA's RFI.

[25] This involves Google's use of content to develop features which appear to reduce click-through rates for publishers, reducing the benefit they receive for contributing relative to functions within traditional organic listings such as index and ranking. Publishers are unable to identify how far their content is used for these new purposes, or exercise controls to address this change in value proposition.

are likely reduced,[26] which risks damage to publishers' brands and may affect end users by eroding trust and transparency.

(e) Insufficient attribution can also undermine consumers' ability to check and verify the underlying sources of AI-generated responses – something that we expect to become increasingly important over time.

1.16 Collectively, these limitations can restrict publishers' abilities to invest in new high-quality content, which leads to detriment for end users.[27] This could be particularly significant for certain sectors, such as news.[28] End users will also be harmed were attribution to be insufficiently prominent or accurate, since this would inhibit trust and transparency in AI-generated responses.

---

[26] For example, one publisher said there is merit in being able to opt out of these features in a scenario in which Google designs those services in a way in which there is limited attribution, or the lack of a return path back to the source website on which the article was first published. Financial Times's response to the CMA's RFI.

[27] For example, one publisher's internal document projects that a hypothetical loss of all organic traffic to their site would cost the organisation approximately £5 million in advertising and subscription revenue per year. [✂] internal document [✂].

[28] A Mediatique report for the Cairncross review identified that declining revenues has in the past resulted, amongst other consequences, in a reduction in frontline journalism, consistent with other reports cited in paragraph 5.36. We set out benefits of news consumption in paragraphs 5.36 to 5.41. Source: Mediatique for Department for Digital, Culture, Media & Sport, Overview of recent dynamics in the UK press market, April 2018, accessed by the CMA on 23 January 2025.

2. The aim of our Publisher CR

2.1 Taking into account the nature of the concerns set out in Section 1 above, we propose to introduce a CR to ensure publishers can make properly informed and meaningful choices over how their Search Content is used in relation to Google's generative AI services and features. To make a properly informed and meaningful choice, publishers require sufficient:

(a) controls that enable them to withhold the use of their content for Google's generative AI services and features (including on the SERP) without affecting their position or display in Google's general search in a way that reduces their search traffic;

(b) transparency to make an informed choice, including regarding:

(i) how Google uses Search Content and the effect those controls will have on such use, so that they can be well-understood, including in terms of the scope of the generative AI services and features to which they apply; and

(ii) how users are engaging with content in Google's search generative AI features, so that they can understand the potential benefits of permitting Google to use their Search Content in those features; and

(c) trust that Google will provide sufficient attribution of Search Content within Google's search generative AI features, so that they can understand the potential risk in how Search Content is used or attributed in those features.

2.2 **The aim of the Publisher CR would therefore be to ensure that publishers have sufficient:**

(a) **controls over Google's use of their Search Content in its generative AI services and features;**

(b) **transparency over Google's use of their Search Content in its generative AI services and features and user engagement with their Search Content in search generative AI features; and**

(c) **trust that where Google uses Search Content in its search generative AI features, that content is sufficiently and accurately attributed,**

to enable them to make properly informed and meaningful decisions about whether and how they interact with Google in respect of general search services.

2.3   Overall, this would enable publishers to make more informed decisions about how their Search Content is used and help address the bargaining power that Google has over publishers because of its SMS in general search. We recognise that this will not address all the issues currently facing publishers. However, we consider that it would be an important precondition for addressing broader concerns about publishers' ability to negotiate reasonable terms with Google for use of their Search Content.

Statutory objective(s)

2.4   As explained in the 'Introduction to the consultation' document, the Act provides that CRs must seek to achieve one or more of three statutory objectives.[29]

2.5   The proposed Publisher CR would pursue:

   (a) The fair dealing objective (section 19(6) of the Act): that publishers (as users of Google's general search services[30]) are treated fairly and able to interact, whether directly or indirectly, with Google on reasonable terms; and

   (b) The trust and transparency objective (section 19(8) of the Act): that publishers have the information they require to understand Google's general search services and the terms on which they are provided, and make properly informed decisions about whether and how they interact with Google in respect of general search services.

Permitted type(s)

2.6   As explained in the 'Introduction to the consultation' document, each CR must fall within an exhaustive list of 'permitted types' set out in the Act.[31]

---

[29] Section 19(5) of the Act.
[30] 'User' includes any person, legal or natural, and, in relation to a digital activity, means any user of the relevant service or digital content (s118(1)); while 'using' includes, 'in relation to a service or digital content, interacting, or carrying out activities that interact, in any way, directly or indirectly, with the service or digital content' (s118(2)(b)). See further the explanatory notes to the Act, paragraph 533(f).
[31] Sections 19(9) and 20 of the Act.

2.7   The components of the proposed Publisher CR would fall under the following
       permitted types:

   (a) Section 20(3)(d): preventing Google from requiring or incentivising
       publishers to use its generative AI services and features alongside its
       traditional search engine (Google Search) – see paragraphs 2 and 3 of
       the draft CR;[32]

   (b) Section 20(2)(c): obliging Google to provide clear, relevant, accurate and
       accessible information to publishers about general search services: in
       particular, how their content, collected for the purpose of general search,
       is used, attributed and engaged with – see paragraphs 4, 5 and 6 of the
       draft CR; and

   (c) Section 20(2)(e): obliging Google to present to publishers any options or
       default settings in relation to general search services in a way that allows
       publishers to make informed and effective decisions in their own best
       interests about those options or settings – see paragraph 4 of the draft
       CR.

Consumer benefits likely to result from this CR

2.8   Before imposing a CR, the CMA must have regard in particular to the benefits
       for consumers that it considers would likely result from the CR.[33]

2.9   We consider that our proposals would represent a positive step towards
       improving publishers and consumer outcomes. In particular, we expect the
       proposals would lead to consumer benefits in several ways, including:

   (a) By supporting publishers facing challenges created by the inclusion of
       AI-generated responses in general search services, consumers would
       be likely to benefit from improved quality and availability of web content.

   (b) By ensuring Google meaningfully attributes content in its AI-generated
       responses, we expect the Publisher CR would contribute to consumers'
       ability to assess and trust content they read on the web.

---

[32] The explanatory notes to the Act (paragraph 191) note that 'This permitted type also encompasses a situation
where a designated undertaking requires or incentivises the use of one part of the designated activity (e.g. a
payment service) alongside another (e.g. an app store)'.
[33] Section 19(10) of the Act.

(c) By facilitating a more level playing field between Google and its competitors, we expect the CR would support improved generative AI service offerings which benefit consumers.

2.10 We provide greater detail of how these consumer benefits would arise (together with benefits to businesses such as publishers) in paragraphs 5.27 to 5.48 below.

3. Our proposed Publisher CR and Interpretative Notes

Publisher CR

3.1 Having identified our aim (see paragraph 2.2) based on the concerns identified in Section 1, we are proposing to impose the following draft Publisher CR on the basis of the effectiveness and proportionality analysis set out in Sections 4 and 5 below respectively.

Definitions

1. For the purposes of this conduct requirement:
    a. **broader generative AI services** means Google's generative AI-dependent products and services offering information retrieval capabilities outside of general search, such as Gemini AI Assistant and the Vertex AI API.
    b. **general search** has the meaning given to it in the SMS Decision Notice dated 10 October 2025, as revised from time to time.
    c. **general search services** has the meaning given to it in the SMS Decision Notice dated 10 October 2025, as revised from time to time.
    d. **generative AI services and features** means Google's search generative AI features and broader generative AI services, together.
    e. **publisher** means any party that makes content available on the web to any natural or legal persons located in the UK using Google's general search services.
    f. **Search Content** means publisher content collected through Googlebot or any Google crawler fulfilling the function of crawling open web content for Google's general search.
    g. **search generative AI features** means Google's generative AI-dependent features offered within general search, such as AI Overviews and AI Mode.

Controls

2. Google shall provide publishers with effective controls to withhold their Search Content from being used in:

    a. the training and grounding of its broader generative AI services; and

    b. the grounding of its search generative AI features.

3. In relation to the controls described in paragraph 2, Google:

---

a. shall ensure that those controls evolve in an appropriate way as generative AI services and features evolve;

b. shall not:

   i. maintain or introduce ranking signals whose purpose is to downrank opted-out Search Content in general search outside of search generative AI features; or

   ii. cause Search Content to be presented or displayed differently in general search outside of search generative AI features depending upon whether it is opted in or out through those controls; and

c. shall not attempt to circumvent any publisher's choice to withhold its Search Content by acquiring that content through other sources.

Transparency

4. Google shall:

   a. publish clear and detailed information explaining how Search Content is used for the training and grounding of its generative AI services and features; and

   b. ensure that the effect and scope of the controls described in paragraph 2 is transparent and can be well-understood by publishers.

5. Google shall provide publishers with clear and detailed metrics on user engagement with their Search Content where it is used in its search generative AI features.

Attribution

6. Google shall:

   a. take reasonable steps to ensure that Search Content is sufficiently attributed when used in search generative AI features; and

   b. publish clear and detailed information explaining the steps it takes to:

      (i) ensure that Search Content is sufficiently attributed in search generative AI features; and

      (ii) ensure and measure the factuality of search generative AI features.

Interpretative Notes

3.2 The CMA may publish interpretative notes to accompany a CR. Interpretative notes will provide greater clarity over the CMA's interpretation of a CR, including how it may apply in particular circumstances, for the benefit of both the SMS firm and other industry participants.[34] It would be open to the SMS firm to take a different approach to the one outlined in the interpretative notes where it is able to demonstrate to the CMA that its approach complies with the terms of the CR.[35]

3.3 We propose that the Publisher CR be accompanied by the following set of interpretative notes.

Effective controls

1. Under the requirement set out in paragraphs 2 and 3 of the conduct requirement, in order to provide effective controls we expect Google to:

   (a) In relation to the training and grounding of broader generative AI services:

   (i) make any changes necessary to the existing Google-Extended control in order for it to enable publishers to opt their Search Content out of the training of generative AI models and grounding of broader generative AI services; and

   (ii) publish clear and detailed information describing the scope of the Google-Extended control. This should include a description of Google-Extended and its overall purpose, the key exceptions and limitations to its scope and an explanation of the publisher content Google-Extended covers.

   (b) In relation to the grounding of search generative AI features:

   (i) introduce a new control to enable publishers to opt out of the use of their Search Content at both directory-level and page-level; and

   (c) not take any actions or omissions that would frustrate the effectiveness of the controls.

---

[34] See Digital Markets Competition Regime Guidance (CMA194), paragraphs 3.59 to 3.60.
[35] See Digital Markets Competition Regime Guidance (CMA194), paragraph 3.61.

2. Paragraph 3.a. of the conduct requirement requires that the controls evolve in an appropriate way. The CMA expects Google to keep the scope and operation of the controls under active review over time and to extend them to new generative AI services and features as they are released. The CMA also expects Google to adopt a product-agnostic definition such that the controls scope in services which produce a substantial amount of generative output, grounded on or otherwise using publisher content, in response to user queries.

3. Pursuant to paragraph 3.b.i. of the conduct requirement, Google should not intentionally rank publishers lower or remove them from organic search results based on their use of the controls; however, the CMA recognises that Google may not be able to control for all possible second order effects of the controls on Google's overall ranking algorithm.

4. Paragraph 3.b.ii. of the conduct requirement encompasses any change to the presentation of publishers' content in general search that could reduce the rate at which users click through to their content. Opting out of their Search Content being included in any of Google's generative AI services and features should therefore not lead to a publisher losing features such as snippets.

5. Paragraph 3.c. of the conduct requirement requires Google not to actively attempt to circumvent a publisher's choice. The CMA expects Google not to, for example, pay a third party to crawl the website of a publisher that has opted out of its Search Content being used by Google through these controls. However, the CMA considers that it would be reasonable for Google to acquire such content through open-source datasets, where these datasets have obtained content legally, given the nature of such sources.

Transparency over Google's use of publisher content and publisher controls

6. The information Google is required to provide pursuant to paragraph 4.a. of the conduct requirement should be sufficiently clear and detailed to assist publishers in deciding whether to give consent to Google for the use of their Search Content for these purposes, including through the controls Google is required to provide under paragraph 2 of the conduct requirement.

7. Paragraph 4.b. of the conduct requirement requires that the effect and scope of the publisher controls is transparent and can be well-understood. The CMA expects Google to make available to publishers clear and comprehensive information about the controls. For example, guidance could be posted on Google Search Central. This guidance should list those

18

Google products, services and other activities to which each control relates. Google should invite representations on whether the information provided is sufficiently clear and take action to respond to any issues raised by publishers.

Transparency over user engagement

8. The information Google is required to provide under paragraph 5 of the conduct requirement should enable publishers to take informed decisions about:

    (a) whether to give consent to Google for the use of their Search Content in Google's search generative AI features, including through the controls Google is required to provide under paragraph 2.

    (b) their Search Content in relation to its use and discoverability in Google's search generative AI features.

9. The CMA expects the data Google is required to provide under paragraph 5 to include the following:

    (a) User impressions – namely data on the display of a publisher's Search Content as part of a Google search generative AI feature in response to a user query. This should include data on where the publisher's Search Content is attributed in a search generative AI feature returned in response to a user query.

    (b) User engagement with the publisher's Search Content – including click-throughs to the publisher's content from links in search generative AI features and data to enable publishers to assess the 'quality' of those clicks.

    (c) Click-through-rate – providing publishers with the percentage of users who click on a link to that publisher within a Google search generative AI feature ('clicks') where that publisher's Search Content has been displayed within a search generative AI feature ('impressions').

10. The data should be provided:

    (a) on a disaggregated basis for each publisher;

    (b) fully disaggregated from other elements of general search, such as organic search results (including web results, images and videos) and search features other than other search generative AI features; and

(c) through a commonly accessible platform, such as Google Search Console.

