A Model of Behavioral Manipulation*

Daron Acemoglu † Ali Makhdoumi ‡ Azarakhsh Malekian § Asuman Ozdaglar ¶

May 19, 2023

Abstract

This paper builds a model of online behavioral manipulation. Our approach is motivated by advances in AI that are massively expanding the amount of information platforms have about users. In our model, platforms dynamically offer one of n products and an associated price to a user who is uncertain about and can slowly learn the quality of the products. The signals about product quality, however, also depend on extraneous factors, such as the appearance of the good or other attributes that make it appear more attractive than it is, at least in the short run — what we refer to as "*glossiness*". AI tools enable platforms to better estimate the glossiness of products and enable them to engage in behavioral manipulation. Formally, we study a continuous-time experimentation problem in which the platform has different beliefs and guides the experimentation choices of the user.

Our analysis provides three main results. First, when glossiness is absent or short-lived, the superior information of the platform benefits consumers, thus validating the default position in the literature that AI can help users and consumers. Second, in sharp contrast, when glossiness is long-lived, superior information of the platform makes users worse off because it enables behavioral manipulation — the platform systematically distorts user choices towards glossy products. Third, as the number of products increases, behavioral manipulation intensifies, and the user is more likely to be confronted with low-quality, glossy products.

Keywords: behavioral economics, behavioral manipulation, behavioral surplus, data, experimentation, online markets, platforms.

JEL Classification: D90, D91, L86, D83.

*We thank participants at various seminars and conferences for comments and feedback. We are particularly grateful to our discussant Liyan Yang. We also gratefully acknowledge financial support from the Hewlett Foundation, Smith Richardson Foundation, and the NSF. This paper was prepared in part for and presented at the 2023 AI Authors' Conference at the Center for Regulation and Markets (CRM) of the Brookings Institution, and we thank the CRM for financial support as well.

†Department of Economics, Massachusetts Institute of Technology. daron@mit.edu

‡Fuqua School of Business, Duke University. ali.makhdoumi@duke.edu

§Rotman School of Management, University of Toronto. azarakhsh.malekian@rotman.utoronto.ca

¶Dept. of Electrical Engineering and Computer Science, Massachusetts Institute of Technology. asuman@mit.edu

1 Introduction

The data of billions of individuals are currently being utilized for personalized advertising, product offerings, and pricing. The collection and mining of individual data are set to grow exponentially in the coming years with more extensive data collection from new online apps and integrated technologies such as the Internet of Things and the more widespread applications of artificial intelligence (AI) and machine learning techniques. This exponentially increasing amount of data in the hands of tech companies such as Google, Facebook, Amazon, Netflix, and others are raising a host of concerns, including those related to privacy.

Many economists, AI experts, and researchers have a fairly optimistic take on this: more data will mean more informative advertising, better product targeting, and much finer service differentiation, benefiting users. Zuboff [2019] in *The Age of Surveillance Capitalism* takes a diametrically opposed position. She argues that the new economic model "claims human experience as free raw material for hidden commercial practice process of extraction, prediction, and sales" (p. 9). She maintains that the economic logic of this new system, dominated by Big Tech, turns on the extraction of "behavioral surplus", created by what she calls "behavioral modification", meaning the ability of tech companies with vast amounts of data to be able to modify the behavior of users and profit from this. A similar perspective has been discussed in Hanson and Kysar [1999] who note: "Once one accepts that individuals systematically behave in non-rational ways, it follows from an economic perspective that others will exploit those tendencies for gain" (p. 105).

Yet these arguments are insufficient since the modification of behavior may be for the user's good (directing her towards better products). The behavior of users can be modified in a way that is consistent with their intrinsic or extrinsic welfare. In this paper, we are interested in the conditions under which behavioral modification harms consumers. In particular, we develop a behavioral model to theoretically explore the "bad" type of modification, which we refer to as *behavioral manipulation*. Some examples include retailers forecasting whether a woman is pregnant and start sending them hidden ads for baby products; various companies estimating "prime vulnerability moments" and send ads for beauty products; marketing strategies targeted at more "vulnerable populations" such as the elderly or children; websites favoring products (including credit cards and subscription programs) with delayed costs and apparent short-term benefits, especially to users that they estimate might be attracted by such offers; or streaming platforms algorithmically estimating more addictive videos for certain consumer types to maximize engagement. This type of behavioral manipulation would not be possible if consumers were fully rational and understood that with vast amounts of data, platforms may know and manipulate relevant information that they do not know about their preferences. However, when the full implications of the superior information of platforms, due to advances in AI and big data, are not properly understood by users, there could be room for much more pernicious types of behavioral modification, which we systematically explore in this paper.

1.1 Our Model and Results

We suppose consumers are sufficiently used to and can thus act rationally in the status quo environment. But they are not "hyper-rational" enough to understand how the environment changes after technology platforms can collect vast amounts of data both on them and on other consumers with similar characteristics, thus having much greater *behavioral predictive power* (meaning prediction of future behavior conditional on current inducements). It is this behavioral predictive power of platforms that opens the door to behavioral manipulation.

More formally, we study a setup in which a platform can offer to a user one of n products and an associated price at any point in time. Each product i has a quality that is either 1 (high) or 0 (low). When the user consumes a product, she receives a signal about its quality. There is an extraneous factor that impacts the signal received by the user: some products may initially be in a state that masks bad news, capturing some systematic tendencies that users have (we tend to be favorably disposed toward products with some properties such as better appearance, salient attractive characteristics, or hidden and delayed costs). We denote this state for product i at time 0 by $\alpha_{i,0}$, and refer to it as the "*glossiness*" of the item. We model the dynamics of $\alpha_{i,t}$ with a continuous-time Markov chain that starts from either 0 or 1. The bad news about the quality of a product will be masked if and only if the initial glossiness state is $\alpha = 1$ — meaning that the product is glossy. However, this state is temporary, indicating that a low-quality product may initially appear good, but eventually, the user will understand this. In particular, the continuous-time Markov chain transitions from an initial glossiness state of 1 to an absorbing glossiness state of $\alpha = 0$ at the rate ρ. Finally, if a low-quality product is not glossy, the user receives bad news about it at the rate γ (here, bad news means the product is of low quality).

We consider two informational environments: (i) A *pre-AI environment* (or status quo) in which the platform has no informational advantage, and the glossiness state is unknown to both the platform and the user. By observing the signals — specifically, whether bad news has arrived or not — the platform and the user update their belief about both the product's underlying quality and the product's state. Put differently, in this environment, there is symmetric information between the user and the platform. (ii) A *post-AI environment* in which the platform can estimate the product's glossiness for this user from the behavior of others with similar characteristics from whom it has collected data.[1] The key (and only) behavioral bias in our model is that users do not understand that the platform knows and can exploit its knowledge of the glossiness state in this new environment. This means, in particular, that the user does not learn from the offering and the price decision of the platform.

After characterizing the learning dynamics and the equilibrium strategy of the platform and the user in both the pre-AI and post-AI environments, we turn to our main results.

First, we prove that for a large enough transition rate ρ — which means that the initially glossy products remain glossy for a short period — the ex-ante expected user utility, welfare, and platform profits in

[1]Specifically, the platform can harness the purchase, webpage visit and review data of users with similar characteristics to estimate which products have strong short-term appeal to the relevant demographic group.

In this context, see Ekmekci [2011] and Che and Hörner [2018] for past analysis of the information content of online recommendation systems. In Acemoglu et al. [2022a], even if users are fully Bayesian, they will learn less from such reviews than the platform, because they see a summary statistic of past scores.

the post-AI environment are larger than in the pre-AI environment. This confirms the common wisdom in the literature that data collected by online platforms will enable better product offerings and higher quality matches for the users, thus ultimately benefiting them. To further understand the intuition for this result, we highlight two forces that impact the user utility in the post-AI environment in opposing directions. There is a *manipulation effect*, resulting from the platform offering a low-quality glossy product ahead of non-glossy products. Counteracting this, there is a *helpfulness effect*, which means that the platform guides the user towards higher-quality products. With a sufficiently large transition rate ρ, the helpfulness effect dominates. This is because, for low-quality products, the "manipulation"state in which bad news is masked is short, making the manipulation effect both less profitable for the platform and also negligible for the utility of the user.

