Generative AI, Information Supply and Quality Online

Kacper Krasowski
Universitat Autònoma de Barcelona
Collegio Carlo Alberto
kacper.krasowski@carloalberto.org

Draft: April 14, 2026

Abstract

This paper develops a model of the internet to study how generative AI affects the supply of online information and the quality of information users consume. In the baseline model, users choose among traditional online sources, while creators decide whether to enter based on expected traffic and fixed entry costs. The extended model introduces generative AI as both a new source of information and a technology that can aggregate existing knowledge and lower creators' production costs. These channels generate opposing effects. On the one hand, better AI diverts users away from traditional websites and weakens incentives to create content online. On the other hand, it can expand information supply by reducing entry costs and improve consumed quality by synthesizing dispersed knowledge. The model therefore predicts a threshold effect. When AI is weak, stronger AI capability and AI-enabled cost reductions can lower equilibrium consumed quality. When AI becomes sufficiently effective, the same forces raise consumed quality and can even outperform the no-AI benchmark. The framework helps reconcile recent empirical evidence of declining activity on online knowledge platforms with the possibility of longer-run gains from AI. It also implies that policies such as source attribution, redirection, and compensation for original creators may be important for sustaining the information ecosystem on which AI itself depends.

Keywords: generative AI, online information, content creation, platform competition, information quality

1 Introduction

The internet has brought many important benefits. It shapes how people learn, communicate, and make decisions. It provides immediate access to news, educational materials, scientific knowledge, and public debate. As a result, individuals can stay informed more quickly than through traditional media alone. It also lowers barriers to information by making knowledge accessible across countries, social groups, and institutions.

For students and researchers, online content is especially valuable because it supports learning, comparison of sources, and access to recent developments in a field. When online content is accurate and accessible, it can promote education, innovation, and social inclusion. For this reason, the internet is not only a source of information, but also a powerful environment in which knowledge is created, shared, and contested.

Since 2022, generative AI has advanced at an extraordinary pace. At first, it was unable to solve simple puzzles or complete basic tasks reliably. Today, it can compete with or even outperform humans in a range of tasks, including selected programming and reasoning settings. However, it still lags behind humans in more difficult mathematics, complex reasoning, and longer-horizon problem-solving tasks as shown in Maslej et al. (2025). More recently, growing attention has focused on the role of AI in creating content and providing information online.

AI makes it easier to produce information and content online by serving as a copywriter, research assistant, and tool for building content infrastructure. This suggests that generative AI could expand the stock of knowledge available online. At the same time, generative AI, and especially chatbots, compete with traditional web-based information sources. Instead of *googling* for information, a user can now *chat* it or *gpt* it. This may reduce the stock of information available online, because AI can draw business away from traditional sources and thereby weaken incentives to create new content.

Because the internet and the information it contains play such an important role in our lives, and because generative AI is rapidly transforming the internet through these channels, it is natural to ask: *How does generative AI (GenAI) affect online content supply and the quality of information users consume?* This paper aims to answer this question and to contribute to a better theoretical (and in the future empirical) understanding of how AI tools affect the online stock and quality of knowledge. To do so, I develop a model of the internet.

On the demand side, internet users choose between traditional online information sources and a chatbot. On the supply side, information creators decide whether to enter the market. The chatbot can access and aggregate knowledge available online and present it as its own output. In this model, generative AI affects (i) users, by serving as an alternative source of information, and (ii) creators, by lowering the cost of market entry.

The model predicts that if AI becomes sufficiently effective, consumed quality can rise; the stock of traditional information rises only if AI-enabled cost reductions dominate business stealing. The stock of knowledge rises because entry becomes cheaper, and the reduction in costs dominates the business-stealing effect. Quality also improves for two reasons: (i) AI increases the stock of information available online, and (ii) AI can improve quality by aggregating existing knowledge.

If AI is not sufficiently effective, both the stock of knowledge and its quality may decline. Users shift away from traditional sources, which reduces incentives for creators to enter the market. As a result, AI has less information on which to train. In this case, AI does not reduce costs enough to encourage creators to enter.

The remainder of the paper proceeds as follows. Section 2 reviews the related literature. Section 3 presents the baseline model, and Section 4 extends it to incorporate generative AI. Section 5 discusses the main results and their implications. Section 6 highlights the limitations of the model. Section 7 concludes.

2 Literature Review

The literature most closely related to this paper studies the impact of generative AI on the internet. This branch of research has focused primarily on empirical analyses of two major platforms: Stack Overflow and Wikipedia.

With a focus on Stack Overflow, Burtch, Lee, and Chen (2024) provide evidence that LLMs affect participation and content production in online knowledge communities. Similarly, Shan and Qiu (2025) find that generative AI is associated with an increase in the number of answers generated by users, and that these answers tend to be shorter and easier to read. In related work, Rio-Chanona, Laurentsyeva, and Wachs (2023) investigate how the introduction of ChatGPT reduced activity on Stack Overflow by comparing activity in countries with access to ChatGPT to activity in countries with more limited access.

Turning to Wikipedia, Reeves, Yin, and Simperl (2025) show that recent advances in generative AI have reduced views and edits on the platform. Likewise, Lyu et al. (2025) find that Wikipedia content overlapping with material generated by ChatGPT 3.5 experienced a larger decline in editing and viewership after the November 2022 launch of ChatGPT than articles with less overlap.

In the context of online news, Zhao and Berman (2025) document a moderate decline in traffic to news publishers. They show that blocking GenAI bots can be associated with a reduction in total website traffic for large publishers relative to not blocking them. On the hiring side, however, they do not find evidence that LLMs are yet replacing editorial or content-production jobs. From a different angle, Koren et al. (2026) study how vibe coding[1] affects the creation of open-source software. They find that AI use may reduce the future availability of open-source software. Their approach has also influenced the modelling choices in this paper.

A related body of work studies the effects of digital technologies on collective information and knowledge formation. For example, Dasaratha and He (2022) and Acemoglu, Ozdaglar, and Siderius (2024) examine how social media affect learning. In the context of the knowledge economy, Ide and Talamàs (2025) show that output may be higher with autonomous AI, which contrasts with much of the empirical evidence cited above. In addition, Agarwal, Moehring, Rajpurkar, et al. (2023) show, through an experiment with radiologists, that humans collaborating with AI can outperform humans working alone. With a focus on productivity, Cullen, Danielle Li, and S. Li (2025) also contribute to this broader literature.