Trust in Google's attribution of Search Content

11. In order for attribution to be sufficient, Google should take reasonable steps to ensure that:

   (a) attribution is accurate; and

   (b) where a publisher's Search Content makes a significant contribution to a search generative AI feature response, end users will be provided with a clear means to access that Search Content.

12. In assessing whether Google has complied with the requirement set out in paragraph 6.a. of the conduct requirement, the CMA expects to take into account the need for attribution to also reflect broader end-user experience, design aesthetics and source diversity.

13. We expect Google's obligation to publish an explanation of its approach to attribution, pursuant to paragraph 6.b.i. of the conduct requirement, to include information about:

   (a) how Google identifies what content to attribute in its search generative AI features;

   (b) the steps it takes to monitor the accuracy of its attributions and what steps are available for publishers to identify content that has not been attributed accurately; and

   (c) how it seeks to provide end users with access to relevant Search Content in order to test the veracity of the result, including how it ensures that such Search Content is sufficiently prominent for end users to identify and access it.

14. Given the way generative AI responses are created, we recognise that search generative AI features cannot always be completely accurate; however, pursuant to paragraph 6.b.ii. of the conduct requirement, we expect Google to outline the steps it takes to improve the factuality of search generative AI features, how it deals with factual inaccuracies when they arise, and metrics on factuality that would help publishers to take a decision on whether to give consent for Google to use their Search Content in search generative AI features.

Publication of information for publishers

15.   In relation to the information Google is required to publish pursuant to paragraphs 4 and 6.b. of the conduct requirement, Google should:

    (a) Provide the information in an accessible location and format. For example, it could be incorporated in Google's information pages (eg What Is Googlebot | Google Search Central | Documentation | Google for Developers);

    (b) Publicise the information, and direct publishers towards it when queries are raised about the use of Search Content;

    (c) Keep the information up to date as Google's product offerings and use of Search Content evolves and provide clarity over the scope by reference to the product-agnostic definition as set out in paragraph 2 above.

    (d) Invite feedback from publishers on areas where the information provided lacks clarity or is incomplete and act on such feedback.

4. Effectiveness of our proposed Publisher CR

4.1 Having identified an aim (see paragraph 2.2 above), the CMA must identify a CR, or combination of CRs, that would likely be effective in achieving this aim. As part of this, the CMA will consider both the content and form of potential CRs.[36]

4.2 This section sets out the analysis we have undertaken to identify the most effective design of our proposed Publisher CR and inform the proposed interpretative notes. It covers:

    (a) Our proposed design choices that resulted in the draft Publisher CR set out in Section 3 above;

    (b) Implementation and compliance; and

    (c) CRs that would be equally effective.

Key design issues we have considered

4.3 To inform and test the shape of the Publisher CR we have identified a series of key design issues to ensure it effectively meets our aim and mitigates potential risks, including unintended consequences. These include how the Publisher CR would enable publishers to make properly informed and meaningful choices about whether and how they interact with Google in respect of general search services by providing publishers with:

    (a) sufficient controls over Google's use of their Search Content in its generative AI services and features;

    (b) sufficient transparency over Google's use of their Search Content in its generative AI services and features and user engagement with their Search Content in search generative AI features; and

    (c) sufficient trust that where Google uses Search Content in its search generative AI features, that content is sufficiently and accurately attributed.

---

[36] See Digital Markets Competition Regime Guidance (CMA194), paragraph 3.20(b).

***How the CR ensures that Google will provide sufficient controls over Google's
use of publishers' Search Content in its generative AI services and features***

4.4 In order to ensure that publishers can make properly informed and meaningful
choices about whether and how they interact with Google in respect of
general search services, the first limb of our aim involves providing publishers
with sufficient controls over Google's use of their Search Content in its
generative AI services and features.

4.5 In designing the choice elements of the Publisher CR (and the underlying
Interpretative Notes), the CMA has had regard to the following key design
parameters relevant to their effectiveness, which are considered further
below:

  (a) the use cases covered by the controls;

  (b) the scope of the controls;

  (c) the granularity of the controls; and

  (d) supplemental requirements needed to ensure the effectiveness of the
  controls.

*The use cases covered by the controls should allow publishers to opt Search
Content out of being used for grounding of search generative AI features and
grounding and training of broader generative AI services*

4.6 As explained above, publishers have no realistic option but to allow their
content to be crawled for Google's general search because of the market
power Google holds in general search. However, Google currently uses that
content in both its search generative AI features and in its broader generative
AI services.

4.7 Accordingly, in order to be able to make meaningful decisions about how
Google uses their Search Content, we consider that publishers need the
ability effectively to opt their Search Content out of both Google's search
generative AI features and Google's broader generative AI services.

4.8 Publishers' Search Content may be used at several points in the development
and operation of Google's generative AI services and features:

  (a) In training the underlying foundation models (eg the Gemini foundation
  model family) from which Google derives the wider suite of models used
  within and outside general search.

(b) In fine-tuning the derivative models used in general search.[37]

(c) In grounding the generated responses produced by its generative AI services and features (ie both within and outside general search), and displaying grounding sources to the user as part of the generated response.[38]

4.9 As described in paragraph 1.11, publishers have expressed concerns that they are not able to exercise effective control over their Search Content being used both to ground Google's generative AI services and features,[39] as well as to train or fine-tune the AI models underpinning these services and features.[40]

4.10 Google has told us that:

(a) Google-Extended already enables publishers to control whether their Search Content is used in training Google's foundation models and grounding Google's broader generative AI services.[41]

(b) It would be technically feasible for Google to develop a control that would allow publishers to opt out their content from being used by search generative AI features for the specific purpose of grounding generative answers.[42] AI Overviews and AI Mode rely heavily on corroboration with grounding sources.[43] The models underlying these features are trained to provide an overview response that closely follows

---

[37] Google confirmed to us that the models underlying AI Overviews derive from the Gemini foundation model family. The Google-Extended control allows publishers to opt out their Search Content from being used to train the Gemini foundational model family but does not enable them to opt-out of Search Content being used to fine-tune the AI models within Search, to the extent that occurs. See Google's response to the CMA's RFI. Google's submission to the CMA.

[38] Google's response to the CMA's RFI.

[39] Several publishers suggested they do not have effective choice over how their content is used for display and grounding in search generative AI features. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. Of these, 8 ([✂]) explicitly mentioned concerns relating to display and grounding in broader generative AI services.

[40] Several publishers said they are concerned how their content is used for fine-tuning and training models powering generative AI services and features [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[41] Google's submission to the CMA.

[42] Google's submission to the CMA.

[43] Google's submission to the CMA.

the search snippets and links returned in response to a user's query.[44] If a publisher were to opt out of grounding, these features would not be able to generate accurate, up-to-date and cited responses using that publisher's content.[45]

(c) There is no realistic prospect of harm to publishers in respect of training/fine-tuning of AI models for search and search generative AI features. Fine-tuning helps the model learn how to process information rather than what current information to display; this internal processing does not create a substitute for publisher websites. A model relying solely on patterns learned during training would be static, often outdated, and prone to hallucinations.[46]

(d) Allowing publishers to opt out of the fine-tuning of models underlying search generative AI features would be counterproductive. Fine-tuned models are used for a variety of purposes across general search including query understanding and ranking. An opt-out would therefore raise the risk of downranking or mis-ranking publisher content in organic search results. Scoping a control to cover the fine-tuning of "generative" models but not "ranking" models presents a false dichotomy that would be impractical and restrict Google's ability to innovate, as models are not static and can evolve beyond their initial purpose. 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.[47]

4.11 As provided for in paragraph 2 of the Publisher CR, our provisional view is therefore that:

(a) Outside Google's general search, effective controls should allow publishers to opt out of their Search Content being used for any generative AI purpose (ie both training and grounding).

(b) Within Google's general search, effective controls need to enable publishers to opt out of their Search Content being used to ground responses in search generative AI features. Responses that are

---

[44] Google's response to the CMA's RFI.
[45] Google's submission to the CMA.
[46] Google's submission to the CMA.
[47] Google has told us that, [✂]. Google's submission to the CMA.

displayed in search generative AI features are closely based on publisher Search Content which has been retrieved through grounding.

4.12 In practice, we expect that this would mean Google maintaining its current approach to Google-Extended as a control covering training and grounding outside of general search, and also implementing a new control allowing publishers to opt out of their Search Content being used in the grounding of content in its search generative AI features.

4.13 We note that Google-Extended currently offers a single control covering both training and grounding use cases. We would like to hear further evidence on the benefits and risks of Google providing separate controls over training and grounding outside of general search.

The scope of the controls should be product-agnostic and consistent within and outside general search

4.14 Google's existing control, Google-Extended, relates to how publisher Search Content is used in broader generative AI services. Google has publicly said that Google-Extended allows publishers to manage whether Search Content may be used for training future generations of Gemini models that power Gemini Apps and Vertex AI API for Gemini for grounding (providing content from the Google Search index to the model at prompt time to improve factuality and relevancy) in Gemini Apps and Grounding with Google Search on Vertex AI.[48]

4.15 As described in paragraph 1.11, several publishers have told us that the scope of Google-Extended is ambiguous. For example, it is unclear which services are excluded and how the control evolves as services develop. Ambiguity is likely to reduce publishers' understanding of the choice they are being offered. We also recognise that Google's general search products are evolving rapidly, and that the scope of controls must be sufficiently flexible to be able to adapt to this evolution.

4.16 We provisionally consider that, by defining this control in relation to specific services, Google's current approach does not provide sufficient certainty to publishers as to the effect of the control in relation to its broader generative AI services and in relation to any future services it may develop and roll out.

---

[48] Google, 'Crawling Infrastructure', dated 21 November 2025, accessed by the CMA on 10 December 2025.

4.17   We have therefore stated in the Interpretative Notes for the Publisher CR that we expect the effective controls to include a product-agnostic definition of their scope and such a definition should be applied to controls operating both within and outside general search.[49]

4.18   We consider this would ensure that the controls are effective by:

(a) preventing the perception that Google might cherry-pick which services fall within scope of the control and would provide a clear expectation as to which future services would be brought within scope; and

(b) scoping in services which produce a substantial amount of generative output, grounded on or otherwise using publisher content, in response to user queries.

4.19   In practice, we expect this would mean Google ensuring and clarifying that its existing Google-Extended control meets this definition of its scope and also applying this definition of scope to a new control covering its search generative AI features.

The granularity of controls should allow publishers to opt out at directory-level and page-level within general search

4.20   There are different methods by which publisher web controls may be implemented. For example, Google-Extended is implemented via a text file – robots.txt – saved in a website's root directory[50] whereas nosnippet is implemented via HTML metatags contained in the source code of individual web pages.[51]

4.21   Our understanding is that different methods may lend themselves to indicating preferences at different levels of granularity. For example, robots.txt preferences are well suited to expressing directory-level preferences, whereas HTML metatags are more suited to page or sub-page level preferences. We are also aware that there are ongoing discussions in the IETF's AI Preferences Working Group on standardising a method for indicating AI preferences.

---

[49] Applying a product-agnostic definition within Search would ensure a control covers any future search features – beyond AI Overviews and AI Mode – which Google may roll out, should they fall within this definition.
[50] Robots Exclusion Protocol is a standard agreed by IETF: RFC 9309: Robots Exclusion Protocol.
[51] Google, 'Meta tags and attributes that Google supports,' accessed by the CMA on 7 January 2026.

4.22   A few publishers have told us that being able to opt out at different levels of granularity would be beneficial. They note that this would allow them to assess the merits of opting in different types of Search Content separately.[52] We have seen evidence in Google's internal documents that suggests [✂].[53]

4.23   Our provisional view is that, within general search, an effective control should allow publishers to opt out at both directory-level and page-level:

(a)   Directory-level control is necessary to ensure that publishers are able to opt out their entire website in a simple and easy manner. Publishers may also wish to opt out only certain directories using the control – for example opting out certain content genres, or paywalled content.

(b)   Page-level control is necessary to ensure publishers can exercise different choices where the costs and benefits of opting out vary across their Search Content. This could include allowing publishers to benefit from their Search Content generally appearing within search, whilst being able to opt out, for example, (a) Search Content where there are particularly high risks or stakes for insufficient attribution quality (eg where inaccuracies or low prominence might erode their brand),[54] (b) relatedly, Search Content where generative summaries most restrict click-throughs, or (c) premium Search Content where a directory-level structure would not be workable.[55]

4.24   We have therefore specified in the Interpretative Notes that, in order for the new control covering search generative AI features to be effective, we would expect Google to allow publishers to opt out at both directory-level and page-level.[56]

---

[52] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[53] Google's internal document. This document also shows that [✂].

[54] We discuss these issues further in paragraphs 4.51 to 4.69.

[55] We recognise that a solution providing page-level controls could be more costly than a solution without, because it would depart from the robots.txt standard. However, Google's current offering (in particular, nosnippet) permits some similar page level control; we infer from this that the costs are manageable and likely lower than the benefits.

[56] Unless it is not technically possible to do so, our provisional view is that effective compliance with our Publisher CR could additionally require augmenting the current functionality in Google Search Console to allow publishers to block quickly a specific page from appearing in search generative AI features, whilst still appearing in non-generative AI search results. This would, for example, support publishers concerned that particular content is inaccurately displayed in search generative AI features causing reputational damage. We would welcome views from publishers and Google on whether this would support the effectiveness and proportionality of the Publisher CR.

4.25   We note that Google-Extended does not currently offer easy page-level control. It would be more onerous to require Google to reimplement this control at page-level, and our view is that the benefits of this control outside of general search would be more limited.[57] We have therefore not specified this in the Interpretative Notes, but would like to hear further evidence on the benefits and risks of Google providing content-level controls for Search Content usage outside of general search.