Second, in contrast with the modal view in the literature, we establish that when the transition rate ρ is small — the initially glossy products remain glossy for a long period — the ex-ante expected user utility and welfare in the post-AI environment will be smaller than in the pre-AI environment. While this happens, platform profits always increase, because platforms benefit from this manipulation. Intuitively, in this case with persistent platform superior information, the manipulation effect dominates, as glossy items remain so for a while, enabling the platform to extract more surplus from the user.

Third, we establish that, in the case with a sufficiently small transition rate ρ, as the number of products increases, the ex-ante expected user utility and welfare in the post-AI environment decreases, while the platform continues to benefit. Intuitively, a greater number of products provides more opportunities for the platform to manipulate user behavior by finding glossy items that are profitable. This result implies that as advances in AI simultaneously expand the ability of platforms to collect more information about users and widen their offerings, there may be a "double whammy" from the viewpoint of the consumers.

1.2 Related Literature

Our paper relates to the literature on the legal and economic aspects of big data owned by digital platforms. For example, Pasquale [2015] and Zuboff [2019] argue that big data gives digital platforms enormous power for actions that can potentially harm users. Similarly, Calo [2013], Lin [2016], and Zarsky [2019] argue that the combination of knowing the consumers' behavior and being able to adapt the product (or content) accordingly enables digital platforms to manipulate consumers which requires regulation (see also Tirole [2020] for a discussion of regulations and policy implications).

More closely related to our paper are a few works investigating whether information harms consumers, either because of surveillance (Tirole [2019]) or because of price discrimination reasons. In particular, Taylor [2004], Acquisti and Varian [2005], Fudenberg and Villas-Boas [2006], Hidir and Vellodi [2019], Bergemann et al. [2015], Bonatti and Cisternas [2020] study how the history of consumers' behavior can be used for estimating their willingness-to-pay and price discrimination. See Acquisti et al. [2016], Bergemann and Bonatti [2019], and Agrawal et al. [2018] for excellent surveys of different aspects of this literature. We depart from this literature by focusing on behavioral manipulation — distortion of product choice and experimentation — rather than price discrimination.[2] Within this literature, our paper is most closely related to independent work by Liu et al. [2023], which examines the impact of data sharing on consumer welfare and algorithmic inequality. This paper argues that data sharing benefits consumers by providing access to desired products and services, but also enables firms to identify and exploit temptation and self-control problems. This differs but is complementary to our emphasis on platforms identifying sources of mis-learning due to glossiness and other extraneous factors.

Our paper also relates to the growing literature on information markets. One branch of this literature focuses on the use of personal data for improved allocation of online resources (e.g., Bergemann and Bonatti [2015], Goldfarb and Tucker [2011], and Montes et al. [2019]). Another branch investigates how information can be monetized either by dynamic sales or optimal mechanisms. For example, Anton and Yao [2002], Babaioff et al. [2012], Eső and Szentes [2007], Horner and Skrzypacz [2016], Bergemann et al. [2018], and Eliaz et al. [2019] consider either static or dynamic mechanisms for selling data, and Admati and Pfleiderer [1986] and Begenau et al. [2018] study markets for financial data. A third branch focuses on optimal collection and acquisition of information, for example, Roesler and Szentes [2017], Zhong [2017], Agarwal et al. [2019], Chen and Zheng [2019], and Chen et al. [2018]. More recently, Kamenica et al. [2011], Fudenberg and Lanzani [2020], Fudenberg et al. [2020], Frostig and Weiss [2016], Heidhues et al. [2016], Heidhues et al. [2018], Gossner et al. [2018], Liang et al. [2021], Bardhi [2020], and Heidhues and Kőszegi [2018] have studied the sale and aggregation of information.

Our work more directly builds on the literature on experimentation and multi-armed bandits. Bandit models, first analyzed in Robbins [1952], have been used for many applications, including, pricing (Rothschild [1974]); the optimal design of clinical trials (Berry and Fristedt [1985]); product search (Bolton and Harris [1999]); research and development problems (Keller and Rady [2010]); and learning and wage setting (Felli and Harris [1996]). See also Berry and Fristedt [1985] and Krylov [2008] for book-length treatments of bandit problems. Even more closely related to our analysis is the exponential bandit framework of Keller et al. [2005] which we build on and extend. The mathematical arguments in our paper are different from those in these past works, however, because we have to keep track of the beliefs of the platform and the user, which follow different laws of motion, and we use structural properties of the Gittins index characterization in order to compare the expected discounted values of different strategies.

Finally, our work is inspired and draws on a growing body of evidence that platforms are acquiring a large amount of information about users and influencing their behavior by various targeting, advertising, and product offering strategies. A recent contribution to this literature that summarizes previous work is Susser and Grimaldi [2021]. A few other contributions include Acquisti et al. [2015] and Martin and Nissenbaum [2016] who empirically study individuals' ability to control their personal information and Calvo et al. [2020] who study how platform algorithms can enable price manipulation.

[2] Our paper also relates to the works that study privacy concerns in data markets such as Choi et al. [2019], Bergemann et al. [2019], Acemoglu et al. [2022b], and Gradwohl [2017], who study data externalities; Fainmesser et al. [2019] and Jullien et al. [2020], who consider the negative effects of leaking user's (private) personalized; and Ichihashi [2020b] and Ichihashi [2020a], who explore the role of information intermediaries and dynamic data collection by platforms. More recent work by Köbis et al. [2021] has empirically investigated broader social implications of extensive data collection and use of machine learning techniques online.

The rest of the paper proceeds as follows. Section 2 presents our model. Section 3 provides the learning trajectory of the user and the platform and preliminary characterization of the equilibrium. Sections 4 and 5 characterize the equilibrium in the pre-AI and post-AI environments, respectively. Section 6 establishes that manipulation occurs — it is possible for the user to be worse off in the post-AI. Section 7 concludes, while the Appendix presents the proofs and additional results.

2 Model

We consider a platform that hosts $n > 1$ products denoted by $\mathcal{N} = \{1, \ldots, n\}$. Each product $i \in \mathcal{N}$ has a quality denoted by $\theta_i \in \{0, +1\}$ and an initial glossiness state $\alpha_{i,0} \in \{0, 1\}$. The quantity θ_i is fixed and we refer to it as the *quality* of product i ($\theta_i = 1$ means a high-quality product). On the other hand, $\alpha_{i,t}$ changes over time and we refer to $\alpha_{i,t}$ as the *glossiness state*, or simply *glossiness*, of product i at time t.

The platform interacts with a user over an infinite time horizon in continuous time, and at each time $t \in \mathbb{R}_+$ offers one of the products and an associated price to the user. We let $x_{i,t} \in \{0, 1\}$ denote the platform's decision about offering product i at time t and $p_{i,t}$ stand for the offered price. If the user purchases the product, she receives a signal about its quality and updates her belief (about both θ_i and $\alpha_{i,t}$). In particular, if the user purchases product i at time t, she observes a signal, described as follows:

1. If $\theta_i = 1$, then there is no "bad" news about product quality.
2. If $\theta_i = 0$ and $\alpha_{i,t} = 0$, then bad news (fully revealing signal about product quality) arrives at the rate $\gamma > 0$.
3. If $\theta_i = 0$ and $\alpha_{i,t} = 1$, then there is no bad news about product quality.