More specifically, part of this literature focuses on the effects of generative AI on information acquisition. Among these papers, Acemoglu, Kong, and Ozdaglar (2026) study how generative AI may lead to a knowledge collapse in society. They show that, in a static sense, AI can help humans, but in a dynamic sense, it can undermine the accumulation of collective knowledge. Similarly, Ide (2025) show that automating entry-level tasks can reduce long-term growth because it disrupts the intergenerational transmission of tacit knowledge, which typically takes place within firms through interactions between novices and experts. Finally, Agarwal, Moehring, and Wolitzky (2025) discuss how optimal collaboration schemes can be designed using a mechanism design approach.

This paper contributes to this literature by proposing a structural model of the internet and, in future work, estimating it empirically This framework makes it possible to disentangle several channels through which generative AI may affect the stock and quality of information available online, which remains the largest source of information. Ideally, these effects can be understood clearly in theory and then validated in the data.

[1]Vibe coding is an AI-driven software development approach where users describe desired applications in plain English, allowing large language models (LLMs) to generate, refine, and debug code.

3 Baseline Model

3.1 Environment

The goal of the model is to describe and characterize the choices people make when searching for information online and when creating it. There is a mass 1 of users who search for and choose information online. There are M information creators[2], each of whom decides whether to enter the market and create a single piece of information.

Each user i receives utility U_{ik} from consuming information source k, which depends on the quality of the source k, denoted by Q_k. In addition, each user is subject to an idiosyncratic taste shock ϵ_{ik}, which is distributed according to an i.i.d. Type-I extreme value distribution with scale parameter θ. This shock captures heterogeneity in user preferences, such as personal feelings toward the author, the appearance of the website, style of writing, etc. Utility is given by

$$U_{ik} = \ln Q_k + \epsilon_{ik}.$$

Using this utility function and the properties of the Type-I extreme value distribution[3], I can characterize the share of users who visit source k:

$$s_k = \frac{Q_k^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}}, \tag{1}$$

where $\mathcal{A}$ is a set of active firms. This expression highlights a natural implication of the model. The higher the quality of the information, the more users will visit the source. The larger is θ, the less users care about quality.

On the supply side, creators draw their quality Q_k from a Pareto distribution, where γ is the shape parameter. The Pareto distribution is chosen to capture the idea that much online information is low quality, while a small number of sources, such as Wikipedia, are exceptionally high quality. but some sources excel, like Wikipedia. Given their quality, they decide whether entering the market is profitable. The profit of creator k is

$$\pi_k = r s_k - F,$$

where r is the revenue per share of users, s_k is the share of users consuming the information, and F is the fixed cost required to enter the market.

Although highly stylized, this framework provides a useful approximation of the real world. Each creator either monetizes content through ads, sponsorships, and related channels, or derives intrinsic value from producing content, as shown in Xu and Dahui Li (2015). Both types of returns increase with audience size. Because the object of analysis is a single piece of information, the fixed cost F represents all costs involved in producing information online, such as typing, copywriting, and research.

A creator will enter the market whenever $r s_k \geq F$. Therefore, there must exist a cutoff quality Q_0 at which a creator is indifferent between entering and not entering. This cutoff can be derived using equation (1), the entry condition, and the properties of the Pareto distribution. The derivations are provided in the appendix. The resulting cutoff is

$$Q_0 = \left(\frac{FM\gamma}{r(\gamma - 1/\theta)} \right)^{1/\gamma}. \tag{2}$$

[2] Which is assumed to be large.

[3] Proof in the appendix.

For the purposes of this paper, it is important to define the supply of information $m_{\mathcal{A}}$. Notice that the supply of information equals the mass of creators M multiplied by the probability of a random creator entering the market, $Pr(Q \geq Q_0)$. This simplifies to

$$m_{\mathcal{A}} \equiv MPr(Q \geq Q_0) = \frac{r(\gamma - 1/\theta)}{F\gamma}. \tag{3}$$

A natural implication of the baseline model is that the stock of information decreases with the fixed cost and increases with revenue per user.

3.2 Equilibrium

A baseline equilibrium is a pair $(Q_0, \mathcal{A})$ such that

$$\mathcal{A} = \{k : Q_k \geq Q_0\},$$

users allocate demand across active sources according to

$$s_k = \frac{Q_k^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}}, \qquad k \in \mathcal{A},$$

and the marginal entrant earns zero profit:

$$r \frac{Q_0^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}} = F.$$

Given $(Q_0, \mathcal{A})$, all other equilibrium objects are induced. In particular, the mass of active creators is

$$m_{\mathcal{A}} = M \Pr(Q \geq Q_0).$$

Thus, equilibrium is fully characterized by the cutoff rule and the demand shares it induces. In the symmetric cutoff equilibrium,

$$Q_0 = \left(\frac{FM\gamma}{r(\gamma - 1/\theta)} \right)^{1/\gamma}, \qquad m_{\mathcal{A}} = \frac{r(\gamma - 1/\theta)}{F\gamma}.$$

3.3 Information Quality

To analyze the quality of information online, I use the quality consumed by users, defined as $\tilde{Q} \equiv \sum_{j \in \mathcal{A}} s_j Q_j$. This can be expressed as a function of the cutoff quality Q_0:

$$\tilde{Q} = \frac{\gamma - 1/\theta}{\gamma - 1/\theta - 1} Q_0.$$

The higher the cutoff, the fewer firms enter the market, and the higher is the consumed quality. This is because low-quality firms do not enter and therefore are not consumed. Because there is no cost of consumption, users always choose to consume information.

Naturally, when users care more about quality, that is, when θ decreases, higher-quality information is consumed. When the quality distribution has a fatter tail, that is, when γ is lower, there is a greater probability of extremely high-quality information. that is, when γ is higher, there is a greater probability of extremely high-quality information. As a result, the consumed quality is higher.

4 Extended Model with Generative AI

Following the rationale from the Introduction, AI should affect online information services in three main ways. First, it can aggregate available information and present it as its own output. Second, it can serve as an alternative to traditional online search. Third, it can reduce production and entry costs for information creators.

