

AI Search Features and Quality Degradation*

Kevin Mi[†]

September 28, 2026

Abstract

Search platforms, such as Google and Bing, increasingly answer queries with artificial intelligence (AI) search features that scrape and synthesize content from online publishers. In this paper, I develop a theoretical model to examine how the introduction of AI search features affects content investment decisions, platform design choices, and welfare. Improvement in AI content quality enhances users' direct benefit from the search platform but can also divert users away from publishers' sites. Consequently, it erodes the publishers' incentives to invest in high-quality content and, in turn, reduces the quality of the AI content. In the face of this tradeoff, I show that introducing an AI search feature improves user surplus but may reduce social welfare if publishers sufficiently value user clicks. Furthermore, I analyze how this tradeoff shapes the platform's choice of AI search quality. Finally, I apply the model to evaluate several proposed policies for regulating AI search features. I find that copyright restrictions can implement the socially optimal outcome; platform competition and pay-per-crawl payments may reduce social welfare; opt-out options have ambiguous effects on welfare; and licensing deals can improve welfare, although the optimal form of licensing depends on market conditions.

Keywords: search engines, quality degradation, AI overviews, AI summaries

JEL Codes: D21, L15, L82

*I would like to thank Jay Pil Choi, Arijit Mukherjee, and Hanzhe Zhang for their guidance and support throughout the development of this project. I also thank Jon Eguia and various audiences at Michigan State University and the 90th Midwest Economics Association meeting. All mistakes are my own.

[†]Department of Economics, Michigan State University, E-mail: kwmi@msu.edu

1 Introduction

Search platforms, e.g., search engines like Google and Bing, play a crucial role in the digital economy by processing billions of queries every day and transforming an overwhelming mass of online information into an organized and accessible resource. Due to recent advances in generative artificial intelligence (AI) and large language model (LLM) technologies, search platforms are increasingly utilizing AI-powered search features (e.g., Google with AI Overviews/Gemini and Bing with Microsoft Copilot).¹ These AI search features synthesize relevant information across the web into AI content, and within the search platform’s interface, such content is typically placed before other webpage links and advertisements.

Because online content creators heavily rely on search platforms for user traffic, the introduction of AI search features has led to two key problems, prompting legal action and scrutiny from competition authorities. First, AI search results are displayed before the organic and paid results, which reduces users’ need to click on original sources.² This is commonly referred to as the “zero-click phenomenon.” In 2025 alone, Penske Media, the parent company of *Rolling Stone*, *Billboard*, and *Variety*, along with Chegg filed lawsuits against Google alleging that the roll-out of AI Overviews has removed the users’ need to visit their websites.

The other key problem with AI models is copyright infringement. Specifically, online content creators and large news publishers are concerned that their copyright protected content is used as training data for generative AI technologies. In 2023, the *New York Times* filed a lawsuit against OpenAI alleging that OpenAI violated copyright laws by training their AI models (e.g. ChatGPT) on millions of their articles. A point of contention for publishers is that when ChatGPT was asked questions regarding topics covered in *Times* articles, ChatGPT could reproduce the articles verbatim.³

Thus, the introduction of AI features on search platforms presents a novel tradeoff: it may enhance users’ benefits from the search platform but divert users away from publishers’ sites, weakening publishers’ incentives to invest in content creation and, in turn, reducing the quality of the AI content itself.

In this paper, I provide a tractable model of a platform with AI search features that captures this tradeoff and examine how these features affect competition between online content creators and their investment in high quality content. I use this framework to deter-

¹According to a report from McKinsey & Company (2025), about 50 percent of Google searches already contain AI summaries and, by 2028, AI search features will result in \$750 billion in US revenue.

²According to data from SimilarWeb, *Forbes* and *HuffPost* each lost 40% of their user traffic, and CNN.com saw a 28% drop in user traffic about a year after the introduction of AI search results.

³However, OpenAI argues that scraping data from the *Times* is protected under “fair use,” as the service is not meant to regurgitate documents but to facilitate research or teaching.

mine the welfare implications of the AI search features, and examine a set of widely debated policy responses, including (i) enforcing fair use copyright policies, (ii) providing publishers with opt-out options, (iii) facilitating licensing deals between platforms and publishers, (iv) implementing pay-per-crawl payments, and (v) promoting platform competition. For each type of response, I explore the tradeoff it creates and determine its welfare consequences.

In the baseline model without AI search features, a monopoly search platform intermediates content consumption between users and two publishers. The publishers compete à la Hotelling and choose how much to invest in content quality. Thus, publishers are horizontally differentiated by how well their content matches users' search queries and vertically differentiated through their quality choices. Crucially, the platform has an informational advantage regarding users' locations on the line. Users do not know their locations and must rely on the platform to allocate them to a publisher. Given this informational advantage, the platform selects an allocation threshold that directs users towards a publisher.

With the incorporation of AI search features, the search platform possesses two additional strategic choices: a mixing rule and a quality-conversion ability. The AI's mixing rule determines how each publisher's content is combined to create each user's AI content. The AI's quality-conversion ability captures how efficiently the AI synthesizes publisher content, and can be affected through external government policies such as scraping barriers or copyright restrictions.⁴ In addition, with AI search features, each user first reads the AI content and then decides whether to continue searching and consume their allocated publisher's content.

I find that the inclusion of AI search features leads to three separate effects. First, reading the AI content provides users with a direct benefit without requiring any additional search. Second, reading the AI content provides users with information of their preferences on the Hotelling line, improving their match quality with publishers. Finally, publishers degrade their content quality as the platform's quality-conversion ability increases, because investing in content quality also bolsters the AI's content quality and induces more users to substitute away from clicking on organic content. Therefore, from the platform's perspective, improving the AI conversion ability increases the direct benefit each user receives but lowers publishers' content quality. This degradation, in turn, indirectly reduces the AI content quality. As a result, I show that the platform should optimally restrain its AI conversion ability.

I also find that although user surplus is higher with AI search features, social welfare may be lower. In fact, I establish that the socially-optimal quality-conversion ability cannot exceed the user-optimal quality-conversion ability. This implies that policies restricting the

⁴In my model, this is also a strategic choice from the platform. For example, Google's AI Overviews uses a lighter version of Gemini and OpenAI can choose between GPT-4o-mini and GPT-4o depending on the user's plan and search query.

platform’s quality-conversion ability, such as an ex-post fair use copyright policy, can be socially beneficial and should be a primary target for policymakers.

If such a policy cannot be implemented, other responses may be considered, but they may be harder to implement, depend on market conditions, or lead to lower social welfare. For example, under an interior symmetric equilibrium with platform competition, platforms inflate their quality-conversion ability above the user-optimal level to capture more users. Therefore, user surplus and social welfare are lower with competition, as the feedback effect that degrades content quality in the baseline model is amplified.

Providing publishers with the option to opt out of AI web scraping leads both publishers to fully block AI scraping, which may reduce user surplus and social welfare when search costs are relatively high. Taking this into account, I then allow the platform to pay publishers to regain access to their content. I focus on two types of licensing: lump-sum and quality-adjusted. I show that quality-adjusted licensing partially solves the quality-degradation problem, but lump-sum licensing is preferred by the platform when the market is very competitive. I also consider pay-per-crawl payments where the platform must pay publishers each time it scrapes one of their pages. Under pay-per-crawl payments, the platform distorts its mixing rule by creating some users’ AI content using only one publisher, magnifying the quality degradation problem and harming users.

The rest of the paper proceeds as follows. Section 2 reviews relevant literature. Section 3 introduces the baseline model without AI search features and presents equilibrium results. Section 4 extends the baseline by incorporating and analyzing the model with AI search features, and then derives welfare consequences. Section 5 analyzes several responses to AI search features and Section 6 extends the baseline model to account for platform competition, providing several policy implications discussed in Section 7. Section 8 concludes. All proofs are found in Appendix A, additional figures are provided in Appendix B, and model robustness is shown in Appendix C.

2 Related Literature

The literature on the effects of AI search features on content investment is quite limited.⁵ Chan (2026) theoretically examines how AI search results may improve users’ search experi-

⁵Agarwal and Sen (2026) provide experimental evidence that AI Overviews divert traffic away from publishers without improving the users’ search experience. Zhao and Berman (2026) provide empirical evidence that many publishers block AI web crawlers which adversely affects their traffic and that content becomes richer as a result of AI search features.

ence while diverting traffic away from content producers.⁶ While he treats AI content quality as exogenous, the AI content quality in my model is endogenously determined by publishers’ content quality. Therefore, even though he also finds that AI search degrades content quality, this paper yields a mechanism absent from his paper: the platform degrades its own AI content quality to preserve publishers’ investment incentives. Lyu and Chen (2026) also study a platform-design tradeoff generated by cannibalization of sponsored search revenue. However, the tradeoff highlighted in my model arises on the supply side: stronger AI lowers the return to publisher investment and thereby degrades the external content inputs used to produce AI answers.

This paper is also related to several papers on news aggregators. Empirical studies on news aggregators include Athey and Mobius (2012), Chiou and Tucker (2017), Calzada and Gil (2020), and Athey et al. (2021). This paper is most relevant to the theoretical studies including Dellarocas et al. (2013) and Jeon and Nasr (2016).⁷ Jeon and Nasr (2016) find that the effect of news aggregators is ambiguous as it consists of two distinct effects: business-stealing and readership-expansion effects. The business-stealing effect represents how the aggregator steals homepage attention from publishers and the readership-expansion effect represents how readership for higher-quality articles increases even though the newspapers’ user-base may remain constant. A key difference between news aggregators and AI search features is that aggregators act as a referral technology directing users to high quality articles while AI search features act as a content technology. Also, my model endogenizes the platform’s strategic choice on the effectiveness of its AI search feature, while the news aggregator in Jeon and Nasr (2016) is purely mechanical.⁸

This paper is also closely related to the literature on search engines. In particular, my model specification is most similar to that of de Cornière and Taylor (2014). Their model is used to analyze vertical integration between a publisher and how it may bias the allocation threshold. Instead, my model is more focused on the competitive effects of AI search features. Taylor (2013) models search engine competition where each search engine provides an advertisement and organic link. He highlights a novel trade-off where providing high quality organic results may lead to revenue loss from advertisements. Thus, search engines are incentivized to deliberately degrade their organic result quality. A similar aspect exists in my model.

This paper naturally builds on the nascent literature on AI content generation and its

⁶Wu et al. (2026) focus on citation and monetary compensation mechanisms to sustain creator incentives. I instead endogenize the platform’s AI design itself to determine how it affects creator incentives.

⁷Jeon (2018) provides a comprehensive survey on the economics of news aggregators.

⁸Other theoretical papers on news aggregators have focused on its impact on versioning decisions (Calzada and Ordóñez-Calafi, 2012) and multi-homing (George and Hogendorn, 2012).

impact on copyright policies. Gans (2024b) models how AI content generation affects user contributions on query-and-answer platforms. He shows that AI answers substitute for answers to easier queries, which gives contributors more incentive to answer more complex and valuable questions. Gans (2024a) analyzes how various copyright regimes on generative AI affect content quality. My model differs from these two works as I introduce competition among content providers in a search platform setting which leads to different effects from AI content generation.

More generally, this paper contributes to the literature on two-sided markets (Caillaud and Jullien, 2003, Anderson and Coate, 2005, Rochet and Tirole, 2006, Armstrong, 2006, Teh and Wright (2024)). Two-sided markets are typically characterized as industries where platforms intermediate transactions or services between two groups. With regard to the media industry, the two sides are often readers and advertisers (Gabszewicz et al., 2001, Anderson and Coate, 2005, Armstrong, 2006, Peitz and Valletti, 2008, Crampes et al., 2009, Anderson and Gans, 2011, Athey et al., 2018, Anderson et al., 2018, Choi and Jeon, 2023, among others). In the baseline environment, I do not explicitly model the advertising side. Rather, I model each publisher’s objectives with a reduced-form to focus more on the strategic interactions in quality choice (similar to Gabszewicz et al., 2001).

Lastly, this paper touches on the literature on licensing contracts (Kamien and Tauman, 1986, Katz and Shapiro, 1986, Bernheim and Whinston, 1998, Sen, 2005). A key takeaway from recent works is that the optimal type of licensing deal, such as exclusive deals or payment type, depends on market conditions. Regarding AI search features, I find that the platform’s preferred licensing structure shifts with market conditions: lump-sum payments that effectively yield singular access can be optimal when publishers are fierce substitutes, whereas quality-adjusted payments for broad access are preferred when publishers are highly differentiated and search costs are high.

3 Baseline model without AI search features

3.1 An Illustrative Example

To better understand how the model will translate to the user search experience without and with AI search features, consider the following illustrative example of a user who has just traveled to Lansing, Michigan and searches “what are things I can do around me” on the search platform. Suppose the user knows they would prefer arts and culture activities over outdoor activities. See Figures 8 and 9 in Appendix B for visual reference.

Example 1a (without AI). The user sees a link directing them to a Lansing-specific

tourist site and must determine whether they want to click on the link to access the page.

Example 1b (with AI). The user sees a block of information that provides general information of arts and culture activities that is clearly created using the Lansing-specific tourist site and outdoor activities that is clearly created using Tripadvisor. The user decides that they would rather see an art museum and must determine whether they want to click on the associated link to learn more.

Even though the user knows their preference between arts and culture activities compared to outdoor activities, they are unaware of how well each page embodies their preference, i.e. users do not know how far away each publisher is on the Hotelling line. Thus, users must rely on the platform to allocate them to the appropriate publisher. As the platform has aggregated data across millions of similar queries, it effectively knows the mapping of user tastes to publishers. Without AI, the less informed user makes their search decision with limited knowledge on how they will enjoy the content in the Lansing-specific page. In contrast, when the user searches with AI search features, the user can better judge how helpful the Lansing-specific page will be after reading the AI search result and, thereby, make a more informed search decision. However, it is also plausible that the user will be fully satisfied with the AI search result’s information and find clicking on the page link a waste of their time and effort.

3.2 The search environment

I consider a market consisting of two publishers on a monopolistic search platform and model their competition for users in a given search query. As much of the backlash on AI search features is related to news and media publishers, such as the *New York Times* and the *Rolling Stone*, I use the term “publisher” throughout the rest of the paper to represent any online content creator.

Users may only directly access one of the two publishers’ sites, i.e. users are single-homing with respect to the publishers’ content. The two publishers are horizontally differentiated and located at each end of a Hotelling segment of length one. Let l_i denote the location of publisher $i \in \{1, 2\}$ where $l_1 = 0$ and $l_2 = 1$. A measure one of users are uniformly distributed along the line. Here, the users’ location represents their implicit preferences for a publisher which may capture political/ideological biases or styling tastes (Mullainathan and Shleifer, 2005, and Gentzkow and Shapiro, 2011). If a user located at $x \in [0, 1]$ consumes publisher i ’s content, then they incur a transportation cost of $t|x - l_i|$, where $t > 0$, due to imperfect preference matching.

Publishers are also vertically differentiated based on their content quality. For publisher i , producing content of quality q_i comes at a cost $C(q_i) = q_i^2/2$. Users derive utility q_i from reading publisher i 's content. They also receive additional utility $v > t$ when accessing a publisher's site which captures additional benefits that are not related to content.⁹

Similar to de Cornière and Taylor (2014), the market is characterized by incomplete information. Users know their own content preferences, but before receiving information from the search platform they do not know how well those preferences match the content offered by each publisher. I represent this informational friction by assuming that users do not initially know their location x on the Hotelling line. The platform, which observes information from users' browsing and search histories, knows this mapping.

Definition 1 *The search platform's allocation threshold is a threshold $\bar{x} \in [0, 1]$, where all users $x < \bar{x}$ are allocated to publisher 1's site and all other users are allocated to publisher 2's site.*

The search platform can set its allocation threshold $\bar{x} \in [0, 1]$ such that it perfectly controls user choice over publishers.¹⁰ Users who access a publisher's site incur a search cost $\theta \sim U[0, \bar{\theta}]$ where $\bar{\theta} > v$. Let F and f denote the associated cdf and pdf, respectively. I provide sufficient conditions for a more general distribution that would give us similar qualitative results in Appendix C.1. Here, θ is a reduced-form friction capturing the time and attention costs of engaging with the results page and arriving at a publisher's page (e.g. locating the result, clicking, and page-load delay).

In summary, users must rely on the search platform to allocate them to a publisher where the utility that a user located at x receives from consuming publisher i 's content after continuing to search is given by

$$u_i(x) - \theta = \underbrace{q_i - t|l_i - x|}_{\text{value of content}} + \underbrace{v}_{\text{additional benefit}} - \underbrace{\theta}_{\text{search cost}} \quad (1)$$

where $u_i(x) := q_i - t|l_i - x| + v$ denotes the utility user x would receive from reading publisher i 's content.

⁹The lower bound on v serves as a full market coverage assumption ensuring that there exists an interior solution. I discuss examples on how publishers can alter v in Section 7.

¹⁰This assumption is consistent with Edelman and Lai (2013) who show that minor changes on the layout of search engines have a large impact on user traffic.

3.3 Player objectives

Publishers earn a per-user return of $r > 0$. While in practice this may be a proxy for net margins from hosting advertisements, selling data, or subscriptions, for the purposes of aggregate welfare analysis, r is best interpreted as the net social benefit or intrinsic value generated when a user consumes the publisher’s content. For example, those who answer coding questions on StackExchange do not earn any revenue; they instead value how much their answer contributes to the community.

I assume that the search platform’s objective is aligned with the users, i.e. the platform maximizes aggregate user surplus. Specifically, the platform selects its strategic instrument(s) to maximize an increasing and differentiable function $\Phi : \mathbb{R}_+ \rightarrow \mathbb{R}_{++}$ of expected user surplus US . In other words, the search platform does not receive any monetary transfers between users, publishers, or outside parties.

I provide two possible justifications for this. First, user surplus serves as a proxy for interactions on the platform which result in richer data. This data can be used by the platform to gain more revenue through other subsidiaries (e.g. Google uses user data to provide more suitable advertisements on YouTube) or the data can simply be sold to other companies. Second, this can also be viewed as a reduced-form way to capture the benefits of growing the platform’s user base including when competing with other search platforms.¹¹ For example, consider a simple microfoundation where information needs for each user arrive according to a Poisson rate $\lambda > 0$. For each information need h , let $\epsilon_h \sim G$ on support $[0, \infty)$ represent the utility of alternatives to searching on the platform. This could capture using a different search engine, visiting a site directly, asking another person, or simply not searching. As a result, the expected number of queries made on the platform is $\lambda G(US)$ and assuming the platform earns expected net margin $m > 0$ for each user search, the platform’s expected profit from a user is $\Phi(US) = m\lambda G(US)$.

In fact, if the platform’s objective function were to instead rely purely on publishers’ user traffic, integrating a substitute AI search feature would not yield a strict benefit. (This will be seen in Section 4.) Therefore, I conjecture that the platform’s main revenue stream from integrating AI search results does not arise from transfers related to organic clicks.

3.4 Timing and discussion of assumptions

The timing of the game proceeds as follows:

- Stage 1: Each publisher i chooses q_i and the platform selects $\bar{x}$ simultaneously.

¹¹A similar argument is seen in Johnson et al. (2023).




Figure 1: Market structure

- Stage 2: Users observe q_1 , q_2 , $\bar{x}$, and their allocated publisher. They decide whether to access their allocated publisher's website by incurring their idiosyncratic search cost.

The assumption that users observe q_1 and q_2 allows demand for publishers' content to change with regard to how much they invest in quality.¹² If I instead assumed users had no prior knowledge of the content on the platform, demand would be fully determined by their prior expectations. The observability of allocation threshold $\bar{x}$ by users can be interpreted as users being aware how well a search engine performs when providing their results page, either through media or prior personal experiences.

The simultaneity of quality choice and selecting the allocation threshold is consistent with the fact that publishers cannot predict how their investment in quality will impact how the platform ranks content for each user. For robustness, I analyze the case where $\bar{x}$ is chosen after the content qualities in Appendix C.3. I assume all users are already on the monopoly platform, so $\bar{x}$ does not affect the platform's user base. I endogenize user's choice of platform when introducing platform competition in Section 6.

3.5 Equilibrium without AI search features

I use (weak) Perfect Bayesian equilibrium as the solution concept. Suppose the search platform has chosen allocation threshold $\bar{x}$ and the publisher qualities are q_1 and q_2 , respectively. Let $Q := (q_1, q_2)$ denote the set of qualities and $\mathcal{X}_i$ be the set of users allocated to publisher i by the platform's allocation threshold.

¹²It may also reflect how consumers learn the quality of content publishers produce over time, perhaps through a repeated game.

Given that the user knows they are allocated to publisher i , the expected utility of accessing publisher i 's webpage and reading the content without the search cost is

$$\rho_i(\bar{x}, Q) := \mathbb{E}[u_i(x)|x \in \mathcal{X}_i] = q_i + v - \frac{t}{2} \cdot |\mathcal{X}_i|. \quad (2)$$

The expected user surplus without AI search features, given Q , is then

$$\sum_{i=1}^2 \int_{\mathcal{X}_i} \mathbb{E} \left[\max\{\rho_i(\bar{x}, Q) - \theta, 0\} \right] dx = \sum_{i=1}^2 \int_{\mathcal{X}_i} \int_0^{\rho_i(\bar{x}, Q)} F(s) ds dx. \quad (3)$$

That is, (3) represents a user's expected value from searching directly when there exists some informational friction, i.e. they are unaware of their location on the line and have updated their belief using $\bar{x}$. The search platform maximizes (3) w.r.t. $\bar{x}$ which gives the optimal allocation threshold

$$\bar{x}^N = \underbrace{\frac{t + q_1 - q_2}{2t}}_{\text{threshold user}} - \underbrace{\frac{q_1 - q_2}{2(4q_1 + 4q_2 - 3t + 8v)}}_{\text{match distortion}}. \quad (4)$$

Unlike the standard Hotelling cutoff, $\bar{x}^N$ is not the location of a user who is indifferent between the two publishers. Rather, it is the platform's optimal partition of users into two informational pools. Changing the threshold both reallocates users and changes the expected match quality—and therefore the search participation—of all users assigned to each publisher. As a result, the platform biases allocation toward the higher-quality publisher, but less strongly than under the classic utility-maximizing Hotelling allocation.