In summary, a high-quality product never generates bad news (the first case), but a low-quality product may (the second case). In the third case, where $\alpha_{i,t} = 1$, the low-quality product is glossy and does not generate bad news. Glossiness is what makes products appear better than they are (at least in the short run) and enables behavioral manipulation. In what follows, we let $S_{i,t} = \mathrm{B}$ to designate the arrival of bad news for product i at time t and $S_{i,t} = \mathrm{NB}$ to designate no bad news for product i at time t. When purchasing product i, the user does not receive any signal about other products $j \in \mathcal{N} \setminus \{i\}$.

We assume that $\alpha_{i,t}$ follows a continuous-time Markov chain that transitions from state 1 to state 0 at the rate $\rho \geq 0$ where state 0 is an absorbing state. This means that a low-quality product appears high-quality for a while, but this type of glossiness is not forever. Formally, $\{\alpha_{i,t}\}_{t\geq 0}$ is a homogeneous continuous-time Markov Chain (CTMC) with two states $\{0, 1\}$ and *transition rate matrix*

$$\begin{bmatrix} 1 & 0 \\ \rho & 1-\rho \end{bmatrix}.$$

2.1 Information Environment

The platform and the user share a common prior about the initial glossiness state and the quality of each product. In particular, they both understand that a low-quality product can initially be glossy ($\alpha_{i,0} = 1$)

or non-glossy ($\alpha_{i,0} = 0$), while high-quality products are always non-glossy. Specifically, for any $i \in \mathcal{N}$ we have the following as common knowledge:

$$\mathbb{P}[\alpha_{i,0} = 0 \mid \theta_i = 1] = 1, \;\; \mathbb{P}[\alpha_{i,0} = 1 \mid \theta_i = 0] = \lambda, \text{ and } \;\; \mu_{i,0} = \mathbb{P}[\theta_i = 1] \tag{1}$$

for some initial belief $\mu_{i,0} \in [0, 1]$ and $\lambda \in (0, 1)$.

At any time t, the information available to the user and the platform is whether bad news for product i has arrived for all $i \in \mathcal{N}$. We let

$$\mathrm{NBN}_{i,t} = \{S_{i,\tau} = \mathrm{NB} \text{ for all } \tau \in [0, t)\}$$

denote the event that there has been no bad news for product i at times $\tau \in [0, t)$. We also use

$$I_{i,t} = \begin{cases} 1 & \text{if } S_{i,\tau} = \mathrm{NB} \text{ for all } \tau \in [0, t) \\ 0 & \text{otherwise} \end{cases}$$

be the indicator that product i's quality at time t is not 0 from the users' perspective (and therefore, the platform can still offer it and make a non-zero payoff).

The information sets of the platform and the user only differ in the initial signals that the platform observes about the glossiness of different products. In particular, we consider two settings:

1. *Pre-AI environment* in which the platform does not have an information advantage and observes the same signals as the user.

2. *Post-AI environment* in which the platform has an information advantage and observes the initial glossiness state $\alpha_{i,0}$ for all $i \in \mathcal{N}$.

The justification for additional information possessed by the platform in the post-AI environment was already discussed in the Introduction, but briefly, we consider this as a simple representation of the fact that the platform has a wealth of data about the behavior and preferences of other users with similar characteristics, and thus can estimate which are the products that will (temporarily) appear to the user in question to be more attractive than their quality warrants. For simplicity, we suppose these data do not inform the platform about the true quality of the product for the user.

In summary, in the pre-AI environment the platform's information at time t is $\{\mathrm{NBN}_{i,t}\}_{i \in \mathcal{N}, t}$ which is the same as the user's information. In the post-AI environment, the platform's information at time t is given by $\{\alpha_{i,0}, \mathrm{NBN}_{i,t}\}_{i \in \mathcal{N}, t}$ which, in addition to the user's information, contains $\alpha_{i,0}$ for all $i \in \mathcal{N}$. Initially, the user's belief about θ_i for all $i \in \mathcal{N}$ is drawn uniformly. Over time, all agents update their belief about the product's true quality using Bayes's rule based on the information they observe (see Figure 1). We characterize the evolution of their belief in Section 3. In what follows, we establish that the information advantage of the platform in the post-AI environment enables behavioral manipulation and can potentially reduce user utility. A few points about the model are worth mentioning. First, as already noted, we assume that in the post-AI environment, the platform can perfectly estimate the

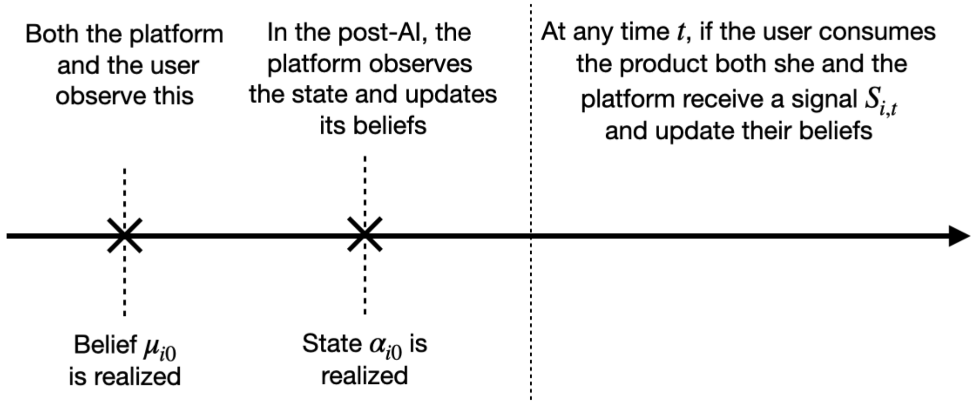

Figure 1: Timing of events and information structure.

initial glossiness state $\alpha_{i,0}$. This is enabled by having access to the data of users with similar preferences and the ability to learn from them. That the platform observes glossiness perfectly is for simplicity, and our main results continue to hold even if there are some errors in the platform's estimate of the initial glossiness state. Second, we assume that the information advantage of the platform is about the initial glossiness state and not the product's truth quality, which could be idiosyncratic across users. Our main results continue to hold even when the platform can estimate the true underlying quality. Third, we are assuming that the only information advantage of the platform is about the initial glossiness state and the platform does not observe when the glossiness state of a product switches from 1 to 0, which is also not directly observed by the user and is likely to be idiosyncratic across users within the same demographic group. This assumption limits the informational advantage of the platform as well. We establish that even with this limited form of informational advantage, the platform can have a very profitable behavioral manipulation strategy, and our qualitative results continue to apply if the platform can perfectly observe $\alpha_{i,t}$. Finally, as we explain below, we will assume that the users will remain unaware of the platform's ability to estimate the initial glossiness state in the post-AI environment. This could be interpreted as a behavioral bias, but our favorite interpretation would be that most users do not understand the informational capabilities of digital platforms in today's economy.

2.2 Utilities

If the platform offers item i at time t, meaning that $x_{i,t} = 1$ and $x_{j,t} = 0$ for all $j \in \mathcal{N} \setminus \{i\}$, at price $p_{i,t}$, the instantaneous user's utility from purchasing the item is

$$\theta_i - p_{i,t},$$

and the instantaneous utility of the platform, if the user purchases, is $p_{i,t}$. We let $b_{i,t} \in \{0, 1\}$ denote the purchase decision of the user if product i is offered to her. We also assume that both the platform and the user discount the future at rate r for some $r \in \mathbb{R}_+$ and update their beliefs using Bayes's rule. The platform's expected profit (utility) at time 0 is given by

$$\mathbb{E}_0^P \left[\int_0^\infty re^{-rt} \sum_{i=1}^n x_{i,t} p_{i,t} b_{i,t} dt \right],$$

and the user's expected utility at time 0 is

$$\mathbb{E}_0^U \left[\int_0^\infty re^{-rt} \sum_{i=1}^n x_{i,t} b_{i,t} \left(\theta_i - p_{i,t} \right) dt \right],$$

where $x_{i,t}$ and $p_{i,t}$ are adapted to the natural filtration of the platform at time t, denoted by $\mathcal{F}_t^P$ and $b_{i,t}$ is adapted to the natural filtration of the user's information set at time t, denoted by $\mathcal{F}_t^U$.