In this spirit, the extended model assumes that generative AI uses all available sources in its production function to produce its own quality Q_A. I assume that the AI production function takes a CES form, with the elasticity of substitution, for tractability, set equal to the shape parameter of the taste shock. The total factor productivity term ϕ captures the ability of AI to convey, aggregate, and synthesize information. It is given by

$$Q_A = \phi \left(\sum_{j \in \mathcal{A}} Q_j^{1/\theta} \right)^{\theta} .$$

The higher is ϕ, the stronger is the ability of generative AI to produce useful information.

Generative AI also serves as an alternative option for users. Users' choices now take the form of a nested logit model. First, they choose whether to use AI or perform a traditional web search. If web search is chosen, the user then selects a source from the available creators. As described above, users choose between traditional sources and an AI chatbot in a nested structure. The utility shock now follows a General Extreme Value distribution, again with shape parameter θ.

This leads to three relevant shares: the share of users consuming AI-generated information s_A, the share of users consuming information from the web s_W, and, conditional on choosing the web, the share of users consuming information from source k, denoted by $s_{k|W}$. These shares are given by

$$\begin{aligned} s_A &= \frac{\phi}{1+\phi}, \\ s_W &= \frac{1}{1+\phi}, \\ s_{k|W} &= \frac{Q_k^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}}. \end{aligned}$$

It is important to note that the total share of users who visit traditional source k is now given by $s_k = s_W s_{k|W}$. The better the AI, the more business it steals from traditional sources.

The fixed cost of entry for creators is reduced by a multiplicative term η, so that profits take the form

$$\pi_k = r s_k - F(1-\eta).$$

The better the AI, the larger is the reduction in costs, and therefore the higher is η. Notice that when there is no cost reduction, so that $\eta = 0$, and AI is unable to aggregate knowledge, so that $\phi = 0$, the extended model collapses to the baseline model.

In addition, as AI improves, both the cost reduction parameter η and the aggregation parameter ϕ should increase. This creates two counteracting forces. Following the same steps as before, we can derive the new cutoff:

$$Q_0^{AI} = \left(\frac{FM\gamma}{r(\gamma - 1/\theta)} (1-\eta)(1+\phi) \right)^{1/\gamma} = Q_0((1-\eta)(1+\phi))^{1/\gamma},$$

which makes these opposing forces explicit.

The supply of information now becomes

$$m^{AI}_{\mathcal{A}} = \frac{r(\gamma - 1/\theta)}{F\gamma(1-\eta)(1+\phi)} = \frac{m_{\mathcal{A}}}{(1-\eta)(1+\phi)}.$$

This expression highlights the two effects. When production becomes cheaper, that is, when η increases, the supply of information rises because more creators enter the market. On the other hand, when AI becomes better at serving and aggregating knowledge, it steals business from traditional sources, thereby reducing the stock of available information.

The overall effect depends on the relative strength of these two parameters. In general, as AI improves, it increases both ϕ and η, so its effect on the stock of knowledge online is ambiguous.

4.1 Equilibrium with Generative AI

An equilibrium in the extended model is a pair $(Q_0^{AI}, \mathcal{A})$ such that

$$\mathcal{A} = \{k : Q_k \geq Q_0^{AI}\},$$

AI quality is

$$Q_A = \phi \left(\sum_{j \in \mathcal{A}} Q_j^{1/\theta} \right)^{\theta},$$

users choose between AI and the web according to

$$s_A = \frac{\phi}{1+\phi}, \qquad s_W = \frac{1}{1+\phi},$$

and, conditional on choosing the web, they allocate demand across traditional sources according to

$$s_{k|W} = \frac{Q_k^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}}, \qquad s_k = s_W s_{k|W}, \qquad k \in \mathcal{A}.$$

The marginal traditional entrant earns zero profit:

$$r\, s_W \frac{(Q_0^{AI})^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}} = F(1-\eta).$$

Given $(Q_0^{AI}, \mathcal{A})$, the mass of active traditional creators is

$$m^{AI}_{\mathcal{A}} = M \Pr(Q \geq Q_0^{AI}).$$

Therefore, equilibrium is characterized by the cutoff rule together with the induced AI quality and user shares. In the symmetric cutoff equilibrium,

$$Q_0^{AI} = \left(\frac{FM\gamma}{r(\gamma - 1/\theta)} (1-\eta)(1+\phi) \right)^{1/\gamma}, \qquad m^{AI}_{\mathcal{A}} = \frac{m_{\mathcal{A}}}{(1-\eta)(1+\phi)}.$$

4.2 Information Quality

When generative AI is available, users now also consume information from AI, thus it should also be included in the consumed quality. In the case of the extended model, the consumed quality is defined as

$$\tilde{Q} \equiv s_A Q_A + s_W \sum_{j \in \mathcal{A}} s_{j|W} Q_j,$$

where users split between AI and traditional web search, AI aggregates the information available online, and traditional creators enter subject to the cutoff rule implied by equilibrium profits. Thus, AI affects consumed quality through two forces. On the one hand, a better AI improves the quality of the AI option and can make use of a larger stock of online information. On the other hand, AI diverts users away from traditional websites and can weaken incentives to create content online.

To state the main results, let $a \equiv \gamma - \frac{1}{\theta}$, and let the follwoing assumptions to be made. Firstly, $a > 1$ rules out the case in which equilibrium outcomes are driven by an infinitesimal set of extremely high-quality sources. Secondly, $\theta - \frac{1}{\gamma} < 1$ requires that users remain sufficiently responsive to quality. Thirdly, $0 \leq \eta < 1$ ensures that the cost remians positive. Lastly, $\phi \geq 0$ ensures that AI produces something valuable.

Proposition 4.1 (Weak AI lowers consumed quality)**.** *Suppose $a > 1$. Fix any $\eta \in [0, 1)$. Then there exists $\underline{\phi}(\eta) > 0$ such that for all*

$$\phi \in [0, \underline{\phi}(\eta)),$$

equilibrium consumed quality satisfies

$$\frac{\partial \tilde{Q}(\eta, \phi)}{\partial \eta} < 0 \qquad \textit{and} \qquad \frac{\partial \tilde{Q}(\eta, \phi)}{\partial \phi} < 0.$$

When AI is still weak, the positive aggregation channel is too small to compensate for the damage AI causes to the traditional information ecosystem. A rise in ϕ immediately shifts users away from traditional websites, reducing their traffic and weakening entry incentives. At the same time, a rise in η lowers entry costs and induces additional entry, but with weak AI this mainly brings in marginal low-quality creators rather than generating enough extra informational value through aggregation. As a result, both stronger AI capability and stronger AI-enabled cost reduction lower the quality users consume when AI is not yet sufficiently effective.