Publisher i 's profit is given by

$$\pi_i(\bar{x}, Q) = r \int_{\mathcal{X}_i} F(\rho_i(\bar{x}, Q)) dx - C(q_i)$$

which is concave in q_i as $\frac{d\pi_i^2}{dq_i^2} = -1 < 0$.

In the subsequent analysis, I focus solely on interior solutions, because corner solutions are not very meaningful given that my aim is to analyze content competition on the platform. Assumption 1 limits the per-user revenue gain from publishers which guarantees that content quality is not so high such that all users will directly access a webpage:

Assumption 1 $r < \bar{\theta} \min\{(\bar{\theta} - v)/2, t\}$.

Solving each publisher's maximization problem and using (4) naturally gives us $q_1^N =$

$q_2^N = q^N \Rightarrow \bar{x}^N = 1/2$ where q^N solves

$$rf(v + q^N - t/4) = 2q^N \Rightarrow q^N = \frac{r}{2\theta} \text{ and } \pi^N = \frac{r(v - \frac{t}{4})}{2\theta} + \frac{r^2}{8\theta^2}.$$

The tradeoffs presented in the baseline model are very simple. When each publisher increases content quality, they capture a larger user base but it comes at an increasing cost.¹³ Intuitively, equilibrium publisher profit π^N is increasing in r , decreasing in t , and increasing in v .

4 Baseline model with AI search features

In this section, the two publishers now compete when the platform implements AI search features. Specifically, in addition to setting an allocation threshold, the search platform can now synthesize its own AI content consisting of a “mixing rule” and “quality-conversion ability.”

Definition 2 *The search platform’s AI mixing rule is characterized as a mapping $\alpha(x) : [0, 1] \rightarrow [0, 1]$ such that its AI content for users is located at $\alpha(x)$ on the Hotelling line.*

Definition 3 *The search platform’s AI quality-conversion ability is characterized as $\gamma \in [0, 1]$ such that the quality of its AI content for users is $\gamma[(1 - \alpha(x))q_1 + \alpha(x)q_2]$.*

When a user x reads the AI content, they receive a utility of

$$u^{AI}(x) = \gamma \left[(1 - \alpha(x))q_1 + \alpha(x)q_2 \right] - t|\alpha(x) - x|. \quad (5)$$

The AI’s mixing rule $\alpha(x)$ represents how the AI can mix the publishers’ content in both vertical and horizontal dimensions. That is, $\alpha(x)$ dictates both user preferences and the quality of the content. The AI’s quality-conversion ability γ represents the possible degradation in quality from using AI. For example, this degradation can be due to hallucinations, inefficiencies in the synthesis process, or limitations when scraping content due to copyright regulations. Given that recent litigation seeks to restrict the platform’s ability to scrape or synthesize publisher content, the quality-conversion parameter, γ , emerges as a critical variable for the subsequent analysis.

¹³Altering $\bar{x}$ ’s timing such that it is chosen after content qualities induces an additional positive effect for each publisher. That is, if a publisher invests in content quality, the allocation would be more biased in their favor.




Figure 2: User search process and associated payoffs

When a user searches on the platform, they must first consume the AI's content at no search cost and then choose to be allocated by the search platform normally with the same search cost θ from before.¹⁴ If a user chooses to both consume the AI content and consume a publisher's content, the value they gain from the content is the maximum of both. In essence, both the AI content and organic publisher content address the same underlying search query. Therefore the user retains only the better of the two pieces of content. I discuss the scenario in which users still retain some of the AI content value when reading organic content in Appendix C. Figure 2 presents the overall user search process and the associated user payoffs.

The timing of the game now proceeds as follows:

- Stage 1: The platform chooses γ .
- Stage 2: Each publisher i chooses q_i and the platform chooses $\alpha(x)$ and $\bar{x}$ simultaneously.
- Stage 3: Users observe Q , $\alpha(x)$, γ , and $\bar{x}$. They then view the AI content and decide whether to continue searching to access a publisher's website by incurring their idiosyncratic search cost.

4.1 Equilibrium with AI search features

As the platform's design parameters are observable by users, users update their beliefs of their location on the line after reading the AI content. In fact, if $\alpha(x)$ is injective, users can

¹⁴Note that I do not increase the search cost when a user is allocated normally. Even if I included some additional search cost, the results still hold if the cost is sufficiently low.

perfectly pin down their location on the line. This is because $u^{AI}(x)$ is comprised of two main expressions: (i) the utility from content which is affected by both γ and $\alpha(x)$; and (ii) the disutility from preference mismatch which is determined solely by $\alpha(x)$. I now show that the optimal mixing rule is perfectly personalized to each user.

Lemma 1 *In any interior equilibrium, AI is perfectly personalized (i.e. $\alpha^I(x) = x$) almost everywhere and $\bar{x}^I = 1/2 + \frac{q_1 - q_2}{2t}$.¹⁵*

When finding the optimal mix of publisher content for a user x , there are two marginal effects the platform faces. First, suppose that one publisher is producing higher quality content than its rival and the mixing rule is $\alpha(x) = x$. Using more of the higher quality content when synthesizing AI content directly raises the quality of AI content, but also increases transportation costs as the location of the AI content is further from x . In any interior equilibrium, Lemma 1 shows that the preference mismatch marginal effect always dominates any possible gains in AI content quality. In addition, the perfectly personalized mixing rule allows the user to learn their exact location on the line and users can, thereby, imitate any decision they would have made for other arbitrary mixing rules. Therefore, the perfectly personalized mixing rule is the platform's optimal mixing rule in any interior equilibrium.

Suppose the publishers have chosen Q and the platform sets γ and $\alpha(x) = x$. A user located at x will access publisher i 's site if and only if

$$\theta \leq u_i(x) - u^{AI}(x).$$

Given $\bar{x}$, publisher i 's maximization problem becomes

$$\max_{q_i} r \int_{\mathcal{X}_i} \underbrace{F(u_i(x) - u^{AI}(x))}_{\Delta_i(x, Q)} dx - C(q_i). \quad (6)$$

where $\Delta_i(x, Q)$ is the difference in value from reading publisher i 's organic content compared to the platform's AI content for user x and given Q . Let $q^I(\gamma)$ be the solution to (6) given some γ in equilibrium. Also, for notational tractability, let $Q^I(\gamma) = (q^I(\gamma), q^I(\gamma))$ and $Q^N = (q^N, q^N)$.

Proposition 1 *With AI search features, there exists a unique interior equilibrium where $q_1 = q_2 = q^I(\gamma) = \frac{r(4-3\gamma)}{8\theta} \leq q^N$ and publisher profit is $\pi^I(\gamma) = \frac{r(v-t/4)}{2\theta} + \frac{r^2(4-3\gamma)(4-5\gamma)}{128\theta^2} \leq \pi^N$. In response to an increase in γ , equilibrium content quality and publisher profits decrease.*

¹⁵Uniqueness of the mixing rule $\alpha(x)$ is understood almost everywhere. In other words, the uniqueness of the equilibrium exists up to changes in the mixing rule on a set of measure zero.

In the equilibrium without any AI search features, publishers faced a simple tradeoff between attracting more users to their site by increasing quality and facing increasing costs of investment. With the incorporation of AI search features, publishers now face an additional conflicting AI quality effect as seen in publisher i 's FOC:

$$r \int_{\mathcal{X}_i} f(\Delta_i(x)) \cdot \left[\underbrace{\frac{\partial \Delta_i}{\partial u_i} \frac{\partial u_i(x)}{\partial q_i}}_{\text{user base effect (+)}} + \underbrace{\frac{\partial \Delta_i}{\partial u^{AI}} \frac{\partial u^{AI}(x)}{\partial q_i}}_{\text{AI quality effect (-)}} \right] dx = C'(q_i). \quad (7)$$

With AI search features, the platform's AI content serves as a new competitor that relies on each publisher's content. The AI quality effect reflects the fact that increasing one's content quality also improves the AI content quality, which indirectly reduces user traffic to their site as more users will be satisfied with the AI search results. This feedback effect is amplified as the AI's quality conversion ability increases. That is, for $q_i > 0$ I have that the AI quality effect is increasing in γ :

$$\frac{d}{d\gamma} \left| \frac{\partial \Delta_i}{\partial u^{AI}} \frac{\partial u^{AI}(x)}{\partial q_i} \right| > 0.$$

This is the underlying reason why equilibrium quality degrades as the AI's quality conversion ability improves.

Now that I have established how equilibrium quality and profits change with γ , how does user surplus with AI compare to the user surplus without AI search features? To determine the equilibrium γ set by the platform, I first characterize expected user surplus with AI search features given $\bar{x}$, Q , and γ :

$$\begin{aligned} \mathcal{U}(\bar{x}, \gamma, Q) &= \sum_{i=1}^2 \int_{\mathcal{X}_i} \mathbb{E} \left[u^{AI}(x) + \max\{\Delta_i(x, Q) - \theta, 0\} \right] dx \\ &= \underbrace{\mathbb{E}[u^{AI}(x)]}_{\text{upfront AI value}} + \underbrace{\sum_{i=1}^2 \int_{\mathcal{X}_i} \int_0^{\Delta_i(x, Q)} F(s) ds dx}_{\text{additional search value}}. \end{aligned} \quad (8)$$

Let $\mathcal{U}^I(\gamma) := \mathcal{U}(\bar{x}^I, \gamma, Q^I(\gamma))$ and $\mathcal{U}^N$ denote the equilibrium user surpluses with and without AI search features, respectively.

Proposition 2 $\mathcal{U}^I$ is concave in γ and $\mathcal{U}^I(0) > \mathcal{U}^N$. This implies that user surplus with AI search features is strictly higher than user surplus without AI search features in equilibrium.

Figure 3 illustrates this finding in more detail. The concavity of user surplus provides a clean interpretation of the tradeoff the platform faces when selecting γ . Consider a marginal

increase in γ . There are two conflicting effects on the upfront AI value each user receives. First, raising γ directly increases each user's upfront valuation for AI content. However, due to the feedback effect of quality degradation, the quality of the AI content will also degrade. I can also easily verify that the additional search value each user gains from accessing a webpage is decreasing in γ . This is because raising the AI's quality conversion ability causes more users to substitute away from organic content (hence the lower additional search value) and reduces organic content quality. When γ is sufficiently low, the direct effect outweighs the indirect feedback effect and the reduction in additional search value. If γ is sufficiently high, the AI search feature used by the platform is eating its own tail to its own detriment.

Although AI search features degrade quality, it still provides users with a key benefit. Recall that upon consuming the AI content, users learn their location on the line. This allows users to make more informed search decisions, resulting in higher quality matches between users and publishers.¹⁶

Going further, I can decompose $\mathcal{U}^I(\gamma) - \mathcal{U}^N$ into three terms:

$$\begin{aligned}
& \mathcal{U}^I(\gamma) - \mathcal{U}^I(0) + \mathcal{U}^I(0) - \mathcal{U}^N \\
= & \underbrace{\mathbb{E}[u^{AI}(x)]}_{\text{direct AI benefit (+)}} + \underbrace{\sum_{i=1}^2 \int_{\mathcal{X}_i^I} \left(\mathbb{E}[\max\{\Delta_i(x, Q^I(\gamma)) - \theta, 0\} - \max\{\Delta_i(x, Q^I(0)) - \theta, 0\}] \right) dx}_{\text{quality degradation effect (-)}} \\
& + \underbrace{\mathcal{U}^I(0) - \mathcal{U}^N}_{\text{match quality effect (+)}}.
\end{aligned}$$

Note at $\gamma = 0$, the direct AI benefit and quality degradation effect both equal 0, leaving only the positive match quality effect.

Next, define equilibrium social welfare with AI search features as

$$SW^I(\gamma) := \underbrace{\mathcal{U}^I(\gamma)}_{\text{user surplus}} + \underbrace{\Phi(\mathcal{U}^I(\gamma))}_{\text{platform profit}} + \underbrace{\sum_{i=1}^2 \pi^I(\gamma)}_{\text{total publisher profits}}.$$

In addition, I define the user-optimal γ as $\gamma^* := \arg \max_{\gamma} \mathcal{U}^I(\gamma)$ and socially-optimal γ as $\gamma^{SO} := \arg \max_{\gamma} SW^I(\gamma)$. As social welfare takes publisher profits into account, and γ^* is the maximizer of user surplus and platform profits, it is quite intuitive that the socially-optimal

¹⁶It is important to note that if users already possessed perfect knowledge of their preferences, or if the baseline platform could perfectly signal users' locations without AI, the strict matching benefit would disappear. In such cases, AI adoption increases user surplus only under specific parameter conditions, as formally analyzed in Appendix C.4.




Figure 3: Numerical example of user surplus without and with AI search features for $\gamma \in [0, 1]$ ($\theta \sim U[0, 4]$, $r = 1$, $v = 1.5$, $t = 1$, $\gamma^* \approx 0.4263$)

γ cannot be greater than the user-optimal γ . The following proposition includes this finding.

Proposition 3 *The following results hold: (i) γ^* is weakly increasing in t and weakly decreasing in v and r ; (ii) $\gamma^{SO} \leq \gamma^* < 1$; (iii) Social welfare with AI search features is greater than without at γ^{SO} ; (iv) Assuming $\Phi(US) = US$ and $t > 16/11$, social welfare is lower at γ^* if r and $\bar{\theta}$ are sufficiently high such that Assumption 1 is satisfied.¹⁷*

With high content differentiation, AI content is more valuable as the platform can personalize the AI search result and reduce transportation costs compared to the organic content. Thus, the platform has stronger incentives to raise its AI's quality-conversion ability, since users will benefit more from just viewing the initial AI content. Similar logic applies to the comparative static regarding v . When the publisher's per-user revenue increases, organic user traffic is more valuable to publishers and strengthens their incentives to invest in content quality. However, this also implies that the quality degradation effect will have a larger bite, which drives the platform to reduce its quality-conversion ability.

Proposition 3 also shows that the platform should always degrade its AI conversion ability, i.e. never set $\gamma^* = 1$. In addition, Proposition 3 implies that from a policy perspective,

¹⁷The result $\gamma^{SO} \leq \gamma^*$ is robust to generalized welfare weights. Suppose a planner maximizes $SW = \omega_1 US + \omega_2 \sum \pi_i$ with $\omega_1, \omega_2 > 0$. Since $\frac{d\pi^I}{d\gamma} < 0$, the marginal change in welfare at γ^* is strictly negative, implying $\gamma^{SO} \leq \gamma^*$. $\Phi(US) = US$ is adopted purely for analytical tractability.




Figure 4: A graphical illustration that social welfare may be lower at γ^* with AI search features ($v = 2, \bar{\theta} = 9, t = 1.5, \Phi(\text{US}) = \text{US}$)

restricting the platform to a lower γ may be socially-optimal. For example, implementing a quality-conversion ceiling $\bar{\gamma}$ that restricts the platform to the socially optimal level γ^{SO} would be (weakly) welfare improving by restoring publisher incentives and reducing quality degradation. Realistically, such a policy could be achieved by enforcing harsher copyright laws or introducing more technical barriers to make scraping less efficient via paywalls and required authentication. I provide one possible mechanism in Section 4.2 that can establish this ceiling.

There is a great need for such a policy when r is very large (see Figure 4). In such a scenario, publishers highly value user traffic for revenue. Thus, the traffic that the platform drives away from publishers has a large effect on publisher profits and social welfare. From a policy standpoint, the last item implies that more prominent publishers, who typically value readership either from a financial, political, etc., standpoint, are the most important targets for regulation. Indeed, it is more likely that AI search features reduce social welfare in equilibrium and that a ceiling on γ greatly improves social welfare for larger publishers.

4.2 Ex-post Fair Use copyright policy

Most AI developers argue that using publicly available content for model training is “transformative,” because it allows the model to generate new outputs rather than redistribute it. For example, OpenAI considers the training of its models as fair-use, while Google similarly describes their training as a transformative use under US copyright law. This becomes more complicated in the AI search context as AI search features retrieve publishers’ content in response to a user’s search query and produce a summary that may satisfy the same informational demand as the publishers’ content. In fact, the US Copyright Office notes that this form of AI content generation is less likely to be transformative as such AI responses may substitute for original works and some uses of copyrighted works in generative AI may qualify as fair-use while others may not, particularly when AI outputs closely reproduce or economically substitute for the underlying works.

This motivates an ex-post approach that permits the platform to use publisher content initially, but imposes liability when the resulting AI content becomes sufficiently similar to—and thereby serves more as an economic substitute with respect to—the publisher’s organic content. Specifically,

1. If $\mathbb{E}[\Delta_i(x)|x \in \mathcal{X}_i] = \mathbb{E}[u_i(x) - u^{AI}(x)|x \in \mathcal{X}_i] \geq d$, the platform is not liable for any damages from using publisher i ’s copyright-protected content.
2. If $\mathbb{E}[\Delta_i(x)|x \in \mathcal{X}_i] = \mathbb{E}[u_i(x) - u^{AI}(x)|x \in \mathcal{X}_i] < d$, the platform is found liable for copyright infringement and must pay a regulatory fine of $\Phi(\mathcal{U}^I(\gamma^*))$ or higher to the social planner.¹⁸

In other words, $d \geq 0$ represents a threshold such that if the expected value of the AI content is too similar to that of publisher i ex-post, then the platform must pay a steep compensation. Suppose a social planner has chosen d and consider the baseline game that then ensues. I can arrive at the following result:

Proposition 4 *The ex-post fair use copyright policy establishes a ceiling on γ and there exists d such that γ^{SO} is implemented in equilibrium. As a result, social welfare weakly increases, but user surplus weakly decreases compared to the baseline.*

To see this, observe that if the liability payment is sufficiently large, then the platform must comply with the copyright policy. Assuming an interior solution, Lemma 1 and Proposition

¹⁸I model this as a fine to prevent publishers from intentionally degrading their content just to trigger a payout. If damages were paid directly to publishers, compensation would need to be dependent on expected equilibrium quality, rather than actual quality, to preserve their baseline investment incentives.

1 do not change, so the fair-use constraint becomes

$$\begin{aligned}\mathbb{E}[\Delta_i(x)|x \in \mathcal{X}_i] \geq d &\iff v - \frac{t}{4} + (1 - \gamma)q^I(\gamma) \geq d \\ &\iff \gamma \leq \bar{\gamma}(d) := \frac{7 - \sqrt{1 + 96\bar{\theta}(d - v + t/4)/r}}{6}.\end{aligned}\tag{9}$$

Therefore, given $d \in [v - t/4, v - t/4 + r/(2\bar{\theta})]$ and by concavity of $\mathcal{U}^I(\gamma)$, the platform's optimal quality-conversion ability is $\gamma^{FU} = \min\{\gamma^*, \max\{0, \bar{\gamma}(d)\}\}$.¹⁹ As $\bar{\gamma}(d)$ is continuous and decreasing in d , there must exist some d^{SO} such that $\gamma^{FU} = \gamma^{SO}$.

Remark (Commitment to γ): If the platform has full control rights over γ and chooses γ *after* qualities are chosen, then the platform would set $\gamma = 1$. This outcome is pareto-dominated by γ^* if $\gamma^* < 1$ as both publisher profits and user surplus would be higher. Thus, setting $d = d^*$ such that $\bar{\gamma}(d^*) = \gamma^*$ would be socially beneficial as it forcefully prevents the platform from ignoring publisher incentives when selecting its quality-conversion ability. This further implies that any $d \in (v - t/4, d^*]$ would be a pareto-dominant outcome over the no commitment case as any d in that interval restricts $\gamma \in [\gamma^*, 1)$.

Remark (Observability of $\Delta_i(x)$): Under $\theta \sim U[0, \bar{\theta}]$, each publisher's expected organic traffic in the symmetric equilibrium is proportional to $\mathbb{E}[\Delta_i(x)|x \in \mathcal{X}_i]$ (specifically, publisher i 's user traffic = $\frac{1}{2\bar{\theta}} \mathbb{E}[\Delta_i(x)|x \in \mathcal{X}_i]$). Furthermore, Proposition 1 establishes that publisher profit is strictly decreasing in γ . Therefore, a planner does not need to observe user utility; they can implement the same conversion ceiling $\bar{\gamma}(d)$ by enforcing a threshold directly on observable publisher traffic or publisher revenue.

In summary, there are two main advantages for an ex-post fair use copyright policy. First, the mechanism restores publisher incentives as long as the AI content is sufficiently different from organic content. Although this reduces user surplus, it improves social welfare. Second, the copyright policy prevents a commitment problem the platform faces when determining its quality-conversion ability. Indeed, without such a policy, the platform can always perfectly imitate the publishers' content to provide users with as much direct benefit as possible. Given that publishers can anticipate this outcome, they would spend less producing their content which hurts all players. In practice, it may be the case that this mechanism is costly to implement or to determine ex-post damages. However, a feature that may make this practical is the fact that the policy need not be established on the user-level and can be implemented more broadly on the market-level by observing user traffic or publisher revenue outcomes.

¹⁹The interval on d comes from the fact that $\bar{\gamma}(v - t/4) = 1$.

5 Responses to AI search features

In this section, I examine several more responses to AI search features: (1) opt-out options, (2) licensing deals, and (3) pay-per-crawl payments. In each scenario, I extend the baseline model and explore the relevant tradeoffs and welfare consequences that arise. For the subsequent analysis in this section, I assume a closed-form $\Phi(\text{US}) = \text{US}$ in order to understand the trade-offs and outcomes more closely.