2.3 Equilibrium Concept

In the pre-AI environment, when the platform does not have an information advantage, neither the user nor the platform knows or receives signals about the initial glossiness $\alpha_{i,0}$ for $i \in \mathcal{N}$. Therefore, the platform's decision at time t, $(x_{i,t}, p_{i,t})$ for $i \in \mathcal{N}$, does not contain any extra information about the quality of the product or the glossiness state. In this environment, both the platform and the user are fully Bayesian, and we adopt *Markov Perfect Equilibrium (MPE)* as the solution concept. This means that the platform's and the user's decisions at each time t are dependent on the game's history through their respective beliefs about the quality of each product and their strategies are conditioned on the payoff-relevant states in this game, which are summarized by these beliefs.

In the post-AI era, the platform possesses superior knowledge about the initial glossiness state $\alpha_{i,0}$. As noted above, the platform does not receive additional information about this variable after the initial date. Nevertheless, the knowledge of $\alpha_{i,0}$ enables it to estimate $\alpha_{i,t}$ more accurately than the user for any t. In this environment, we again use the notion of *Markov Perfect Equilibrium (MPE)*, except that in this case each player evaluates payoffs according to their beliefs, which differ because they have access to different information. In particular, the platform's decision at time t, $(x_{i,t}, p_{i,t})$ for $i \in \mathcal{N}$, contains additional information about the quality of the product and the glossiness state, and the only behavioral aspect of our model is that the user does not recognize this informational superiority and does not update her beliefs on the basis of the offers of the platform. This behavioral assumption is in line with the equilibrium concepts proposed in Eyster and Rabin [2005] and Esponda [2008]. In particular, our equilibrium concept can be viewed as the equivalent of the Berk-Nash equilibrium notion in Esponda [2008].[3]

[3](Non-behavioral) Bayesian Nash equilibrium would involve the user drawing inferences from the product offering and pricing decisions of the platform. It is straightforward to show that one (simple) equilibrium in this case would be that the user interprets any deviation from the offering of the highest-belief product as a signal of manipulation, and thus the platform would be unable to use any of its post-AI information.

3 Belief Dynamics and Equilibrium Characterization

We next characterize the learning trajectories of the user and the platform and then provide some preliminary characterization of the equilibrium.

3.1 Belief Dynamics

We now provide belief dynamics in the pre-AI and post-AI environments.

Pre-AI environment: At any given time t, if bad news has been received for a product $i \in \mathcal{N}$, both the user and the platform believe with certainty that the product is of low quality. Otherwise, their Bayesian belief will be given by the posterior probability:

$$\mu_{i,t} = \mathbb{P}\left[\theta_i = 1 \mid \mathrm{NBN}_{i,t}\right],$$

which represents the probability of the product being of high quality given that no bad news has been received up to that point in time. In forming this expectation, both the user and the platform recognize that the reason why no bad news may have been received so far is that the product is glossy ($\alpha_{i,t} = 1$). For this reason, they also have to update the probability that the product is glossy, which is

$$\lambda_{i,t} = \mathbb{P}\left[\alpha_{i,t} = 1 \mid \theta_i = 0, \mathrm{NBN}_{i,t}\right].$$

We next characterize the dynamics of the user's belief in the pre-AI environment, which are identical to the platform's belief. These belief dynamics depend on the platform's strategy given by $\{x_{i,t}, p_{i,t}\}_{t\in\mathbb{R}_+, i\in\mathcal{N}}$ and the user's purchasing strategy given by $\{b_{i,t}\}_{t\in\mathbb{R}_+, i\in\mathcal{N}}$. Throughout, when writing the continuous-time belief updating dynamics, we follow the standard practice of dropping terms of order $(dt)^2$ and smaller.

Lemma 1. *For any $\{x_{i,t}, p_{i,t}, b_{i,t}\}_{i\in\mathcal{N},t}$, in the pre-AI environment, the beliefs of the user and the platform evolve as follows:*

$$\begin{aligned}
d\mu_{i,t} &= x_{i,t} b_{i,t} \mu_{i,t}(1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt \quad \textit{with } \mu_{i,0} = \mathbb{P}[\theta_i = 1] \\
d\lambda_{i,t} &= x_{i,t} b_{i,t} \lambda_{i,t}\left((1-\lambda_{i,t})\gamma - \rho\right) dt \qquad \textit{with } \lambda_{i,0} = \lambda.
\end{aligned}$$

Moreover, the probability of receiving bad news for product i at time t, given that bad news has not arrived during $[0, t)$, is:

$$\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1] = x_{i,t} b_{i,t}\left(1-\mu_{i,t}\right)\left(1-\lambda_{i,t}\right)\gamma dt.$$

The proof of this lemma as well as other omitted proofs are given in the Appendix.

The evolution of these beliefs highlights that $\mu_{i,t}$ moves more rapidly when it is closer to $1/2$ (here, $1/2$ means that the user's current belief is uninformative) and when γ is larger (which means bad news about low-quality products are more frequent). We emphasize again that the user is fully Bayesian and knows about the (existence of) glossiness state $\alpha_{i,t}$. The updating equation also shows that when $\lambda_{i,t}$ is smaller, $\mu_{i,t}$ moves faster, because it is recognized that no bad news is more likely to correspond to a high-quality product in this case (whereas with higher $\lambda_{i,t}$ it may be the product's glossiness that hides bad news).

Post-AI environment: Again, at time t, if bad news about product $i \in \mathcal{N}$ has arrived, then the belief of the platform and the user is that the product is of low quality with probability 1. Otherwise, the user's posterior Bayesian belief is

$$\mu_{i,t} = \mathbb{P}\left[\theta_i = 1 \mid \mathrm{NBN}_{i,t}\right].$$

In the post-AI environment, the main difference is that the platform knows the initial glossiness state $\alpha_{i,0}$ of all products and therefore has a different belief about product qualities. We let

$$\mu^{(P)}_{i,t} = \mathbb{P}\left[\theta_i = 1 \mid \mathrm{NBN}_{i,t}, \alpha_{i,0}\right]$$

stand for the belief of the platform (hence, the superscript "P") about product i's quality. Notice that if the initial glossiness state is $\alpha_{i,0} = 1$, then the platform knows the product is of low quality with probability 1, and hence $\mu^{(P)}_{i,t} = 0$. If the initial glossiness state is $\alpha_{i,0} = 0$, then the platform's belief evolves as we characterize in the next lemma. We also denote by

$$\lambda^{(P)}_{i,t} = \mathbb{P}\left[\alpha_{i,t} = 1 \mid \mathrm{NBN}_{i,t}, \alpha_{i,0}\right]$$

the belief of the platform about the product being in state $\alpha_{i,t} = +1$ given the initial glossiness state and that no bad news has arrived at time t. Notice that if the initial glossiness state is $\alpha_{i,0} = 0$, then the state remains at $\alpha_{i,t} = 0$ at time t, and therefore $\lambda^{(P)}_{i,t} = 0$. If the initial glossiness state is $\alpha_{i,0} = 1$, then $\lambda^{(P)}_{i,t}$ evolves as we show next.