Proposition 4.2 (Sufficiently effective AI raises consumed quality)**.** *Suppose $a > 1$ and $\theta < 1 + \frac{1}{\gamma}$.*

Fix any $\eta \in [0, 1)$. Then there exists $\overline{\phi}(\eta)$ such that for all

$$\phi > \overline{\phi}(\eta),$$

equilibrium consumed quality satisfies

$$\frac{\partial \tilde{Q}(\eta, \phi)}{\partial \eta} > 0 \qquad \textit{and} \qquad \frac{\partial \tilde{Q}(\eta, \phi)}{\partial \phi} > 0.$$

Hence, once AI is sufficiently effective, further improvements in AI capability and stronger AI-induced cost reductions both increase the quality users consume.

When AI becomes sufficiently capable, the balance of forces reverses. A larger ϕ no longer just steals users from traditional sources; it also substantially improves the quality of the AI option. Likewise, a larger η no longer merely lowers the entry cutoff; it also expands the stock of online information enough to make

AI significantly more useful. In this region, the gains from aggregation and expanded information supply dominate the negative business-stealing and selection effects, so that improvements in AI raise consumed quality.

Taken together, the two propositions 4.1 and 4.2 imply that the effect of AI on consumed quality depends on how effective AI has become. When AI is weak, both more AI capability and more AI-enabled cost reduction lower equilibrium consumed quality. When AI is sufficiently effective, both forces raise consumed quality instead. The model therefore predicts a threshold effect: AI first degrades the quality of information consumed online, but after it becomes capable enough, the same channels can reverse sign and improve quality.

Proposition 4.3 (When AI improves consumed quality relative to the baseline)**.** *Suppose that a >1 and* $\theta < 1 + \frac{1}{\gamma}$. *Fix any* $\eta \in [0, 1)$. *Then there exists a threshold* $\bar{\phi}(\eta) > 0$ *such that for all*

$$\phi > \bar{\phi}(\eta),$$

equilibrium consumed quality in the AI economy exceeds equilibrium consumed quality in the no-AI baseline:

$$\tilde{Q}^{AI}(\eta, \phi) > \tilde{Q}^{B}.$$

Equivalently, AI improves consumed quality whenever

$$(1-\eta)^{1/\gamma}(1+\phi)^{1/\gamma-1}\left[1+\frac{\gamma-1/\theta-1}{\gamma-1/\theta}\phi^2\left(\frac{r}{F(1-\eta)(1+\phi)}\right)^{\theta}\right] > 1.$$

This proposition compares the AI economy directly to the no-AI benchmark. The expression on the left-hand side can be read as the product of two forces. The first term, $(1-\eta)^{1/\gamma}(1+\phi)^{1/\gamma-1}$, captures the deterioration of the traditional web ecosystem relative to the baseline: AI steals users from traditional websites, and lower entry costs reduce the quality cutoff for entry. On its own, this term is below one and therefore pushes consumed quality downward.

The second term inside the brackets captures the gain from AI aggregation. As AI becomes more capable, it synthesizes a larger stock of online information into a higher-quality output, and users benefit directly from consuming that output. The proposition 4.3 states that when AI capability is high enough, this aggregation force becomes strong enough to outweigh the weakening of the traditional web sector. In that region, the AI economy delivers higher consumed quality than the no-AI baseline.

5 Discussion

The main takeaway from the model is that the recent rise of generative AI does not affect the internet and its content in a monotonic way. Instead, two counteracting forces are at work. On the one hand, AI systems can help aggregate knowledge, present it more effectively, and reduce the cost of producing information online. On the other hand, they attract users away from traditional sources, thereby cannibalizing the ecosystem on which they depend. The model predicts that if AI is not sufficiently effective, the second channel dominates, leading to worse outcomes overall. If AI cannot aggregate knowledge well or reduce costs sufficiently, the business-stealing effect prevails. However, if AI becomes strong enough, the overall quality of information may exceed the one in a world without AI. Even though it steals business from traditional sources, it may still be able either to aggregate existing knowledge effectively or to reduce costs enough to keep market entry profitable for creators.

Recent empirical work suggests that the amount of information both created and consumed has declined, based on analyses of Wikipedia, Stack Overflow, and news platforms. The model shows that these effects are not necessarily unambiguously bad in the long run. They may instead indicate that we are in a transitional period, after which further improvements in AI could lead to a larger supply of information sources.

The model also has clear implications for policymakers and platform design. If the business-stealing channel is the dominant force, as the empirical literature currently suggests, then it becomes important to focus on mechanisms such as redirecting users from AI-generated content to original sources, or compensating original creators. Such interventions would reduce the extent to which AI cannibalizes its own information ecosystem.

To my knowledge, this is the first economic model of the online information ecosystem. It provides a framework for thinking about the internet as an equilibrium system in which users, creators, and AI interact. It shows not only that AI can either help or harm the internet, but also identifies the specific channels through which these effects arise. In particular, it highlights the importance of information aggregation, creators' dependence on traffic, and cost reduction. These predictions create a natural bridge to the empirical work proposed in the following section.

6 Limitations and Avenue for Future Work

AI's Impact on the Creators

In reality, AI affects the production of information in more complex ways than simply lowering production costs. It not only makes content easier to produce, but also affects the availability of new information for creators themselves. In addition, the assumption that creators randomly draw the quality of their information may seem unrealistic. Although this assumption may still capture the heterogeneous nature of online information reasonably well, a fuller understanding of AI's effects requires a richer treatment of the production side.

To address these issues, I plan to extend the model in future work with a more realistic approach. In particular, creators would be equipped with an explicit information production function that is itself affected by the introduction of AI.

Quality of AI

In the model, AI quality improves through exposure to all available sources of information, regardless of their quality. In reality, this may not be accurate. The best generative AI models are typically designed to avoid low-quality training data. In addition, for simplicity and tractability, I assume that the elasticity of substitution is equal to the shape parameter of the taste shock distribution. Relaxing this assumption would substantially reduce the tractability of the model.

The assumptions on the CES function and the substitution parameter are extremely convenient. Unfortunately, they do not fully capture the richness of the real setting. In future work, I plan to relax both assumptions to test whether the model's results remain robust under alternative functional forms and parametric assumptions.