Supplemental requirements necessary to ensure the effectiveness of the controls

4.26   We consider that the controls would likely not be effective if:

(a) they fail to reflect the way in which generative AI services and features will continue to evolve in the future;[58] and

(b) Google were to be able to undermine a publisher's expressed choice by:

(i) intentionally affecting how a publisher's Search Content appears in organic search results if they opt out of that Search Content being used in generative AI services and features (eg by downranking or changing how it is presented);[59] or

(ii) circumventing publisher choices by acquiring opted-out content through other means. For example, in principle Google could pay a third party to crawl an opted-out website.[60]

4.27   We have therefore also included a high-level requirement addressing these risks at paragraph 3 in the Publisher CR.

---

[57] In particular, the high costs to publishers of opting Search Content out of Google Search may not apply to the same extent to all of Google's wider generative AI services and features, for example where the importance of appearing in such features is lower (eg because Google does not have a position of strategic significance). This implies that Google-Extended is more likely to benefit publishers even where publishers are forced through lack of granularity to opt some Search Content out which they would rather opt in, than were this lack of granularity also evident in the control.

[58] See also our analysis of the need for a product-agnostic description of the scope of the controls at paragraphs 4.14 to 4.19.

[59] However, we recognise that Google may not be able to control or prevent all possible second order effects of these controls on its overall ranking algorithm.

[60] We would not consider Google attempting to acquire content opted out of Google-Extended via open-source datasets to amount to circumvention, where these datasets have obtained content legally, given the open nature of that data source.

How the CR ensures that Google will provide transparency over its use of publishers' Search Content in its generative AI services and features and user engagement with their Search Content in search generative AI features

4.28　The above analysis relates to ensuring that Google provides publishers with effective controls to enable them to opt their Search Content out of Google's generative AI services and features. As part of that analysis, we considered the information required to allow publishers to understand the extent and scope of those controls.[61] However, in order for publishers to be able to make informed decisions, including in relation to those controls, we consider it would also be necessary to ensure that Google provides publishers with sufficient transparency more broadly, as set out in the second limb of our aim.

4.29　In designing the transparency elements of the Publisher CR obligations (and the underlying Interpretative Notes), the CMA has had regard to the following key design parameters relevant to their effectiveness, which are considered further below:

　　(a)　providing publishers with sufficient transparency over how Google uses publisher Search Content in generative AI services and features;

　　(b)　transparency in relation to publisher controls; and

　　(c)　providing publishers with sufficient transparency in relation to user engagement and performance in search generative AI features.

Google should provide publishers with sufficient transparency over how it uses publishers' Search Content in generative AI services and features

4.30　In order for publishers to be able to make properly informed and meaningful decisions in relation to how they interact with Google, they need to understand how Google uses Search Content in its generative AI services and features.

4.31　Currently, there is a general information asymmetry between Google and publishers over the functioning of generative AI services and features given the pace of development of these new technologies, which has reduced

---

[61] See paragraphs 4.14 to 4.19 and paragraphs 4.34 to 4.35.

publishers' understanding of the commercial exchange between them and Google.[62]

4.32 We have therefore included, at paragraph 4 of the Publisher CR, a requirement on Google to publish detailed information enabling publishers to understand how their Search Content may be used in the training and grounding of generative AI services and features.

4.33 We would expect Google to provide this information and guidance in a public and accessible manner to ensure that it can be used by publishers and have reflected this in the Interpretative Notes.

Google should provide transparency over the effect and scope of its publisher controls

4.34 We are concerned that Google does not offer sufficient transparency around the scope and operation of its existing publisher controls. As described in paragraph 1.11, publishers have told us that the definition of Google-Extended is ambiguous.

4.35 To support publishers in deciding whether to exercise its controls, we think Google should publicly provide clarity over:

(a) The scope of its controls (via a product-agnostic definition);

(b) Any exceptions to the scope of its controls, with reasoned explanations; and

(c) Examples of current use cases falling within and outside the scope of its controls.

Google should provide publishers with sufficient transparency over user engagement and performance within search generative AI features

4.36 We have also considered what information publishers need on user engagement for them to make a properly informed and meaningful choice over how Google uses their Search Content. Where Google uses Search Content in its broader generative AI services outside general search (for example in its Gemini AI assistant) we would expect competition from other

---

providers to incentivise Google to provide appropriate information to publishers. However, Google's search generative AI features form part of its general search services in which Google has SMS and therefore does not face such competitive discipline. We therefore consider that, in order to meet our aim, an effective CR should require Google to provide publishers with information on user engagement and performance in search generative AI features (but not broader generative AI services).

4.37 Publishers make use of various metrics to understand the performance of their content within Google Search, provided via Google Search Console. The most fundamental of these metrics are 'Clicks' (how many clicks a publisher receives from Google Search), 'Impressions' (the number of times a publisher link is present in the SERP) and 'Click-Through Rate' (the ratio of clicks to impressions) (**CTR**).

4.38 Currently, publishers are unable to view disaggregated metrics showing the performance of their content within search generative AI features. For example, Clicks and Impressions from these features are aggregated across Google Search as a whole within Google Search Console. Publishers have told us that they would value greater insight into the performance of user engagement with their content in search generative AI features such as 'Clicks', 'Impressions' and 'CTR' provided separately from performance data from Google Search as a whole.[63]

4.39 Our provisional view is that Google should provide, at a minimum, performance data such as 'Clicks', 'Impressions' and 'CTR' in aggregate for search generative AI features as separate metrics from those provided in relation to non-generative AI search results. We consider that this information should:

(a) inform publishers' choice of whether to allow their Search Content to be used in Google's search generative AI features, by providing them with an insight into the potential value of opting in; and

---

[63] [✂] response to the CMA's RFI. Candr Media Group's response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. Iconic Media's response to the CMA's RFI. [✂] response to the CMA's RFI. Financial Times's response to the CMA's RFI. BBC's response to the CMA's RFI. The Independent's response to the CMA's RFI. [✂] response to the CMA's RFI.

(b) allow publishers to optimise their business for performance within search generative AI features by understanding which content performs well and drives traffic to their websites.

4.40 We have therefore included a requirement to provide publishers with metrics on user engagement in the Publisher CR and further clarified this in the Interpretative Notes.

Access to performance data on a per-feature basis

4.41 We have further considered to what extent this data needs to be disaggregated per-feature to provide the benefits outlined above, ie whether one metric covering all search generative AI features would be sufficient, or whether publishers require access to data on a 'per-feature' basis. Providing data on a per-feature basis would mean providing publishers with separate data regarding specific search generative AI features such as AI Overviews and AI Mode.

4.42 Several publishers explicitly stated they wished to receive performance and user engagement data on a 'per-feature' basis.[64] For example, one publisher informed us that it would enable them to better assess the impact of each feature independently and optimise their content strategy to what users were seeking in either AI Overviews or AI Mode.[65]

4.43 Google have told us that providing data disaggregated on a 'per-feature' basis to publishers would not be an appropriate level of granularity to meet the stated aim, as (among other reasons):

(a) Providing data on a feature-by-feature basis would not be future proof because individual features are likely to evolve over time.

(b) Google has not historically provided engagement metrics broken out at a per feature level of its non-generative AI search results.[66]

4.44 Our provisional view is that it would not be necessary for Google to provide performance data on a 'per-feature' basis, so we have not reflected this in our Publisher CR or Interpretative Notes. This is because it is currently not clear to what extent (if any) there is an actionable benefit to the provision of

---

[64] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.
[65] [✂] response to the CMA's RFI.
[66] Google's response to the CMA's RFI.

performance data such as 'Clicks' and 'Impressions' on a 'per-feature' basis. An approach based on the provision of performance data from specific search generative AI features may also be less flexible as those features continue to evolve and develop.

4.45 We would like to receive further evidence on the benefits and risks of Google providing performance and engagement information on a 'per-feature' basis within general search.

Transparency over 'Click Quality' for traffic referred from search generative AI features

4.46 The introduction of search generative AI features has led to concerns from publishers regarding the volume of traffic they are receiving from Google, with reporting of significant drops in organic traffic. For many publishers, reductions in traffic volumes reduce opportunities to monetise their businesses whether via advertising or subscriptions.

4.47 Google has stated that AI Overviews have sent higher quality traffic to websites (meaning traffic which is more likely to convert into a positive commercial outcome such as a subscription), as users use search generative AI features as starting point for further in-depth research, and these follow-on clicks are of more value. Alongside this, Google stated that AI Overviews and other search generative AI features improve the quality of the user experience, and whilst there may be shifts in traffic to individual sites, those that are adapting to changing user habits, are receiving higher quality traffic.[67]

4.48 Publishers have told us that currently they cannot accurately assess whether referral traffic from search generative AI features is of higher value than that referred to publishers via traditional search links.[68] We expect that this information on click quality would materially impact a publisher's decision as to whether to opt in to use of their content in search generative AI features, by enabling them to understand the value of referral traffic from these sources.

4.49 Our provisional view is that publishers should have access to greater information regarding the 'Click Quality' of traffic referred from search

---

[67] Google, The Keyword, 'AI in Search is driving more queries and higher quality clicks', dated 6 August 2025, accessed by the CMA on 16 December 2025.
[68] [✕] response to the CMA's RFI. [✕] response to the CMA's RFI. [✕] response to the CMA's RFI. [✕] response to the CMA's RFI. [✕] response to the CMA's RFI.

generative AI features to be able to make an informed decision on the use of their Search Content in those search generative AI features.

4.50   As such we have clarified in the Interpretative Notes that Google should also provide publishers with a clear and actionable insight into the quality of clicks provided via search generative AI features. This could involve directly providing this metric or facilitating publishers to calculate this metric themselves (for example, by indicating which referrals arrive via a search generative AI feature). We welcome submissions on the most effective method of providing publishers with this information.

How the Publisher CR ensures that Search Content will be sufficiently attributed in search generative AI features

4.51   To meet our aim, the Publisher CR also needs to ensure that publishers have sufficient trust that where Google uses Search Content in its search generative AI features, that content is sufficiently attributed. Without such trust, publishers will not be able to make properly informed and meaningful decisions about whether and how they interact with Google in respect of general search services. Effective attribution is also important for enabling users to verify the AI-generated responses that are provided, by clicking through to the underlying content if they wish.[69]

4.52   In designing the attribution elements of the Publisher CR (and the underlying Interpretative Notes), the CMA has had regard to the following key design parameters relevant to their effectiveness, which are considered further below:

   (a)   the factors relevant to how Search Content is attributed within search generative AI features;

   (b)   the further steps Google needs to take to ensure that Search Content is sufficiently attributed within search generative AI features; and

   (c)   whether publishers require any additional means to raise complaints about the way Search Content is attributed within search generative AI features.

---

[69] Where Google uses Search Content in its broader generative AI services outside general search (for example in its Gemini AI assistant) we would expect competition from other providers to incentivise Google to attribute content effectively. We do not have the same expectation in relation to Google's search generative AI features, which form part of its general search services in which Google has SMS. We have therefore focused our assessment on attribution of Google's content in search generative AI features.

Factors relevant to how Search Content is attributed within search generative AI features

4.53   As set out at paragraph 1.14 above, while we recognise that Google has taken steps to improve the attribution of content, we consider that sufficient attribution is, and will continue to be, important. This is not only because of Google's scale, and its significance as a source of traffic for publishers, but also because effective attribution helps consumers understand that information is automated and derivative and may need checking, both now, and as services continue to develop.

4.54   In relation to accurate attribution within search generative AI features, several publishers have expressed concerns that they do not have sufficient information from Google to monitor how their Search Content is attributed and identify factual inaccuracies. Although publishers may receive anecdotal reports of inaccurate citations from their audience and internal teams, they would benefit from additional transparency and tools, and the Search Console is not sufficient on its own.[70] Publishers told us that accurate attribution is important, for example because:

   (a) In instances when AI Overviews summarise information inaccurately, this has damaging implications for publisher brands;[71]

   (b) Increased access to real-time, granular data from Google would reduce publishers' need to rely on third-party tools, lower monitoring costs and resources, and improve a publisher's ability to track content performance and leverage effective strategies.[72]

4.55   In relation to the prominence of attribution, several publishers have also argued for specific changes to how content is displayed in search generative AI features.[73]

---

[70] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[71] [✂] response to the CMA's RFI.

[72] [✂] response to the CMA's RFI.

[73] As an indication: one publisher asked for expanded use of in-line anchor text citations rather than primarily requiring the user to tap/click on the link icon and find the source in the citation pane. [✂] response to the CMA's RFI. Another publisher asked for a prominent logo to provide branding that clearly identifies the publisher to users, and a clickable publisher URL along with a clickable headline link which is more prominent than the answer text provided in AI Overviews. [✂] response to the CMA's RFI. Some publishers said that the attribution

4.56 We recognise that:

(a) There is a link between lack of prominent attribution and lower click-through rates, as subtle and poor attribution means users are less likely to click through to the original source.[74]

(b) Prominent attribution within search generative AI features can lead to better protection of publisher brands, since it helps promote users' awareness of which publisher(s) are generating the content that supports the search generative AI features.[75] This is likely to drive user engagement with publisher content, facilitating them in maintaining flows of traffic, both from Google and from users visiting their site directly.

(c) Prominence and recognition of publisher brands is also important for user engagement and content monetisation, and helps maintain the commercial viability of publisher business models, ultimately supporting the plurality of web content for end users.[76]

The further steps Google needs to take to ensure that Search Content is sufficiently attributed in search generative AI features

4.57 Giving publishers a control to opt out of their Search Content being used in search generative AI features should help provide Google with stronger incentives to ensure attribution of Search Content is sufficient.