Lemma 2. *For any $\{x_{i,t}, p_{i,t}, b_{i,t}\}_{i\in\mathcal{N}, t\in\mathbb{R}_+}$, in the post-AI environment, the user's beliefs are the same as in Lemma 1, and the platform's beliefs evolve as follows:*

- *If $\alpha_{i,0} = 1$, then*

$$\mu^{(P)}_{i,t} = 0 \text{ and } d\lambda^{(P)}_{i,t} = x_{i,t} b_{i,t} \lambda^{(P)}_{i,t}((1-\lambda^{(P)}_{i,t})\gamma - \rho)dt \text{ with } \lambda^{(P)}_{i,0} = 1.$$

- *If $\alpha_{i,0} = 0$, then*

$$d\mu^{(P)}_{i,t} = x_{i,t} b_{i,t} \mu^{(P)}_{i,t}(1-\mu^{(P)}_{i,t})\gamma dt \text{ with } \mu^{(P)}_{i,0} = \frac{\mu_{i,0}}{\mu_{i,0} + (1-\mu_{i,0})(1-\lambda)} \text{ and } \lambda^{(P)}_{i,t} = 0.$$

Moreover, from the platform's perspective, the probability of receiving bad news for product i at time t, given *that bad news has not arrived during* $[0, t)$*, is:*

$$\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1, \alpha_{i,0} = 0] = x_{i,t} b_{i,t} (1 - \mu_{i,t}^{(P)}) \gamma dt \quad \textit{for } \alpha_{i,0} = 0$$
$$\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1, \alpha_{i,0} = 1] = x_{i,t} b_{i,t} (1 - \lambda_{i,t}^{(P)}) \gamma dt \quad \textit{for } \alpha_{i,0} = 1.$$

3.2 Preliminary Equilibrium Characterization

We first prove a simple lemma that characterizes the equilibrium pricing strategy of the platform and the equilibrium purchasing strategy of the user. We then use this result to pin down the platform's equilibrium strategy.

Lemma 3. *The time-t utility of the user from purchasing product i for which no bad news has arrived ($I_{i,t} = 1$) is*

$$\mathbb{E}\left[\theta_i \mid \mathrm{NBN}_{i,t}\right] - p_{i,t} = \mu_{i,t} - p_{i,t}.$$

Thus, the time-t equilibrium price of product i for which no bad news has arrived is

$$p_{i,t} = \mathbb{E}[\theta_i \mid \mathrm{NBN}_{i,t}] = \mu_{i,t}.$$

For a product i at time t for which bad news has arrived ($I_{i,t} = 0$), the equilibrium price is zero.

Intuitively, the continuation game gives zero utility to the user according to her own expectations, because of the platform's take-it-or-leave-it offers. Put differently, the platform can extract the full surplus by making the user indifferent between purchasing and not purchasing a product. But importantly, this is surplus as estimated by the user's expectation, which may differ from the platform's information (and the platform can accurately estimate the user's expectation, because it has access to the same information as the user as well).

We next provide expressions for platform and user utilities in the pre-AI and post-AI environments.

Pre-AI environment: In the pre-AI environment, the platform's equilibrium strategy is a solution to a stochastic dynamic optimization problem whose state is given by beliefs about the product's quality. Since both the platform and the user share the same information in this setting, we let $\{\mu_i\}_{i \in \mathcal{N}}$ denote their initial belief. With this notation, starting from any initial belief $\{\mu_i\}_{i \in \mathcal{N}}$, the platform's offering strategy in equilibrium denoted by $\{x_{i,t}^{\mathrm{pre-AI}}\}_{i \in \mathcal{N}, t \in \mathbb{R}_+}$ solves

$$\Pi^{\mathrm{pre-AI}}(\{\mu_i\}_{i=1}^n) = \max_{\{x_{i,t}\}_{i \in \mathcal{N}, t \in \mathbb{R}_+}} \mathbb{E}_0 \left[\int_0^\infty r e^{-rt} \sum_{i=1}^n x_{i,t} \mu_{i,t} dt \right] \tag{2}$$
$$d\mu_{i,t} = I_{i,t} x_{i,t} \mu_{i,t} (1 - \mu_{i,t})(1 - \lambda_{i,t}) \gamma dt \text{ with } \mu_{i,0} = \mu_i$$
$$d\lambda_{i,t} = I_{i,t} x_{i,t} \lambda_{i,t} \left((1 - \lambda_{i,t})\gamma - \rho\right) dt \quad \text{with } \lambda_{i,0} = \lambda$$
$$\mu_{i,t} = 0 \quad \text{if } I_{i,t} = 0$$
$$\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1] = x_{i,t}\left(1 - \mu_{i,t}\right)\left(1 - \lambda_{i,t}\right)\gamma dt \text{ if } I_{i,0} = 1.$$

The platform's equilibrium pricing strategy in this problem is given by $p_{i,t}^{\text{pre}-\text{AI}} = \mu_{i,t}$ and the user's purchasing decision $b_{i,t}$ does not appear in this expression, because at this price she always purchases the offered product (Lemma 3). Notice that in this pre-AI environment the platform and the user have the same information set and hence the expectation $\mathbb{E}_0$ is with respect to this information set. In particular, this expectation takes into account that $\mu_i \sim \text{Unif}[0,1]$, $\theta_i \sim \text{Bern}(\mu_i)$, and α_i drawn according to (1) for all $i \in \mathcal{N}$.

The expectation of the user's utility can then be computed as

$$U^{\text{pre}-\text{AI}}(\{\mu_i\}_{i=1}^n) = \mathbb{E}_0\left[\int_0^\infty re^{-rt}\sum_{i=1}^n x_{i,t}^{\text{pre}-\text{AI}}\left(\theta_i - \mu_{i,t}\right)dt\right]. \tag{3}$$

Notice also that the user's utility in the pre-AI, when the platform has no information advantage, does not depend on products' glossiness, $\{\alpha_{i,0}\}_{i\in\mathcal{N}}$, because the platform does not observe these quantities and naturally its equilibrium strategy does not depend on them.

The expected utilitarian welfare is then the sum of the user's utility given in (3) and the platform's payoff given in (2), which is equal to

$$W^{\text{pre}-\text{AI}}(\{\mu_i\}_{i=1}^n) = \mathbb{E}_0\left[\int_0^\infty re^{-rt}\sum_{i=1}^n x_{i,t}^{\text{pre}-\text{AI}}\theta_i dt\right].$$

Post-AI environment: In the post-AI environment, again, the platform's equilibrium strategy is a solution to a stochastic dynamic optimization problem. In contrast with the pre-AI setting, here, the platform knows the initial glossiness state of the product and therefore has a different belief from the user. Starting from any user's belief $\{\mu_i\}_{i\in\mathcal{N}}$ and the initial glossiness state $\{\alpha_i\}_{i\in\mathcal{N}}$, the platform's equilibrium offering strategy $\{x_{i,t}^{\text{post}-\text{AI}}\}_{i\in\mathcal{N},t\in\mathbb{R}_+}$ solves

$$\begin{aligned}
\Pi^{\text{post}-\text{AI}}(\{\mu_i,\alpha_i\}_{i=1}^n) = \max_{\{x_{i,t}\}_{i\in\mathcal{N},t\in\mathbb{R}_+}} & \mathbb{E}_0^P\left[\int_0^\infty re^{-rt}\sum_{i=1}^n x_{i,t}\mu_{i,t}dt\right] \qquad (4)\\
& d\mu_{i,t} = I_{i,t}x_{i,t}\mu_{i,t}(1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt \text{ with } \mu_{i,0} = \mu_i\\
& d\lambda_{i,t} = I_{i,t}x_{i,t}\lambda_{i,t}\left((1-\lambda_{i,t})\gamma - \rho\right)dt \text{ with } \lambda_{i,0} = \lambda\\
& d\mu_{i,t}^{(P)} = I_{i,t}x_{i,t}\mu_{i,t}^{(P)}(1-\mu_{i,t}^{(P)})\gamma dt \quad \text{and} \quad \lambda_{i,t}^{(P)} = 0 \text{ if } \alpha_i = 0\\
& \mu_{i,t}^{(P)} = 0 \quad \text{and} \quad d\lambda_{i,t}^{(P)} = I_{i,t}x_{i,t}\lambda_{i,t}^{(P)}\left((1-\lambda_{i,t}^{(P)})\gamma - \rho\right)dt \text{ if } \alpha_i = 1\\
& \mu_{i,t} = \mu_{i,t}^{(P)} = 0 \quad \text{if } I_{i,t} = 0\\
& \mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1, \alpha_i = 0] = x_{i,t}b_{i,t}\left(1-\mu_{i,t}^{(P)}\right)\gamma dt\\
& \mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1, \alpha_i = 1] = x_{i,t}b_{i,t}\left(1-\lambda_{i,t}^{(P)}\right)\gamma dt, \text{ if } I_{i,0} = 1.
\end{aligned}$$

The platform's equilibrium pricing strategy is again $p_{i,t}^{\text{post}-\text{AI}} = \mu_{i,t}$ (Lemma 3). However, we are now evaluating the platform's expected profits according to its finer information set, and hence the expectation is denoted by $\mathbb{E}_0^P$.