Utility Function

In the utility function, the main characteristic of an online source that matters to users is the quality of the information source. All other considerations are assumed to be captured by the taste shock. In reality, however, users make decisions based on more than this broadly defined notion of quality. This simplification is made for tractability.

To account for this, one could enrich the utility function by adding further terms, such as AI aversion or other user-specific preferences.

Distributional Dependence

The model's results depend heavily on the assumption that the quality of creators, and therefore of the content itself, follows a Pareto distribution. This assumption implies that there is almost always some very high-quality content online from which AI can learn.

Because this assumption is so important, it should be verified empirically, even if the actual distribution of online content quality may plausibly resemble a Pareto distribution.

Redriection and Google AIO

Chatbots often provide users with links to original sources. As a result, users may still visit the original content creators, even when AI is their primary point of access. This would imply less business stealing from online creators.

Another issue concerns Google AI Overviews. When users search on Google, they are often shown an AI-generated answer even when using traditional web search. Thus, even if users choose the traditional web, some of them still end up being effectively exposed to AI. This may steal even more business from online information creators.

To account for these two features, I would introduce exogenous redirection rates. The first rate would capture the share of AI users who are redirected to the original content creators. The second rate would capture the share of web users who are redirected to AI.

Model Estimation and the Data

The ultimate goal of this project is to estimate the model empirically. To do so, I plan to use the GESIS Panel.dbd Digital Behavioral Data (currently pending approval). This dataset tracks the online activity of more than 2,500 users from February 2023 to November 2023. Each observation corresponds to a single online activity and includes a user ID, the exact timestamp, the visited website and URL, and, most importantly, an HTML snapshot of the page. This would allow me to observe the precise content of the visited website. In addition, the dataset links each user to socio-economic characteristics as well as to their knowledge and opinions about AI, the internet, and politics.

This dataset would allow me to estimate the demand side of the model using a nested logit framework. To do so, I would need to define the utility function more precisely, which I plan to do once I gain access to the data. On the supply side, the HTML snapshots of visited websites would be especially valuable. However, without a well-specified production function, I am not yet able to say much more about estimation on that side of the model.

7 Conclusions

This paper studies how generative AI affects the internet as an information ecosystem. I develop a model in which users choose between traditional online sources and an AI chatbot, while creators decide whether to enter the market and produce information. In this setting, generative AI changes the equilibrium through two main channels: it competes with traditional websites for users, and it reduces the cost of producing online content. At the same time, it can aggregate and synthesize existing information into a new source of value for users.

The main result is that the effect of generative AI on online information is not monotonic. When AI is still weak, the business-stealing effect dominates. Users shift away from traditional websites, creators lose traffic, entry incentives fall, and the quality of information consumed can decline. In that region, even AI-enabled cost reductions may worsen outcomes by encouraging entry by relatively low-quality creators without generating enough value through aggregation. By contrast, when AI becomes sufficiently effective, the balance reverses. Aggregation improves the quality of the AI option, lower costs support entry, and the resulting expansion in the stock of information can raise the quality users consume. If AI is strong enough, the model predicts that the AI economy can outperform the no-AI baseline.

These results suggest that recent declines in activity on platforms such as Wikipedia, Stack Overflow, and online news sites should not necessarily be interpreted as definitive evidence that AI will reduce social welfare in the long run. Instead, they may reflect a transitional phase in which AI has already begun to divert attention away from traditional sources but has not yet become effective enough to compensate through better aggregation or sufficiently strong cost reductions. More broadly, the paper highlights that the future effects of AI depend not only on how much traffic it displaces, but also on how much new value it creates from the information ecosystem on which it relies.

The model also has clear policy implications. If the business-stealing channel is quantitatively important, then platform design and regulation may need to focus on preserving incentives for original content creation. Mechanisms such as redirecting users to source material, improving attribution, or compensating creators could help reduce the extent to which AI undermines the very ecosystem from which it learns. In this sense, the paper suggests that sustaining the supply of online information is not separate from improving AI quality; it is one of its preconditions.

Finally, the framework provides a foundation for future empirical work. While the model is intentionally stylized, it offers a tractable way to organize the relationship between users, creators, and AI within a single equilibrium system. Extending the production side, relaxing assumptions on AI aggregation, and estimating the model with digital behavioral data are natural next steps. More generally, the paper argues that understanding generative AI requires studying not only its direct usefulness to users, but also its equilibrium effects on the production and preservation of knowledge online.

References

Acemoglu, Daron, Dingwen Kong, and Asuman Ozdaglar (2026). *AI, Human Cognition and Knowledge Collapse*. Tech. rep. National Bureau of Economic Research.

Acemoglu, Daron, Asuman Ozdaglar, and James Siderius (2024). "A model of online misinformation". In: *Review of Economic Studies* 91(6), pp. 3117–3150.

Agarwal, Nikhil, Alex Moehring, Pranav Rajpurkar, et al. (2023). *Combining human expertise with artificial intelligence: Experimental evidence from radiology*. Tech. rep. National Bureau of Economic Research.

Agarwal, Nikhil, Alex Moehring, and Alexander Wolitzky (2025). *Designing human-ai collaboration: A sufficient-statistic approach*. Tech. rep. National Bureau of Economic Research.

Burtch, Gordon, Dokyun Lee, and Zhichen Chen (2024). "The consequences of generative AI for online knowledge communities". In: *Scientific Reports* 14(1), p. 10413.

Cullen, Zoë, Danielle Li, and Shengwu Li (2025). *Labor as Capital: AI and the Ownership of Expertise*. Tech. rep. Technical report, Working Paper.

Dasaratha, Krishna and Kevin He (2022). "Learning from viral content". In: *arXiv preprint arXiv:2210.01267*.

Ide, Enrique (2025). "Automation, AI, and the intergenerational transmission of knowledge". In: *arXiv preprint arXiv:2507.16078*.

Ide, Enrique and Eduard Talamàs (2025). "Artificial Intelligence in the Knowledge Economy". In: *Journal of Political Economy* 133(12), pp. 3762–3800.

Koren, Miklós et al. (2026). "Vibe Coding Kills Open Source". In: *arXiv preprint arXiv:2601.15494*.

Lyu, Liang et al. (2025). "Wikipedia contributions in the wake of chatgpt". In: *Companion Proceedings of the ACM on Web Conference 2025*, pp. 1176–1179.