5.1 Opt-out options

As it currently stands, publishers can use several Web Application Firewalls (WAFs) that prevent content extraction. Google also claims publishers can opt out of their AI web scraping if they follow the “robots.txt” protocol which is famously used by BBC and *The New York Times*.²⁰ Simply put, the robots.txt protocol is a file that websites can use that indicates web crawlers which portions of the site cannot be visited. However, this is a gentleman’s agreement, i.e. there is no technical capability that prevents web crawlers from scraping their content and the protocol is based on voluntary compliance.

The need for having opt-out options has been widely discussed by policymakers and competition authorities. For example, the Competition and Markets Authority (CMA) and the US Department of Justice (DOJ) proposed that publishers have the option to opt out of their content being used to power Google’s AI features.²¹ The CMA and DOJ also require that Google cannot penalize or downrank sites that choose the opt-out option. The Independent Publishers Alliance also set out a complaint to the European Commission alleging that Google’s AI search feature has caused significant harm as publishers have no way to opt out from their content being used for large language models.²²

To determine how opt-out options affect content quality and welfare outcomes, I now allow each publisher to block the platform’s AI web scraping simultaneously to when γ is chosen. Their choice to opt out is observable by all players. I also assume that if a publisher chooses to block, it is perfect and costless, because I want to capture a rigorous opt-out option such as the current robots.txt protocol with full compliance.

Formally, publisher i can choose $b_i \in [0, 1]$ such that a user now derives utility

$$u^{\text{AI}}(x) = \gamma \left[(1 - b_1) \cdot (1 - \alpha(x))q_1 + (1 - b_2) \cdot \alpha(x)q_2 \right] - t|\alpha(x) - x|$$

²⁰Other search engines also follow this protocol such as AOL, Bing, DuckDuckGo, Yahoo!, and more.

²¹See <https://www.gov.uk/government/news/cma-proposes-package-of-measures-to-improve-google-search-services-in-uk>.

²²See <https://www.reuters.com/legal/litigation/googles-ai-overviews-hit-by-eu-antitrust-complaint-independent-publishers-2025-07-04/>.

from the platform’s AI content. Observe that the opt-out decisions only affect how content quality is converted into AI content quality, and not the AI’s ability to form content along the line. That is, even if publisher i chooses to block the AI web scraping, the platform can still set $\alpha(x) = l_i$ and form AI content at location l_i . One may interpret this as the platform still being able to scrape “small” publishers who provide low quality content or the AI’s ability to hallucinate content that is of no direct value to users.

Proposition 5 *In the simultaneous opt-out game, the equilibrium outcome is always equivalent to full blocking. Specifically, if the platform chooses $\gamma > 0$,*

- (i) *Full blocking is a strictly dominant strategy for both publishers.*
- (ii) *User surplus is lower with opt-out options if and only if $2r < t\bar{\theta} + \frac{16}{7}\bar{\theta}^2 - 4v\bar{\theta}$.*
- (iii) *Social welfare is lower with opt-out options if $2r < \bar{\theta}(\bar{\theta} - \frac{8r}{7}) + t\bar{\theta} - 4v\bar{\theta}$.*²³

Notice that if the platform optimally sets $\gamma = 0$, publishers are indifferent to blocking. Because any blocking choice in this boundary case yields the exact same payoffs as full blocking, the strategic outcome is equivalent.

When both publishers partially block the platform from scraping their content, they each have an incentive to deviate and enforce harsher blocking. This is because blocking now removes the AI quality effect such that the AI content relies solely on the other publisher. Therefore, the publisher with more blocking will have greater incentive to raise their content quality and attract more users. Now, given there is one publisher who is blocking more, the rival publisher can “even the odds” by deviating and blocking. Again, this removes the AI quality effect and the equilibrium outcome is effectively when $\gamma = 0$. Allowing opt-out options, in turn, results in degenerate AI content, but still provides a matching benefit to users.²⁴

Notice then that user surplus is reduced by providing opt-out options if and only if $\gamma^* > 0$ giving us the condition in item (ii). To interpret the condition, suppose $\bar{\theta}$ is sufficiently high relative to v , i.e. the additional benefits users gain from accessing direct websites are relatively low. In such a case, visiting publishers directly is more costly on average, so users must rely on the platform’s AI search feature for content. The platform would have set some $\gamma^* > 0$ to provide value to users, implying that opt-out options are harmful. Similar logic

²³A sufficient and necessary condition can be found in the proof, but is omitted because it is algebraically complex.

²⁴It is worth noting this welfare comparison assumes the platform retains an independent means to personalize AI content even when premium publishers opt out. If blocking access to publishers’ content also destroyed the AI’s horizontal matching capability, the baseline would revert to the no-AI environment, potentially making opt-outs even more harmful to users.

applies when considering social welfare consequences. The only additional consideration is that publishers obtain a higher profit with opt-out options despite the increasing production cost when $\gamma^* > 0$. If $\bar{\theta}$ is sufficiently high, this effect is dominated by the loss of user surplus from no longer consuming AI content, i.e. $u^{AI} = 0$. Figure 5 graphically illustrates the condition.




Figure 5: A graphical representation of how opt-out options affect social welfare

5.2 Licensing deals

Now that I have established that publishers opt-out of AI scraping when given the option, I move on to options for the platform to regain access to organic content. AI platforms have begun making licensing deals with several large news publishers for a couple of reasons. Some first-order reasons, perhaps the most obvious, are to avoid copyright litigation risk and to secure rights to use the content for training AI models. Although the structure of these licensing deals is still unclear, there seem to be two main types of licensing deals as it stands. First, there are lump-sum payments that serve as AI training licensing deals. These were more common when popular large language models, such as OpenAI's ChatGPT, became public around 2023-2024. For example, OpenAI formed fixed-fee deals with the Associated Press, News Corp, Financial Times, etc., in 2023 and early 2024. Another type of licensing deal that has started to gain traction in the video content industry is quality-adjusted payments. Several AI companies such as Treveo, Protégé Media, and Corpus are now sourcing video content from creators for \$1-4 per minute depending on resolution, clarity for accurate object identification, and frame rates.

Policymakers have also advocated for licensing deals. In June 2025, the European Parliament called on the European Commission to introduce a remuneration obligation on AI providers. Another example is Australia which introduced a News Bargaining Incentive that imposes a 2.25% levy on platform revenues unless they form licensing deals with local news publishers.²⁵

5.2.1 Lump-sum licensing

Suppose the platform sets a lump-sum licensing payment T simultaneously with its quality-conversion ability γ . A publisher has two distinct choices when choosing to opt out. They can either block the platform from scraping or accept the payment T . As a result, there are two types of pure strategy equilibria that can arise (in addition to both publishers blocking): both publishers accept the payment or only one publisher accepts. Also, while the baseline model assumes the platform maximizes gross user surplus, under a licensing regime the platform must fund these transfers. Consequently, the platform's objective here is to maximize user surplus net of any licensing payments.

First, consider a licensing equilibrium where each publisher sets $b_1 = b_2 = 0$. Notice that the payment does not affect publishers' investment choice in content quality, so both publishers will set $q^I(\gamma)$ as shown in Section 4. In turn, T must satisfy each publisher's participation constraint:

$$\pi^I(\gamma) + T \geq \pi_R^L(\gamma) \quad (PC^L)$$

where $\pi_R^L(\gamma)$ denotes the publisher's deviation profit from rejecting payment T .

Alternatively, consider a singular acceptance equilibrium where only one publisher accepts the payment, say publisher 1 without loss of generality. The AI content for user x is valued at $\gamma(1-x)q_A$ such that each publisher's content quality choice and the platform's allocation in equilibrium is characterized by the system of equations:

$$\bar{x}^L = \frac{1}{2} + \frac{q_A^L - q_R^L}{2t}, \quad q_A^L = \frac{r\bar{x}^L}{\theta} \left[1 - \gamma(1 - \bar{x}^L/2) \right], \quad \text{and} \quad q_R^L = \frac{r(1 - \bar{x}^L)}{\theta}.$$

On the publisher side, a singular acceptance equilibrium occurs if both the accepting and rejecting publisher's constraint is satisfied:

$$\pi_A^L(\gamma) + T \geq \pi^I(0) \quad (PC_A^L)$$

$$\pi_R^L(\gamma) \geq \pi^I(\gamma) + T \quad (PC_R^L)$$

²⁵The News Bargaining Incentive proposal applies to Google but excludes pure AI companies like OpenAI.

where $\pi_A^L(\gamma)$ denotes the accepting publisher's equilibrium profit without the payment.

To summarize, for lump-sum payment T ,

$$\begin{cases} \text{Both publishers block} & \text{if } T \leq \pi^I(0) - \pi_A^L(\gamma) \\ \text{Only one publisher blocks} & \text{if } \pi^I(0) - \pi_A^L(\gamma) \leq T \leq \pi_R^L(\gamma) - \pi^I(\gamma) \\ \text{Both publishers accept} & \text{otherwise} \end{cases}$$

where the platform maximizes

$$\left\{ \mathcal{U}^I(0), \mathcal{U}(\bar{x}^L, \gamma, (q_1^L, q_2^L)) - (\pi^I(0) - \pi_A^L(\gamma)), \mathcal{U}^I(\gamma) - 2(\pi_R^L - \pi^I(\gamma)) \right\}.$$

Note that the above characterization applies when $\pi^I(0) - \pi_A^L(\gamma) \leq \pi_R^L(\gamma) - \pi^I(\gamma)$ as it is possible that both publishers accept and both publishers block can occur in equilibrium given $T \in (\pi_R^L(\gamma) - \pi^I(\gamma), \pi^I(0) - \pi_A^L(\gamma)]$ if $\pi^I(0) - \pi_A^L(\gamma) > \pi_R^L(\gamma) - \pi^I(\gamma)$. In the proceeding results, I restrict attention to pure strategies and focus on the pure-strategy licensing equilibrium that yields the highest payoff for the platform, as Section 5.1 already covers the case where both publishers block.

5.2.2 Quality-adjusted licensing

Consider a quality-based payment scheme where the platform offers a payment rate $\phi > 0$. Similar to before, a publisher can reject the offer by blocking the AI scraping. Crucially, if a publisher accepts ϕ , they earn additional profit that is dependent on the quality of their content. That is, publisher i 's profit when accepting the offer is now

$$\pi_i = r \int_{\mathcal{X}_i} F(\Delta_i(x)) dx + \phi q_i - C(q_i). \quad (10)$$

Let $\pi^{qual}(\gamma, \phi)$ denote the equilibrium publisher profits given ϕ without the actual payment from the platform and q^{qual} be the equilibrium content quality when both publishers accept. For the payment to deter blocking for both publishers, it must be that

$$\pi^{qual}(\gamma, \phi) + \phi q^{qual} \geq \pi_R^{qual}(\gamma, \phi) \quad (PC^{qual})$$

where $\pi_R^{qual}(\gamma, \phi)$ is the publisher's deviation profit from rejecting the payment.

Under a singular acceptance equilibrium, I can again characterize the platform's alloca-

tion threshold and the accepting and rejecting publisher's content quality as

$$\bar{x}^{qual} = \frac{1}{2} + \frac{q_A^{qual} - q_R^{qual}}{2t}, \quad q_A^{qual} = \frac{r\bar{x}^{qual}}{\bar{\theta}} \left[1 - \gamma(1 - \bar{x}^{qual}/2) \right] + \phi, \text{ and } q_R^{qual} = \frac{r(1 - \bar{x}^{qual})}{\bar{\theta}}.$$

The accepting and rejecting publisher's participation constraint is:

$$\pi_A^{qual}(\gamma, \phi) + \phi q_A^{qual} \geq \pi^I(0) \quad (PC_A^{qual})$$

$$\pi_R^{qual}(\gamma, \phi) \geq \pi^{qual}(\gamma, \phi) + \phi q^{qual} \quad (PC_R^{qual})$$

where $\pi_A^{qual}(\gamma, \phi)$ is the accepting publisher's profit without the quality-adjusted payment. In contrast to the lump-sum payment, $\pi_A^{qual}(\gamma, \phi)$ and $\pi_R^{qual}(\gamma, \phi)$ both depend on ϕ as ϕ affects content quality choice. This will play an important role in the following results.

Given ϕ ,

$$\begin{cases} \text{Both publishers block} & \text{if } \phi \leq \frac{\pi^I(0) - \pi_A^{qual}(\gamma, \phi)}{q_A^{qual}} \\ \text{Only one publisher blocks} & \text{if } \frac{\pi^I(0) - \pi_A^{qual}(\gamma, \phi)}{q_A^{qual}} \leq \phi \leq \frac{\pi_R^{qual}(\gamma, \phi) - \pi^{qual}(\gamma, \phi)}{q^{qual}} \\ \text{Both publishers accept} & \text{otherwise} \end{cases}$$

Similar to the lump-sum case, if these participation cutoffs cross, it is possible for both joint acceptance and joint blocking to occur in equilibrium for the same payment rate; in the proceeding results, I continue to focus on the equilibria where licensing actually occurs.

Proposition 6 *In an interior equilibrium,*

- *There exists a region near the lower boundary $t = r/\bar{\theta}$ in which the platform prefers the lump-sum payment scheme and only one publisher accepts the offer; in this case, social welfare is higher than in the baseline model.*
- *There exists a region where $t, v, \bar{\theta}$ are jointly large and $\bar{\theta}$ is sufficiently large relative to v and t in which the platform prefers the quality-adjusted payment scheme and both publishers accept the offer; in this case, social welfare is higher than in the baseline model.*

Proposition 6 states that the platform's preferred payment scheme is dependent on how competitive the content market is (see Figure 10 for a graphical example). In the proof, I first establish that a lump-sum payment where both publishers accept is never the platform's optimal contract choice. To see this, consider an equilibrium where both publishers accept under lump-sum and its associated minimum per-publisher payment T^* such that (PC^L) is

binding.²⁶ The platform can simply set $\phi = \frac{T^*}{q^I(\gamma)}$ to also satisfy each publisher's participation constraint. In addition, using a quality-adjusted payment raises equilibrium content quality (when fixing γ), because it acts as a subsidy on quality investment, partially offsetting the negative AI quality effect.²⁷ In contrast, lump-sum payments do not affect publishers' investment choices, so the equilibrium content quality remains the same as in Section 4. This means that the user surplus is higher with quality-adjusted payments. Although implementing the quality-adjusted payment scheme is more costly as $\phi q^{qual} = \frac{T^* q^{qual}}{q^I(\gamma)} > T^*$, I show that the improvement in user surplus from quality-adjusted payments dominates the increase in costs when the platform's budget constraint is satisfied under lump-sum.

Next, consider when publishers are not very differentiated, i.e. consider when $t \rightarrow \frac{r}{\bar{\theta}}$. Under such a scenario, the platform pays one publisher T to remove them from the market such that almost every user is allocated to the rejecting publisher. This provides the rejecting publisher with higher incentive to produce high-quality content which is beneficial only when t and $\bar{\theta}$ is low. Instead, under a singular acceptance equilibrium for the quality-adjusted payment, there are two opposing economic forces when selecting $\bar{x}^{qual}$. Allocating more users to the accepting publisher provides the accepting publisher with more incentive to let the AI scrape and use its content. However, this simultaneously reduces the rejecting publisher's market share increasing their incentive to deviate and also accept the payment. When $\bar{\theta}$ is also low, this user surplus gain is much more costly compared to the lump-sum case as the platform has no need to generate a high quality AI search result. Therefore, the gain from a quality-adjusted payment is not relatively valuable. A similar comparison for an equilibrium where both publishers accept the quality-adjusted payment can be made.

Now, I turn to the second item in Proposition 6. Under higher content differentiation, AI content is more valuable because it can be personalized to remove transportation costs. Coupled with high search costs, any singular acceptance agreement would be less preferred, as AI content is synthesized through one publisher's content and users are less willing to continue searching for organic content. For each of the scenarios discussed in Proposition 6, social welfare is higher with licensing than in the baseline. While the payments between the platform and publisher have no net effect on social welfare, licensing alters each publisher's (accepting and rejecting) incentives and results in higher quality content on average for users who continue to search organically.

²⁶In the proof of Proposition 6, I also prove this result for singular acceptance equilibria, but it is much more complex in terms of analysis, so I do not focus on it in the main section.

²⁷In fact, when both publishers accept the quality-adjusted payment, the equilibrium content quality is $q^{qual} = q^I(\gamma) + \phi$.

5.3 Pay-per-crawl

A shortcoming of the baseline model is that publishers cannot monetize platform access to their webpage. There are several services that allow publishers to charge micro-payments when AI bot crawlers scrape their content. For instance, Cloudflare and TollBit grant publishers three options for AI web scrapers. Publishers can choose to block, freely allow, or charge a flat-fee payment. Similarly, IAB Technology Lab, a non-profit research consortium, proposed a pay-per-crawl framework that allows publishers to monetize their content. Platforms have also begun to use a similar type of payment. Perplexity’s Publisher Program compensates publishers for a percentage of ad revenue Perplexity makes when one of the publisher’s webpage is cited as a source for that response.

To model such an option, suppose publishers are paid a pay-per-crawl payment rate $R > 0$ that is chosen by an outside policy maker or a pay-per-crawl service. That is, the platform must pay $R > 0$ to a publisher if they scrape and use their content for AI content.²⁸ To isolate the direct effect of this payment structure on mixing and investment, I hold the platform’s quality-conversion ability γ fixed and analyze the resulting continuation game.

Given allocation threshold $\bar{x}$, only using publisher i ’s content for user $x \in \mathcal{X}_i$ improves user x ’s utility if and only if

$$u_i^{AI}(x) := \underbrace{\gamma q_i - t|l_i - x|}_{u^{AI} \text{ using only publisher } i} \geq u^{AI}(x) = \underbrace{\gamma[(1-x)q_1 + xq_2]}_{u^{AI} \text{ using both publishers}}.$$

The gain in expected user surplus for user x when allocated to pub. i for mixing versus just one publisher is then

$$\begin{aligned} \Delta \mathcal{U}_i(x) &= u^{AI}(x) - u_i^{AI}(x) + \int_0^{\Delta_i(x)} F(s)dx - \int_0^{\Delta_i^{corner}(x)} F(s)ds \\ &= u^{AI}(x) - u_i^{AI}(x) + \frac{(\Delta_i(x))^2 - (\Delta_i^{corner})^2}{2\bar{\theta}} \end{aligned} \quad (11)$$

where $\Delta_i^{corner} = q_i(1 - \gamma) + v$. Note that $\Delta_i(x) - \Delta_i^{corner} = \gamma|l_i - x|(q_i - q_{-i}) - t|l_i - x| = -(u^{AI}(x) - u_i^{AI}(x))$. Then, (11) can then be reduced to

$$\Delta \mathcal{U}_i(x) = \left[\gamma|l_i - x|(q_{-i} - q_i) + t|l_i - x| \right] \underbrace{\left(1 - \frac{\Delta_i(x) + \Delta_i^{corner}}{2\bar{\theta}} \right)}_{\geq 0}.$$

²⁸We assume the platform must access a publisher’s site to retrieve information for each user. Perplexity AI follows this approach, while other AI services like ChatGPT and Gemini follow a hybrid approach. Modeling a hybrid approach is beyond the scope of this paper.

Notice that $\Delta\mathcal{U}_1(x)$ is increasing in x and $\mathcal{U}_2(x)$ is decreasing in x if $\gamma|q_1 - q_2| < t$. Under any interior equilibrium, then the optimal choice by the platform is

$$\alpha^{PPC}(x) = \begin{cases} 0 & \text{for } x \in [0, \bar{a}_1] \\ x & \text{for } x \in [\bar{a}_1, 1 - \bar{a}_2] \\ 1 & \text{for } x \in [1 - \bar{a}_2, 1] \end{cases}$$

where $\bar{a}_i$ solves $\Delta\mathcal{U}_1(\bar{a}_1) = R$ and $\Delta\mathcal{U}_2(1 - \bar{a}_2) = R$.²⁹ Note that even though $\alpha^{PPC}(x)$ pools users near the endpoints, users still perfectly infer their location x .³⁰

For the publisher's problem, each publisher's FOC now integrates over two regions. Consider publisher 1's problem given that R_1 and R_2 were chosen:

$$\max_{q_1} r \int_0^{\bar{a}_1} F(\Delta_1^{corner}) dx + r \int_{\bar{a}_1}^{\bar{x}} F(\Delta_1(x)) dx + R_1(1 - \bar{a}_2) - C(q_1).$$

Solving for publisher 1's FOC gives

$$q_1 = \frac{r}{\bar{\theta}} \left[(1 - \gamma)\bar{x} + \frac{\gamma}{2}(\bar{x}^2 - \bar{a}_1^2) \right]. \quad (12)$$

Because $\alpha(x)$ is chosen simultaneously with the content qualities, changing q_1 does not affect $\bar{a}_i$ in the FOC. In the unique symmetric equilibrium where $\bar{x} = 1/2$ and $\bar{a}_1 = \bar{a}_2 = \bar{a}$, this condition neatly simplifies to $q^{PPC} = q^I(\gamma) - \frac{r\gamma\bar{a}^2}{2\bar{\theta}}$.

Proposition 7 *Suppose $\gamma > 0$. With pay-per-crawl payments, there exists a unique symmetric equilibrium where*

- (i) *Content quality and user surplus when publishers use pay-per-crawl payments is lower than in the baseline model.*³¹
- (ii) *More mixing occurs as t and $\bar{\theta}$ increases.*
- (iii) *Given γ , social welfare is lower with pay-per-crawl if search costs are sufficiently high relative to publishers' per-user revenue (specifically $r < \frac{3(\bar{\theta}-v)}{2}$).*

At first glance, it seems that publishers should invest more in content quality as they have an additional revenue outlet similar to the quality-adjusted payment scenario in Section

²⁹I assume R is sufficiently low such that $\bar{a}_i \leq 1/2$ and omit the case where $\bar{a}_i > 1/2$. In such a case, $\alpha^{PPC}(x) = 0$ if $x \leq 1/2$ and $\alpha^{PPC}(x) = 1$ otherwise.