Note also that if the initial glossiness state is $\alpha_{i,0} = 1$, then the platform knows that

$$\mu_{i,0}^{(P)} = 0 \text{ and } \lambda_{i,0}^{(P)} = 1,$$

and if the initial glossiness state is $\alpha_{i,0} = 0$, then Bayes's rule implies that the platform's beliefs are

$$\mu_{i,0}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0} + (1-\mu_{i,0})(1-\lambda)} \text{ and } \lambda_{i,0}^{(P)} = 0.$$

In what follows, whenever the platform has more information than the user, we will evaluate the user's utility and utilitarian welfare according to this finer information set (or according to the "true events" that incorporate information about the glossiness of products). Hence:

$$U^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^n) = \mathbb{E}_0^P \left[\int_0^\infty re^{-rt} \sum_{i=1}^n x_{i,t}^{\text{post}-\text{AI}} \left(\theta_i - \mu_{i,t}\right) dt \right]. \tag{5}$$

The expected utilitarian welfare is the sum of user utility in (5) and platform profits in (4), and thus equal to

$$W^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^n) = \mathbb{E}_0^P \left[\int_0^\infty re^{-rt} \sum_{i=1}^n x_{i,t}^{\text{post}-\text{AI}} \theta_i dt \right].$$

With these expressions, we can provide a full characterization of equilibrium in the pre-AI and post-AI environments.

4 Equilibrium Characterization in the Pre-AI Environment

We first characterize the evolution of user beliefs in the pre-AI environment.

Lemma 4. *In the pre-AI environment, for any product $i \in \mathcal{N}$ that has been offered for $[0, t)$ with $\mathrm{NBN}_{i,t}$ (no bad news), we have*

$$\lambda_{i,t} = \frac{(\gamma - \rho)\lambda}{\lambda\gamma + e^{(\rho-\gamma)t}((1-\lambda)\gamma - \rho))},$$

and

$$\mu_{i,t} = \frac{\mu_{i,0}(\gamma - \rho)}{\mu_{i,0}(\gamma - \rho) + (1-\mu_{i,0})\left(e^{-\rho t}\lambda\gamma + e^{-\gamma t}\left((1-\lambda)\gamma - \rho\right)\right)}.$$

Moreover, $\mu_{i,t}$ is increasing in $\mu_{i,0}$ and t and converges to 1 as $t \to \infty$. Additionally:

- *If $\rho > \gamma$, then $\lambda_{i,t}$ is monotonically decreasing and limits to zero as $t \to \infty$.*
- *If $\rho \leq \gamma$, then:*
  - *If $(1-\lambda)\gamma - \rho \geq 0$, $\lambda_{i,t}$ is monotonically increasing and converges to $\frac{\gamma-\rho}{\gamma}$ as $t \to \infty$.*
  - *If $(1-\lambda)\gamma - \rho < 0$, $\lambda_{i,t}$ is monotonically decreasing and converges to $\frac{\gamma-\rho}{\gamma}$ as $t \to \infty$.*

Lemma 4 characterizes the evolution of user beliefs about the quality of a product, conditioning on no bad news realization until the time in question. This lemma enables us to compare the belief trajectory of different products based on the user's initial belief and their history of experimentation. Building on this property, we next characterize the platform's equilibrium decision in the pre-AI environment.

Theorem 1. *In the pre-AI environment, the platform's equilibrium decision is to offer the product with the highest initial user belief until bad news occurs and then offer the product with the second highest belief, and so on.*

We establish this theorem by leveraging the Gittins index theorem (see, e.g., [Gittins et al., 2011, Chapter 2]), which asserts that the platform always chooses the product with the highest Gittins index. The theorem then follows by using two key properties of the Gittins index in our context: (i) products with higher beliefs have initially higher Gittins indices, and (ii) in the absence of bad news, the Gittins index for a product increases. As a result, the platform begins by offering the product with the highest belief. As long as there is no bad news for this product, its Gittins index continues to grow, maintaining its status as the product with the maximum Gittins index. Thus, the platform continues to offer this product until bad news occurs, at which point it switches to the product with the second highest belief, and so on.

5 Equilibrium Characterization in the Post-AI Environment

We start with the evolution of user and platform beliefs in the post-AI environment.

Lemma 5. *In the post-AI environment, for any product $i \in \mathcal{N}$ that has been offered for $[0,t)$ with $\mathrm{NBN}_{i,t}$ (no bad news), the dynamics of $\mu_{i,t}$ and $\lambda_{i,t}$ are the same as Lemma 4, and additionally*

- *If $\alpha_{i,0} = 0$, then*

$$\mu_{i,t}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0} + e^{-\gamma t}(1-\lambda)(1-\mu_{i,0})},\ \lambda_{i,t}^{(P)} = 0,\ \mathbb{P}[\mathrm{NBN}_{i,t} \mid \alpha_{i,0} = 0] = \frac{\mu_{i,0} + e^{-\gamma t}(1-\lambda)(1-\mu_{i,0})}{1-\lambda(1-\mu_{i,0})},$$

and

$$\mathbb{P}[I_{i,t+dt} = 0, I_{i,t} = 1 \mid \alpha_{i,0} = 0] = \frac{e^{-\gamma t}(1-\lambda)(1-\mu_{i,0})}{1-\lambda(1-\mu_{i,0})}\gamma dt.$$

- *If $\alpha_{i,0} = 1$, then*

$$\mu_{i,t}^{(P)} = 0,\ \lambda_{i,t}^{(P)} = \frac{\gamma - \rho}{\gamma - e^{(\rho-\gamma)t}\rho},\ \mathbb{P}[\mathrm{NBN}_{i,t} \mid \alpha_{i,0} = 1] = \frac{e^{-\rho t}\gamma - e^{-\gamma t}\rho}{\gamma - \rho},$$

and

$$\mathbb{P}[I_{i,t+dt} = 0, I_{i,t} = 1 \mid \alpha_{i,0} = 1] = \frac{\left(e^{-\rho t} - e^{-\gamma t}\right)\rho}{\gamma - \rho}\gamma dt.$$

Lemma 5 characterizes the evolution of user and platform beliefs, conditioning on no bad news realization in the post-AI environment. This lemma enables us to compare the belief trajectories of different products based on the user's initial beliefs about them and the initial glossiness state $\alpha_{i,0}$. Building on this lemma, we next determine the platform's equilibrium decision in the post-AI environment.