(a) As set out above, we are proposing that Google should provide a new control allowing publishers to opt their content out of search generative AI features, whilst remaining in traditional search links and so continuing to attract at least some click-through.[77] This would reduce the costs to

---

must appear at the top of the AI-generated summary. For example: [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[74] For example, one publisher said that the main issue flowing from search generative AI features is the dramatic decrease in traffic through their website pages because the user's query is answered without the user needing to click through to their websites. This is being exacerbated by poor attribution. The same publisher reported that attribution occurs only sporadically and inconsistently, and when it does, the link is not prominent. [✂] response to the CMA's RFI. Another publisher mentioned that attribution should be at least as good as that currently supplied in non-AI search, to ensure that the click-through rates on AIOs are no worse than those on non-AI search. [✂] response to the CMA's RFI. Another publisher said that while a citation link may appear, it is often subtle, lacks prominence, and fails to provide the user with a clear reason to click through to the original source. [✂] response to the CMA's RFI.

[75] See also paragraph 5.31(d).

[76] See also paragraph 5.31(d).

[77] In order for this to be an effective tool for publishers, it is important that the control operates at a page level (as does noindex today) – otherwise publishers may have to opt out entire segments or even potentially all their content in response to a factual issue relating to an individual page.

publishers of opting content out of search generative AI features, relative to the status quo.[78]

(b) We would expect that reducing the costs to publishers from exercising the opt-out in these ways would provide Google with stronger incentives to ensure prominence and improve accuracy in attribution. This is because Google would have a greater need to encourage publishers to continue to opt their Search Content into its search generative AI features. We set out further evidence on this benefit in paragraph 5.31(d).

4.58 However, we consider that, although introducing a new publisher control would incentivise more effective attribution, given the evolving nature of search generative AI features, Google's scale and the importance of the issues for publishers and users alike, this alone would not ensure effective attribution.

4.59 Our provisional view is therefore that, in order to be effective in achieving the aim, Google should also be required to:

(a) take reasonable steps to ensure that Search Content is sufficiently attributed when used in a search generative AI feature; and

(b) publish further information describing the steps it takes to ensure Search Content is attributed sufficiently in search generative AI features and ensure and measure the factuality of such responses.[79]

4.60 We expect 'sufficient attribution' to relate to both the accuracy and the prominence of attribution and to reflect Search Content that makes a significant contribution to a search generative AI feature response. We also recognise that decisions on attribution require Google to balance this against end-user experience; design aesthetics; and source diversity. We have proposed clarifying these points as part of our Interpretative Notes.[80]

4.61 We consider that, consistent with our aim, these additional obligations would:

---

[78] Because publishers will continue to attract click-through from organic listings, even where they opt their content out of Search generative AI features.

[79] We would expect this information to include an explanation of how Google identifies what Search Content to attribute, how it monitors the factuality of attribution, how it seeks to ensure search generative AI features are as accurate as possible and how it ensures that attribution is sufficiently prominent.

[80] We would expect to consult Ofcom, which has duties with respect to media plurality in the UK, when assessing how Google has weighed up these factors.

(a) improve publishers' trust that Search Content will be sufficiently attributed in search generative AI features and

(b) improve publishers' ability to decide whether they wish to opt any Search Content out of search generative AI features on the basis of the risk of poor factuality, without expensive monitoring.

Whether publishers require additional means to raise complaints about the way Search Content is attributed within search generative AI features

4.62   Publishers have also raised concerns about their lack of ability to deal with issues relating to incorrect attribution of their content.

(a) For example, a local news publisher suggested that Google should allow publishers to report wrong or misleading uses of their content quickly and have a rapid correction process.[81] The same respondent said that they cannot check referrals from AI Overviews or AI Mode reliably, and manual checking is impossible at scale, and that they need a dashboard or alerts showing how often and where their content appears in AI answers, and tools to flag incorrect or outdated information linked to their brand.[82] Other publishers made a related point that a mechanism in Google Search Console to report issues with AI Overviews would help address concerns about inaccurate or misleading AI results.[83]

(b) We also heard that publishers need further data, insights and information about how content is surfaced in response to user queries to assist with identifying risks associated with instances of incorrect or out-of-date information appearing in Google's search generative AI features, and that a two-way communication and a feedback loop between relevant Google personnel and publishers is essential to ensure that any identified issues can be promptly rectified, noting the immediate impact and reputational/user trust damage to a publisher as a result of incorrect information being surfaced and attributed to that publisher.[84]

4.63   We agree that inaccuracies within search generative AI features can have a significant reputational impact on publishers, and can also harm end users who need to verify the veracity of the response, and that, in order to be

---

[81] Iconic Media's response to the CMA's RFI.
[82] Iconic Media's response to the CMA's RFI.
[83] For example, see [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.
[84] [✂] response to the CMA's RFI.

effective, a Publisher CR would need to ensure that publishers are able to take appropriate action if these issues arise.

4.64 We also recognise that Google has taken steps in relation to the accuracy of search generative AI features. In particular, Google told us that:

(a) It tries to optimise towards factuality in AI Overviews by using human Search Quality raters to review and rate example AI Overviews to determine how well the AI Overview response has grounded to inputs. Specifically, [✂].[85]

(b) It also offers a mechanism by which publishers and users can report instances of search generative AI features – namely AI Overviews and AI Mode – containing inaccurate information. This is accessed via the 'thumbs up / thumbs down' buttons which are displayed below any generative AI response. These icons allow publishers and users to flag if something is wrong. If the 'thumbs down' icon is clicked, a user can then select 'Not factually correct'.[86] Google told us it receives approximately [hundreds of thousands of] pieces of thumbs up / thumbs down feedback on AI Overviews and AI Mode globally each day.[87] We also understand, based on Google's submissions, that Google considers this feedback on an aggregated basis and [✂].[88]

4.65 Our initial view is that Google's existing thumbs up/down mechanism provides some ability for publishers to identify concerns, which can help Google to improve accuracy of its search generative AI features. However, on its own it has limitations: it is not a two-way publisher-facing mechanism for publishers to raise complaints; and there is a lack of transparency between the data that Google aggregates internally and the information available to publishers.

4.66 We also consider that the publisher choice proposals set out above should provide additional control for publishers to address these concerns. In particular, as set out in paragraphs 4.20 to 4.25 above, we expect our Publisher CR to provide publishers with the ability to block their content from AI-generated responses at page level. Having a page-level control would enable publishers to withdraw Search Content from search generative AI

---

[85] Google's response to the CMA's RFI.
[86] Google's response to the CMA's RFI.
[87] Google's response to the CMA's RFI.
[88] Google's response to the CMA's RFI.

features, taking action themselves to address factuality issues when they arise.

4.67   In addition to our attribution proposals set out above, we consider that there could be benefits in Google providing further tools and information to publishers. This could include, for example, a mechanism through which publishers can more easily communicate their reasons for blocking content from appearing in search generative AI features.[89] This could give Google helpful data for the purposes of ensuring the factuality of its search generative AI features and its compliance reporting. We welcome views from Google and publishers whether the addition of such a mechanism would enhance the effectiveness of our CR whilst ensuring it remains proportionate.

4.68   Our initial view is that these measures represent the most effective and proportionate way to reduce the negative effects of inaccuracies in search generative AI features. At this stage, we therefore do not consider it necessary to specify a new and separate two-way publisher facing complaints process as part of the Publisher CR.[90]

4.69   In particular, our initial view is that a bespoke complaints process which goes further and includes requirements on Google to assess whether to take action is likely to make the CR more onerous than necessary. This is because we consider such a requirement would provide relatively little additional benefit beyond other measures we have set out above,[91] but involve expense from Google reading, actioning and prioritising complaints. Such complaints could in principle stem from a material fraction of its index of over 20 billion websites,[92] relating to search generative AI features which are difficult to replicate (because these responses exhibit high variability across the same

---

[89] As set out in paragraphs 4.20 to 4.25, we would expect Google's compliance with our Publisher CR to extend to offering page-level control. This could include augmenting the current functionality in Google Search Console to allow publishers to block quickly a specific page from appearing in search generative AI features, whilst still appearing in non-generative AI search results. We consider that, as part of this mechanism, Google could allow publishers to note their rationale for withdrawing a page, with 'Not factually correct' as an option. Because Google Search Console requests can only be submitted by the owner of websites, this would allow Google to differentiate between views of users – collected through the thumbs up/down mechanism – and perceptions of the content creator itself.

[90] Publishers would already be able to raise concerns if they consider that Google has breached the attribution requirements in the Publisher CR.

[91] Our proposed Publisher CR design allows publishers to restrict their content from appearing in Google's search generative AI features, providing immediate redress. Further, we consider simpler mechanisms to report inaccurate summaries will support the production and monitoring of sufficient attribution more widely.

[92] Strategic Market Status investigation into Google's general search services: Final Decision (SMS Decision), 10 October 2025, paragraph 5.208.

query),[93] and are generated probabilistically from a model (implying that individual complaints would be difficult to address).

Implementation and compliance

4.70 A CR comes into force at a time determined by the CMA.[94] Once in force, Google would be required to provide the CMA with a compliance report in relation to that CR[95] and the CMA would be required to keep under review the extent to which Google is complying with the Publisher CR.[96] This section sets out our proposed approach to ensure that any final Publisher CR is implemented effectively and to monitoring compliance.

Approach to monitoring and compliance

4.71 We propose that the conduct requirement would come into force within six months following imposition. During this period Google should, within one month of imposition, submit an implementation plan and engage constructively with the CMA and third parties to develop and implement changes to comply with the requirement.

4.72 We propose to monitor compliance through four main mechanisms: an initial baseline compliance audit, a six-monthly compliance report from Google, reporting of key information and data from Google; and ongoing stakeholder engagement and feedback.

4.73 To assure publishers that Google's controls work as described, Google should carry out, submit and publish a baseline compliance audit. We recognise that appointing a 'skilled person'[97] may introduce some further costs that Google would not have to bear if carrying out its own internal compliance audit. However, given that one of the core issues we are seeking to address is a lack of trust by publishers that Google implements its controls in a way consistent with its public statements, we provisionally consider that there is an intrinsic value in a baseline report being carried out by an independent third

---

[93] See for example the analysis in Wellows, How AI Answer Variability Impacts SEO Results, 23 December 2025.
[94] Section 19(11)(a) of the Act.
[95] Section 84(1) of the Act.
[96] Section 25(b) of the Act.
[97] Section 79 of the Act, pursuant to which the CMA may appoint a 'skilled person' to provide it with a report about matters relevant to its digital markets functions in relation to a designated undertaking. See CMA194, paragraphs 5.68-5.80.

party. We would expect this audit to include adversarial testing to ensure that publisher content does not appear where it has been opted out.

4.74 Google's six-monthly compliance report would include:

(a) An explanation of how it has complied with the Publisher CR over the relevant period. Including:

(i) any updates to the implementation of its publisher controls;

(ii) how it has applied its publisher controls to new generative AI services and features or changed the application to existing services or features;

(iii) how it has continued to monitor that its controls function as intended and described to publishers, including internal adversarial testing;

(iv) any updates to public transparency offered in relation to its publisher controls, its use of publisher Search Content in generative AI services and features, and its attribution of publishers in search generative AI features;

(v) an overview of any changes it has made to its attribution of publisher content in search generative AI features, including steps to ensure accurate attribution and that end users are provided with a clear means to access Search Content; and

(vi) any updates on the steps Google takes to measure and ensure the factuality of its search generative AI features.

(b) A summary of the most frequent stakeholder feedback received with respect to the above.

4.75 As part of compliance reporting, we propose to require Google to supply the following information and data to enable us to assess the impact and effectiveness of the Publisher CR and Google's compliance with it:

(a) The number of publishers making use of its controls, including details of any cases where a publisher's decision to exercise or otherwise these controls become subject to contractual agreement with Google;

(b) An aggregation of the user engagement metrics it provides publishers in relation to its search generative AI features;[98] and

(c) Information and data with respect to its search generative AI features, including for example the feedback collected through Google's thumbs up/thumbs down mechanism.[99]

4.76 Beyond this reporting, we would maintain regular communication with stakeholders on this Conduct Requirement. This would enable them to raise issues with us if they believe Google is failing to comply with the requirement.

4.77 In the interest of transparency, the CMA considers that Google should prepare a non-confidential version (alongside the confidential version) of each compliance report and related performance metrics and publish this at the same time as submitting it to the CMA. This would improve confidence in Google's compliance with the Publisher CR and enable third parties to provide further views on Google's compliance.

Consistency with existing and expected laws and standards

4.78 This CR would sit alongside existing legal and regulatory frameworks governing the use of online content in AI systems. This section outlines how it interacts with existing legal regimes, regulatory guidance and emerging technical standards.

Existing Copyright Framework & robots.txt debate

4.79 The Government is currently considering responses to its consultation on the UK AI Copyright framework, including what publisher content can be used for AI model training and on what terms.[100] The focus of our CR is on addressing issues of Google's market power, rather than cross-cutting copyright issues. The Publisher CR applies to Google, as it has SMS, rather than the market as a whole.

---

[98] This could include average (a) click-through rate and (b) click-quality at breakdowns which support evaluation.
[99] See paragraph 4.64(b) above.
[100] See Intellectual Property Office, Department for Science Innovation and Technology, and Department for Culture, Media and Sport, 'Copyright and Artificial Intelligence', dated 17 December 2024, accessed by the CMA 16 December 2025.

Internet Engineering Task Force

4.80 In parallel with legislative developments, the Internet Engineering Task Force (**IETF**) has initiated work to address the technical dimension of content governance in the context of AI.[101] The newly chartered AI Preferences (**AIPREF**) Working Group is tasked with developing standardised mechanisms for expressing how internet content may be used in the development, deployment, and operation of AI systems.

4.81 The AIPREF Working Group does not engage with the legal enforceability of these preferences but seeks to improve clarity and interoperability across platforms. Its work responds to growing concerns among content creators and publishers about the opaque and inconsistent AI preferences signalling practices currently used by AI developers. The CMA is mindful that our CR on publisher control needs to be flexible to reflect possible future changes in this area.