³⁰For example, for $x \in [0, \bar{a}_1]$, $\alpha(x) = 0$ and the user experiences AI utility $u^{AI}(x) = \gamma q_1 - tx$. Because $u^{AI}(x)$ is strictly monotonic in x , experiencing the AI content perfectly reveals x .

³¹If the platform were to endogenously re-optimize γ in the first stage, it might select a lower γ to partially offset this quality degradation. Thus, this result highlights the policy's direct, partial-equilibrium distortion.

5.2. This does not occur because when pay-per-crawl payments are in effect, the platform’s optimal mixing rule is no longer perfectly personalized. Instead, the platform is incentivized to distort the mixing rule for users on the ends of the Hotelling line towards the publisher closer to them. Indeed, mixing now becomes too costly compared to the net user surplus gain, given that the user is already sufficiently close to a publisher. This distortion implies that the negative AI quality effect established in the baseline model is amplified for end users. Therefore, publishers produce lower quality content as the AI content cannibalizes user traffic at a higher rate than in the baseline.




Figure 6: Distorted mixing rule under pay-per-crawl

Naturally, the mixing rule becomes less distorted as the market becomes less competitive, i.e. t and $\bar{\theta}$ increases. As transportation costs or search costs increase, using both publishers’ content to synthesize AI content becomes more valuable to users. To see the welfare result, recall that any transfers between the publishers and platform does not affect net social welfare. While publisher profits are indeed higher with pay-per-crawl payments, the user surplus reduction due to the reduced content quality and less personalized AI content dominates given that publishers’ per-user revenue is sufficiently low.

6 Platform competition

I extend the monopolist platform model by allowing for platform competition. To keep things as simple as possible, I employ a double Hotelling model with a hierarchical market structure where each platform $j \in \{A, B\}$ chooses an allocation threshold $\bar{x}_j$, AI mixing rule $\alpha_j(x)$, and AI quality-conversion ability γ_j . To keep the analysis tractable, I continue to assume $\Phi(\text{US}) = \text{US}$.

This extended environment is similar to that of Chi et al. (2026), consisting of two hierarchical layers where each platform is located at the end of the Hotelling line of length one ($l_A = 0$ and $l_B = 1$) in the upper layer (refer to Figure 7). Users are uniformly distributed along the upper layer and first select a platform to join. Let $\tau > 0$ be the transportation cost parameter that represents the horizontal differentiation between both platforms. Then, a user located at $y \in [0, 1]$ in the upper layer who accesses platform j receives $\mathcal{U}_j - \tau|y - l_j|$ where $\mathcal{U}_j$ denotes the expected utility from joining platform j . The marginal user that is indifferent between both platforms is, thereby, located at $\bar{y} := \frac{1}{2} + \frac{\mathcal{U}_A - \mathcal{U}_B}{2\tau}$.

Once a user joins a platform, then the search and allocation process is that of the monopoly case described in Sections 3 and 4. The lower layer in Figure 7 represent users' tastes for publishers and the allocation threshold each platform sets. Crucially, the users' location on the upper layer is independent of their location on the lower layer.

Before, each user had to access content through the monopoly platform. With competition, platform j 's objective is to maximize $|\mathcal{Y}_j| \cdot \Phi(\mathcal{U}_j)$ where $\mathcal{Y}_j$ is the set of users who joined platform j in the upper layer. Observe that, as in Sections 3 and 4, each platform j maximizes expected user surplus (and hence platform profits) by setting $\alpha_j(x) = x$ and $\bar{x}_j = \frac{1}{2} + \frac{q_1 - q_2}{2t}$ in a symmetric equilibrium.

Given $(\bar{x}_A, \bar{x}_B)$, publisher i 's maximization problem now becomes

$$\max_{q_i} r \sum_{j=1}^2 \int_{\mathcal{X}_{ij}} |\mathcal{Y}_j| \cdot F(u_i(x) - u_j^{AI}(x)) dx - C(q_i)$$

where $\mathcal{X}_{ij}$ denotes the set of users allocated to publisher i in the lower layer given they joined platform j , and $u_j^{AI}(x)$ represents the utility of platform j 's AI content.

The timing of the game is as follows:

- Stage 1: Each platform j chooses γ_j simultaneously.
- Stage 2: Users observe (γ_A, γ_B) and decide which platform to join.
- Stage 3: Publishers observe how many users are on each platform. Each publisher chooses q_i and each platform chooses $\alpha_j(x)$ and $\bar{x}_j$ simultaneously.
- Stage 4: Users at platform j observe Q , $\alpha_j(x)$, and $\bar{x}_j$. They then view the AI content at their platform and decide whether to continue searching to access a publisher's website by incurring their idiosyncratic search cost. Outcomes are realized.

Proposition 8 *Under platform competition, there exists a symmetric equilibrium if τ is sufficiently large such that the equilibrium content quality $q^c \leq q^I$ and each platform's quality-conversion ability $\gamma^c \geq \gamma^*$ with strict inequality if γ^* is interior. In addition, the equilibrium*




Figure 7: Modified market structure with platform competition

publisher profit, user surplus, and aggregate platform welfare are weakly lower compared to the baseline model.

Each platform's FOC when maximizing w.r.t. γ can be isolated into two terms:

$$\underbrace{|\mathcal{Y}_j| \cdot \Phi'(\mathcal{U}_j) \frac{d\mathcal{U}_j}{d\gamma_j}}_{\text{user surplus reduction } (-)} + \underbrace{\Phi(\mathcal{U}_j) \cdot \frac{d|\mathcal{Y}_j|}{d\gamma_j}}_{\text{platform user base effect } (+)} = 0.$$

The user surplus reduction effect captures the fact that raising $\gamma_j > \gamma^*$ lowers user surplus as discussed previously in the monopoly model. The platform's user base effect represents the fact that increasing γ entices more users to join the platform given that each publisher sets the same quality and each platform designs its AI content from the same publisher pool. Observe that when platform j marginally increases its quality-conversion ability, it also reduces user surplus for the other platform which is another contributing factor to this user base effect.

The user surplus reduction effect is greater than zero when evaluating the FOC at γ^* . A competing platform internalizes only half of the publisher-quality degradation effect and imposes a negative externality onto the rival platform with a unilateral increase in γ_j . As a result, each platform wants to deviate upwards in a “race to the top.” Indeed, the intensive margin of raising γ must be negative to offset the user-stealing effect which gives $\gamma^c > \gamma^*$. Therefore, introducing more competition among platforms causes the AI quality effect by the publishers to hold less weight when deciding γ , and worsening the content quality degradation problem relative to the monopoly platform case. The decrease in content quality and increase

in γ implies that expected user surplus and publisher profits are now lower.

Remark: As $\tau \rightarrow 0$ and competition becomes more fierce, platforms race all the way to the top and maximize their quality-conversion ability. This essentially captures the scenario where platforms are Bertrand competing for users. In such a scenario, the intense user-stealing incentive drives platforms in a race to the top, pushing conversion ability toward $\gamma = 1$ where AI content fully substitutes for publishers’ organic content.

Another important implication from this result is that implementing the ex-post fair use copyright policy from Section 4.2 can restrict γ to γ^* such that social welfare does not change from the monopoly model, minus the upper-layer transportation costs. This is because along a symmetric equilibrium path where $\gamma_A = \gamma_B = \gamma$, the platforms’ equal market shares yield a common publisher quality $q^c(\gamma, \gamma) = q^I(\gamma)$. Thus, the expected fair-use gap coincides perfectly with the monopoly benchmark, meaning a regulator can target the socially optimal symmetric conversion ability using the same threshold d^{SO} . Since $\gamma^c \geq \gamma^* \geq \gamma^{SO}$, the range of d that improves welfare expands and the socially optimal threshold d^{SO} implementing γ^{SO} still holds under platform competition.

7 Policy Implications

The central policy challenge is not whether AI search features should be permitted or prohibited as the effect of implementing AI search features on social welfare is ambiguous. Rather, the need for policy intervention is particularly salient for markets where publishers’ expected per-user revenue is high relative to user search costs, provided the platform still chooses to deploy a functional AI substitute. In general, the best policy target is the platform’s effective ability to transform publisher content into substitute AI content. This can take the form of copyright rules, technical restrictions, attribution requirements, or liability standards. For example, an ex-post fair-use mechanism that evaluates whether AI output is sufficiently differentiated from publisher content can implement an upper bound on the platform’s effective quality-conversion ability. In practice, regulators might approximate this mechanism by targeting empirical proxies, such as platform-attributed user leakage or the resulting profit loss to publishers.³²

Second, if restricting the platform’s quality-conversion ability is not possible, then providing opt-out options for publishers may be another solution. That is, providing opt-out options results in publishers blocking AI webscraping which restores publisher incentives, but also eliminates the direct content value of the AI search feature. A caveat is that opt-out rules may protect publishers at the cost of reducing user surplus and social welfare, espe-

³²This is similarly proposed in Gans (2024a), although in a different context.

cially when users rely heavily on the AI result due to high search costs. An opt-out policy is more attractive in markets where publishers derive higher per-use revenue or readership value which also happens to be when the need for policy intervention is greatest. In addition, opt-out options should be paired with mechanisms that encourage licensing agreements, so that publishers have credible control over how their content is used but platforms can still obtain access.

This leads to the third point on the design of licensing payments. Lump-sum payments can compensate publishers for allowing access, but they do not directly change the marginal return to investing in quality. Instead, quality-adjusted payments raise the marginal return to quality and partially offset the AI quality effect. That being said, lump-sum licensing may be attractive in practice because it is simple and avoids the need to measure usage or contribution. In addition, quality-adjusted payments depend on the platform being truthful to publishers on measures that proxy the publisher’s contribution to AI quality, such as article usage, citation frequency, retrieval frequency, topic authority, freshness, user engagement with cited sources or other quality-weighted measures. This may be the reason why only the video content creation market has seen this type of licensing deal as it is easier to track reliable quality measures.

The optimal licensing scheme also depends on market conditions. Under extreme market conditions, licensing is socially beneficial. When competition between publishers is very intense and users face relatively low search costs, a targeted lump-sum arrangement with one publisher can be attractive for the platform. In that case, the platform can obtain content access from one publisher while allocating more users to the rival publisher, strengthening the rival’s incentive to produce quality. However, when the market is more segmented, licensing both publishers becomes more attractive because the value of maintaining high-quality AI and organic content is greater. Thus, policies should not focus on a single type of contract for all markets.

Fourth, while pay-per-crawl systems may seem appealing at first glance, they distort the way AI content is mixed for users. Rather than using the best combination of publisher content for each user, the platform relies too heavily on a single publisher’s content in order to reduce payment costs. This further reduces publishers’ incentives to invest in high quality content. Hence, holding the platform’s AI capabilities fixed, pay-per-crawl results in socially worse outcomes if search costs are sufficiently high. Therefore, policymakers should not assume pay-per-crawl serves as an effective substitute for licensing.

There have been several proposals from competition authorities that may indirectly create more competition among search platforms with respect to AI content. For instance, UK’s CMA has advocated for choice screens on Android and Chrome, and in June 2025, explicitly

stated that choice screens should cover AI assistants. This would make it easier for users to choose and switch search services. The US’s DOJ also barred Google from entering exclusive contracts related to the distribution of Gemini and ordered Google to make user-interaction data available to rival platforms.

Platform competition may be harmful to users and socially inefficient. That is, the model implies that competition can amplify, rather than solve, the quality-degradation problem. Because both platforms draw from the same publisher pool, a platform can attract users by unilaterally increasing its AI conversion ability while externalizing part of the resulting quality-degradation cost onto its rival. As a result, competing platforms collectively choose an excessively high quality-conversion ability. This does not mean that competition policies should never be implemented. Rather, the model suggests that if competition remedies intensify the race to capture users, they risk exacerbating this quality-degradation problem. Therefore, policymakers should consider pairing competition remedies with content-access rules, licensing obligations, or copyright standards that constrain the degree to which platforms can appropriate publisher content.

While the analysis has primarily focused on characterizing how AI search features affect content investment decisions and subsequent policy responses, it is also useful to establish some practical solutions publishers can employ to combat the introduction of AI search features.

a) Investment in webpage-specific features

Many webpages include interactive tools, calculators, data visualizations, comment sections, personalized recommendations, multimedia content, games, live updates, expert databases, community features, and other forms of user engagement that are specific to the publisher. Such features increase the incremental value of clicking through after seeing the AI answer as they are less vulnerable to AI appropriation.

b) Ability to access content directly and facilitating convenient search

If direct webpage access is very convenient, AI search becomes a much weaker substitute. Users are more likely to click through when a publisher’s page loads fast, is not cluttered with advertisements, or is easy to navigate on their devices. While paywalls can effectively block web scraping, publishers must carefully choose when to implement them, as restrictive access may alienate their user base and significantly reduce organic traffic from the search platform.

c) Build direct user relationships and foster user trust

Web pages can mitigate the zero-click problem by building direct relationships with users, which is essential in growing a user base that does not rely on a search platform. For example, publishers can provide subscriber-exclusive rewards such as newsletters and briefings, com-

munity forums, behind-the-scenes reporting, webinars, etc. In addition, publishers should signal to users how trustworthy their content is by providing methodology notes, discussing editorial standards, and being transparent on where they source their information.

8 Concluding Remarks

In this paper, I studied how the incorporation of AI search features on search platforms affects the investment of content quality from publishers. I show that while AI content improves match quality between users and publishers by helping users learn their preferences, the platform’s appropriation of publisher content degrades investment in quality. Thus, it is not necessarily optimal for the platform to maximize its AI quality-conversion ability (even when maximizing user-surplus), as restraining this capability preserves publisher incentives.

This paper is a first step towards modeling AI content creation on search engines, opening other directions for future work. The presented results provide several research questions that empiricists can look into. I highlighted a mechanism at which introducing AI search features and improving AI content leads to content degradation, how this mechanism is affected by various underlying factors (such as content differentiation), and several predictions depending on how different players respond. These predictions can and should be tested as more data becomes readily available. For example, an empirical study could use a difference-in-differences design to estimate how AI Overviews affects content quality measures by exploiting variation in the timing of its roll-out across countries and regions. Researchers could also leverage within-publisher variation when the same publisher operates separate regional divisions (e.g. *The Guardian* and *The Independent* have separate US and UK teams) that were exposed to AI Overviews at different times.³³

One can also extend my theoretical framework in many ways. For instance, I restricted users to single-home with respect to choosing a publisher and the model only captures competition in a singular topic. It would be interesting to enrich the environment by allowing users to multi-home or extending competition among multiple topics as in Jeon and Nasr (2016) with news aggregators. Also, I allowed the platform to have full control rights over several design parameters such as the allocation threshold, AI mixing rule, and AI quality-conversion ability. One can allow for bargaining or other negotiating tactics to endogenize publisher’s position on the search platform. The analysis of policy responses can also be further extended. For example, I did not allow publishers or the platform to select the pay-per-crawl rate and the results did not characterize all the possible equilibria when the

³³One would have to account for the fact that publishers who produce in different regions face different topics and audiences.

platform licenses content. Lastly, one can extend my environment by vertically integrating one of the publishers with the platform.

References

- Agarwal, Saharsh and Ananya Sen (2026) “The Impact of Google AI Overviews on Publisher Traffic and User Experience: Evidence from a Field Experiment,” *TSE Working Paper*.
- Anderson, Simon P and Stephen Coate (2005) “Market provision of broadcasting: A welfare analysis,” *The Review of Economic Studies*, 72 (4), 947–972.
- Anderson, Simon P, Øystein Foros, and Hans Jarle Kind (2018) “Competition for advertisers and for viewers in media markets,” *The Economic Journal*, 128 (608), 34–54.
- Anderson, Simon P and Joshua S Gans (2011) “Platform siphoning: Ad-avoidance and media content,” *American Economic Journal: Microeconomics*, 3 (4), 1–34.
- Armstrong, Mark (2006) “Competition in two-sided markets,” *RAND Journal of Economics*, 37 (3), 668–691.
- Athey, Susan, Emilio Calvano, and Joshua S Gans (2018) “The impact of consumer multi-homing on advertising markets and media competition,” *Management science*, 64 (4), 1574–1590.
- Athey, Susan and Markus Mobius (2012) “The impact of news aggregators on internet news consumption: The case of localization.”
- Athey, Susan, Markus Mobius, and Jeno Pal (2021) “The impact of aggregators on internet news consumption,” Working Paper 28746, National Bureau of Economic Research, 10.3386/w28746.
- Bernheim, B Douglas and Michael D Whinston (1998) “Exclusive dealing,” *Journal of political Economy*, 106 (1), 64–103.
- Caillaud, Bernard and Bruno Jullien (2003) “Chicken & egg: Competition among intermediation service providers,” *RAND Journal of Economics*, 309–328.
- Calzada, Joan and Ricard Gil (2020) “What do news aggregators do? Evidence from Google News in Spain and Germany,” *Marketing Science*, 39 (1), 134–167.

- Calzada, Joan and Guillem Ordóñez-Calafi (2012) “Competition in the news industry: Fighting aggregators with versions and links,” *Available at SSRN 2164399*.
- Chan, Alex (2026) “AI and the Collapse of the www,” Working Paper 35344, National Bureau of Economic Research, 10.3386/w35344.
- Chi, Chang-Koo, Jay Pil Choi, Jong-Hee Hahn, and Seongkyun Kim (2026) “Regulating Platform Bias: Self-Preferencing and Transaction Fees as Strategic Substitutes.”
- Chiou, Lesley and Catherine Tucker (2017) “Content aggregation by platforms: The case of the news media,” *Journal of Economics & Management Strategy*, 26 (4), 782–805.
- Choi, Jay Pil and Doh-Shin Jeon (2023) “Platform design biases in ad-funded two-sided markets,” *The RAND Journal of Economics*, 54 (2), 240–267.
- de Cornière, Alexandre and Miklos Sarvary (2023) “Social media and news: Content bundling and news quality,” *Management Science*, 69 (1), 162–178.
- de Cornière, Alexandre and Greg Taylor (2014) “Integration and search engine bias,” *The RAND Journal of Economics*, 45 (3), 576–597, <https://doi.org/10.1111/1756-2171.12063>.
- Crampes, Claude, Carole Haritchabalet, and Bruno Jullien (2009) “Advertising, competition and entry in media industries,” *The Journal of Industrial Economics*, 57 (1), 7–31.
- Dellarocas, Chrysanthos, Zsolt Katona, and William Rand (2013) “Media, aggregators, and the link economy: Strategic hyperlink formation in content networks,” *Management science*, 59 (10), 2360–2379.
- Edelman, Benjamin and Zhenyu Lai (2013) *Exclusive preferential placement as search diversion: Evidence from flight search*: Harvard Business School.
- Gabszewicz, Jean J, Dider Laussel, and Nathalie Sonnac (2001) “Press advertising and the ascent of the ‘Pensée Unique’,” *European Economic Review*, 45 (4-6), 641–651.
- Gans, Joshua S. (2024a) “Copyright Policy Options for Generative Artificial Intelligence,” *Journal of Law and Economics*, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4707911.
- (2024b) “Will user-contributed AI training data eat its own tail?” *Economics Letters*, 242, 111868, <https://doi.org/10.1016/j.econlet.2024.111868>.
- Gentzkow, Matthew and Jesse M Shapiro (2011) “Ideological segregation online and offline,” *The Quarterly Journal of Economics*, 126 (4), 1799–1839.

- George, Lisa M and Christiaan Hogendorn (2012) “Aggregators, search and the economics of new media institutions,” *Information Economics and Policy*, 24 (1), 40–51.
- Jeon, Doh-Shin (2018) “Economics of news aggregators,” *TSE Working Paper*.
- Jeon, Doh-Shin and Nikrooz Nasr (2016) “News Aggregators and Competition among Newspapers on the Internet,” *American Economic Journal: Microeconomics*, 8 (4), 91–114, 10.1257/mic.20140151.
- Johnson, Justin P., Andrew Rhodes, and Matthijs Wildenbeest (2023) “Platform Design When Sellers Use Pricing Algorithms,” *Econometrica*, 91 (5), 1841–1879, <https://doi.org/10.3982/ECTA19978>.
- Kamien, Morton I and Yair Tauman (1986) “Fees versus royalties and the private value of a patent,” *The Quarterly Journal of Economics*, 101 (3), 471–491.
- Katz, Michael L and Carl Shapiro (1986) “How to license intangible property,” *The Quarterly Journal of Economics*, 101 (3), 567–589.
- Lyu, Gaoyan and Ying-Ju Chen (2026) “Designing AI-Powered Search Platforms: Advertiser Incentives, Quality Paradox, and AI Sequencing,” *Quality Paradox, and AI Sequencing (January 31, 2026)*.
- McKinsey & Company (2025) “New front door to the internet: Winning in the age of AI search,” McKinsey’s Growth, Marketing, & Sales Practice.
- Mullainathan, Sendhil and Andrei Shleifer (2005) “The market for news,” *American economic review*, 95 (4), 1031–1053.
- Peitz, Martin and Tommaso M Valletti (2008) “Content and advertising in the media: Pay-tv versus free-to-air,” *international Journal of industrial organization*, 26 (4), 949–965.
- Pew Research Center (2025) “What Web Browsing Data Tells Us About How AI Appears Online,” Pew Research Center, <https://pewrsr.ch/4kDb6I5>.
- del Rio-Chanona, R Maria, Nadzeya Laurentsyeva, and Johannes Wachs (2024) “Large language models reduce public knowledge sharing on online Q&A platforms,” *PNAS nexus*, 3 (9), pgae400.
- Rochet, Jean-Charles and Jean Tirole (2006) “Two-sided markets: a progress report,” *RAND Journal of Economics*, 37 (3), 645–667.