Maslej, Nestor et al. (2025). "Artificial intelligence index report 2025". In: *arXiv preprint arXiv:2504.07139*.

Reeves, Neal, Wenjie Yin, and Elena Simperl (2025). "Exploring the impact of ChatGPT on Wikipedia engagement". In: *Collective Intelligence* 4(3), p. 26339137251372599.

Rio-Chanona, Maria del, Nadzeya Laurentsyeva, and Johannes Wachs (2023). "Are large language models a threat to digital public goods? evidence from activity on stack overflow". In: *arXiv preprint arXiv:2307.07367*.

Shan, Guohou and Liangfei Qiu (2025). "Examining the impact of generative ai on users' voluntary knowledge contribution: Evidence from a natural experiment on stack overflow". In: *Information Systems Research*.

Xu, Bo and Dahui Li (2015). "An empirical study of the motivations for content contribution and community participation in Wikipedia". In: *Information & management* 52(3), pp. 275–286.

Zhao, Hangcheng and Ron Berman (2025). "The Impact of LLMs on Online News Consumption and Production". In: *arXiv preprint arXiv:2512.24968*.

Appendix

.1 Additional derivations omitted from the main text

Throughout, assume that creator quality follows a Pareto distribution with tail

$$\Pr(Q \geq q) = q^{-\gamma}, \qquad q \geq 1,$$

so the density is

$$f(q) = \gamma q^{-\gamma-1}, \qquad q \geq 1.$$

.1.1 Baseline demand share s_k

Users choose the source that gives the highest utility,

$$U_{ik} = \ln Q_k + \epsilon_{ik},$$

where ϵ_{ik} is i.i.d. Type-I extreme value with scale parameter θ. By the standard multinomial logit formula,

$$s_k = \Pr(U_{ik} \geq U_{ij}\ \forall j \in \mathcal{A}) = \frac{\exp\left(\frac{\ln Q_k}{\theta}\right)}{\sum_{j\in\mathcal{A}} \exp\left(\frac{\ln Q_j}{\theta}\right)} = \frac{Q_k^{1/\theta}}{\sum_{j\in\mathcal{A}} Q_j^{1/\theta}}.$$

.1.2 Baseline cutoff Q_0

In a cutoff equilibrium, the marginal entrant has quality Q_0 and earns zero profit:

$$F = r\, s(Q_0) = r\frac{Q_0^{1/\theta}}{\sum_{j\in\mathcal{A}} Q_j^{1/\theta}}.$$

Under the cutoff rule, $\mathcal{A} = \{j : Q_j \geq Q_0\}$. Hence, by the law of large numbers,

$$\sum_{j\in\mathcal{A}} Q_j^{1/\theta} = M\,\mathbb{E}\left[Q^{1/\theta}\mathbb{1}\{Q \geq Q_0\}\right] = M\int_{Q_0}^{\infty} q^{1/\theta}\gamma q^{-\gamma-1}\,dq.$$

Evaluating the integral,

$$\sum_{j\in\mathcal{A}} Q_j^{1/\theta} = M\frac{\gamma}{\gamma - 1/\theta} Q_0^{1/\theta-\gamma},$$

provided that $\gamma > 1/\theta$. Substituting back into the zero-profit condition,

$$F = r\frac{Q_0^{1/\theta}}{M\frac{\gamma}{\gamma-1/\theta}Q_0^{1/\theta-\gamma}} = r\frac{\gamma - 1/\theta}{M\gamma}Q_0^{\gamma}.$$

Therefore,

$$Q_0 = \left(\frac{FM\gamma}{r(\gamma - 1/\theta)}\right)^{1/\gamma}.$$

.1.3 Baseline mass of active creators $m_{\mathcal{A}}$

The mass of active creators is

$$m_{\mathcal{A}} = M \Pr(Q \geq Q_0) = M Q_0^{-\gamma}.$$

Substituting the expression for Q_0,

$$m_{\mathcal{A}} = M \left(\frac{FM\gamma}{r(\gamma - 1/\theta)} \right)^{-1} = \frac{r(\gamma - 1/\theta)}{F\gamma}.$$

.1.4 Extended-model shares s_A, s_W, and $s_{k|W}$

Conditional on choosing the web, users allocate demand across traditional sources according to

$$s_{k|W} = \frac{Q_k^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}}.$$

The inclusive value of the web nest is

$$V_W = \left(\sum_{j \in \mathcal{A}} Q_j^{1/\theta} \right)^{\theta}.$$

By assumption, AI quality is

$$Q_A = \phi V_W.$$

Hence the upper-level shares are

$$s_A = \frac{Q_A}{Q_A + V_W} = \frac{\phi V_W}{\phi V_W + V_W} = \frac{\phi}{1 + \phi}, \qquad s_W = \frac{V_W}{Q_A + V_W} = \frac{1}{1 + \phi}.$$

Therefore, the total share visiting traditional source k is

$$s_k = s_W s_{k|W} = \frac{1}{1 + \phi} \frac{Q_k^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}}.$$

.1.5 Extended-model cutoff Q_0^{AI}

For the marginal traditional entrant,

$$F(1 - \eta) = r\, s(Q_0^{AI}) = \frac{r}{1 + \phi} \frac{(Q_0^{AI})^{1/\theta}}{\sum_{j \in \mathcal{A}} Q_j^{1/\theta}}.$$

Under the cutoff rule, $\mathcal{A} = \{ j : Q_j \geq Q_0^{AI} \}$, so

$$\sum_{j \in \mathcal{A}} Q_j^{1/\theta} = M \int_{Q_0^{AI}}^{\infty} q^{1/\theta} \gamma q^{-\gamma - 1}\, dq = M \frac{\gamma}{\gamma - 1/\theta} (Q_0^{AI})^{1/\theta - \gamma}.$$

Therefore,

$$F(1 - \eta) = \frac{r}{1 + \phi} \frac{(Q_0^{AI})^{1/\theta}}{M \frac{\gamma}{\gamma - 1/\theta} (Q_0^{AI})^{1/\theta - \gamma}} = \frac{r}{1 + \phi} \frac{\gamma - 1/\theta}{M\gamma} (Q_0^{AI})^{\gamma}.$$