CRs that could be equally effective

Crawler separation

4.82 Alongside considering what controls and transparency Google should provide to publishers, we have carefully considered whether Google should be required to operate separate crawlers to gather data for its search and AI products and services. This could be an alternative to a requirement on Google to improve its existing publisher controls, outlined at paragraphs 4.4 to 4.27 4.27 above. We refer to this option as 'crawler separation'.

4.83 As noted at paragraph 1.11, publishers do not at present have sufficient choice over whether their content, crawled for Google's general search, is used in Google's generative AI services and features. Because Google Search is such a critical route to market for most publishers, they have no choice but to grant access to Google's crawler to ensure they appear in search results. Crawler separation could potentially address this issue by enabling publishers to accept or reject access to Google's AI crawler independently from the search crawler.

4.84 We have considered a range of arguments in favour of crawler separation. For example, one proposal is that there should be distinct crawlers for:

---

[101] See IEFT Datatracker 'AI Preferences (aipref)', accessed by the CMA 16 January 2026.

(a) search indexing,

(b) AI training/fine-tuning, and

(c) AI grounding or summarisation.

4.85   This separation would allow publishers to block or allow each type of crawler independently, without needing to rely on controls provided by Google within its existing search crawler.

4.86   Our initial assessment is that crawler separation could be effective in meeting the aim set out at paragraph 2.2 above. This is because it would enable publishers to make decisions separately over whether their content is crawled for different purposes. We note that crawler separation would need to be operated alongside other requirements in order to ensure Google was not able to recombine data from the different crawlers or use data for wider purposes.

4.87   However, our view is that crawler separation would bring more costs and risks of unintended consequences to both Google and third parties than our proposed CR. As set out in the proportionality assessment at paragraph 5.7 to 5.12, crawler separation would require Google to create and maintain a separate crawling infrastructure. For publishers, an additional crawler would increase the load imposed on their servers, increasing costs.

4.88   For this reason, our assessment is that improving Google's controls would be a more proportionate way of meeting our aim.

5. Provisional proportionality assessment for the Publisher CR

5.1 The CMA may only impose a CR if it considers that it would be proportionate to do so for the purposes of one or more of the statutory objectives, having regard to what the CR is intended to achieve (as set out in Section 2 above).[102]

5.2 This section sets out our provisional proportionality analysis for the proposed Publisher CR. A proportionate CR is one that:

   (a) Is effective in achieving its intended aim;

   (b) Is no more onerous than it needs to be to achieve its intended aim;

   (c) Is the least onerous CR, where the CMA has identified multiple equally effective options that would achieve the intended aim; and

   (d) Does not produce disadvantages that are disproportionate to its aim.[103]

5.3 We have provisionally considered each of these four criteria below.

The CR is effective at achieving its intended aim

5.4 As set out in the effectiveness section, improved controls would enable publishers to make a meaningful choice over how their content, crawled for Google's general search, is used in Google's generative AI services and features. That decision would be supported by improved transparency over how their content is used and engaged with, and measures to ensure effective attribution of content. For example, additional transparency and effective attribution would further empower publishers in making decisions about how Google uses their content, what types of content they produce, and how they optimise that content for display in Google Search.

5.5 In the sections below, we set out evidence on the significance of the benefits and costs which would flow from achieving this aim, given our proposed CR design set out above.

---

[102] Section 19(5) of the Act.
[103] Digital markets competition regime guidance, December 2024 (CMA194), paragraph 3.33.

The CR is no more onerous than necessary

5.6   The effectiveness section sets out the key policy design questions we have considered and explains why the provisions of the CR are no more onerous than necessary. In particular, it demonstrates why:

    (a) The proposed elements of the CR are necessary to meet the aims we have articulated in Section 2.

    (b) The CR does not go beyond what is necessary. We have set out in the effectiveness section a number of points where we have decided not to include elements of the CR because they could impose more onerous burdens. For example, we have set out why we would not require Google to provide greater granularity over the different uses for which publishers can opt out content outside of general search (which would require potentially onerous changes to Google-Extended).

    (c) We have also sought to ensure the supporting transparency aspects of our remedy mitigate costs whilst maintaining the effectiveness of the remedy. Specifically: we have minimised the data scope required, are providing Google with the flexibility to use existing infrastructure and reporting frequencies and have ensured the attribution reporting requirement is practicable.

The CR is the least onerous of equally effective measures

5.7   As outlined above, we have considered crawler separation as a potentially equally effective intervention in relation to providing sufficient choice over the use of publishers' content. We received submissions on the costs of crawler separation, which we compare to the costs of publisher content controls.

5.8   We received mixed submissions on the relative costs to Google of crawler separation and publisher content controls. Google submitted that crawler separation would lead it to incur costs significantly above those presented by a new search generative AI features content control. These stem from direct upfront and ongoing costs to Google of [over £150 million] per year;[104] further

---

[104] Google said upfront costs including developing a separate index, implementing complex data pipelines, enforcing access control policies to limit which products can read the index, and defining recrawl policies, and ongoing costs of running and maintaining the separate crawler and crawler index, including crawling, compute and storage costs. We annualise costs over 5 years, assuming the upfront costs occur in year 1 and the ongoing costs are evenly distributed across years 2-5, discounting in line with the Green Book rate of 3.5%. Google's response to the CMA's RFI. Google notes that it is very difficult to accurately forecast the costs of potential

significant costs associated with Google retraining its generative models and adjusting functionality to use only data collected by the separate crawler;[105] and further opportunity costs from being diverted from undertaking work to advance its models' capability in competition with AI assistants of the order of £[✂] billion per year.[106] A third party submitted that it would not be technically difficult for Google to operate separate search and AI crawlers, arguing that Google already operating multiple crawlers undermines any efficiency arguments and gives it experience.[107]

5.9   We also received mixed submissions around the relative third-party costs of crawler separation and publisher content controls. Google identified various costs to third parties (content delivery networks, website hosts and publishers) totalling [hundreds of millions] per year when annualised and discounted, from serving content multiple times, increasing use of bandwidth and server requests from Google, and costs from a chilling effect on innovation.[108] A third party content delivery network submitted that crawler separation would address problems of server load by enabling website operators to block Google's crawler for purposes that they object to.[109]

5.10   Our analysis indicates that third parties would be likely to bear significantly greater costs under separate crawlers compared to improved publisher content controls. Separation is likely to lead to duplication of crawling for traditional search functions and generative AI features.[110] While third-party costs are relatively small per website,[111] the overall impact is significant accounting for impacts across Google's index of hundreds of billions of

---

interventions with a high degree of confidence and so Google's cost estimates across all the potential CRs should be regarded as directional only. Google's response to the CMA's RFI.

[105] For the costs of model training, Google refers to estimates of hundreds of millions of dollars [✂]. Google's response to the CMA's RFI.

[106] For opportunity costs, Google estimates [✂]; and notes that this could represent an opportunity cost of £[✂] billion. Google's response to the CMA's RFI.

[107] Cloudflare also argued that if Google were to split out its crawlers, website operators would likely see a reduction in load if they prevent crawling for certain purposes, and even if they permit all purposes are unlikely to see a significant increase in activity. Cloudflare's submission to the CMA.

[108] CMA analysis of Google's cost data. We assume Google's estimate of £[✂] billion in third party costs is evenly distributed across years 2-5 in line with Google's treatment of its own ongoing costs, and discount these costs in line with the Green Book rate of 3.5%. Google's response to the CMA's RFI.

[109] Cloudflare's submission to the CMA.

[110] Publishers that exercise improved controls can choose to opt out of crawling for one purpose but not another, providing Google with incentives to reduce any crawling it would do to use content only for the opted-out purpose. Where publishers wish to be crawled for both purposes, under a single crawler Google has no incentive to crawl multiple times; under multiple crawlers it would be forced to.

[111] Cloudflare data implies around 4% of web traffic is today attributable to Google's crawlers. Cloudflare, 'Cloudflare Radar – Bot Traffic', accessed by the CMA on 2 December 2025.

webpages.[112] Illustratively, a 30%[113, 114] increase in crawling would imply incremental ongoing data transfer costs alone to third parties of keeping Google's 100 million gigabyte index up to date would be at least £25 to £50 million annually.[115]

5.11   Further, we consider the costs to Google of publisher content controls would likely be materially lower than the costs of crawler separation.

(a)   In relation to Google's upfront costs, a new control would largely work within Google's existing system architecture: Google already provides granular controls to publishers, such as the nosnippet control which offers sub-page-level granularity. By contrast, a separate crawler would require Google to incur higher upfront costs developing new infrastructure duplicating that which already exists for Googlebot.[116] Further, Google launched Google-Extended (an entirely new control at the time) after more than [✂].[117] Whilst Google submitted that developing a new search generative AI features content control would involve "massive complexity", Google indicated that the control could

---

[112] Strategic Market Status investigation into Google's general search services: Final Decision (SMS Decision), 10 October 2025, paragraph 5.208.

[113] We calibrate our illustrative figure based on OpenAI's crawlers (GPTbot and GPT-user) which appeared to operate at around 30% of the scale of GoogleBot's in July 2025. Cloudflare, 'The Crawl-to-Click Gap', 29 August 2025, accessed by the CMA on 2 December 2025.

[114] Our analysis shows that Google in June 2025 showed a similar scale of queries showing AI Overviews ([✂]) as OpenAI responded to ChatGPT queries ([✂]); so we did not scale for usage intensity. Third-party ([✂]) data for November 2024 and July 2025 shows that Googlebot's crawling in this period was broadly in line with that before AI Overviews were rolled out at scale; so we did not adjust GoogleBot data to incorporate any reduced activity of that crawler. Sources: OpenAI's response to the CMA's RFI. CMA analysis of data submitted in Google's response to the CMA. Cloudflare's submission to the CMA.

[115] Carrying forward the crawling increase, we assume Google ingests 30% of its index monthly, in line with research showing that Googlebot finds the median new site under 30 days (some SEO research suggests this is cautious, and as a sensitivity we assume twice as much data is consumed). We assume a price of $0.12 per GB in line with GCP pricing for egress from Europe to North America (0-1TiB), assuming these costs occur in years 2 to 5, converting to GBP and discounting at the Green Book rate of 3.5%. This gives around £25m. Sources: Google, 'Organizing Information – How Google Search Works', accessed by the CMA on 2 December 2025. IndexCheckr, 'Google Indexing Study: Insight from 16 Million Pages', 28 February 2025, accessed by the CMA on 2 December 2025. Fatrank, 'How Often Does Google Crawl your Site?', accessed by the CMA on 9 December 2025. Google, 'Announcement of pricing change for egress traffic', accessed by the CMA on 2 December 2025. Bank of England, 'XUAAUSS,' accessed by the CMA on 9 December 2025.

[116] These costs potentially include developing data ingestion pipelines, and associated infrastructure. There is extensive research demonstrating the savings that can be made from reducing duplication and reusing shared infrastructure. This is captured for example in the UK Government's Technology Code of Practice, 7 July 2025.

[117] Google's response to the CMA's RFI.

potentially go live in [✀].[118] Indeed Google appears [✀].[119] Contemporaneous documents to the decision [✀].[120]

(b) In relation to Google's ongoing costs, just as for third parties, we also expect reduced crawling activity under the publisher content controls option to reduce Google's costs relative to crawler separation. Using publicly available data, we calculate illustratively that Google's electricity use from the additional web crawling would be significantly greater under crawler separation,[121] implying our chosen measure would avoid additional costs in at least the tens of millions annually.[122]

5.12 We therefore consider that crawler separation would be the more onerous proposal.[123] Since improved controls is an equally effective remedy (as described above), we do not intend to proceed with crawler separation.

---

[118] Google presentation to the CMA.

[119] Google's documents show it has been evaluating more granular controls for how Search uses content since at least April 2024, [✀]. Sources: Google internal document, 'Search (including SGE) Publisher Controls', published on Antitrust Division, U.S. Department of Justice, dated 12 April 2024, accessed by the CMA 4 December 2025. Google's internal document.

[120] A Google document from [✀]. Sources: Google's internal document. Google's internal document.

[121] CommonCrawl have published data on their emissions (320KgCO2e in a single month from electricity use) and assumptions on US grid intensity of 30KgCO2e/kWh), which imply annual electricity use of around 10,000kWh/year. We scale CommonCrawl's implied electricity use for relative crawling intensity between CommonCrawl and Google's crawlers (0.04% compared to 45% of crawl requests, based on Cloudflare data scaling CCbot figure for the denominator excluding non-AI crawling by GPTbot data presented with and without non-AI crawling) and for Googlebot gathering images, JavaScript and CSS content which is wider than CommonCrawl's focus on text (2% of page content, based on data for the median page from HTTP archive). On a highly illustrative basis, this method implies 0.15 TWh of electricity use per year, representing 0.5% of Google's overall usage. Sources: Common Crawl and Tailpipe. 'Case Study: Measuring the Carbon Cost of Crawling Five Billion Web Pages', dated 8 May 2025, accessed by the CMA 2 December 2025; Cloudflare, 'The Crawl-to-Click Gap', 29 August 2025, accessed by the CMA on 2 December 2025; HTTP Archive, 'Part III Chapter 15 - Page Weight' 11 November 2024, accessed by the CMA on 2 December 2025; Google, 'Google 2025 Environmental Report' published on Google Sustainability, dated June 2025, accessed by the CMA 2 December 2025, page 107.

[122] We proxy the cost of the volume of electricity above using UK Green Book data for the commercial sector. We consider a range between the low and high estimates given uncertainty around Google's purchasing arrangements. This gives £30-45m per year. Department for Energy Security and Net Zero, 'Green Book supplementary guidance: valuation of energy use and greenhouse gas emissions' dated 30 November 2023, accessed by the CMA on 6 January 2026, table 9.