- Sen, Debapriya (2005) “On the coexistence of different licensing schemes,” *International Review of Economics & Finance*, 14 (4), 393–413.
- Taylor, Greg (2013) “Search Quality and Revenue Cannibalization by Competing Search Engines,” *Journal of Economics & Management Strategy*, 22 (3), 445–467, <https://doi.org/10.1111/jems.12027>.
- Teh, Tat-How and Julian Wright (2024) “Competitive bottlenecks and platform spillovers,” *Available at SSRN*.
- Wu, Yihang, Jiajun Tang, Jinfei Liu, Haifeng Xu, and Fan Yao (2026) “Do AI Overviews Benefit Search Engines? An Ecosystem Perspective,” <https://arxiv.org/abs/2601.22493>.
- Zhao, Hangcheng and Ron Berman (2026) “Strategic Response of News Publishers to Generative AI,” <https://arxiv.org/abs/2512.24968>.

A Omitted Proofs

The derivation of Proposition 4 is established in the main text. All other omitted proofs are presented here.

Proof of Lemma 1

Proof. Fix Q and $\bar{x}$. Consider an arbitrary mixing rule $\alpha(x)$ and let a denote the realized AI-content location observed by user x . Let $d_\alpha(a, \theta) \in \{0, 1\}$ denote the user's optimal continuation decision under $\alpha(x)$ where $d_\alpha = 1$ if they continue search.

First, note that $q_i \leq r/\bar{\theta}$. To see this, observe that holding the platform's choices fixed, a marginal increase in q_i raises the continuation value for a user allocated to publisher i by at most one. Because the search cost θ is distributed uniformly on $[0, \bar{\theta}]$, the density is $1/\bar{\theta}$. Thus, a marginal increase in q_i increases a user's probability of clicking by at most $1/\bar{\theta}$. Since the total mass of users allocated to publisher i cannot exceed one, the derivative of publisher i 's total expected user traffic with respect to q_i is bounded above by $1/\bar{\theta}$. Setting the marginal revenue $r/\bar{\theta}$ equal to the marginal cost $C'(q_i) = q_i$ implies that any interior best response satisfies $q_i \leq r/\bar{\theta}$. Assumption 1, hence, implies $\gamma|q_1 - q_2| < t$.

Next, compare the arbitrary mixing rule $\alpha(x)$ with $\hat{\alpha}(x) = x$. For a user located at x , let the AI-content utility generated by the content being located at a be

$$u^{AI}(x; a) = \gamma[(1 - a)q_1 + aq_2] - t|a - x|.$$

Comparing this to when AI content is located at exactly x , we have

$$\begin{aligned} u^{AI}(x; x) - u^{AI}(x; a) &= \gamma(x - a)(q_2 - q_1) + t|a - x| \\ &\geq (t - \gamma|q_2 - q_1|) \cdot |a - x| \geq 0. \end{aligned}$$

Thus, perfectly personalized AI provides weakly greater user utility state by state.

In addition, under $\hat{\alpha}(x) = x$, the user perfectly learns their location and can imitate the decision they would have taken under the mixing rule $\alpha(x)$ by choosing $\hat{d}(x, \theta) = d_\alpha(\alpha(x), \theta)$.

There are two cases:

1. If $\hat{d}(x, \theta) = 0$, the payoff under $\hat{\alpha}(x) = x$ is $u^{AI}(x; x) \geq u^{AI}(x; \alpha(x))$. So, the user is weakly better off when $\hat{\alpha}(x) = x$.
2. If $\hat{d}(x, \theta) = 1$ and the user is allocated to publisher i , the payoff under $\hat{\alpha}(x) = x$ is $\max\{u^{AI}(x; x), u_i(x)\} - \theta \geq \max\{u^{AI}(x; \alpha(x)), u_i(x)\} - \theta$. So, the user is again weakly better off when $\hat{\alpha}(x) = x$.

Therefore, the mixing rule $\hat{\alpha}(x)$ weakly dominates every alternative mixing rule. To establish that $\hat{\alpha}(x) = x$ is uniquely optimal almost everywhere, note that for any alternative location $a \neq x$, the AI utility gain is strictly positive. For every user x , there exists a positive-measure interval of high search costs near $\bar{\theta}$ where the user chooses not to continue searching and receives exactly the AI utility such that any alternative mixing rule differing from $\hat{\alpha}(x) = x$ on a positive-measure set of locations yields a strictly lower expected user surplus. Thus, because the platform's objective is strictly increasing in user surplus, $\alpha^I(x) = x$ almost everywhere.

Finally, under $\alpha^I(x) = x$, users know their true location before deciding whether to continue search, so the optimal allocation threshold can be routinely shown as $\bar{x}^I = 1/2 + \frac{q_1 - q_2}{2t}$. ■

Proof of Proposition 1

Proof. Given $\alpha^I(x) = x$ and $\bar{x}^I$, taking the difference of publisher FOCs from (7) results in

$$(q_1 - q_2) \left[1 - \frac{r(2 - \gamma)}{2t\bar{\theta}} \right] = 0 \iff q_1 = q_2.$$

Thus, the interior equilibrium qualities are the same for both publishers implying that $\bar{x}^I = 1/2$.

We next solve for $q^I(\gamma)$ and prove that $q^I(\gamma)$ is decreasing in γ . Solving publisher 1's FOC (wlog) for (6) gives us

$$r \int_{\mathcal{X}_i} f(v + q_1(1 - \gamma(1 - \alpha(x))) - \gamma q_2 \alpha(x) - tx) \cdot [1 - \gamma(1 - \alpha(x))] dx = q_1.$$

In a symmetric equilibrium, $\bar{x} = 1/2$ and $\alpha(x) = x$ which implies that q^I solves

$$r \int_0^{1/2} f(v + q^I(1 - \gamma) - tx) \cdot [1 - \gamma(1 - x)] dx = q^I.$$

Under uniform distribution, this simplifies to $q^I = \frac{r(4-3\gamma)}{8\bar{\theta}}$ and $v + q^I(\gamma) \leq \bar{\theta} \forall \gamma \in [0, 1]$ so the solution is interior. This is the unique equilibrium as each publisher's profit function is strictly concave given $\alpha(x)$ and $\bar{x}$:

$$\frac{d^2 \pi_i}{dq_i^2} = -1 < 0.$$

By the implicit function theorem,

$$\frac{dq^I}{d\gamma} = -\frac{r \int_0^{1/2} q^I f'(\Delta^I)(1 - \gamma(1 - x))dx + r \int_0^{1/2} (1 - x)f(\Delta^I)dx}{1 - r(1 - \gamma) \int_0^{1/2} f'(\Delta^I)(1 - \gamma(1 - x))dx}$$

where $\Delta^I = v + q^I(1 - \gamma) - tx$. Observe that the denominator is positive, so we just need to verify that the numerator is non-negative. We prove that the numerator is non-negative with a more general distribution mentioned in Section C.1 using Assumption 2 (which uniform distribution satisfies). Note that $q^I \leq \frac{r\bar{f}}{2}(1 - \gamma/2)$. Using Assumption 2, we have for $x \in [0, 1/2]$,

$$\begin{aligned} r \cdot \sup f \cdot \sup \left(-\frac{f'}{f} \right) &\leq 1 \Rightarrow \frac{r}{2} \cdot \sup f \cdot \sup \left(-\frac{f'}{f} \right) (1 - \frac{\gamma}{2}) \leq \frac{1}{2} \\ &\Rightarrow q^I \cdot \sup \left(-\frac{f'}{f} \right) \leq \frac{1}{2} \\ &\Rightarrow q^I \cdot \sup \left(-\frac{f'}{f} \right) \frac{1 - \gamma(1 - x)}{2(1 - x)} \leq \frac{1}{2} \\ &\Rightarrow q^I \cdot \sup \left(-\frac{f'}{f} \right) (1 - \gamma(1 - x)) \leq 1 - x \end{aligned}$$

which is a sufficient condition for the numerator being non-negative. Next, given γ , equilibrium publisher profits with AI search features are

$$\pi^I(\gamma) = r \int_0^{1/2} F(\Delta^I)dx - \frac{(q^I)^2}{2}$$

where $\Delta^I = v + q^I(1 - \gamma) - tx$. Differentiating w.r.t. γ gives us

$$\begin{aligned} \frac{d\pi^I}{d\gamma} &= -r \int_0^{1/2} q^I f(\Delta^I)dx + \left[r \int_0^{1/2} f(\Delta^I) \frac{\partial \Delta^I}{\partial q^I} dx - q^I \right] \frac{dq^I}{d\gamma} \\ &= -r \int_0^{1/2} \left[\gamma x \frac{dq^I}{d\gamma} + q^I \right] f(\Delta^I)dx \\ &= -rq^I \mathcal{A} - r\gamma \mathcal{B} \frac{dq^I}{d\gamma} \end{aligned}$$

where $\mathcal{A} := \int_0^{1/2} f(\Delta^I) dx$ and $\mathcal{B} := \int_0^{1/2} x f(\Delta^I) dx$. Recall,

$$\begin{aligned} \frac{dq^I}{d\gamma} &= -\frac{r[q^I \int_0^{1/2} f'(\Delta^I)(1 - \gamma(1 - x))dx + \int_0^{1/2} f(\Delta^I)(1 - x)dx]}{1 - r(1 - \gamma) \int_0^{1/2} f'(\Delta^I)(1 - \gamma(1 - x))dx} \\ &= -\frac{r(q^I \mathcal{C} + \mathcal{A} - \mathcal{B})}{1 - r(1 - \gamma)\mathcal{C}} \end{aligned}$$

where $\mathcal{C} := \int_0^{1/2} f'(\Delta^I)(1 - \gamma(1 - x))dx$. We can also easily rewrite $q^I = r[(1 - \gamma)\mathcal{A} + \gamma\mathcal{B}]$. Finally, we have

$$\begin{aligned} \frac{d\pi^I}{d\gamma} &= -r \frac{r[(1 - \gamma)\mathcal{A}^2 + \gamma\mathcal{B}^2] - (q^I)^2 \mathcal{C}}{1 - r(1 - \gamma)\mathcal{C}} < 0 \iff (q^I)^2 \mathcal{C} < r[(1 - \gamma)\mathcal{A}^2 + \gamma\mathcal{B}^2] \\ &\iff \mathcal{C} < \frac{r[(1 - \gamma)\mathcal{A}^2 + \gamma\mathcal{B}^2]}{(q^I)^2} \end{aligned}$$

which is always satisfied as $\mathcal{C} < 0$ and the RHS is greater than 0. ■

Proof of Proposition 2

Proof. By (8), we have that

$$\begin{aligned} \mathcal{U}^I(\gamma) &= \gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \int_0^{1/2} [v + q^I(\gamma)(1 - \gamma) - tx]^2 dx \\ &= \gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \int_0^{1/2} [B(\gamma) + t/4 - tx]^2 dx \\ &= \gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \left[\frac{(B(\gamma))^2}{2} + \frac{t^2}{96} \right] \end{aligned}$$

where $B(\gamma) := v + q^I(\gamma)(1 - \gamma) - t/4$. Using Proposition 1,

$$\begin{aligned} \frac{d^2 \mathcal{U}^I}{d\gamma^2} &= -\frac{3r}{4\bar{\theta}} + \frac{1}{\bar{\theta}} \left[(B'(\gamma))^2 + B''(\gamma)B(\gamma) \right] = -\frac{3r}{4\bar{\theta}} + \frac{1}{\bar{\theta}} \left[\frac{r^2(6\gamma - 7)^2}{64\bar{\theta}^2} + \frac{3rB(\gamma)}{4\bar{\theta}} \right] \\ &= \frac{3r}{4\bar{\theta}} \left(\frac{B(\gamma)}{\bar{\theta}} - 1 \right) + \frac{r^2(6\gamma - 7)^2}{64\bar{\theta}^3}. \end{aligned}$$

Both terms are maximized when $\gamma = 0$. Therefore, it suffices to show that

$$\frac{49r}{16\bar{\theta}} \leq 3(\bar{\theta} - B(0)) = 3(\bar{\theta} - v - \frac{r}{2\bar{\theta}} + t/4).$$

By Assumption 1, we have

$$\frac{49r}{16\bar{\theta}} < \frac{49(\bar{\theta} - v)}{32} < \frac{72(\bar{\theta} - v) + 24t}{32} < 3(\bar{\theta} - B(0))$$

which gives the result. Next, we have

$$\begin{aligned} \mathcal{U}^I(0) &= \mathcal{U}(1/2, 0, Q^N) = \frac{1}{\bar{\theta}} \int_0^{1/2} (v + q^N - tx)^2 dx \\ &> \frac{1}{\bar{\theta}} \int_0^{1/2} (v + q^N - t/4)^2 dx = \mathcal{U}^N \end{aligned}$$

where the rest of the proposition is immediate. ■

Proof of Proposition 3

Proof. I first prove the comparative statics results in item (i). Suppose $\gamma^* \in (0, 1)$ is the γ such that

$$\frac{d\mathcal{U}^I(\gamma)}{d\gamma} = \underbrace{q^I(\gamma) + \gamma \frac{q^I(\gamma)}{d\gamma} + \frac{1}{\bar{\theta}}(v - t/4 + \frac{r(1-\gamma)(4-3\gamma)}{8\bar{\theta}}) \frac{r(6\gamma-7)}{8\bar{\theta}}}_{=H} = 0.$$

Solving the above gives

$$v^* = \frac{-18(\gamma^*)^3 r + 63(\gamma^*)^2 r - 73\gamma^* r + 12\gamma^* t \bar{\theta} + 48\gamma^* \bar{\theta}^2 + 28r - 14t\bar{\theta} - 32\bar{\theta}^2}{8\bar{\theta}(6\gamma^* - 7)}$$

when $\gamma = \gamma^*$. By the implicit function theorem,

$$\left. \frac{d\gamma^*}{dr} = -\frac{H_r}{H_\gamma} \right|_{\gamma=\gamma^*} < 0, \left. \frac{d\gamma^*}{dt} = -\frac{H_t}{H_\gamma} \right|_{\gamma=\gamma^*} > 0, \left. \frac{d\gamma^*}{dv} = -\frac{H_v}{H_\gamma} \right|_{\gamma=\gamma^*} < 0$$

as $H_\gamma < 0$ by concavity of user surplus and

$$\begin{aligned} H_r|_{\gamma=\gamma^*} &= -\frac{r(1-\gamma^*)(4-3\gamma^*)(7-6\gamma^*)}{64\bar{\theta}^3} < 0, H_t|_{\gamma=\gamma^*} = \frac{r(14-12\gamma^*)}{64\bar{\theta}^2} > 0, \\ H_v|_{\gamma=\gamma^*} &= \frac{r(48\gamma^* - 56)}{64\bar{\theta}^2} < 0. \end{aligned}$$

For item (ii), it is immediately true if γ^* is 0 or 1. At interior γ^* ,

$$SW^{II'}(\gamma^*) = \underbrace{\frac{d\mathcal{U}^I(\gamma^*)}{d\gamma} + \Phi'(\mathcal{U}^I(\gamma^*)) \frac{d\mathcal{U}^I(\gamma^*)}{d\gamma}}_{=0} + \underbrace{\sum_{i=1}^2 \frac{d\pi_i^I(\gamma^*)}{d\gamma}}_{\leq 0}.$$

Note publisher profits cannot be increasing in γ by Proposition 1. In addition, $\frac{d\mathcal{U}^I}{d\gamma} \leq 0$ for all $\gamma \geq \gamma^*$ by concavity established in Proposition 2 and Φ is non-decreasing. Therefore, SW^I is non-increasing on $[\gamma^*, 1]$, and increasing γ above γ^* reduces welfare, so $\gamma^{SO} \leq \gamma^*$. From Proposition 2, $\mathcal{U}^I(\gamma)$ is concave in γ , so

$$\left. \frac{d\mathcal{U}^I}{d\gamma} \right|_{\gamma=1} = \frac{-2r}{8\bar{\theta}} + (v - t/4) \cdot \frac{-r}{8\bar{\theta}^2} < 0 \Rightarrow \gamma^* < 1.$$

For the third item, social welfare at γ^{SO} being higher than without AI search features is immediately true when considering that social welfare is higher when $\gamma = 0$ than without AI search features by using Proposition 2. The last part is established given that social welfare at γ^* is lower than without AI iff

$$\begin{aligned} D &:= \pi^N - \pi^I(\gamma^*) - [\mathcal{U}^I(\gamma^*) - \mathcal{U}^N] > 0 \\ \iff \frac{t^2}{96\bar{\theta}} + \frac{r\gamma^*(4 - 3\gamma^*)}{8\bar{\theta}} - \frac{r\gamma^*(7 - 3\gamma^*)}{8\bar{\theta}^2} \left[v - \frac{t}{4} + \frac{r(8 - 7\gamma^* + 3(\gamma^*)^2)}{16\bar{\theta}} \right] &< \frac{r^2\gamma^*(32 - 15\gamma^*)}{128\bar{\theta}^2}. \end{aligned}$$

From the proof of Proposition 2,

$$\lim_{\bar{\theta} \rightarrow \infty} \lim_{r \rightarrow \bar{\theta}t} \frac{d\mathcal{U}^I}{d\gamma} \rightarrow t \frac{(2 - 3\gamma)}{4} = 0 \iff \gamma = 2/3.$$

Therefore, γ^* approaches 2/3 as $\bar{\theta}$ goes to infinity and $r/\bar{\theta}$ is sufficiently close to t . This implies that

$$\lim_{\bar{\theta} \rightarrow \infty} \lim_{r \rightarrow \bar{\theta}t} D \rightarrow \frac{t(11t - 16)}{96} > 0$$

By continuity, there must exist some threshold $\hat{\theta}$ and ratio threshold $\hat{k} < t$ such that social welfare is lower with AI search features than without for $\bar{\theta} > \hat{\theta}$ and $r \in (\hat{k}\bar{\theta}, \bar{\theta}t)$ where γ^* is also interior. ■

Proof of Proposition 5

Proof. While the functional form of $u^{AI}(x)$ is different now, it is routine to check that the logic in Lemma 1 remains the same such that $\alpha(x) = x$ is preferred by the platform. I first prove item (i) by showing that blocking is a dominant strategy. Let $R_i(q_{-i}; b_i)$ denote publisher i 's best response function given b_i . Without loss of generality, consider publisher 2's problem and the associated FOC:

$$\max_{q_2} r \int_{\mathcal{X}_2} F(\Delta_2) dx - C(q_2) \Rightarrow \frac{r}{\theta} \int_{\bar{x}}^1 \frac{\partial \Delta_2}{\partial q_2} dx = q_2$$

where $\frac{\partial \Delta_2}{\partial q_2} = 1 - \gamma(1 - b_2)x$. By the above FOC, and recalling that $\bar{x}^I = \frac{1}{2} + \frac{q_1 - q_2}{2t}$ in equilibrium, we know that

$$\frac{dR_2}{dq_1} = \frac{\frac{r}{2t\theta}[-1 + \gamma(1 - b_2)\bar{x}^I]}{1 - \frac{r}{2t\theta}[1 - \gamma(1 - b_2)\bar{x}^I]} < 0$$

In fact, $|\frac{dR_2}{dq_1}| < 1$ by Assumption 1. Note that publisher 2's reaction function is unaffected by publisher 1's decision to block. Now, consider publisher 1's FOC as

$$q_1 = \frac{r}{\theta} \int_0^{\bar{x}} [1 - \gamma(1 - b_1)(1 - x)] dx$$

which means that publisher 1's best response quality is highest when $b_1 = 1$, i.e. $R_1(q_2; 1) > R_1(q_2; b_1) \forall b_1 \neq 1$. Let q_1^k be the quality that solves $R_1(R_2(q_1, k), k) = q_1$. We know that $q_1^1 > q_1^j \forall j < 1$ because R_2 is unaffected by b_1 which also implies $R_2(q_1^1; b_2) < R_2(q_1^j; b_2) \forall j < 1$. Also, for each $(b_1, b_2) \in [0, 1]^2$, there is a unique sub-game equilibrium for publishers' qualities. Define $\Psi(q_1) = R_1(R_2(q_1))$ wlog. By the chain rule, $|\Psi'(q_1)| = |\frac{dR_1}{dq_2}| \cdot |\frac{dR_2}{dq_1}| < 1$, thereby Ψ is a contraction mapping. By the Banach fixed-point theorem, Ψ has a unique fixed point q_1^* and $q_2^* = R_2(q_1^*)$ is also unique.

For the sake of notation, let $\Pi_1^*(\bar{x}, b_1) = \max_{q_1} r \int_{\mathcal{X}_1} F(\Delta_1) dx - C(q_1)$ denote publisher 1's profit evaluated exactly at the continuation equilibrium for blocking choice b_1 given $\bar{x}$. Therefore,

$$\frac{d\Pi_1^*}{db_1} = \underbrace{\frac{\partial \Pi_1^*}{\partial b_1}}_{>0} + \underbrace{\frac{\partial \Pi_1^*}{\partial q_1} \frac{\partial q_1}{\partial b_1}}_{=0} + \underbrace{\frac{\partial \Pi_1^*}{\partial q_2}}_{<0} \underbrace{\frac{\partial q_2}{\partial b_1}}_{<0} + \underbrace{\frac{\partial \Pi_1^*}{\partial \bar{x}}}_{>0} \underbrace{\frac{d\bar{x}}{db_1}}_{>0} > 0$$

so full blocking is a dominant strategy. By symmetry, the same holds for publisher 2 so $(b_1, b_2) = (1, 1)$ is the unique equilibrium.