Solving for the cutoff yields

$$Q_0^{AI} = \left(\frac{FM\gamma}{r(\gamma - 1/\theta)} (1 - \eta)(1 + \phi) \right)^{1/\gamma} = Q_0 \big((1 - \eta)(1 + \phi) \big)^{1/\gamma}.$$

.1.6 Extended-model mass of active creators $m_{\mathcal{A}}^{AI}$

The mass of active traditional creators is

$$m_{\mathcal{A}}^{AI} = M \Pr(Q \geq Q_0^{AI}) = M(Q_0^{AI})^{-\gamma}.$$

Substituting the expression for Q_0^{AI},

$$m_{\mathcal{A}}^{AI} = M\left(\frac{FM\gamma}{r(\gamma - 1/\theta)}(1-\eta)(1+\phi)\right)^{-1} = \frac{r(\gamma - 1/\theta)}{F\gamma(1-\eta)(1+\phi)} = \frac{m_{\mathcal{A}}}{(1-\eta)(1+\phi)}.$$

.2 Preliminaries

Define

$$a \equiv \gamma - \frac{1}{\theta}, \qquad B \equiv \frac{a}{a-1}, \qquad \bar{C} \equiv \left(\frac{FM\gamma}{ra}\right)^{1/\gamma}.$$

In the extended model, the equilibrium cutoff is

$$Q_0^{AI} = \left(\frac{FM\gamma}{ra}(1-\eta)(1+\phi)\right)^{1/\gamma} = \bar{C}\big((1-\eta)(1+\phi)\big)^{1/\gamma}.$$

The consumed quality object is

$$\tilde{Q} \equiv s_A Q_A + s_W \sum_{j \in A} s_{j|W} Q_j,$$

with

$$s_A = \frac{\phi}{1+\phi}, \qquad s_W = \frac{1}{1+\phi}.$$

Lemma .1 (Closed-form expression for consumed quality)**.** *Suppose $a > 1$. Then equilibrium consumed quality can be written as*

$$\tilde{Q}(\eta, \phi) = \bar{C}\left[B(1-\eta)^{1/\gamma}(1+\phi)^{1/\gamma - 1} + \left(\frac{r}{F}\right)^{\theta} \phi^2 (1-\eta)^{1/\gamma - \theta}(1+\phi)^{1/\gamma - 1 - \theta}\right].$$

Proof. First, the traditional web component equals

$$\sum_{j \in A} s_{j|W} Q_j = B Q_0^{AI},$$

which is the same expression as in the baseline model, evaluated at the new cutoff Q_0^{AI}.

Second, the AI quality term is

$$Q_A = \phi\left(\sum_{j \in A} Q_j^{1/\theta}\right)^{\theta}.$$

Under the cutoff rule, the mass of active traditional creators is

$$m_A^{AI} = M \Pr(Q \geq Q_0^{AI}) = \frac{ra}{F\gamma(1-\eta)(1+\phi)}.$$

For a Pareto distribution with shape parameter γ,

$$\mathbb{E}[Q^{1/\theta} \mid Q \geq Q_0^{AI}] = \frac{\gamma}{\gamma - 1/\theta}(Q_0^{AI})^{1/\theta} = \frac{\gamma}{a}(Q_0^{AI})^{1/\theta}.$$

Hence

$$\sum_{j \in A} Q_j^{1/\theta} = m_A^{AI} \cdot \frac{\gamma}{a}(Q_0^{AI})^{1/\theta} = \frac{r}{F(1-\eta)(1+\phi)}(Q_0^{AI})^{1/\theta}.$$

Therefore,

$$Q_A = \phi \left(\frac{r}{F(1-\eta)(1+\phi)} \right)^{\theta} Q_0^{AI}.$$

Substituting into

$$\tilde{Q} = s_A Q_A + s_W \sum_{j \in A} s_{j|W} Q_j$$

gives

$$\tilde{Q} = \frac{\phi}{1+\phi} \cdot \phi \left(\frac{r}{F(1-\eta)(1+\phi)} \right)^{\theta} Q_0^{AI} + \frac{1}{1+\phi} \cdot B Q_0^{AI}.$$

Using

$$Q_0^{AI} = \bar{C}\big((1-\eta)(1+\phi)\big)^{1/\gamma}$$

yields

$$\tilde{Q}(\eta, \phi) = \bar{C} \left[B(1-\eta)^{1/\gamma}(1+\phi)^{1/\gamma - 1} + \left(\frac{r}{F}\right)^{\theta} \phi^2 (1-\eta)^{1/\gamma - \theta}(1+\phi)^{1/\gamma - 1 - \theta} \right].$$

□

.3 Useful derivative formulas

Differentiating the expression in Lemma .1 with respect to η gives

$$\frac{\partial \tilde{Q}}{\partial \eta} = \bar{C}(1-\eta)^{1/\gamma - \theta - 1}(1+\phi)^{1/\gamma - 1 - \theta} \left[\left(\theta - \frac{1}{\gamma}\right)\left(\frac{r}{F}\right)^{\theta} \phi^2 - \frac{B}{\gamma}(1-\eta)^{\theta}(1+\phi)^{\theta} \right]. \tag{4}$$

Differentiating with respect to ϕ gives

$$\frac{\partial \tilde{Q}}{\partial \phi} = \bar{C}(1-\eta)^{1/\gamma - \theta}(1+\phi)^{1/\gamma - 2 - \theta} \left[\left(\frac{r}{F}\right)^{\theta} \phi\Big(2 + \big(1 + \tfrac{1}{\gamma} - \theta\big)\phi\Big) - B\left(1 - \frac{1}{\gamma}\right)(1-\eta)^{\theta}(1+\phi)^{\theta} \right]. \tag{5}$$

.4 Proof of Proposition 1

Proof. Fix $\eta \in [0, 1)$. Evaluate the two derivatives at $\phi = 0$.

From (4),

$$\left.\frac{\partial \tilde{Q}}{\partial \eta}\right|_{\phi=0} = -\bar{C}\frac{B}{\gamma}(1-\eta)^{1/\gamma - 1} < 0.$$

From (5),

$$\left.\frac{\partial \tilde{Q}}{\partial \phi}\right|_{\phi=0} = -\bar{C}B\left(1 - \frac{1}{\gamma}\right)(1-\eta)^{1/\gamma} < 0.$$

Both inequalities are strict because $\bar{C} > 0$, $B > 0$, $\gamma > 1$, and $1 - \eta > 0$.