Theorem 2. *In the post-AI environment, the platform's equilibrium decision is to offer either the product with the highest initial user belief among the ones in the initial glossiness state $\alpha = 0$ for which no bad news has arrived, or the product with the highest initial user belief among the ones in the initial glossiness state $\alpha = 1$ for which no bad news has arrived.*

Notice that the conditioning here is on the initial glossiness state, $\alpha_{i,0}$, not on the current glossiness state, since the platform does not observe the latter. This theorem's proof uses the Gittins index theorem and the same properties of Gittins indices as in the proof of Theorem 1, discussed above.

Theorem 2 characterizes the structure of the platform's equilibrium decision, but it does not pin down the comparison between a product with $\alpha_{i,0} = 0$ and another product with $\alpha_{i',0} = 1$. We next provide a characterization of the platform's equilibrium decision between two products with different initial glossiness states.

In the pre-AI environment, the platform always offers the product with the highest belief. In the post-AI environment, this is no longer the case, as the platform may prefer to opt for a glossy product ($\alpha_{i,0} = 1$) with a lower belief, rather than a higher belief non-glossy product ($\alpha_{i,0} = 0$). We next formalize this observation.

Proposition 1 (Manipulation Effect)**.** *Consider two products $i, i' \in \mathcal{N}$ in the post-AI environment with $\alpha_{i',0} = 1$ and $\alpha_{i,0} = 0$. There exists a function $\mu_{i,0} \mapsto f(\mu_{i,0}; \lambda, \gamma, r)$ parameterized by λ, γ, r that is everywhere below $\mu_{i,0}$ and $\underline{\rho}$ such that for $\rho \leq \underline{\rho}$ when*

$$\mu_{i',0} > f(\mu_{i,0}; \lambda, \gamma, r),$$

the platform prefers to offer product i' rather than offering product i.

We call this *manipulation effect* because the platform is offering a low-quality product with a lower user belief. To understand this effect and why it requires ρ to be small, observethat there are two forces that determine the platform's choice: (i) when $\alpha_{i',0} = 1$, the platform knows that the quality is low, and therefore, once $\alpha_{i',t}$ becomes 0, the user will receive bad news at the rate γ. This force incentivizes the platform to not offer product i'; (ii) when $\alpha_{i',0} = 1$, the platform knows that the user's belief increases for a while because bad news will not arrive. This force incentivizes the platform to offer product i'. For sufficiently small ρ, the second force dominates, because the glossy state is very persistent and thus bad news for product i' will not arrive for a long time. In this case, the platform prefers to offer product i' rather than product i.

Proposition 1 thus establishes that for small enough ρ, the platform prefers to offer a product with lower user belief but $\alpha_{i,0} = 1$ rather than a product with a higher user belief but $\alpha_{i,0} = 0$. We next show that the opposite phenomenon arises when ρ is sufficiently large. In this case, the platform actually stays away from products with $\alpha_{i,0} = 1$, because it recognizes that these are low-quality products, and the user will quickly learn this as their glossiness wears out. We refer to this as the *helpfulness effect* to contrast it with the manipulation effect.

Proposition 2 (Helpfulness Effect). *Consider two products $i, i' \in \mathcal{N}$ in the post-AI environment with $\alpha_{i',0} = 1$ and $\alpha_{i,0} = 0$. There exists a function $\mu_{i',0} \mapsto g(\mu_{i',0}; \lambda, \gamma, r)$ parametrized by λ, γ, r that is everywhere below $\mu_{i',0}$ and $\bar{\rho}$ such that for $\rho \geq \bar{\rho}$ when*

$$\mu_{i,0} > g(\mu_{i',0}; \lambda, \gamma, r),$$

the platform prefers to offer product i relative to offering product i'.

In this configuration, the platform prefers not to offer a low-quality with a moderately higher user belief than an alternative product with unknown quality with a lower user belief. This is helpful to the user as it makes her avoid the low-quality product. The platform recognizes the low-quality product simply from the fact that it has an initial glossiness state $\alpha_{i,0} = 1$, and deploys this information for the user's benefit.

To see the intuition of why the platform now moves away from the low-quality product, recall that when ρ is large, the platform understands that the glossiness of the product will not last, and bad news will start arriving for product i' (which has initial glossiness state $\alpha_{i',0} = 1$). Therefore, even if $\mu_{i',0} > \mu_{i,0}$, so long as the belief $\mu_{i,0}$ is not too small (ensured by the condition $\mu_{i,0} > g(\mu_{i',0}; \lambda, \gamma, r)$), the platform prefers offering product i instead of product i'. We conclude this section by noting that both functions $f(\cdot)$ and $g(\cdot)$ characterized in Propositions 1 and 2, can be written explicitly in terms of the primitives $(\mu, \lambda, \gamma, r)$.

6 Who Benefits from the Platform's Superior Information?

In this section, we discuss how the information advantage of the platform in the post-AI environment can have either positive or negative effects on the user, while the platform itself always benefits from this informational advantage. We then investigate the implications of expanding the set of products available to the platform. Throughout, as noted above, we evaluate all expectations using the finer information set of the platform, but to simplify the notation we denote these expectations simply by $\mathbb{E}$.

6.1 When Users Benefit from the Platform's Superior Information

Here, we show that the platform's information advantage in the post-AI environment can benefit the user. In particular, our next theorem establishes that for sufficiently large ρ, the information advantage of the platform always benefits the user and the platform.

Theorem 3. *Suppose the initial beliefs $\{\mu_i\}_{i=1}^n$ are i.i.d. and uniform over $[0, 1]$. For any r, γ, λ, there exists ρ_h such that for $\rho \geq \rho_h$ we have*

$$\begin{aligned}
\mathbb{E}\left[\Pi^{\text{post-AI}}(\{\mu_i, \alpha_i\}_{i=1}^n)\right] &> \mathbb{E}\left[\Pi^{\text{pre-AI}}(\{\mu_i\}_{i=1}^n)\right] \\
\mathbb{E}\left[U^{\text{post-AI}}(\{\mu_i, \alpha_i\}_{i=1}^n)\right] &> \mathbb{E}\left[U^{\text{pre-AI}}(\{\mu_i\}_{i=1}^n)\right] \\
\mathbb{E}\left[W^{\text{post-AI}}(\{\mu_i, \alpha_i\}_{i=1}^n)\right] &> \mathbb{E}\left[W^{\text{pre-AI}}(\{\mu_i\}_{i=1}^n)\right].
\end{aligned}$$

The informational advantage of the platform always increases its own profits, as it enables the platform to modify the user's behavior. Theorem 3 establishes that this informational advantage also increases the expected user's utility and welfare for large enough ρ. This theorem confirms what might be viewed as the conventional wisdom in the literature: more data enables better allocation of products and therefore benefits the users and society as a whole. This theorem follows from the fact that for large enough ρ, the helpfulness effect, characterized in Proposition 2, dominates the manipulation effect of Proposition 1.

In the next subsection, we will see that this intuition does not hold in general and the platform's information advantage may harm users.

6.2 When Behavioral Manipulation Harms Users

Here, we show that in the post-AI environment, the user's utility decreases because of behavioral manipulation — since the manipulation effect dominates the helpfulness effect.