Next, to prove item (ii), observe that including an opt-out option lowers user surplus if and only if $\frac{d\mathcal{U}^I(\gamma)}{d\gamma}\Big|_{\gamma=0} > 0$ because we know $\mathcal{U}^I$ is concave. Recall that when θ is uniform,

$$\begin{aligned}\mathcal{U}^I(\gamma) &= \gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \int_0^{1/2} (v + q^I(\gamma)(1 - \gamma) - tx)^2 dx \\ &= \frac{\gamma r(4 - 3\gamma)}{8\bar{\theta}} + \frac{1}{\bar{\theta}} \left(\frac{A(\gamma)^2}{2} - \frac{A(\gamma)t}{4} + \frac{t^2}{24} \right)\end{aligned}$$

where $A(\gamma) = v + (1 - \gamma)\frac{r(4-3\gamma)}{8\bar{\theta}}$. Now,

$$\begin{aligned}\gamma^* > 0 &\iff \frac{d\mathcal{U}^I(\gamma)}{d\gamma}\Big|_{\gamma=0} > 0 \\ &\iff \frac{r}{2\bar{\theta}} + \frac{1}{\bar{\theta}}(A'(0)(A(0) - t/4)) > 0 \\ &\iff \frac{r}{2\bar{\theta}} - \frac{1}{\bar{\theta}}\frac{7r}{8\bar{\theta}}(v + \frac{r}{2\bar{\theta}} - \frac{t}{4}) > 0 \\ &\iff \underbrace{\frac{r}{32\bar{\theta}^3}}_{>0}(16\bar{\theta}^2 + 7t\bar{\theta} - 28v\bar{\theta} - 14r) > 0.\end{aligned}$$

Thus, the statement holds iff $\frac{16}{7}\bar{\theta}^2 + t\bar{\theta} - 4v\bar{\theta} - 2r > 0$.

For item (iii), social welfare given γ and content qualities $Q = (q_1, q_2)$, where $q_1 = q_2 = q$, is derived as

$$\begin{aligned}SW(\gamma, Q) &= 2\mathcal{U}(\bar{x}^I, \gamma, Q) + 2\frac{r}{\bar{\theta}} \left[\frac{v + (1 - \gamma)q}{2} - \frac{t}{8} \right] - q^2 \\ &= 2\gamma q + \frac{(v + (1 - \gamma)q)^2}{\bar{\theta}} + \frac{r - t/2}{\bar{\theta}}(v + (1 - \gamma)q) + \frac{t^2}{12\bar{\theta}} - \frac{rt}{4\bar{\theta}} - q^2.\end{aligned}$$

After some algebra,

$$\begin{aligned}SW(0, q^N) &< SW^I(\gamma^*) \\ \iff 2v - \frac{t}{2} + r + \underbrace{\frac{r(8 - 7\gamma^* + 3\gamma^{*2})}{8\bar{\theta}}}_{\leq \frac{r}{\bar{\theta}}} &< \underbrace{\frac{2\bar{\theta}(4 - 3\gamma^*)}{7 - 3\gamma^*}}_{\geq \bar{\theta}/2} + \underbrace{\frac{3r(8 - 3\gamma^*)}{8(7 - 3\gamma^*)}}_{\geq 3r/7}.\end{aligned}$$

Therefore, a sufficient condition is $2v - t/2 + r + r/\bar{\theta} < \frac{\bar{\theta}}{2} + \frac{3r}{7}$. ■

Proof of Proposition 6

Proof of item (i)

We separate the proof into two parts. First, we show that for any equilibrium such that both publishers accept the offer for a lump-sum scheme, there exists some ϕ such that the quality-adjusted scheme satisfies each publisher's participation constraint and is preferred by the platform. Second, we show that when t is sufficiently close to $\frac{r}{\theta}$, such that Assumption 1 is still satisfied, a lump-sum payment that induces only one publisher to accept dominates any type of equilibrium under quality-adjusted payments. Note that some algebra is omitted for brevity and can be easily verified through reliable software such as Mathematica.

Lemma A.1 *Under any equilibrium such that both publishers accept the offer for a lump-sum scheme and is preferred by the platform over the full-blocking benchmark, there exists some ϕ such that both publishers accept the payment and the outcome preferred by the platform.*

Proof. Let $T^*(\gamma)$ be the minimum lump sum payment T to sustain an acceptance equilibrium where both publishers accept. We first show that for ϕ such that $\phi q^I(\gamma) = T^*(\gamma)$, each publisher's participation constraint under the quality-adjusted payment scheme is also satisfied. Solving for the publisher's FOC in (10) and using a similar symmetry argument gives us the equilibrium quality as $q^{qual} = q^I(\gamma) + \phi > q^I(\gamma)$. This implies that $\pi^{qual}(\gamma, \phi) + \phi q^{qual} = \pi^I(\gamma) + \frac{\phi^2}{2} + \frac{r(1-\gamma)\phi}{2\theta}$.

For payment scheme $\mathcal{P}$, the rejecting publisher's profit is

$$\begin{aligned}\pi_R^{\mathcal{P}} &= r \cdot \int_0^{s_R^{\mathcal{P}}} F(\Delta_1) dx - \frac{(q_R^{\mathcal{P}})^2}{2} \\ &= \frac{r}{\theta} \left[v + q_R^{\mathcal{P}} - (t + \gamma q_A^{\mathcal{P}}) \frac{s_R^{\mathcal{P}}}{2} \right] s_R^{\mathcal{P}} - \frac{(q_R^{\mathcal{P}})^2}{2} \\ &= \frac{r}{\theta} \left[v s_R^{\mathcal{P}} + \left(\frac{r}{2\theta} - \frac{t}{2} - \frac{\gamma q_A^{\mathcal{P}}}{2} \right) (s_R^{\mathcal{P}})^2 \right]\end{aligned}$$

where $s_R^{\mathcal{P}}$ ($s_A^{\mathcal{P}}$) denotes the rejecting (accepting) publisher's market share and $q_R^{\mathcal{P}}$ ($q_A^{\mathcal{P}}$) denotes the rejecting (accepting) publisher's quality in equilibrium. Solving each publisher's maximization problem where publisher 1 accepts and publisher 2 rejects, wlog, gives us $\bar{x}^{\mathcal{P}} = s_A^{\mathcal{P}} = \frac{1}{2} + \frac{q_A^{\mathcal{P}} - q_R^{\mathcal{P}}}{2t}$. For the lump-sum payment, we have

$$q_A^L = \frac{r\bar{x}^L}{\theta} \left[1 - \gamma(1 - \bar{x}^L/2) \right], q_R^L = \frac{r(1 - \bar{x}^L)}{\theta}$$

and for the quality-based payment scheme, we have

$$q_A^{qual} = \frac{r\bar{x}^{qual}}{\bar{\theta}} \left[1 - \gamma(1 - \bar{x}^{qual}/2) \right] + \phi, q_R^{qual} = \frac{r(1 - \bar{x}^{qual})}{\bar{\theta}}.$$

By substituting the equilibrium qualities into the allocation threshold $\bar{x}^{\mathcal{P}}$ to define a fixed-point map $g^{\mathcal{P}} : [0, 1] \rightarrow [0, 1]$:

$$g^L(\bar{x}) := \frac{1}{2} + \frac{1}{2t} \left[\frac{r\bar{x}}{\bar{\theta}} (1 - \gamma(1 - \bar{x}/2)) - \frac{r(1 - \bar{x})}{\bar{\theta}} \right] \text{ and} \\ g^{qual}(\bar{x}) := g^L(\bar{x}) + \frac{\phi}{2t}.$$

To ensure this mapping supports an interior equilibrium, we restrict attention to feasible payment rates where $g^{qual}(1) \leq 1$, which requires $\phi \leq t - \frac{r}{\bar{\theta}}(1 - \gamma/2)$. Under this condition, $g^{\mathcal{P}}$ maps $[0, 1]$ into itself. We also have

$$\frac{dg^L}{d\bar{x}} = \frac{dg^{qual}}{d\bar{x}} = \frac{r}{2t\bar{\theta}} [2 - \gamma + \gamma\bar{x}] \leq \frac{r}{t\bar{\theta}}.$$

By the mean value theorem, for $\bar{x}, \bar{x}' \in [0, 1]$, $|g^{\mathcal{P}}(\bar{x}) - g^{\mathcal{P}}(\bar{x}')| \leq \frac{r}{t\bar{\theta}} |\bar{x} - \bar{x}'|$ which means $g^{\mathcal{P}}$ is a contraction mapping. By the Banach fixed point theorem, there exists a unique fixed point which implies that $\bar{x}^L < \bar{x}^{qual} \Rightarrow s_R^{qual} < s_R^L \Rightarrow q_R^{qual} < q_R^L$ and $q_A^{qual} > q_A^L$. Note that $\bar{x}^{qual} - \bar{x}^L = g^{qual}(\bar{x}^{qual}) - g^L(\bar{x}^L) \geq \frac{\phi}{2t}$ since g^L is increasing in $\bar{x}$. Also,

$$\pi_R^{qual} = \frac{r}{\bar{\theta}} \int_{\bar{x}^{qual}}^1 (v + q_R^{qual} - (t + \gamma q_A^{qual})(1 - x)) dx - \frac{(q_R^{qual})^2}{2} \\ = \frac{r}{\bar{\theta}} \left[v s_R^{qual} - \frac{t(1 + \gamma) - r/\bar{\theta}}{2} (s_R^{qual})^2 + \frac{\gamma(2t - r/\bar{\theta})}{2} (s_R^{qual})^3 \right].$$

Next, we have that

$$\frac{d\pi_R^{qual}}{ds_R^{qual}} = \frac{r}{\bar{\theta}} \left[v - (t(1 + \gamma) - r/\bar{\theta}) s_R^{qual} + \frac{3\gamma}{2} (2t - r/\bar{\theta}) (s_R^{qual})^2 \right] \geq \frac{r\gamma t}{4\bar{\theta}}.$$

To see the last inequality, let $a := t(1 + \gamma) - r/\bar{\theta}$ and let $b := \frac{3\gamma}{2} (2t - r/\bar{\theta})$. We have two cases:

1. If $a \geq 2b$, then $\frac{d\pi_R^{qual}}{ds_R^{qual}}$ is minimized at $s_R^{qual} = 1$ such that $\frac{d\pi_R^{qual}}{ds_R^{qual}} = \frac{r}{\bar{\theta}}(v - a + b) = \frac{r}{\bar{\theta}}(v + t(2\gamma - 1) + \frac{r}{\bar{\theta}}(1 - \frac{3\gamma}{2})) > \frac{r\gamma t}{4\bar{\theta}}$ under Assumption 1 and $v > t$.
2. If $a < 2b$, the minimum value of $\frac{d\pi_R^{qual}}{ds_R^{qual}}$ is $\frac{r}{\bar{\theta}}(v - \frac{a^2}{4b})$. Therefore, because $v > t$, it suffices

to show that

$$1 - \frac{\gamma}{4} \geq \frac{a^2}{4bt} = \frac{(1 + \gamma - r/(\bar{\theta}t))^2}{6\gamma(2 - r/(\bar{\theta}t))}.$$

Note the RHS is decreasing in $r/(\bar{\theta}t)$, so the RHS cannot be greater than $\frac{(1+\gamma)^2}{12\gamma}$ which is less than $1 - \gamma/4$ for $\gamma > (5 - \sqrt{21})/4$. If $\gamma < (5 - \sqrt{21})/4$, we know that $r/(\bar{\theta}t) > \frac{1-5\gamma}{1-3\gamma}$ from the assumption that $a < 2b$. At $r/(\bar{\theta}t) = \frac{1-5\gamma}{1-3\gamma}$, the RHS is $\frac{3\gamma(1-\gamma)}{2(1-3\gamma)} \leq 1 - \frac{\gamma}{4}$ for $\gamma \in [0, (5 - \sqrt{21})/4]$.

Because π_R^L is the same as $\pi_R^{qual}|_{\phi=0}$ and $\frac{d\pi_R^{qual}}{ds_R^{qual}}$ is bounded below by a positive value,

$$\pi_R^L - \pi_R^{qual} \geq \frac{r\gamma t}{4\bar{\theta}} \cdot \frac{\phi}{2t} = \frac{\gamma r \phi}{8\bar{\theta}} > 0$$

by the mean value theorem.

Therefore, at $\phi = T^*/q^I(\gamma)$,

$$\begin{aligned} \pi^{qual}(\gamma, \phi) + \phi q^{qual} - \pi_R^{qual} &= \pi^I(\gamma) + \frac{\phi^2}{2} + \frac{r(1-\gamma)\phi}{2\bar{\theta}} - \pi_R^L + \pi_R^L - \pi_R^{qual} \\ &= \frac{\phi^2}{2} + \frac{r(1-\gamma)\phi}{2\bar{\theta}} - T^* + \pi_R^L - \pi_R^{qual} \\ &\geq \frac{\phi^2}{2} + \frac{r(1-\gamma)\phi}{2\bar{\theta}} + \frac{\gamma r \phi}{8\bar{\theta}} - T^* \\ &= \frac{\phi^2}{2} + q^I(\gamma) \frac{T^*}{q^I(\gamma)} - T^* = \frac{\phi^2}{2} > 0. \end{aligned}$$

so whenever the lump-sum payment scheme is feasible, the quality-based scheme is also feasible. Finally, we need to show that the platform is better off for a $\phi = T^*(\gamma)/q^I(\gamma)$. we have

$$\begin{aligned} \Gamma^{qual} - \Gamma^L &= \left[\mathcal{U}(\bar{x}^I, \gamma, Q^{qual}(\gamma, \phi)) - \mathcal{U}^I(0) - 2\phi q^{qual} \right] \\ &\quad - \left[\mathcal{U}^I(\gamma) - \mathcal{U}^I(0) - 2T^*(\gamma) \right] \\ &= \mathcal{U}(\bar{x}^I, \gamma, Q^{qual}(\gamma, \phi)) - \mathcal{U}^I(\gamma) - 2\phi q^{qual} + 2T^*(\gamma) \\ &= \mathcal{U}(\bar{x}^I, \gamma, Q^{qual}(\gamma, \phi)) - \mathcal{U}^I(\gamma) - 2\phi^2. \end{aligned}$$

The above expression is greater than 0 if

$$\mathcal{U}(\bar{x}^I, \gamma, Q^{qual}(\gamma, \phi)) - \mathcal{U}^I(\gamma) > \gamma(q^{qual} - q^I(\gamma)) = \gamma\phi > 2\phi^2 \iff \phi < \gamma/2.$$

But, assuming that $\Gamma^L > 0$, we have

$$\mathcal{U}^I(\gamma) - \mathcal{U}^I(0) > 2T^*(\gamma) = 2\phi q^I(\gamma).$$

Lastly, we have

$$\begin{aligned} \mathcal{U}^I(\gamma) - \mathcal{U}^I(0) &= \gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \int_0^{1/2} \left[(v + q^I(\gamma)(1 - \gamma) - tx)^2 - (v + q^N - tx)^2 \right] dx \\ &< \gamma q^I(\gamma) \end{aligned}$$

which implies that $\phi < \gamma/2$ completing the proof. ■

Lemma A.2 *There exists an admissible parameter region where $\bar{\theta}$ is sufficiently low ($v < \bar{\theta}$) and t is sufficiently close to the lower boundary $r/\bar{\theta}$ such that a lump-sum payment scheme dominates any quality-adjusted payment scheme.*

Proof. Consider an admissible sequence of parameter vectors satisfying $\bar{\theta} > v > t$ and Assumption 1 that approaches the lower boundary $t = r/\bar{\theta}$, with r , v , and $\bar{\theta}$ varying jointly whenever necessary to preserve these restrictions. One such example is $\bar{\theta}_\epsilon = \epsilon$, $v_\epsilon = \frac{9}{10}\epsilon$, $r_\epsilon = \frac{1}{25}\epsilon^2$, $t_{\epsilon,\eta} = (\frac{1}{25} + \eta)\epsilon$ where $\epsilon, \eta > 0$ approaches 0. We first solve for the equilibrium if only one publisher accepts, say publisher 1 without loss of generality, under a quality-adjusted payment. By Lemma A.1, we know that the accepting and rejecting publisher's equilibrium content quality is

$$q_A^{qual} = \frac{r\bar{x}^{qual}}{\bar{\theta}} \left[1 - \gamma(1 - \bar{x}^{qual}/2) \right] + \phi \text{ and } q_R^{qual} = \frac{r}{\bar{\theta}}(1 - \bar{x}^{qual}),$$

respectively. Given that the allocation threshold is $\bar{x}^{qual} = \frac{1}{2} + \frac{q_A^{qual} - q_R^{qual}}{2t}$, we can rewrite ϕ as

$$\phi = t(2\bar{x}^{qual} - 1) + \frac{r}{\bar{\theta}} \left[1 - (2 - \gamma)\bar{x}^{qual} - \gamma \frac{(\bar{x}^{qual})^2}{2} \right].$$

At $t = \frac{r}{\bar{\theta}}$, $\phi = \frac{r}{\bar{\theta}} \gamma \bar{x}^{qual} (1 - \frac{\bar{x}^{qual}}{2}) \Rightarrow q_A^{qual} = \frac{r}{\bar{\theta}} \bar{x}^{qual}$ and $q_R^{qual} = \frac{r}{\bar{\theta}} (1 - \bar{x}^{qual})$. The accepting

publisher's participation constraint gives

$$\begin{aligned}
\pi_A^{qual} + \phi q_A^{qual} &\geq \pi^I(0) \\
\iff \frac{r}{\bar{\theta}} \int_0^{\bar{x}^{qual}} (v + q_A^{qual}(1 - \gamma(1 - x)) - tx) dx - \frac{(q_A^{qual})^2}{2} + \phi q_A^{qual} &\geq \frac{r}{\bar{\theta}} \left(\frac{v}{2} - \frac{t}{8} \right) + \frac{r^2}{\bar{\theta}^2} \\
\iff \frac{r}{\bar{\theta}} \left[v \bar{x}^{qual} - \frac{t(\bar{x}^{qual})^2}{2} + q_A^{qual}(1 - \gamma) \bar{x}^{qual} + \gamma q_A^{qual} (\bar{x}^{qual})^2 / 2 \right] - \frac{(q_A^{qual})^2}{2} + \phi q_A^{qual} &\geq \frac{rv}{2\bar{\theta}} \\
\iff \frac{rv}{\bar{\theta}} \bar{x}^{qual} \geq \frac{rv}{2\bar{\theta}} \iff \bar{x}^{qual} \geq \frac{1}{2}.
\end{aligned}$$

Recall the rejecting publisher's equilibrium quality after deviating is

$$q^I(\gamma) + \phi = \frac{r(4 - 3\gamma)}{8\bar{\theta}} + \frac{r}{\bar{\theta}} \left(\gamma \bar{x}^{qual} - \frac{\gamma(\bar{x}^{qual})^2}{2} \right) = \frac{r}{8\bar{\theta}} N(\bar{x}^{qual})$$

where $N := 4 - 3\gamma + 8\gamma\bar{x}^{qual} - 4\gamma(\bar{x}^{qual})^2$. The rejecting publisher's participation constraint gives

$$\begin{aligned}
\pi_R^{qual} &\geq \frac{r}{\bar{\theta}} \int_0^{1/2} (v + (q^I(\gamma) + \phi)(1 - \gamma) - tx) dx - \frac{(q^I(\gamma) + \phi)^2}{2} + \phi(q^I(\gamma) + \phi) \\
&= \frac{r}{\bar{\theta}} \left[\frac{v}{2} - \frac{r}{8\bar{\theta}} + \frac{(1 - \gamma)rN}{16\bar{\theta}} \right] - \frac{r^2 N^2}{128\bar{\theta}^2} + \frac{r^2 N \gamma \bar{x}^{qual} (1 - \bar{x}^{qual}/2)}{8\bar{\theta}^2}.
\end{aligned}$$

The rejecting publisher's payoff in a singular acceptance equilibrium is

$$\begin{aligned}
\pi_R^{qual} &= \frac{r}{\bar{\theta}} \int_{\bar{x}^{qual}}^1 (v + q_R^{qual} - t(1 - x) - \gamma q_A^{qual}(1 - x)) dx - \frac{(q_R^{qual})^2}{2} \\
&= \frac{r}{\bar{\theta}} \left[(v + q_R^{qual})(1 - \bar{x}^{qual}) - (t + \gamma q_A^{qual}) \left(\frac{1 + (\bar{x}^{qual})^2}{2} - \bar{x}^{qual} \right) \right] - \frac{(q_R^{qual})^2}{2} \\
&= \frac{r}{\bar{\theta}} \left[v(1 - \bar{x}^{qual}) - \gamma \bar{x}^{qual} \frac{r}{2\bar{\theta}} (1 - \bar{x}^{qual})^2 \right].
\end{aligned}$$

After strenuous algebra, we can show the participation constraint is satisfied if and only if $\bar{x}^{qual} \leq 1/2$. Therefore, combining both constraints yields $\bar{x}^{qual} = 1/2 \Rightarrow q_A^{qual} = q_R^{qual} = q^N$ and $\phi = \frac{3\gamma r}{8\bar{\theta}}$.

Consider the singular acceptance equilibrium under a lump-sum payment where publisher 1 accepts wlog. In equilibrium, the accepting and rejecting publisher's equilibrium content quality is

$$q_A^L = \frac{r\bar{x}^L}{\bar{\theta}} \left[1 - \gamma(1 - \bar{x}^L/2) \right] \text{ and } q_R^L = \frac{r}{\bar{\theta}} (1 - \bar{x}^L),$$

respectively. Given that the allocation threshold is $\bar{x}^L = \frac{1}{2} + \frac{q_A^L - q_R^L}{2t}$, $\bar{x}^L$ must satisfy

$$\frac{r\gamma}{2\bar{\theta}}(\bar{x}^L)^2 + \left(\frac{r(2-\gamma)}{\bar{\theta}} - 2t\right)\bar{x}^L + \left(t - \frac{r}{\bar{\theta}}\right) = 0.$$

At $t = \frac{r}{\bar{\theta}}$, it must be that $\bar{x}^L = 0 \Rightarrow q_A^L = 0$ and $q_R^L = t$. The minimum payment for the platform is hence

$$T^* = \pi^I(0) - \pi_A^L = \pi^I(0) = \frac{rv}{2\bar{\theta}}$$

which means the utility of the platform under lump sum is

$$\begin{aligned} \mathcal{U}(\bar{x}^I, \gamma, 0, q_R^L) - T^* &= \frac{1}{2\bar{\theta}} \int_0^1 (v + t - t(1-x))^2 dx - \frac{rv}{2\bar{\theta}} \\ &= \frac{(v^2 + \frac{r}{\bar{\theta}}v + \frac{r^2}{3\bar{\theta}^2})}{2\bar{\theta}} - \frac{rv}{2\bar{\theta}}. \end{aligned}$$

This is larger than that of quality-adjusted iff

$$12v(1 + \gamma - 2\bar{\theta}) + \frac{r}{\bar{\theta}}(6 + 3\gamma - 2\gamma^2) + 9\gamma r - 12\gamma\bar{\theta} > 0.$$

Because the equilibrium outcomes and the relevant payoff differences are continuous in the primitives, these strict inequalities continue to hold in a nonempty open neighborhood within the admissible parameter space. Notice the LHS is decreasing in θ and its limit is positive as $\bar{\theta} - v$ becomes small jointly with r . So for sufficiently small $\bar{\theta}$ such that Assumption 1 still holds, this condition holds.