[123] Whilst we recognise that Google's estimates of direct costs were not made with a high degree of confidence, we consider it appropriate at this stage to place weight on the broad magnitudes implied by these estimates (taking into account the factors Google accounted for in its estimates, its broad approach and the wider evidence we have set out above to sense-check the overall conclusions we have drawn). We also recognise Google made additional representations around indirect costs raised by Google in its response to the CMA's RFI; given our conclusion on the direct costs we have not needed to consider these indirect cost estimates in detail.

The CR does not produce disadvantages which are disproportionate to the aim

5.13   In this section we assess the costs and benefits of our proposed CR design. We assess these against an alternative in which Google maintains the controls and transparency which exist in the status quo, as outlined in the issues section paragraphs 1.8 to 1.16. Whilst Google has (as described in paragraph 5.11(a)) [✂]. Further, without certainty provided by compliance mechanisms, any controls introduced would not resolve issues stemming from confidence in Google's practices above. We therefore treat the costs and benefits of the proposed new control as additional.[124]

Potential costs

5.14   In this section we assess in turn the potential costs that Google, web publishers and end users may experience from our proposed Publisher CR.[125]

Costs to Google

5.15   We assess the requirements that Google develop a new search generative AI features content control and provide supporting transparency measures in turn.

5.16   Google said developing a new search generative AI features content control would be technically complex, for example because it would require direct integration with Google Search's indexing, crawling and serving systems. Google also submitted this work would lead it to incur opportunity costs, and that these would need to be considered by the CMA in its proportionality assessment alongside the costs to users and broader impact on the Search experience more generally.[126] Google said the overall costs are inherently uncertain, but less than the estimated costs of complying with a crawler

---

[124] Less any sunk costs Google has already incurred in exploring these options.
[125] We have considered the costs of interventions over a five-year period. Costs incurred in the future have been discounted using the Social Time Preference Rate of 3.5% based on guidance issued by HM Treasury in the Green Book. Annual costs are reported as the total discounted cost over the five years divided by five. Source: HM Treasury, 'Guidance – The Green Book (2022)' dated 16 May 2024, accessed by the CMA 6 January 2026.
[126] Google's response to the CMA's RFI.

separation remedy (which Google estimate to be [over £150 million] per year).[127, 128]

5.17   In relation to the publisher content controls, for the reasons set out in paragraphs 5.7 to 5.12 whilst we consider the cost of providing improved publisher content controls would still be material, we consider that publisher content controls would be substantially less costly than crawler separation over the course of the designation.[129]

5.18   Google estimated the costs of certain additional transparency measures as follows:

   (a) Providing public information on how content is used in its generative AI services would cost [up to £2.5 million] [✂].[130]

   (b) Providing publishers with additional data on user engagement with the content in search generative AI features would cost [up to £1 million] on average per year.[131]

   (c) Providing publishers with a public report on the factuality of Google's search generative AI features would cost [✂] per year. Providing publishers with information about how publisher content is attributed in Search generative AI features would cost [✂] per year. In total, these aspects could cost [up to £2.5 million] per year. [132]

5.19   There are some reasons these estimates could overstate the costs arising from our intervention: the Google's estimates include additional costs associated with identifying the set of UK-based publishers to which the factuality assessment would be restricted, implying that the costs of providing a CR tailored to UK publishers would be greater than the costs of providing

---

[127] CMA analysis of Google data, discounting at the Green Book rate of 3.5%. See Google's response to the CMA's RFI.

[128] As set out above, the scope of Google-Extended may comply with our design (see paragraph 4.19). Whilst Google would be required to provide additional information on Google-Extended than it makes available today, these costs are unlikely to be high and may to an extent overlap with the costs set out in paragraph 5.18(a) below.

[129] As set out in paragraph 5.52 below, we consider the costs would be materially lower than the benefits.

[130] Google's response to the CMA's RFI. The CMA considers these cost estimates to be on the high side relative to the nature of the activities involved, and will welcome further evidence from Google to support the magnitude of these claims.

[131] Specifically, Google estimated that this intervention would incur [✂]. See Google's response to the CMA's RFI.

[132] Specifically, Google estimated that this intervention would incur [✂]. Google also submitted that [✂]. All figures in this paragraph are discounted and annualised in line with the Green Book rate of 3.5%. See Google's response to the CMA's RFI.

this on a global basis. Since our CR design allows Google to provide the factuality report on a global basis to capture all publishers that make content available to UK users, we have excluded costs that Google would incur were it to choose to implement our CR by providing a CR tailored to the UK. Google's submissions suggest this could materially reduce the costs above.[133]

5.20   In relation to the requirement for Google to ensure their Search Content is sufficiently attributed, we consider the costs will be low. Google today provides attribution and is taking steps to improve this further.[134] We expect larger costs would only be incurred if Google were to provide insufficient attribution.[135]

Costs to web publishers

5.21   We asked a range of publishers whether they foresaw any unintended consequences or costs to their businesses as a result of improving their choices in the way proposed above. Very few publishers mentioned such costs. One publisher explicitly said it was unaware of any unintended consequences or costs.[136] Another discussed potential costs of implementing new systems to manage and monitor controls and risks of their content being downranked if they exercise the control.[137]

5.22   We consider that publishers would only incur these costs and risks if they decided to exercise Google's new controls, and therefore we take them into account as potential limitations to the benefits in our analysis below (see paragraph 5.29). Further, our CR would prevent Google intentionally downranking the websites of publishers that use the proposed control in non-generative AI search features.[138] In any event, we consider that any such costs could to a material extent be mitigated through the precise implementation and monitoring of the Publisher CR.

---

[133] See Google's response to the CMA's RFI.

[134] As set out at paragraph 1.14 we recognise that Google has taken steps to improve attribution.

[135] Further, in paragraph 4.67 we consider there could be benefits in Google providing publishers with a way to communicate their reasons for blocking content from appearing in AI-generated responses. We consider this likely to impose small costs, since it broadly aligns with current practices within the 'thumbs up/thumbs down' mechanism and because it appears possible to implement it within an existing offering (Google Search Console).

[136] [✂] response to the CMA's RFI.

[137] [✂] response to the CMA's RFI.

[138] There may be second-order effects as a result of opting out which induce downranking. For example, Google gave the example of the scenario where a publisher opts out of AI Overviews and as a result, it reduces its own opportunities to attract clicks. In doing so, this may in turn have an impact on user traffic. See Google's submission to the CMA.

Costs to consumers

5.23   Google did not identify any particular costs to users from the new search generative AI features content control applying to grounding and display.[139] Whilst we have identified some possible negative effects, we consider that the risk they are significant is small and Google has the ability and incentive to mitigate them.

5.24   Specifically, a reduction in the amount of content used to ground generative AI features could risk lowering their quality.[140] This could lead to an increase in search friction, eroding trust in generative AI services and features as a source of reliable information.[141] However, Google's index is very significant in scale, which would likely provide some robustness to opt-outs because it is more likely Google's index contains content it can substitute for the opted-out content. Insofar as opt-outs increase the risk of inaccuracies, Google's current approach would be [✂].[142] Further, in our view, Google would have the resources and be incentivised to strike deals with key publishers to maintain access to their content and allow Google to maintain the quality of their generative AI services in Search. Indeed, [✂].[143, 144]

5.25   Whilst increased pressure to provide payment for content might to some extent increase Google's costs and be passed through to end-consumers in the short run, we expect the value consumers derive from payments to publishers supporting their profitability would more than offset any such impacts.[145] As set out in paragraph 5.45 below, our Publisher CR would be likely to increase competitive pressure on Google; it may through this mechanism effect offsetting cost reductions.

---

[139] Google said the costs it raises of including training and fine-tuning in a search control would be avoided, were the control only to provide choice over display and grounding. See Google's response to the CMA.

[140] [✂].

[141] Given the scale of Google Search's usage by UK consumers, the costs to users of even a few seconds of additional search time on affected AI Overviews could be in the millions of pounds per year.

[142] Google submitted that it may [✂] if it receives feedback from users that flags that AIO Overviews generated in response to that query are 'Not factually correct'. See Google's response to the CMA's RFI.

[143] In the context of Google-Extended, Google's internal documents indicate [✂]. However, the same document notes that [✂]. Another Google internal document indicates [✂]. Source: Google's internal document. Google's internal document.

[144] Further, as explained below, our interventions could support AI assistants and other GSEs in serving a wider range of user queries. This could strengthen Google's incentives to mitigate any such frictions, and allow consumers to avoid additional search frictions by using an alternative service.

[145] Passthrough in less competitive industries (such as general search services) is typically below passthrough in more competitive industries (such as web content development).

Overall view on costs

5.26   Overall, we recognise that our Publisher CR would lead Google to incur some direct costs. We consider however that the costs of the choice aspects of our CR would be significantly less than the [over £150 million] per year in annualised and discounted terms which Google estimated for a crawler separation remedy.[146] Google estimated additionally [up to £3.5 million] per year in annualised and discounted terms for the transparency and attribution aspects of our CR.[147] Whilst there is some risk that consumers would experience additional search frictions if a significant number of key publishers restrict Google from using their content, we assess this risk to be low. As set out in in paragraphs 5.49 to 5.52 below, we consider that the costs are lower than the benefits we expect.

Potential benefits

5.27   As set out in the effectiveness section above, we expect the design of our Publisher CR would lead to publishers being able to make better decisions for their businesses, and consumers being able to make better decisions about the content they interact with. As set out in paragraph 5.13, we measure these benefits relative to an alternative in which Google maintains its existing practices and so the issues set out in paragraphs 1.8 to 1.16 persist.

5.28   In particular, the evidence suggests that our CR would:

   (a) Empower publishers to exercise existing controls: several of the publishers that responded to our RFI stated that improvements to control and transparency would enable them to have greater choice over how their content is accessed and what it is used for by Google.[148] A few publishers responding to our RFI highlighted the particular role of transparency in providing them with clearer information on which to base opt-out decisions in relation to grounding in AI-generated responses.[149] Opt-out rates for Google-Extended are below rates for other similar

---

[146] Applying discounting at the Green Book rate of 3.5%. Google's response to the CMA's RFI.

[147] Applying discounting at the Green Book rate of 3.5%. Google's response to the CMA's RFI.

[148] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[149] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

crawlers (such as GPTbot),[150] and a few publishers stated their decision to opt-in was related to uncertainty around its functioning.[151]

(b) Enable publishers to make better decisions in running their businesses, such as in content production: Most publishers responding to our RFI explained more transparency would support decisions around content (such as editorial allocation across content type).[152]

5.29 There could be some factors that limit the extent to which publishers exercise the control in practice. Whilst our control would reduce the costs of publishers opting their content out from search generative AI features, there may still be material costs of doing so. These costs could increase were search generative AI features to become more central to general search services.[153] On the other hand, a few publishers told us that they would expect to exercise the control or rely on it in negotiations with Google.[154] Furthermore, even if it were infrequently used, we would expect it to play a disciplining role on Google. The provision of transparency measures and requiring Google to attribute would also have benefits beyond enabling choice. For these reasons, we expect the CR design above allows for substantial benefits, as we evaluate further below.

5.30 We have identified three main categories of benefits which flow from these improvements. These are improved viability of web publishers' business models; benefits flowing from resulting continued consumption of web content; and increased competition in general search and adjacent activities.

---

[150] Google's internal document.

[151] [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI.

[152] Sources: [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI.

[153] Google's internal document on the impact of [✁]. The same document indicates Google [✁]. Source: Google's internal document.

[154] [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. [✁] response to the CMA's RFI. Further, at least for the foreseeable future, we also expect consumers to rely in significant part in their search journeys on organic non-generative listings, which reduces the prospect that publishers are unable to exercise the control on any content because of the need to appear in AI-generated responses. [✁] Google's data showing that [✁] of clicks stem from organic listings [✁] where AI Overviews are shown. For example, in June 2025, on queries that surfaced an AI Overview, organic links below the AI Overview received approximately [✁] million clicks, compared to links inside the AI Overview, which received [✁] million. See Google's response to the CMA's RFI.

Supporting the viability of web publishers' business models

5.31   Increased empowerment of publishers would lead to publishers producing better products and therefore increasing their revenue. Publishers advanced a range of ways in which the proposed improved choice, transparency and attribution could improve their ability to monetise their content.

    (a) Improved ability to use controls may allow publishers to improve their click-through rate by removing certain content from AI Overviews or wider generative AI features. As set out in paragraph 1.15(a) this could include opting out of AIOs where the publisher feels it is not benefiting from being included or where it is not happy with the way it is attributed. Two publishers we spoke to told us that a choice remedy would help them improve traffic,[155] including the ability to test the effect of AI Overviews on traffic and subscriptions and adjust their controls as a result.[156] For example, a publisher submitted to us that with controls they would be able to withhold higher-value content from appearing in AOIs.[157]

    (b) Better user engagement data would support publishers' strategies around content development, subscriptions and advertising revenue.[158] Several publishers mentioned other commercial benefits, including highlighting that transparency would help publishers take actions to increase the effectiveness of advertising on their content and being able to price their subscriptions in a way that maximises revenue.[159] The evidence shows that higher-quality publisher articles improve ad performance and conversions for advertisers.[160] Evidence also shows that optimisation of web pages for selection in grounding in generative AI features can lead to direct benefits for consumers such as content

---

[155] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[156] [✂] said reliable reporting would allow them to test the effect of AI Overviews on traffic and subscriptions (for example, incremental discovery vs cannibalisation) and amend their use of controls. [✂] response to the CMA's RFI.

[157] [✂] response to the CMA's RFI.