Since $\tilde{Q}(\eta,\phi)$ is continuously differentiable on $[0,1)\times[0,\infty)$, both partial derivatives are continuous in ϕ. Therefore, there exists $\underline{\phi}(\eta) > 0$ such that for all $\phi \in [0,\underline{\phi}(\eta))$,

$$\frac{\partial \tilde{Q}}{\partial \eta} < 0 \qquad \text{and} \qquad \frac{\partial \tilde{Q}}{\partial \phi} < 0.$$

This proves Proposition 1. □

.5 Proof of Proposition 2

Proof. Fix $\eta \in [0,1)$ and suppose

$$\theta < 1 + \frac{1}{\gamma}.$$

Because $\gamma > 1$, this implies in particular that $\theta < 2$.

Consider first (4). The prefactor outside the square brackets is strictly positive. Hence the sign of $\partial\tilde{Q}/\partial\eta$ is determined by

$$G_\eta(\phi) \equiv \left(\theta - \frac{1}{\gamma}\right)\left(\frac{r}{F}\right)^\theta \phi^2 - \frac{B}{\gamma}(1-\eta)^\theta(1+\phi)^\theta.$$

Since $\theta < 2$, the positive term grows like ϕ^2 while the negative term grows like ϕ^θ. Therefore,

$$G_\eta(\phi) \to +\infty \qquad \text{as} \qquad \phi \to \infty.$$

Hence there exists ϕ^1_η such that for all $\phi > \phi^1_\eta$,

$$\frac{\partial \tilde{Q}}{\partial \eta} > 0.$$

Now consider (5). Again, the prefactor outside the square brackets is strictly positive, so the sign of $\partial\tilde{Q}/\partial\phi$ is determined by

$$G_\phi(\phi) \equiv \left(\frac{r}{F}\right)^\theta \phi\Big(2 + \big(1 + \tfrac{1}{\gamma} - \theta\big)\phi\Big) - B\left(1 - \frac{1}{\gamma}\right)(1-\eta)^\theta(1+\phi)^\theta.$$

Because $\theta < 1 + 1/\gamma$, the coefficient

$$1 + \frac{1}{\gamma} - \theta$$

is strictly positive. Thus the leading positive term in $G_\phi(\phi)$ grows like ϕ^2, while the negative term grows like ϕ^θ with $\theta < 2$. Hence

$$G_\phi(\phi) \to +\infty \qquad \text{as} \qquad \phi \to \infty.$$

Therefore there exists ϕ^2_η such that for all $\phi > \phi^2_\eta$,

$$\frac{\partial \tilde{Q}}{\partial \phi} > 0.$$

Let

$$\overline{\phi}(\eta) \equiv \max\{\phi^1_\eta, \phi^2_\eta\}.$$

Then for every $\phi > \overline{\phi}(\eta)$,

$$\frac{\partial \tilde{Q}}{\partial \eta} > 0 \qquad \text{and} \qquad \frac{\partial \tilde{Q}}{\partial \phi} > 0.$$

This proves Proposition 2. □

.6 Proof of Proposition 3

Proof. Let

$$a \equiv \gamma - \frac{1}{\theta}, \qquad B \equiv \frac{a}{a-1}.$$

In the no-AI baseline, consumed quality is

$$\tilde{Q}^B = BQ_0,$$

where

$$Q_0 = \left(\frac{FM\gamma}{ra}\right)^{1/\gamma}.$$

In the AI economy, consumed quality is

$$\tilde{Q}^{AI} = Q_0^{AI} \frac{1}{1+\phi} \left[B + \phi^2 \left(\frac{r}{F(1-\eta)(1+\phi)}\right)^{\theta}\right],$$

where

$$Q_0^{AI} = Q_0\big((1-\eta)(1+\phi)\big)^{1/\gamma}.$$

Therefore,

$$\frac{\tilde{Q}^{AI}}{\tilde{Q}^B} = (1-\eta)^{1/\gamma}(1+\phi)^{1/\gamma-1} \left[1 + \frac{1}{B}\phi^2 \left(\frac{r}{F(1-\eta)(1+\phi)}\right)^{\theta}\right].$$

Since $1/B = (a-1)/a = (\gamma - 1/\theta - 1)/(\gamma - 1/\theta)$, this becomes

$$\frac{\tilde{Q}^{AI}}{\tilde{Q}^B} = (1-\eta)^{1/\gamma}(1+\phi)^{1/\gamma-1} \left[1 + \frac{\gamma - 1/\theta - 1}{\gamma - 1/\theta}\phi^2 \left(\frac{r}{F(1-\eta)(1+\phi)}\right)^{\theta}\right].$$

Hence $\tilde{Q}^{AI} > \tilde{Q}^B$ if and only if the inequality in the proposition holds.

It remains to show that this inequality holds for all sufficiently large ϕ. Expanding the ratio,

$$\frac{\tilde{Q}^{AI}}{\tilde{Q}^B} = (1-\eta)^{1/\gamma}(1+\phi)^{1/\gamma-1} + \frac{\gamma - 1/\theta - 1}{\gamma - 1/\theta}\left(\frac{r}{F}\right)^{\theta}(1-\eta)^{1/\gamma-\theta}\phi^2(1+\phi)^{1/\gamma-1-\theta}.$$

As $\phi \to \infty$, the first term is asymptotically negligible relative to the second, while the second behaves like

$$C(\eta)\phi^{1+1/\gamma-\theta},$$

for some constant

$$C(\eta) = \frac{\gamma - 1/\theta - 1}{\gamma - 1/\theta}\left(\frac{r}{F}\right)^{\theta}(1-\eta)^{1/\gamma-\theta} > 0.$$

Under the assumption

$$\theta < 1 + \frac{1}{\gamma},$$

the exponent $1 + 1/\gamma - \theta$ is strictly positive. Therefore,

$$\frac{\tilde{Q}^{AI}}{\tilde{Q}^B} \to \infty \qquad \text{as} \qquad \phi \to \infty.$$

By continuity of the ratio in ϕ, there exists $\bar{\phi}(\eta) > 0$ such that for all $\phi > \bar{\phi}(\eta)$,

$$\frac{\tilde{Q}^{AI}}{\tilde{Q}^{B}} > 1.$$

Equivalently,

$$\tilde{Q}^{AI}(\eta, \phi) > \tilde{Q}^{B}.$$

$\square$