Now, suppose both publishers accept some quality-adjusted payment rate. We know from the singular acceptance case that the rejecting publisher's payoff in a singular acceptance is not satisfied iff $\bar{x} \geq 1/2$. Thus, $\phi = \gamma \frac{r}{\bar{\theta}} \bar{x}^{qual} (1 - \bar{x}^{qual}/2) \in [\frac{3\gamma r}{8\bar{\theta}}, \frac{r\gamma}{2\bar{\theta}}] \Rightarrow q^{qual} = q^I(\gamma) + \phi \in [\frac{r}{2\bar{\theta}}, \frac{r(4+\gamma)}{8\bar{\theta}}]$. When both publishers accept, the platform's net utility is

$$\mathcal{U}(\bar{x}^I, \gamma, Q^{qual}) - 2\phi q^{qual} = \gamma q^{qual} + \frac{1}{\bar{\theta}} \left[0.5A(q^{qual})^2 - \frac{r}{4\bar{\theta}}A(q^{qual}) + \frac{r^2}{24\bar{\theta}^2} \right] - 2(q^{qual} - q^I(\gamma))q^{qual}$$

where $A(q^{qual}) = v + (1 - \gamma)q^{qual}$. This does not exceed the platform's net utility under the lump sum case iff

$$\begin{aligned} r^2 + 4q^{qual}\bar{\theta}^2((\gamma - 1)(2v + q^{qual}(1 - \gamma)) - 2(\gamma - 2q^{qual})\bar{\theta}) + 2r\bar{\theta}(q^{qual} \\ + 3v - 4q^{qual}\bar{\theta} - 2v\bar{\theta} + q^{qual}\gamma(3\bar{\theta} - 1)) \geq 0 \end{aligned}$$

which holds for sufficiently small $\bar{\theta}$; substituting the specific sequence $\bar{\theta} = \epsilon$, $v = 0.9\epsilon$, and $r = 0.04\epsilon^2$, the left-hand side reduces to an expression of order ϵ^4 whose leading coefficient is strictly positive for all $\gamma \in [0, 1]$. ■

By continuity, these strict comparisons continue to hold in a nonempty open neighborhood of the admissible limiting sequence.

Finally, we compare the social welfare under singular lump-sum to the baseline. As $t \rightarrow \kappa := \frac{r}{\bar{\theta}}$, social welfare under singular lump-sum is

$$\frac{v^2 + \kappa v + \kappa^2/3}{\bar{\theta}} + \kappa \int_0^1 (v + \kappa - t(1-x))dx - \kappa^2/2 \rightarrow \frac{v^2 + \kappa v + \kappa^2/3}{\bar{\theta}} + \kappa v.$$

Social welfare under the baseline as $t \rightarrow \kappa$ is

$$2\gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \left[(v + (1-\gamma)q^I(\gamma))^2 - \kappa \frac{v + (1-\gamma)q^I(\gamma)}{2} + \frac{\kappa^2}{12} \right] + \kappa(v + (1-\gamma)q^I(\gamma) - \kappa^2/4) - (q^I(\gamma))^2.$$

After some tedious algebra, the difference in social welfare under singular lump-sum and baseline is

$$\frac{\kappa}{64\bar{\theta}} \left[\underbrace{16\kappa + 32v}_{\geq 48\kappa} + \underbrace{112\gamma v - 64\gamma\bar{\theta} + 48\gamma^2(\bar{\theta} - v)}_{\geq 0 \text{ for } \bar{\theta} \text{ sufficiently close to } v} + \underbrace{\kappa\gamma(28 + 32\bar{\theta} - 61\gamma - 15\gamma\bar{\theta} + 42\gamma^2 - 9\gamma^3)}_{\geq -33\kappa} \right].$$

Thus, the difference as $t \rightarrow r/\bar{\theta}$ must be greater than 0.

Proof of item (ii)

Consider an admissible sequence along which t , v , and $\bar{\theta}$ become jointly large while the ratios $v/t > 1$ and $\bar{\theta}/t > v/t$ remain fixed and $\bar{\theta} + t/2 > 2v$. Assumption 1 is then satisfied for all sufficiently large parameter values. As $t \rightarrow \infty$, $\bar{x}^{qual} \rightarrow 1/2$ under singular acceptance. Therefore, it suffices to show that for $\bar{x}^{qual} = 1/2$, an equilibrium where both publishers accept the payment is preferable to a singular acceptance equilibria. From Lemma A.2, we established that when $\bar{x}^{qual} = 1/2$, $q_1 = q_2 = q^N$ and $\phi = \frac{3\gamma r}{8\bar{\theta}}$, such that the rejecting publisher's participation constraint is binding. Then, the platform's net benefit if both publishers accept is

$$\begin{aligned} (\gamma q^N - \frac{\gamma q^N}{2}) - \phi q^N + \frac{1}{\bar{\theta}} \int_0^{1/2} (A - tx)^2 dx - \frac{1}{2\bar{\theta}} \int_0^{1/2} (A + \gamma q^N x - tx)^2 dx \\ - \frac{1}{2\bar{\theta}} \int_0^{1/2} (A + \gamma q^N (1-x) - tx)^2 dx \end{aligned} \quad (13)$$

where $A := q^N(1 - \gamma) + v$. After some tedious algebra, (13) reduces to $(\frac{\gamma}{2} - \phi)q^N + \frac{\gamma q^N}{2\bar{\theta}} \left[\frac{t}{4} + \frac{q^N(4\gamma-6)}{6} - v \right]$. Note that this expression is increasing in t . In addition, the expression equals zero if and only if

$$t = 4v - 2\bar{\theta}(2 - 4\phi/\gamma) - \frac{2q^N(4\gamma - 6)}{3} = 4v + 3r - 4\bar{\theta} + \frac{2r(3 - 2\gamma)}{3\bar{\theta}}.$$

So, if $\bar{\theta}$ is sufficiently greater than v , then the platform's net benefit from both publishers accepting is greater as $t \rightarrow \infty$.

Using similar logic, we now compare the singular acceptance outcome of lump-sum to that of quality-adjusted where both accept in equilibrium. Assuming an interior solution as t and $\bar{\theta}$ grow large, $\bar{x}^L \rightarrow 1/2$. In this limit, the difference between the publishers' continuation qualities and $q^I(\gamma)$ shrinks to zero (vanishing at a rate of $1/\bar{\theta}$). Therefore, we can asymptotically approximate each publisher's content quality as $q^I(\gamma)$ and the minimum transfer as $T^* \approx \pi^N - \pi^I(\gamma)$. Now, consider $\phi = \frac{3\gamma r}{8\bar{\theta}}$ such that each publisher's constraint is satisfied (again recall Lemma A.2). The platform's net benefit from the quality-adjusted payment as opposed to lump-sum is then,

$$\begin{aligned} \gamma q^N - \frac{\gamma q^I(\gamma)}{2} - 2\phi q^N + T^* + \frac{2}{\bar{\theta}} \int_0^{1/2} (A - tx)^2 dx - \frac{1}{\bar{\theta}} \int_0^{1/2} (B + \gamma q^I(\gamma)x - tx)^2 dx \\ - \frac{1}{\bar{\theta}} \int_0^{1/2} (B + \gamma q^I(\gamma)(1 - x) - tx)^2 dx \end{aligned}$$

where $B := q^I(\gamma)(1 - \gamma) + v$. Observe that the first four terms do not depend on t and the last three terms are weakly greater than the last three terms in (13) as $q^I(\gamma) \leq q^N$. Thus, it must also be that for sufficiently large $\bar{\theta}$, the platform's net benefit from quality-adjusted is preferred over the lump-sum as $t \rightarrow \infty$. The last item can be easily shown by proving, for symmetric quality q , social welfare is increasing in q .

$$\begin{aligned} SW(\gamma, Q) &= 2\mathcal{U}(\bar{x}^I, \gamma, Q) + 2\frac{r}{\bar{\theta}} \int_0^{1/2} (v + (1 - \gamma)q - tx)dx - q^2 \\ &= 2\gamma q + \frac{2}{\bar{\theta}} \int_0^{1/2} (v + (1 - \gamma)q - tx)^2 dx + 2\frac{r}{\bar{\theta}} \int_0^{1/2} (v + (1 - \gamma)q - tx)dx - q^2 \\ \Rightarrow \frac{dSW}{dq} &= 2\gamma + \frac{2(1 - \gamma)}{\bar{\theta}}(v + (1 - \gamma)q - t/4) + \frac{r(1 - \gamma)}{\bar{\theta}} - 2q. \end{aligned}$$

Note that in this scenario $q \in [q^I(\gamma), q^N]$ depending on ϕ and $q^N = \frac{r}{2\bar{\theta}} \rightarrow 0$ as $\bar{\theta} \rightarrow \infty$. Therefore, for sufficiently high t and $\bar{\theta}$, $\frac{dSW}{dq} > 0$ which yields the result as any transfers between the platform and publishers cancel out. Continuity implies that the comparison

remains strict throughout a nonempty open region of sufficiently large admissible parameter vectors.

Proof of Proposition 7

Proof. Assuming symmetry and given γ , equation (12) becomes

$$q^{PPC} = q^I(\gamma) - \frac{r\gamma\bar{a}^2}{2\bar{\theta}}.$$

Again, the publisher's profit function is strictly concave so this is the unique best response conditional on $\bar{a}$ where $\bar{a}$ uniquely solves $\Delta\mathcal{U}_1(\bar{a}) = \Delta\mathcal{U}_2(\bar{a}) = R$. Also, because $q^{PPC} \leq q^I(\gamma)$, this solution is non-negative. The difference in user surplus in equilibrium with pay-per-crawl compared to the baseline with AI search features given γ and $Q^{PPC} = (q^{PPC}, q^{PPC})$ is

$$\begin{aligned} & 2 \int_0^{\bar{a}} \left[\gamma q^{PPC} - tx + \frac{(v + (1 - \gamma)q^{PPC})^2}{2\bar{\theta}} \right] dx + 2 \int_{\bar{a}}^{1/2} \left[\gamma q^{PPC} + \frac{(v + (1 - \gamma)q^{PPC} - tx)^2}{2\bar{\theta}} \right] dx \\ & - \mathcal{U}(1/2, \gamma, Q^{PPC}) = 2 \int_0^{\bar{a}} \left[-tx \left(1 - \frac{v + (1 - \gamma)q^{PPC}}{\bar{\theta}} \right) - \frac{t^2 x^2}{2\bar{\theta}} \right] dx \\ & = - \left[t\bar{a}^2 \left(1 - \frac{v + (1 - \gamma)q^{PPC}}{\bar{\theta}} \right) + \frac{t^2 \bar{a}^3}{3\bar{\theta}} \right] < 0 \end{aligned}$$

under Assumption 1 and given $\bar{a} \in (0, 1/2)$ (equivalently $R > 0$). Also, note that $\mathcal{U}^I(\gamma) \geq \mathcal{U}^I(1/2, \gamma, Q^{PPC})$ because $q^I \geq q^{PPC}$ completing the first item.

For the second item,

$$\frac{d\bar{a}}{dt} = - \frac{\frac{\partial \Delta\mathcal{U}_i(x)}{\partial t}}{\frac{\partial \Delta\mathcal{U}_i(\bar{a})}{\partial a}} < 0 \text{ and } \frac{d\bar{a}}{d\theta} = - \frac{\frac{\partial \Delta\mathcal{U}_i(x)}{\partial \theta}}{\frac{\partial \Delta\mathcal{U}_i(\bar{a})}{\partial a}} < 0$$

as $\Delta\mathcal{U}_i$ is increasing in x , $\frac{\partial \Delta\mathcal{U}_i(x)}{\partial t} \Big|_{x=\bar{a}} = \bar{a} \left[1 - \frac{v + (1 - \gamma)q^{PPC}}{\bar{\theta}} + \frac{t\bar{a}}{2\bar{\theta}} \right] + \frac{t\bar{a}^2}{2\bar{\theta}} > 0$, and $\frac{\partial \Delta\mathcal{U}_i(x)}{\partial \theta} = t\bar{a} \left[\frac{v + 2(1 - \gamma)q^{PPC}}{\bar{\theta}^2} - \frac{t\bar{a}}{2\bar{\theta}^2} \right] > 0$.

For the social welfare comparison, social welfare is better under the baseline if the user surplus gain is greater than the publisher profit difference. The publisher profit difference is

$$\begin{aligned} & \frac{r}{\bar{\theta}} \int_0^{\bar{a}} (v + (1 - \gamma)q^{PPC}) dx + \frac{r}{\bar{\theta}} \int_{\bar{a}}^{1/2} (v + (1 - \gamma)q^{PPC} - tx) dx - C(q^{PPC}) \\ & - \frac{r}{\bar{\theta}} \int_0^{1/2} (v + (1 - \gamma)q^I(\gamma) - tx) dx + C(q^I(\gamma)) = \frac{r^2 \gamma^2 \bar{a}^2 (1 - 2\bar{a}^2)}{16\bar{\theta}^2} + \frac{rt\bar{a}^2}{2\bar{\theta}} \leq \bar{a}^2 \left[\frac{r^2 \gamma^2}{16\bar{\theta}^2} + \frac{rt}{2\bar{\theta}} \right]. \end{aligned}$$

Using the work in item (i), the user surplus difference is

$$\begin{aligned}
& \mathcal{U}^I(\gamma) - \mathcal{U}(1/2, \gamma, Q^{PPC}) + t\bar{a}^2(1 - \frac{v + (1-\gamma)q^{PPC}}{\bar{\theta}}) + \frac{t^2\bar{a}^3}{3\bar{\theta}} \\
&= \frac{r\gamma^2\bar{a}^2}{2\bar{\theta}} + \frac{r\gamma(1-\gamma)\bar{a}^2}{2\bar{\theta}^2} \left[v - \frac{t}{4} + (1-\gamma)\frac{q^I(\gamma) + q^{PPC}}{2} \right] + t\bar{a}^2(1 - \frac{v + (1-\gamma)q^{PPC}}{\bar{\theta}}) + \frac{t^2\bar{a}^3}{3\bar{\theta}} \\
&\geq \bar{a}^2 \left[\frac{r\gamma^2}{2\bar{\theta}} + t(1 - \frac{v + (1-\gamma)q^{PPC}}{\bar{\theta}}) \right] \geq \bar{a}^2 \left[\frac{r\gamma^2}{2\bar{\theta}} + t(1 - \frac{v}{\bar{\theta}} - \frac{r}{2\bar{\theta}^2}) \right] > \bar{a}^2 \left[\frac{r\gamma^2}{2\bar{\theta}} + \frac{3t(\bar{\theta} - v)}{4\bar{\theta}} \right]
\end{aligned}$$

where last inequality comes from Assumption 1. Finally, we have

$$\bar{a}^2 \left[\frac{r\gamma^2}{2\bar{\theta}} + \frac{3t(\bar{\theta} - v)}{4\bar{\theta}} \right] > \bar{a}^2 \left[\frac{r^2\gamma^2}{16\bar{\theta}^2} + \frac{rt}{2\bar{\theta}} \right].$$

which can be satisfied if $r < \frac{3}{2}(\bar{\theta} - v)$. ■

Proof of Proposition 8

Proof. First, observe that conditional on the user masses $|\mathcal{Y}_j|$, each platform j 's continuation objective is increasing in the expected surplus of its existing users. Following Lemma 1, each platform j will set $\alpha_j^I(x) = x$ and $\bar{x}_j^I = \frac{1}{2} + \frac{q_1 - q_2}{2t}$ in equilibrium, as it maximizes user surplus and the number of users that will join the platform. Let $\mathcal{U}_j(\bar{x}_j^I, \gamma_j, Q)$ denote the expected user surplus for platform j given this fact. Publisher i 's FOC is

$$r \sum_j |\mathcal{Y}_j| \int_{\mathcal{X}_{ij}} f(u_i(x) - u_j^{AI}(x)) \frac{\partial(u_i(x) - u_j^{AI}(x))}{\partial q_i} dx = q_i.$$

Combining each publisher's FOC gives

$$\begin{aligned}
q_1 - q_2 &= \frac{r(2 - |\mathcal{Y}_A|\gamma_A - |\mathcal{Y}_B|\gamma_B)}{\bar{\theta}} (\bar{x} - \frac{1}{2}) \iff (q_1 - q_2) \left[1 - \frac{r(2 - |\mathcal{Y}_A|\gamma_A - |\mathcal{Y}_B|\gamma_B)}{2t\bar{\theta}} \right] = 0 \\
&\iff q_1 = q_2.
\end{aligned}$$

Therefore, publisher qualities are symmetric in equilibrium where $q_1^c = q_2^c = q^c(\gamma_A, \gamma_B) = \frac{r}{8\bar{\theta}} \left[4 - 3(|\mathcal{Y}_A|\gamma_A + |\mathcal{Y}_B|\gamma_B) \right]$. The threshold user that is indifferent between joining platform A and B is then

$$y^c = \frac{1}{2} + \frac{\mathcal{U}_A(\bar{x}_A^I, \gamma_A, Q^c) - \mathcal{U}_B(\bar{x}_B^I, \gamma_B, Q^c)}{2\tau}$$

where $Q^c = (q^c, q^c)$ implying $|\mathcal{Y}_A| = [0, y^c]$ and $|\mathcal{Y}_B| = [y^c, 1]$. This gives

$$\mathcal{U}_j(\bar{x}_j^I, \gamma_j, Q^c) = \gamma_j q^c + \frac{(v + (1 - \gamma_j)q^c - t/4)^2}{2\bar{\theta}} + \frac{t^2}{96\bar{\theta}}.$$

As $\tau \rightarrow \infty$, $|\mathcal{Y}_j| \rightarrow 1/2 \Rightarrow q^c \rightarrow q^\infty = \frac{r(8-3(\gamma_A+\gamma_B))}{16\bar{\theta}}$. Next, we show that the platform's objective is concave as $\tau \rightarrow \infty$:

$$\begin{aligned} \lim_{\tau \rightarrow \infty} \frac{d^2 |\mathcal{Y}_A| \mathcal{U}_A}{d\gamma_A^2} &= \lim_{\tau \rightarrow \infty} \left\{ |\mathcal{Y}_A| \frac{d^2 \mathcal{U}_A}{d\gamma_A^2} + \frac{\mathcal{U}_A}{2\tau} \left(\frac{d^2 \mathcal{U}_A}{d\gamma_A^2} - \frac{d^2 \mathcal{U}_B}{d\gamma_A^2} \right) + \frac{\frac{d\mathcal{U}_A}{d\gamma_A}}{\tau} \left(\frac{d\mathcal{U}_A}{d\gamma_A} - \frac{d\mathcal{U}_B}{d\gamma_A} \right) \right\} \\ &= \frac{1}{2} \left\{ \frac{(q^\infty)^2}{\bar{\theta}} - \frac{3r}{8\bar{\theta}} \left[1 - \frac{v + 2(1 - \gamma_A)q^\infty - t/4}{\bar{\theta}} \right] + \frac{9r^2(1 - \gamma_A)^2}{256\bar{\theta}^3} \right\} < 0 \end{aligned}$$

by Assumption 1. To formally establish equilibrium existence for finite but large τ , note that the user share $y = |\mathcal{Y}_A|$ solves the fixed-point equation $y = 1/2 + (\mathcal{U}_A(\bar{x}_A^I, \gamma_A, Q^c) - \mathcal{U}_B(\bar{x}_B^I, \gamma_B, Q^c))/(2\tau)$ (note Q^c depends on y). Because the derivative of the RHS w.r.t. y is uniformly bounded on the compact domain $(\gamma_A, \gamma_B) \in [0, 1]^2$, the right-hand side is a contraction mapping for sufficiently large τ . This guarantees a unique fixed point that is smooth in the platforms' strategies by the IFT. Consequently, the platform objectives and their second derivatives converge uniformly to their limits. As the limiting second derivative is continuous and strictly negative on a compact domain, it has a uniform negative upper bound. Therefore, $\exists \bar{\tau}$ s.t. $\tau \geq \bar{\tau} \Rightarrow$ the platform's objective function is concave.