[158] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[159] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[160] Iizuka, K., Seki, Y. and Kato, M.P., 2021, October. The effect of news article quality on ad consumption. In *Proceedings of the 30th ACM International Conference on Information & Knowledge Management* (pp. 3107-3111).

diversity,[161] and that improving transparency can support publisher revenues.[162]

(c) Improving the bargaining position of publishers. Several publishers told us the CR would support their negotiating position by providing more information on the value of their content and how it is used.[163] A few publishers also said that a choice remedy would improve their bargaining position with Google relating to access to their content,[164] with one of these describing withholding consent as their central piece of leverage.[165] Evidence from a Google internal document shows [✂].[166] Indeed, Google has already entered into deals and we consider that our CR would increase Google's incentives to continue with and expand this programme.[167]

(d) Incentivising Google to improve accurate and prominent attribution through increased choice and transparency.[168] A publisher told us that improved choice in respect of search generative AI features would put pressure on Google to improve attribution, referral of traffic and the factuality of results in its AI interfaces, to encourage publishers to opt in.[169] Two publishers suggested that high quality attribution drives user engagement with publisher content, helping maintain flows of traffic from

---

[161] Ma, L., Qin, J., Xu, X. and Tan, Y., 2025. When Content is Goliath and Algorithm is David: The Style and Semantic Effects of Generative Search Engine. *arXiv preprint arXiv:2509.14436*.

[162] A study by the Incorporated Society of British Advertisers found improvements in data materially increased publisher revenues. The study shows the proportion of advertiser spend reaching publishers increasing from 51% to 65% in the period between the 2020 and 2022 as a result of increased AdTech data (e.g., matching impression data, reducing "unknown delta" of ad-spend). Whilst the data differs from precisely that made available in our intervention, it illustrates the potential value of data to publisher revenues. This principle is reflected also the inclusion of transparency elements (including audience data) in the ACCC bargaining code, which led to voluntary agreements. Sources: Incorporated Society of British Advertisers, 'UK programmatic audit suggests improvements in online advertising transparency', published on 20 April 2023, accessed by CMA 10 December 2025. Australian Competition and Consumer Commission, 'Mandatory news media bargaining code', 19 May 2020.

[163] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[164] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[165] [✂] response to the CMA's RFI.

[166] Google's internal document.

[167] For example, the deal between Google and Reddit for access to Reddit's content for use in AI models - Exclusive: Reddit in AI content licensing deal with Google | Reuters

[168] As set out in paragraphs 4.20 to 4.25 our CR compels Google to offer page level opt-outs, which increases publishers' ability to exercise choice since they can restrict use only of content most likely to be poorly attributed.

[169] See [✂] response to the CMA's RFI.

Google and from users visiting their site directly.[170] In addition, several publishers told us that prominent attribution was important for brand recognition.[171] Such attribution helps maintain the commercial viability of publisher business models, and supports social benefits associated with that web content.[172]

5.32   We expect that requiring Google to ensure sufficient attribution of Search Content (see paragraphs 4.51 to 4.69) would further increase the benefits of accurate and prominent attribution for publishers set out above. In particular, recognising the limitations set out in paragraph 5.29 on the extent to which publisher choice alone can drive benefits, we consider that an effective backstop would help assure these benefits. This was supported by some third parties. One publisher told us that prominence should be required to protect the value of a publisher's content and to ensure that users are exposed to this attribution at critical junctions.[173] Further, one association proposed conduct requirements which sought (in addition to ensuring publisher choice around how its content is used) to ensure appropriate prominence and visibility of attribution.[174]

5.33   Collectively, the mechanisms above could enable publishers to mitigate or reverse adverse click-through rate impacts on a certain proportion of queries showing AI Overviews.[175] Even if benefits were felt only across a small fraction of queries where search generative AI features are shown, the total value to publishers would be significant. This is because AI Overviews are triggered on a material number of queries: over [✂] billion searches by UK users returned an AI Overview in June 2025 alone, corresponding to [a minority] of all searches made on Google Search.[176]

---

[170] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. As set out in paragraphs 5.36 to 5.41 below, end users could benefit directly from improved incentives to attribute.
[171] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.
[172] The click-through rate of links that are attributed inside of AI Overviews has generally increased over time. For example, the click-through rate more than doubled (from [✂]% to [✂]%) between August 2024 and June 2025. This improvement in click-through rate may be partially attributable to Google's experimentation with and changes to the attribution style of AI Overviews, as discussed above. Google's response to the CMA's RFI.
[173] [✂] response to the CMA's RFI.
[174] News Media Association's submission to the CMA.
[175] As set out above in paragraph 1.15, traffic on queries showing AI Overviews sees [✂] reduced click through across the whole SERP (by around [✂]% in the UK). Google's internal document.
[176] Data indicates from June 2025 indicates [a minority] of Searches returned an AI Overview. Source: Google's response to the CMA's RFI. More recent data from Google's documents indicates that the proportion of searches made in the UK that surface an AI Overview has increased [slightly] to approximately [✂]. Google's internal document.

5.34　To illustrate the potential impact, we draw on estimates of digital revenues accruing to UK publishers. Based on data from DAMS on UK open display advertising revenues and Ofcom data on the scale of non-advertising revenue such as subscriptions, publishers' digital revenues could be around £3.5 billion per year.[177, 178, 179] An alternative methodology incorporating non-open display ad revenue based on data from the news sector implies web publisher revenues of up to £9 billion.[180] Supporting even 0.2% of these publisher revenues through the mechanisms above would imply £6 million to £17 million in revenue to UK web publishers alone each year.[181] A 0.2% revenue impact for publishers is consistent with publishers being able to offset negative impacts of AI Overviews in [✄] queries currently showing AI Overviews.[182]

5.35　We consider this calculation to be cautious. There are reasons to expect the benefits would be greater (or else that benefits of this scale would in practice be achievable were our intervention to benefit a smaller proportion of AI Overviews). Google is rolling out AI-generated responses across an increasingly broad range of queries, which implies that publishers would

---

[177] In DAMS, we estimated that UK advertisers spent around £2 billion per year on open display advertising, representing around £2.6 billion in today's prices. We also estimated that publishers receive around 65% of ad spend, implying £1.7 billion per year in undiscounted terms. This excludes exclude direct deals and other forms of advertising which also contribute significant revenue for publishers. We undertake a sensitivity to include such revenues explained below. Source: CMA, Digital Advertising Market Study, Appendix C Table C.6; CMA, Digital Advertising Market Study, Appendix R, paragraph 46; ONS, 'CPI Index', November 2025.

[178] We assume that publishers have other revenue sources in equal proportion to advertising revenues, broadly in line with data from the news sector showing subscriptions comprise 72% (Ofcom) to 89% (AOP) of advertising revenues generalises to wider web content, both sources also recognise that there are other revenue sources in addition to subscriptions. Sources: Ofcom, 'Media plurality and online news: discussion document', page 14; Source: Association of Online Publishers, 'AOP and Deloitte Data Reveals Reduction in Display Advertising Revenue for Publishers in Q1 2024', as Subscriptions Continue to Rise, 9 July 2024.

[179] A document from one large online news publisher shows that reduced traffic from search would be expected to result in a significant loss in subscriptions: [✄] internal document.

[180] Based on data from Ofcom, UK news publishers earned advertising revenues of around £592m in 2023, from all sources of digital advertising (not just open display). We express this in 2025 prices and maintain the assumption that non-advertising sources of revenue are of the same magnitude. We scale the revenue data to all web publishers, based on an estimate using Ofcom view time data that news could represent around 13% of relevant content consumed. Specifically, Ofcom's 2025 Online Nation report found that users spend an average of 10 minutes visiting an online news service, of 1 hours 17 minutes they spent online and not using large platforms. Source: Ofcom, 'Online Nation 2024', 28 November 2024, page 57; Ofcom, 'Online Nation 2025', 10 December 2025, pages 23 and 45.

[181] The low end of the range takes the lower revenue estimate above, the high end the higher estimate. We assume constant benefits over a 4-year period beginning in year 2, discounted at the Green Book rate of 3.5%. We report annualised figures using the simple average of the discounted benefit.

[182] Assuming news publisher click data is representative of all web publisher data, clicks from Google Search (excluding discover) represent only around 28% of publisher traffic. If publisher revenue is broadly proportionate to traffic, this implies supporting 0.2% of publisher revenue would require supporting around 0.7% of Google traffic. AI Overviews currently trigger on around [✄]% of queries, implying that supporting this level of revenue requires supporting around [✄]% of AI Overview query traffic. This is equivalent to avoiding a [✄]% reduction in click-through on [✄]% of AI Overviews. Source: CMA analysis of 13 responses to the CMA's RFI from: [✄].

experience greater benefits than set out above (for a constant share of AI Overview impacts avoided). Google is also testing the inclusion of summaries in Google Discover (although this feature is not currently available in the UK), a material source of traffic for news and magazine publishers, with similar impacts for the value of our CR.[183] Support of digital revenues may also help assure non-digital revenues by covering in part fixed costs, or providing indirect benefits to business dependent on the production of web content.[184] The estimates also do not include wider sources of value set out above, such as supporting new payment for content deals.

Supporting the consumption of web content

5.36   The evidence shows that risks faced by web content publishers in making an adequate return may cause content creators to invest less in their coverage.[185] This may lead to gaps in the information that web users can access. Supporting the financial viability of online content would therefore deliver benefits deriving from the consumption of a wider range of web content.

5.37   These include private benefits associated with online content consumption, represented by the value consumers place on access to content over and above what they actually pay. Drawing on academic estimates of the value of leisure time spent online on the open web,[186] were our interventions to

---

[183] Discover represents around 9% of news publisher traffic. CMA analysis of 13 responses the CMA's RFI from: [✂]; also see Google's, The Keyword, 'New AI-powered features help you connect with web content in Search and Discover', dated 13 October 2025, accessed by the CMA on 23 January 2026.

[184] For example, various forms of online content (including news and subject specific expertise) support the development of AI services and decision making by businesses. News for example is valuable in both cases because its content is generally novel and because it supports real time decisions based on current affairs.

[185] There is a volume of literature on the impact of financial pressures on news output in the UK. There is also evidence that industry players perceive generative AI features to be inhibiting content production across a wide range of content: businesses in the open display advertising sector have publicly identified risks to the volume of ad-inventory available to them. See for example Economic Insight, 'Press Sector Financial Sustainability: A report for the Department of Culture, Media and Sport,' May 2021, page 89; Ofcom, 'Media plurality and online news: discussion document,' November 2022, page 35; The Trade Desk, 'As AI Search Threatens Open Web Ad Supply DSPs face a reckoning,' 20 August 2025; and Magnite, Form 10-K, filing for the fiscal year ended December 31 2024.

[186] An academic study estimates UK consumer surplus of online content to be between £33 billion and £121 billion per year. We adjust this surplus to remove the largest online platforms, since we expect our CR would primarily support smaller and UK-specific publishers. Specifically, we multiply this consumer surplus by the fraction of online time spent outside of the 12 largest platforms derived using Ofcom analysis of Ipsos iris data (specifically, we multiply the estimated surplus by approximately 29%). We then rebase the consumer surplus figure to October 2025 prices. Some large platforms remain and could lead this illustrative calculation to overestimate benefits; on the other hand, we conservatively do not account for the increased size of the UK population and consumers spending more of their time on online-leisure activities in 2025 compared to 2011.

support the viability of 0.2% of web content viewed by UK consumers, the benefits would be around £26 million to £94 million per year.[187] Potential benefits from individual types of web content are significant in their own right: one survey indicates that on average consumers pay much less than they value their consumption of entertainment and news content.[188]

5.38   As set out in paragraph 4.57 increased choice and transparency would support consumer benefits and will support prominence and accuracy in attribution. This would facilitate users in identifying sources, and also in verifying the information or understanding its original context more easily by clicking through if they desire.[189] These impacts could help build greater trust in the information consumers are reading:[190] indeed, a BBC study found that when it comes to AI Overviews, users benefit from being able to easily verify content and expect summaries to be accurate, unbiased, and clear.[191] The impacts will allow consumers to get direct and concise answers to their queries faster which produces a consumer benefit in terms of time saved.[192] Our CR would therefore support end-consumers in capitalising on the potential for AI-generated responses to improve their search experience. Provided AI-generated responses are implemented well, this potential is material. For example, data Google presented at its I/O conference suggests AI Overviews reduced time spent per search by 15% but increased query satisfaction rates by 23%.[193]

5.39   As set out in paragraph 4.58 to 4.61, we expect some limits to the extent to which publisher choice could drive these benefits by itself. Effective

---

Sources: Pantea, Smaranda; Martens, Bertin (2014): The Value of the Internet for Consumers, Institute for Prospective Technological Studies Digital Economy Working Paper No.2014/08 p12 and p16; Ofcom, 'Online Nation Report 2025,' p23; ONS, 'CPI Index,' November 2025.

[187] We assume constant benefits over a four-year period beginning in year 2, discounted at the Green Book rate of 3.5%. We report annualised figures using the simple average of the discounted benefit

[188] The survey showed that consumers pay an average of £2.70 for news per month, but would spend up to an average of £9.90 per month. World Economic Forum, 'Understanding Value in Media: Perspectives from Consumers and Industry,' 2020, April 2020, page 28.

[189] These two specific mechanisms imply differential impacts on click-through rates to publishers; more widely we expect our CR's effect on attribution to still lead to benefits to publishers for the reasons set out in paragraph 5.31(d).

[190] The BBC submitted evidence from a study they conducted which showed that attribution strengthens trust and makes the link between the content and the publisher visible at the moment of reading. BBC's response to the CMA's RFI.

[191] BBC's response to the CMA's RFI.

[192] If greater trust leads to even a small reduction in the time spent searching, the cumulative effect across a large number of queries can generate a substantial overall benefit.

[193] Boston Institute of Analytics. 2025. "Will Google AI Overviews Kill Traditional SEO? An In-Depth Analysis (citing HubSpot, 2024), Boston Institute of Analytics, July 29, 2025.

application of our CR's obligation in paragraph 6 of the CR for sufficient attribution would underscore the benefits set out in the paragraph above.