We now show that $\gamma^c \geq \gamma^*$. If $\gamma^* = 0$, then the result is immediate. Therefore, we restrict attention to $\gamma^* \in (0, 1)$. If $\gamma^c = 1$, the strict inequality $\gamma^c > \gamma^*$ follows immediately since Proposition 3 establishes $\gamma^* < 1$. Furthermore, $\gamma^c \neq 0$; because $\gamma^* > 0$, the right-derivative of the monopoly user surplus at zero is strictly positive, and the strictly positive user-stealing effect under competition ensures the competitive objective's right-derivative is even larger. Thus, we can focus on $\gamma^c \in (0, 1)$. Platform j 's FOC is

$$|\mathcal{Y}_j| \cdot \Phi'(\mathcal{U}_j) \cdot \frac{d\mathcal{U}_j}{d\gamma_j} + \underbrace{\Phi(\mathcal{U}_j) \cdot \frac{d|\mathcal{Y}_j|}{d\gamma_j}}_{>0} = 0.$$

Observe that if $\gamma^* > 0$ then the γ that solves the FOC is also greater than 0. The right underbrace is derived by the following:

$$\frac{d|\mathcal{Y}_j|}{d\gamma_j} = \frac{1}{2\tau} \left[\frac{d\mathcal{U}_A}{d\gamma_A} - \frac{d\mathcal{U}_B}{d\gamma_A} \right] = \frac{1}{2\tau} \left[\underbrace{\frac{\partial \mathcal{U}_A}{\partial \gamma_A}}_{>0} \Big|_{q^c} + \underbrace{\frac{\partial \mathcal{U}_A}{\partial q^c} \frac{\partial q^c}{\partial \gamma_A} - \frac{\partial \mathcal{U}_B}{\partial q^c} \frac{\partial q^c}{\partial \gamma_A}}_{=0 \text{ when } \gamma_A = \gamma_B} \right].$$

Note that it must be that at the equilibrium γ_j^c , the left term in the above FOC is negative. This implies that $\left. \frac{d\mathcal{U}_j^I}{d\gamma_j} \right|_{\gamma_j=\gamma_{-j}=\gamma} < 0 \Rightarrow \frac{d\mathcal{U}^I(\gamma)}{d\gamma} < 0 \Rightarrow \gamma_j^c > \gamma^*$ by Proposition 2. The first implication comes from

$$\begin{aligned} \left. \frac{d\mathcal{U}_j}{d\gamma_j} \right|_{\gamma_j=\gamma_{-j}=\gamma} &= \frac{\partial \mathcal{F}}{\partial \gamma} + \frac{\partial \mathcal{F}}{\partial q^I} \frac{\partial q^I}{\partial \gamma} \cdot \frac{1}{2} \\ &> \frac{\partial \mathcal{F}}{\partial \gamma} + \frac{\partial \mathcal{F}}{\partial q^I} \frac{\partial q^I}{\partial \gamma} = \frac{d\mathcal{U}}{d\gamma} \end{aligned}$$

where $\mathcal{F} = \gamma q^I(\gamma) + \frac{(v+(1-\gamma)q^I(\gamma)-t/4)^2}{2\theta} + \frac{t^2}{96\theta}$ is the expected per-user surplus condition on joining a platform. Note that

$$\frac{\partial \mathcal{F}}{\partial q^I} = \gamma + (1-\gamma) \frac{v + (1-\gamma)q^I(\gamma) - t/4}{\bar{\theta}} > 0 \text{ and } \frac{\partial q^I}{\partial \gamma} < 0.$$

Finally, we have that $q^c(\gamma^c, \gamma^c) = q^I(\gamma^c) \leq q^I(\gamma^*)$. Because the symmetric competitive equilibrium induces the exact same lower-layer allocation and quality response as the monopoly model at γ^c , it follows immediately that $\mathcal{U}^I(\gamma^c) \leq \mathcal{U}^I(\gamma^*)$. Since $\Phi(\text{US}) = \text{US}$, aggregate platform welfare is also $\mathcal{U}^I(\gamma^c) \leq \mathcal{U}^I(\gamma^*)$. Furthermore, aggregate publisher profit across the two platforms is $2\pi^I(\gamma^c)$. By Proposition 1, $\gamma^c \geq \gamma^*$ implies that publisher profit is weakly lower under competition than under the monopoly benchmark. ■

B Additional Figures




Figure 8: e.g. “what are things I can do around me” without AI search feature

Google

what are things i can do around me

All Images Short videos News Forums Videos More Tools

East Lansing, MI 48824 Use precise location

AI Overview

If you are in the East Lansing/Greater Lansing area, you can easily explore Michigan State University landmarks, paddle or bike the Lansing River Trail, or check out interactive spots like the Impression 5 Science Center. Choose Lansing +4

Here is a quick guide to what you can do nearby:

Arts & Culture

- MSU Broad Art Museum:** A stunning, ultra-modern building on campus designed by Zaha Hadid, featuring contemporary art.
- Michigan History Center:** Dive deep into the state's past with 3-story interactive exhibits and displays.
- Wharton Center for Performing Arts:** Check the schedule on the Wharton Center website for Broadway shows, concerts, and touring acts. Choose Lansing +4

Outdoors & Nature

- Lansing River Trail:** More than 20 miles of paved trails perfect for a summer bike ride or a long walk.
- Beal Botanical Garden:** A peaceful, beautifully curated spot on the MSU campus featuring thousands of plant species.
- Potter Park Zoo:** A great 102-acre zoo in Lansing featuring over 160 animal species, ideal for a morning walk. Tripadvisor +4

6 sites

Things to Do in Lansing | Activities, Museums, Events ...
Lansing, Michigan, offers a mix of hands-on science at Impression 5 Science Center,...

Choose Lansing

What is some cool stuff around the Lansing area to do - Reddit
Feb 25, 2025 — Active Entertainment: Options for active entertainment include...

Reddit

THE 15 BEST Things to Do in East Lansing (2026) - Tripadvisor
Speciality Museums. Michigan State Capitol. 4.5. (359) 3.7 mi away. Lansing. Governmen...

Tripadvisor

Show all

Figure 9: e.g. “what are things I can do around me” with AI search feature

Platform preference across licensing equilibria




Figure 10: A graphical example of licensing equilibria across varying γ 's ($r = v = 2$)

C Extensions

C.1 Generalized distribution of θ

Previously, I assumed that user search cost θ was uniformly distributed. I classify a more general distribution such that several of the results in Sections 3 and 4 still hold. Suppose users who access a publisher's site incur a search cost $\theta \sim F$ on support $[0, \bar{\theta}]$ where $\bar{\theta} > v$. Note that I allow for $\bar{\theta} = \infty$. Assume the associated density f is log-concave with the mode at 0.

I now require the following assumption:

Assumption 2 $r\bar{f} < \min\{t, 1/\bar{\lambda}\}$ where $\bar{f} = \sup f$ and $\bar{\lambda} = \sup(-f'/f)$.

Assumption 2 is a regularity condition that ensures the search cost's pdf is never too steep which allows more interesting effects to be present in the model. If the assumption is violated, then publishers have too strong of an incentive in quality investment and the AI's effect in the model does not have sufficient bite. This assumption is satisfied in the baseline model and satisfied if $r \leq \lambda \min\{t, \lambda\}$ given θ is exponential where λ is the scale parameter.

If this improvement in matching directly increases expected user traffic, then publishers will be induced to increase their content quality. In other words, even when the AI content provides no utility to users, i.e. $\gamma = 0$, it may still increase equilibrium quality compared to the case without AI search features. One thing to note is that improving matches between publishers and users can also degrade quality if expected user traffic decreases. This can occur if the average demand across users exceeds that of the demand coming from $x = 1/4$. A simple sufficient condition for $q^I(0) \geq q^N$ is that f is convex which is satisfied by both uniform and exponential distribution.³⁴ This implies at the equilibrium quality-conversion ability γ^* , content quality may be higher with AI search features, but the quality degradation result still holds (i.e. $q^I(\gamma)$ is decreasing in γ). Lastly, concavity of user surplus with AI search features is difficult to establish with a generalized distribution without further restrictions, but the conflicting tradeoffs that the platform still holds.

If f were to instead be strictly concave, then $q^I(0) < q^N$ by Jensen's Inequality and $q^I(\gamma)$ is still decreasing in γ . Therefore, the quality degradation result in Proposition 1 still holds. However, I cannot invoke the same argument for Proposition 2. Now, it may be the case that user surplus is lower with AI search features than without as the improvement in matching between users and publishers results in $q^I(0) < q^N$. The result in Proposition 3 that $\gamma^{SO} \leq \gamma^*$ which comes from an argument through the envelope theorem is robust to

³⁴Note that there is no matching effect on content quality when θ is uniform, since the expected utility given some equilibrium quality q is the same regardless of the knowledge of one's x .

this case. Thus, I can still draw the same conclusion that an optimal regulating policy is one that restricts the platform's quality-conversion ability.

C.2 Information retention

In my model, I assume that users' search decision is dictated by comparing $u^{AI}(x)$ and $u_i(x)$. One possible critique is that users may still retain information from the publisher they are not allocated to when they continue search. I show that my results in the baseline are robust if search costs are sufficiently high.

If the user accesses publisher i after viewing the AI content located at a , they receive

$$u_i^R(x; a) = v + q_i - t|l_i - x| + \gamma|l_i - a|q_{-i}.$$

Then, the continuation value of searching is

$$\Delta_i^R(x; a) := u_i^R(x; a) - u^{AI}(x; a) = v + q_i - t|l_i - x| - \gamma|l_{-i} - a|q_i + t|a - x|$$

such that the user clicks if and only if $\theta \leq \Delta_i^R(x; a)$. Thus, for user $x \in \mathcal{X}_i$, the platform's maximization problem is

$$\begin{aligned} \max_{a \in [0,1]} \mathcal{W}_i(x; a) &:= u^{AI}(x; a) + \int_0^{\Delta_i^R(x; a)} F(s) ds = u^{AI}(x) + \frac{(\Delta_i^R(x; a))^2}{2\bar{\theta}} \\ \Rightarrow \frac{d\mathcal{W}_i}{da} &= \gamma(q_2 - q_1) - t \cdot \text{sign}\{a - x\} + \frac{\Delta_i^R(x; a)}{\bar{\theta}} \left[(1 - 2l_i)\gamma q_i + t \cdot \text{sign}\{a - x\} \right]. \end{aligned}$$

For $\alpha(x) = x$ to be optimal, it must be that $\frac{d\mathcal{W}_i}{da} \leq 0$ for $a \geq x$.³⁵ The above condition then reduces to

$$\begin{aligned} \gamma(q_2 - q_1) - t + \frac{\Delta_i^R(x; a)}{\bar{\theta}} \left[(1 - 2l_i)\gamma q_i + t \right] &< 0 \text{ and} \\ \gamma(q_2 - q_1) + t + \frac{\Delta_i^R(x; a)}{\bar{\theta}} \left[(1 - 2l_i)\gamma q_i - t \right] &> 0. \end{aligned}$$

As $\max_{x,a} \Delta_i^R(x; a) = v + t + q_i$, a sufficient condition for the above expressions to hold under a symmetric equilibrium is

$$t\bar{\theta} > (t + \gamma q_i)(v + t + q_i). \quad (14)$$

³⁵Note we use the same decision imitation logic in Lemma 1 that shows if $\alpha(x) = x$ is optimal for every user over any other possible location than it is weakly dominant over any other arbitrary mixing rule.

The publisher's FOC is similar to the baseline:

$$q_i = \frac{r}{\bar{\theta}} \int_{\mathcal{X}_i} (1 - \gamma |l_{-i} - \alpha(x)|) dx.$$

Under a symmetric equilibrium such that $q_1 = q_2$ where $\alpha(x) = x$ is optimal, I then have

$$q_1 = q_2 = q^I(\gamma) = \frac{r(4 - 3\gamma)}{8\bar{\theta}} \Rightarrow \bar{x} = 1/2.$$

(14) reduces to

$$t\bar{\theta} > (t + \frac{r}{6\bar{\theta}})(v + t + \frac{r}{2\bar{\theta}}). \quad (15)$$

The optimal mixing rule is perfectly personalized if (14) is satisfied. If it isn't perfectly personalized, the mixing rule $\alpha(x) > x$ for a subset of users $x \in \mathcal{X}_1$ and $\alpha(x) < x$ for a subset of users $x \in \mathcal{X}_2$. This is because retaining information has a sufficiently strong bite such that the platform focuses more on providing users with the other publisher's content because providing the publisher they are allocated to does not provide any continuation value.

Now, suppose (15) holds, such that $\alpha(x) = x$. Then,

$$\Delta_i^R(x) = v + q_i(1 - \gamma |l_{-i} - x|) - t |l_i - x|.$$

Solving the two publisher's FOCs and the platform's optimal allocation threshold gives us the system that q_1^R, q_2^R , and $\bar{x}^R$ must satisfy:

$$q_1 = \frac{r}{\bar{\theta}} \left[(1 - \gamma)\bar{x} + \frac{\gamma}{2}\bar{x}^2 \right], q_2 = \frac{r}{\bar{\theta}} \left[(1 - \bar{x}) - \frac{\gamma}{2}(1 - \bar{x}^2) \right], \bar{x} = \frac{t + (1 - \gamma)q_1 - q_2}{2t - \gamma(q_1 + q_2)}.$$

Substituting the two publisher's FOCs into the optimal allocation threshold gives

$$(2\bar{x} - 1) \underbrace{\left[2t - \frac{r}{\bar{\theta}}(2 - \gamma) + \frac{r}{\bar{\theta}}\gamma^2\bar{x}(1 - \bar{x}) \right]}_{>0 \text{ by Assumption 1}} = 0 \Rightarrow \bar{x}^R = 1/2, q_1^R = q_2^R = q^I(\gamma) = \frac{r(4 - 3\gamma)}{8\bar{\theta}}.$$

Consequently, $\pi^R = \pi^I(\gamma) + \frac{r^2\gamma(4-3\gamma)}{64\bar{\theta}^2}$ and $\frac{d\pi^R}{d\gamma} < 0$. So, all the results from Proposition 1 still hold.

$$\mathcal{U}^R(\bar{x}^R, \gamma, Q^I) = \gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \int_0^{1/2} [v + (1 - \gamma)q^I(\gamma) - (t - \gamma q^I(\gamma))x]^2 dx \geq \mathcal{U}^I(\gamma)$$

and it can be verified that $\mathcal{U}^R$ is concave, thereby, all the results in Proposition 2 remain

robust. Using the same logic in Proposition 3, the same results can be shown with this extension. The only caveat is that for item (iv), $t > 16/9$ is sufficient for social welfare at the optimal γ^R to be lower than without AI when r and $\bar{\theta}$ are sufficiently high.

C.3 Timing of $\bar{x}$

Suppose $\bar{x}$ was chosen after content qualities instead of simultaneously. Then, $\bar{x}$ represents an allocation rule that follows the same as in (4). Consider publisher 1's maximization problem without AI search features:

$$\max_{q_1} r \int_0^{\bar{x}^N(Q)} F(\rho_1(\bar{x}^N(Q), Q)) dx - C(q_1).$$

Solving for the F.O.C. gives us

$$\underbrace{r \frac{d\bar{x}^N}{dq_1} F(\rho_i)}_{\text{allocation effect}} + \underbrace{r \bar{x}^N f(\rho_i) \frac{d\rho_i}{dq_1}}_{\text{user base effect}} = \underbrace{C'(q_1)}_{\text{MC}}$$

and at a symmetric equilibrium $q_1 = q_2 = q^N \Rightarrow \bar{x}^N = 1/2$ I get that q^N implicitly solves the following

$$q^N = \frac{r}{\bar{\theta}} \left[\frac{1}{4} + \frac{q^N + v}{2t} - \frac{q^N + v - t/2}{16q^N + 16v - 6t} \right].$$

Altering the timing introduces an additional benefit in investing in content quality. When a publisher increases content quality, it incentivizes the platform to bias their allocation rule towards them. Note that instead of Assumption 1, I need $r < \min\{t\bar{\theta}, \frac{4t\bar{\theta}(\bar{\theta}-v)}{2\bar{\theta}+t}\}$ instead to ensure a well-defined and interior solution.

Moving on to the model with AI search features, perfect personalization remains optimal under any symmetric equilibrium characterized in this extension.³⁶ Publisher 1's maximization problem can be seen as

$$\max_{q_1} r \int_0^{\bar{x}^I(Q)} F(\Delta_1(x, Q)) dx - C(q_1)$$

³⁶Suppose $q_1 = q_2 = q$. For any AI location a , $u^{AI}(x; x) - u^{AI}(x; a) = t|a - x| \geq 0$ and we can use the same imitation decision logic established in Lemma 1.

which gives us the F.O.C. as

$$\underbrace{r \int_0^{\bar{x}^I(Q)} \frac{du_1}{dq_1} f(\Delta_1(x, Q)) dx}_{\text{user base effect}} + \underbrace{r \frac{d\bar{x}^I}{dq_1} F(\Delta_1(\bar{x}^I(Q), Q))}_{\text{allocation effect}} - \underbrace{r \int_0^{\bar{x}^I(Q)} \frac{du^{AI}}{dq_1} f(\Delta_1(x, Q)) dx}_{\text{AI quality effect}} = \underbrace{C'(q_1)}_{\text{MC}}.$$

Solving for the symmetric equilibrium gives us

$$q^I(\gamma) = \frac{\frac{r}{\bar{\theta}} \left(\frac{v}{2t} + 1/4 - \frac{3\gamma}{8} \right)}{1 - \frac{r(1-\gamma)}{2t\bar{\theta}}}.$$

Observe that $q^I(0) > q^N$ and simple comparative statics gives $\frac{dq^I}{d\gamma} < 0$. Therefore, $\mathcal{U}^I(0) > \mathcal{U}^N$ still holds. It remains to be shown that $\mathcal{U}^I(\gamma)$ is concave in γ . Recall that

$$\begin{aligned} \mathcal{U}^I(\gamma) &= \gamma q^I(\gamma) + \frac{1}{\bar{\theta}} \int_0^{1/2} [v + (1 - \gamma)q^I(\gamma) - tx]^2 dx \\ \Rightarrow \frac{d^2 \mathcal{U}^I}{d\gamma^2} &= 2 \frac{dq^I(\gamma)}{d\gamma} + \gamma \frac{d^2 q^I(\gamma)}{d\gamma^2} + \frac{1}{\bar{\theta}} \left[(B')^2 + (B - \frac{t}{4})B'' \right] \end{aligned}$$

where $B := v + (1 - \gamma)q^I(\gamma)$. A simple sufficient condition for concavity is $\bar{\theta} > 2v$ and $r < \frac{t^2 \bar{\theta}}{16v}$. Under these sufficient conditions, the marginal savings in investment costs from a reduction in quality are strictly bounded and never offset the corresponding loss in user traffic, ensuring $\pi^{I'}(\gamma) < 0$. Therefore, the social welfare result regarding $\gamma^{SO} \leq \gamma^*$ in Proposition 3 is similarly proven using the concavity of $\mathcal{U}^I$ and strictly decreasing publisher profits.

C.4 User knowledge of x

In the baseline model, I assumed users know their preference between the publishers, but cannot map their preference onto the line. This captures the fact that although users know their tastes for content, they are unaware which publisher best fits their needs.

If users know their location on the line, the platform's off-equilibrium allocation rule without AI changes to the standard Hotelling cutoff. Nevertheless, the equilibrium allocation threshold, content qualities, and publisher profits remain unchanged. With AI search features, the equilibrium allocation rule and publisher outcomes remain as in Section 4. However, the user surplus without and with AI search features when users know their location

is now

$$\begin{aligned}\tilde{\mathcal{U}}^N &= 2 \int_0^{1/2} \int_0^{v+q^N-tx} F(s) ds dx > \mathcal{U}^N \text{ and} \\ \tilde{\mathcal{U}}^I(\gamma) &= \gamma q^I(\gamma) + 2 \int_0^{1/2} \int_0^{v+(1-\gamma)q^I(\gamma)-tx} F(s) ds dx.\end{aligned}$$

Notice now that there is no matching benefit from incorporating AI search features as $\tilde{\mathcal{U}}^I(0) = \tilde{\mathcal{U}}^N$. It can be shown that $\tilde{\mathcal{U}}^I(\gamma)$ is still concave in γ . Therefore, user surplus is strictly higher with AI search features if and only if

$$\begin{aligned}\left. \frac{d\tilde{\mathcal{U}}^I}{d\gamma} \right|_{\gamma=0} > 0 &\iff q^I(0) + \frac{v + q^I(0) - t/4}{\bar{\theta}} \left(\left. \frac{dq^I}{d\gamma} \right|_{\gamma=0} - q^I(0) \right) > 0 \\ &\iff \frac{r}{2\bar{\theta}} + \frac{v + \frac{r}{2\bar{\theta}} - t/4}{\bar{\theta}} \left(-\frac{3r}{8\bar{\theta}} - \frac{r}{2\bar{\theta}} \right) > 0 \\ &\iff 16\bar{\theta}^2 + 7\bar{\theta}t > 28\bar{\theta}v + 14r.\end{aligned}$$

In other words, if the above condition is satisfied, then the platform will set some $\gamma > 0$ increasing user surplus. If the above cannot be satisfied, then the platform would set $\gamma = 0$ which is effectively the same as not incorporating AI search features at all. The first two items in Proposition 3 remain robust with this change. However, the social welfare comparison is different now. Observe that social welfare in equilibrium remains the same with or without AI when $\gamma^* = 0$. Assume $\gamma^* > 0$ and define the new social welfare with AI search features (when $\Phi(\text{US}) = \text{US}$) as

$$\widetilde{SW}^I(\gamma) = 2\tilde{\mathcal{U}}^I(\gamma) + 2\pi^I(\gamma).$$

Note that $\tilde{\mathcal{U}}^I(\gamma) - \tilde{\mathcal{U}}^N \leq \gamma q^I(\gamma)$ because the continuation value with AI cannot be greater when without AI. A sufficient condition for social welfare being lower with AI than without, in equilibrium, is

$$\gamma^* q^I(\gamma^*) < \pi^N - \pi^I(\gamma^*) \iff \frac{r\gamma^*(4 - 3\gamma^*)}{8\bar{\theta}} < \frac{r^2\gamma^*(32 - 15\gamma^*)}{128\bar{\theta}^2} \iff r > \frac{16\bar{\theta}(4 - 3\gamma^*)}{32 - 15\gamma^*}.$$

A simpler sufficient condition is $r > 2\bar{\theta}$. Therefore, a similar intuition holds for the last item in Proposition 3.