5.40 There are also social benefits, where consumption of web content would lead to benefits beyond those consumers account for. For example, these could include news content (Supporting a better-informed public), health content (reducing costs to the taxpayer and/or avoidable illness spread),[194] and open-access educational resources (increasing economic productivity through knowledge spillover effects).[195]

5.41 Given the importance of these issues, and the scale of content affected, these benefits would likely be significant. However, they are harder to quantify. Taking news alone, a few publishers told us there are significant wider societal benefits of ensuring continued viability of a well-functioning and diverse news media.[196] Evidence from public inquiries has identified that news plays a vital role in informing the public, enabling better decision-making, and enabling active civic engagement.[197, 198] Studies show consuming news boosts knowledge, belief accuracy and trust,[199] and there is some evidence that consumption of news through traditional media outlets (rather than digital intermediaries) leads to measurable further benefits.[200]

Supporting contestability of general search and adjacent activities

5.42 Our CR could have benefits for the competitive outcomes in adjacent activities (such as AI assistants) and general search services. We discuss each in turn.

---

[194] Research suggests that online health information assists people in making health related decisions such as improved adherence to the advice of a physician and improved self-care. Source: Thapa DK, Visentin DC, Kornhaber R, West S, Cleary M., 'The influence of online health information on health decisions: A systematic review. Patient Educ Couns', 2021 Apr;104(4):770-784.

[195] Research by Khan Academy indicates online educational resources provide significant educational benefits to individuals. Source: Khan Academy, 'Research on how Khan Academy drives learning outcomes', accessed by the CMA on 1 December 2025.

[196] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. For example, one publisher [✂] said that intervention could reduce the risk of reaching a point where Google refers no traffic at all to news publishers. This publisher explained that many publishers, especially smaller ones, do not receive substantial direct traffic and there is a serious risk they will simply collapse absent intervention. This would have consequences for media diversity and plurality, which in turn would be hugely damaging to democracy.

[197] House of Lords Communications and Digital Committee: 1st Report of Session 2024-5, 'The future of news,' page 9.

[198] Australian Competition and Consumer Commission, 'Mandatory news media bargaining code', 19 May 2020.

[199] Altay, S., Hoes, E. & Wojcieszak, M., 'Following news on social media boosts knowledge, belief accuracy and trust' *Nat Hum Behav* 9, 1833–1842 (2025).

[200] Ofcom, 'News Consumption in the UK: 2024,' page 10; Ofcom, Media plurality and online news, 'Annex 3 survey analysis: news consumption habits and media plurality outcomes', 16 November 2023, page 1.

5.43   The evidence suggests concerns that Google's position in general search and its design of its publisher content controls may provide its AI assistant services with advantages over its competitors.

(a) As set out in paragraph 1.06 the evidence shows that Google's position in general search services implies publisher websites have no realistic alternative but to allow Google to crawl their content for search. As set out in paragraphs 1.11 to 1.13 Google's design of its controls within general search and limited transparency around the functioning of Google-Extended imply publishers are in practice compelled (due to concerns about Search ranking) to allow Google to use their content for generative AI services and features.

(b) For these reasons, AI firms suggested Google has an unfair competitive advantage. One firm said that any AI model developer with preferential access to publisher content may have an advantage over other model developers, potentially resulting in reduced competition in AI models,[201] and [✂].[202, 203]

(c) This is consistent with third party data, which suggests publishers opt their Search Content (in full or in part) out of appearing in Google's wider generative AI services through Google-Extended materially less frequently than they do for rival AI assistant services.[204]

5.44   The improved control provided by our proposed CR could help address concerns about these advantages in adjacent activities, such as AI assistant services.

(a) A few AI assistant firms' submissions suggested that competition in wider AI services may be negatively impacted by Google's preferential access to publishers' websites by virtue of the must have nature of being indexed for Google Search.[205]

---

[201] [✂] response the CMA's RFI.
[202] [✂]. [✂] submission to the CMA.
[203] OpenAI [✂] referenced testimony at the US Google Search hearing that Google's Search Index is an important input to training Google's generative models, including Gemini. Open AI's submission to the CMA.
[204] Cloudflare's submission to the CMA. Cloudflare's data also shows that publishers permit Google's use of their data in Google Search much more frequently than in other services.
[205] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂]response to the CMA's RFI.

(b) Several publisher submissions highlighted concerns around Google's ability to bundle crawling for AI services with crawling for search.[206] These submissions highlighted the importance of clarity that would be provided through our CR on the impact of exercising Google-Extended, and the positive impact that reducing barriers to exercising the control would have on competition.[207]

5.45 Within general search services, Google's ability to incorporate search generative AI features directly into its existing products contrasts with AI assistants which must encourage users to switch to their products.[208] Models powering Search AI features derive from Google's Gemini foundation model family.[209] Google's access to publisher content often provided for the purposes of appearing in traditional search links to train its models reduces its costs of Search product development, relative to those its competitors would all else equal be able to achieve (for similar reasons as set out above). Google's current publisher content controls may therefore have the effect of maintaining its market position in general search services, reducing the prospect that AI assistants develop into a more sustained and significant competitive constraint.[210]

5.46 Our CR's effect of clarifying the role of Google-Extended and ensuring a product agnostic scope could reduce the extent to which Google can offer features that other general search engines and AI assistants are unable to match. Insofar as our CR would lead to improved payment terms to publishers or investment in improved attribution in search generative AI features (through such mechanisms as set out in paragraph 5.30(d) and 5.31(d), this would reduce the extent to which Google faces cost advantages derived from its SMS position in general search services over an important input when compared to rivals.

---

[206] [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI. [✂] response to the CMA's RFI.

[207] One publisher [✂] said this bundling allowed Google to 'effectively [compel] publishers to provide content to fuel its AI products'. [✂] response to the CMA's RFI. Another [✂] told us that with Google's market power, the risk of unintended consequences from applying controls is too great for them to take. [✂] response to the CMA's RFI.

[208] Strategic Market Status investigation into Google's general search services: Final Decision (SMS Decision), 10 October 2025, paragraph 5.262.

[209] Google's submission to the CMA.

[210] Half of consumers are actively seeking out AI-powered search engines according to October 2025 McKinsey & Co consumer polling (Source: Winning in the age of AI search | McKinsey). Also see CMA, Strategic market status investigation into Google's general search services Final Decision, 10 October 2025, paragraph 5.262.

5.47   Even if there would only be small benefits to competition in these activities, the impact would be substantial given their widespread use and role in facilitating everyday activities.[211, 212] End users would benefit from improved choice and higher quality or lower cost services from supporting competition in AI assistant development, and improved attribution.[213]

Overall view on benefits

5.48   Overall, we consider that our CR would lead to substantial benefits over and above how we expect the status quo to evolve, through supporting web publisher business models, supporting consumption of web content by end users; and supporting competition in general search and adjacent activities (with benefits for end users of these services).

Comparing the costs and the benefits

5.49   In this section we bring together the evidence on the costs and benefits and consider their relative size.[214]

5.50   The analysis above indicates substantial benefits of the proposed CR. Improved publisher content controls (enabled by transparency measures facilitating their use) would support web publishers' financial viability and incentives to produce web content. Given the scale of web content consumption, the issues facing web content developers and the significant volume of qualitative evidence linking our intervention to the mitigation of these issues, we consider these benefits significant. To support our evaluation of their approximate order of magnitude, in paragraph 5.34, we set out an illustrative calculation suggesting annualised benefits of the order of £6 million

---

[211] As set in the CMA's final SMS decision, Google has had a market share greater than 90% in general search services for several years and also has tens of millions of UK users. See Strategic Market Status investigation into Google's general search services: Final Decision (SMS Decision), 10 October 2025, paragraphs 5.26 and 5.276.

[212] In AI assistants, high levels of capital expenditure by Google and its competitors supporting generative AI features and services indicates significant expected returns to investments. Google have estimated their capital expenditure to be approximately $91-93 billion, a figure that has been twice revised upwards this year. See AI's capital expenditure shows no sign of cooling, published 9 November 2025. OpenAI have publicised they expected their annual recurring revenue of $20 billion for 2025, Sam Altman says OpenAI has $20B ARR and about $1.4 trillion in data center commitments | TechCrunch.

[213] As set out above, publishers benefit from improved brand reputation. End-consumers benefit from more confidence in answers to their queries with lessened to click through multiple websites. Results from a Google experiment indicate that AI Overviews materially reduce search friction (measured by the number of manual refinements made by users). [✄]. See Google's internal document.

[214] We have not identified relevant notable impacts of this intervention for people with protected characteristics.

to £17 million.[215] Assuring publisher viability would preserve direct consumers benefits from consuming that content: in paragraph 5.37, a similar illustrative calculation suggests annualised benefits of the order of £26 million to £94 million.[216] As set out in paragraph 5.35 and 5.41, the calculations are cautious and we expect additional significant wider benefits flowing from increased consumer trust in web content and from supporting the consumption of socially valuable content such as news. Further, as set out in paragraphs 5.42 to 5.47, our Publisher CR could also support competition in general search and adjacent activities (such AI assistants), where even small impacts could be significant.

5.51 The proposed Publisher CR would also lead to some costs. Whilst the precise magnitude remains uncertain, the largest cost to Google is the design and implementing the new search generative AI features control for publishers. The upfront costs of this are at most £[✂] million in annualised terms; whilst likely to be substantially lower (for the reasons set out in paragraph 5.11(a)), at this stage we take a cautious approach and assume this full amount.[217] The ongoing costs are unknown, and we assume they are comparable to the transparency measures.[218] Based on the costings set out in paragraph 5.26, Google could incur annualised costs of [up to £3.5 million] from supporting transparency measures.[219] Altogether, we expect a reasonable upper bound to the direct costs of our CR to be [up to £40] million per year.[220] Whilst we would not expect publishers to opt out a significant fraction of web content from AI-generated responses in Search or wider AI services (and consider Google to have strong incentives to address the risks of significant opt outs),

---

[215] We assume benefits in years 2-5, annualise over 4 years and discount at the Green Book rate of 3.5%.

[216] As above, we assume benefits in years 2-5, annualise over 4 years and discount at the Green Book rate of 3.5%.

[217] As set out in paragraphs 5.7 to 5.12, Google submitted that the cost of building a new search generative AI features content control would be less than the costs of complying with a crawler separation remedy, which Google estimated to be [at least £150 million] per year (in annualised and discounted terms). As set out in paragraph 5.17, we consider publisher content controls will not incur such significant upfront costs. We assume all upfront costs are incurred in year 1, but annualise the costs over 5 years for comparability with ongoing costs (and there is no need therefore to discount). Source: Google's response to the CMA's RFI. Google's response to the CMA's RFI.

[218] As set out in paragraph 5.11, we consider crawler control costs will incur substantially lower ongoing costs than for publisher content controls and (in each year at least) lower also than the upfront costs. The high end of Google's transparency estimates for years 2 to 5 of the transparency and attribution measures imply costs of [up to £1.5] million when discounted and annualised over a 4 year period, using the Green Book discount rate of 3.5%). Source: Google's response to the CMA's RFI.

[219] Figures are annualised over 5 years and discounted at the Green Book rate of 3.5%.

[220] As above, figures are annualised over 5 years and discounted at the Green Book rate of 3.5%.

we take into account at least some risk that consumers may incur some indirect costs through additional search frictions.

5.52 Given the nature and breadth of the benefits we identify above, we expect the benefits to outweigh the costs. In particular, we have taken account of the substantial benefits consumers derive from web content, the evidence linking our proposed CR to the preservation of these benefits, and the scale of the wider activities in which we expect our intervention to support competition. Whilst our quantification is illustrative, taking the low end of both benefit calculations summarised in paragraphs 5.34 and 5.37 gives a range of £32 million to £111 million per year.[221] Although the cost estimate in paragraph 5.51 falls within that range, it does so at the lower end.[222] Further, the cost estimates represent a reasonable upper bound, whilst the benefit estimates are cautiously estimated and exclude significant sources of benefits, such as social value derived from consumption of web content. Overall therefore we are satisfied that the CR would not produce disadvantages disproportionate to the aim.

Provisional conclusion on proportionality

5.53 Our overall provisional assessment is that the proposed Publisher CR is effective and proportionate. We will continue to refine our understanding of the costs and benefits of the measure as part of the final proportionality assessment, to inform our decision about whether to impose this CR.

---

[221] As above, cost figures are annualised over 5 years and discounted at the Green Book rate of 3.5%. Benefit figures are annualised over 4 years and discounted at the Green Book rate of 3.5%.
[222] Over a 5 year horizon, so comparing costs and benefits over an equal length of time, the cost estimate still falls below the midpoint of the estimated range.

6.   Questions for consultation

6.1   We welcome views on any aspect of the Publisher CR design or analysis set out above. We are particularly interested in stakeholder feedback on the following questions.

6.2   In relation to ensuring publishers have sufficient choice:

    (a) As noted in paragraph 4.13, we would welcome further evidence on the benefits and risks of Google providing separate controls over training and grounding outside of general search.

    (b) As noted in paragraph 4.25, we would like to receive further evidence on the benefits and risks of Google providing page-level controls outside of general search (ie for Google-Extended).

6.3   In relation to greater transparency for publishers:

    (a) As noted in paragraph 4.45, we would like to receive further evidence on the benefits and risks of Google providing performance and engagement information on a 'per-feature' basis within general search.

    (b) As noted in paragraph 4.50, we would welcome views on what would be the most effective way(s) for Google to provide publishers with information to enable them to understand the quality of clicks referred from search generative AI features.

6.4   In relation to attribution:

    (a) As noted in paragraph 4.67, we welcome views on whether a mechanism for publishers to more easily communicate the reasons for blocking content from appearing in Google's search generative AI features would enhance the effectiveness of our proposed Publisher CR whilst ensuring it remains proportionate.

    (b) We welcome further examples of information and metrics that help explain how Search Content is attributed in search generative AI features and the factuality of those features, and views on how these data would be best disseminated.

    (c) We welcome further views on the extent to which our proposed Publisher CR 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.