Theorem 4. *Suppose the initial beliefs $\{\mu_i\}_{i=1}^n$ are i.i.d. and uniform over $[0, 1]$. For any r, γ, λ, there exists ρ_l such that for $\rho \leq \rho_l$ we have*

$$\mathbb{E}\left[\Pi^{\mathrm{post-AI}}(\{\mu_i, \alpha_i\}_{i=1}^n)\right] > \mathbb{E}\left[\Pi^{\mathrm{pre-AI}}(\{\mu_i\}_{i=1}^n)\right]$$
$$\mathbb{E}\left[U^{\mathrm{post-AI}}(\{\mu_i, \alpha_i\}_{i=1}^n)\right] < \mathbb{E}\left[U^{\mathrm{pre-AI}}(\{\mu_i\}_{i=1}^n)\right]$$
$$\mathbb{E}\left[W^{\mathrm{post-AI}}(\{\mu_i, \alpha_i\}_{i=1}^n)\right] < \mathbb{E}\left[W^{\mathrm{pre-AI}}(\{\mu_i\}_{i=1}^n)\right].$$

In this case with low ρ, the platform's informational advantage enables it to engage in behavioral manipulation: the user is pushed towards products with $\alpha = 1$ (or more accurately, towards products that have initial glossiness state $\alpha_{i,0} = 1$). Because these products do not generate bad news in the short run, the user's belief will become more positive for a while, and this will enable the platform to charge higher prices and increase its profits. However, because glossy products are low quality, this behavioral manipulation is bad for user utility and utilitarian welfare. That ρ is small here is important. As we saw, when ρ is large, the platform expects the glossiness of the production to wear off quickly, and thus it is not worthwhile to push the user towards glossy products. But from a welfare point of view, it is more costly to have users consume glossy products when ρ is small, because they will not discover for quite a while that the product is actually not high-quality. It is this feature of behavioral manipulation that reduces user utility and utilitarian welfare.

6.3 Big Data Double Whammy: More Products Negatively Impact User Welfare

The availability of big data provides platforms with valuable insights into predictable patterns of user behavior, which can be leveraged for behavioral manipulation, as we have established thus far. Moreover, the same advances in AI also enable digital platforms to expand the range of products and services they offer. Next, we demonstrate that this combination of greater choice and more platform information may be particularly pernicious — as the number of products increases, the potential for behavioral manipulation increase as well. This result highlights that multiple aspects of the new capabilities of digital platforms closely interact in affecting user welfare.

Theorem 5. *Suppose the initial beliefs* $\{\mu_i\}_{i=1}^{n+1}$ *are i.i.d. and uniform over* $[\frac{1}{2}-\Delta, \frac{1}{2}+\Delta]$*. For any* r, γ, λ*, there exist* $\bar{\Delta}$ *and* $\tilde{\rho}_l$ *such that for* $\Delta \leq \bar{\Delta}$*,* $\rho \leq \tilde{\rho}_l$*, and* $n \geq \lceil 1 - \frac{\log \lambda}{\log(2-\lambda)/(1-\lambda)} \rceil$ *in the post-AI environment we have:*

$$\mathbb{E}\left[\Pi^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^{n+1})\right] > \mathbb{E}\left[\Pi^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^{n})\right]$$
$$\mathbb{E}\left[U^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^{n+1})\right] < \mathbb{E}\left[U^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^{n})\right]$$
$$\mathbb{E}\left[W^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^{n+1})\right] < \mathbb{E}\left[W^{\text{post}-\text{AI}}(\{\mu_i, \alpha_i\}_{i=1}^{n})\right].$$

This result thus shows that as the number of products increases in the post-AI environment and glossiness remains persistent (ρ is small), platform profits increase but user welfare and utilitarian welfare decline. Hence, greater information acquisition by platforms is particularly costly for users when platforms are able to host more products and services.

Before providing an intuition for this result, we first explain the role of the various assumptions we impose in this theorem. The assumption that $\rho \leq \tilde{\rho}_l$ is adopted for the same reasons as before — this is the range where the platform's superior information can be costly to the user. The assumptions $\Delta \leq \bar{\Delta}$ and large n are added in this proposition to ensure that the initial beliefs of the user are not very informative about product quality and that the platform has enough products to engage in behavioral manipulation. This combination then leads to a configuration where the presence of more products creates greater opportunities for the platform to select initially glossy items (with initial beliefs that are sufficiently close to the beliefs of non-glossy products). This enables, and makes profitable, behavioral manipulation. An alternative way of expressing this intuition is that, while in standard choice theory, greater choices are good for the consumer, in the presence of behavioral manipulation, they may be bad precisely because they enable the platform to engage in more intense behavioral manipulation.

Because new AI and big data technologies are simultaneously intensifying data collection by platforms and expanding their capabilities to host a larger number of products, this result suggests that they may be creating a costly "double whammy" for consumers.

7 Conclusion

With breakneck-paced advances in large language models (LLMs) and other new applications, the capabilities and reach of AI algorithms are expanding. LLMs will likely automate various tasks and play a more central role in making recommendations and helping users retrieve access and process information. While the debate on the consequences of these new algorithms is intensifying, there has been less attention on how the growing role of these systems in recommendations will impact consumer choice, experimentation and learning. The default position among economists and AI researchers is that new AI tools will ultimately benefit humans by providing them with better information and enabling more personalized services.

This paper has explored the possibility that the vast amounts of data collected by online platforms may also enable behavioral manipulation, where these platforms modify user behavior in a way that benefits platforms but not necessarily the users. While this possibility has been discussed informally in the legal and sociology literatures, there has been little work on the theoretical foundations of such behavioral manipulation in economics. Our model focuses on the possibility that the superior information of a platform may enable it to push users towards products that will temporarily appear to the user more attractive than they truly are and in the process enable the platform to charge higher prices (and also slow down learning, which is hampered by glossiness).

To formalize these ideas, we introduce an economy in which a platform offers one of n products with an associated price to a user, who is uncertain about the quality of the goods but can slowly learn about them. The new element of our model is that signals depend not just on the quality of a product but also on whether it is "glossy", meaning that it has appearance or other features (such as packaging, good advertisement, hidden costs, or exaggerated benefits) that make it seem more attractive in the short run than it truly is. We conceptualize online data collection as enabling the platform to estimate which types of products are glossy for a user by collecting data from other users with similar characteristics.

While before the arrival of the AI-based informational advantage of the platform, users are assumed to be completely Bayesian and understand the true data generating process given the platform's strategies, they become "behavioral" after the advent of new advances in AI and big data, because they do not fully understand the informational advantage that these confer on the platform.

Our analysis characterizes the learning dynamics of the platform and the user in this environment. From a modeling point of view, the main innovation of our approach is to develop an experimentation model in which the two agents involved in this decision — the platform and the user — have different beliefs. Specifically, because the platform has a finer information partition than the user, the user's beliefs do not follow a martingale process according to the platform's information set. This fundamentally modifies the dynamic programming problem of the platform, and we develop a different line of analysis to determine how these beliefs evolve and what they imply for the equilibrium strategies of the platform and the users.

Our analysis establishes three main results. First, when the effects of glossiness are short-lived (for example, such glossiness wears off rapidly), the superior information of the platform is beneficial to users. The platform can identify glossy products and can lead users away from them. This is profitable for the platform because the profits it can make from this glossiness would be similarly short-lived.

Second, and more importantly for our main focus, when glossiness is not so ephemeral, the platform systematically distorts user choices towards glossy products. Intuitively, the platform knows that users that start consuming glossy products will, at first, improve their assessment of the quality of these products (by definition of their glossiness), and this will create room for the platform to charge higher prices. This distortion is costly for users precisely because it pushes them towards glossy, low-quality products and this choice does not get reversed rapidly. As a result, we show that, in this case, advances in AI increase platform profits but reduce user utility and utilitarian welfare in this model.

Third, as the number of products increases the platform has more opportunities to engage in behavioral manipulation, further reducing user utility and overall welfare.

We view our paper as a first step in the analysis of product choice and consumer experimentation in online markets in which platforms have extensive and growing information about user preferences (and biases). There are several interesting directions for future work. First, in addition to the experimentation distortions we have focused on, there may be simpler forms of distortions, in which users make static mistakes in online platforms, and whether AI can be a tool for avoiding these mistakes or will exacerbate them and enable their more systematic exploitation is a question for future theoretical and empirical work. Second, we simplified the analysis by assuming that the platform can perfectly estimate the initial glossiness of the product. A more challenging but interesting direction for theoretical work is to model the dynamic evolution of the platform's superior information together with the evolution of the beliefs of the user. For example, the platform can dynamically estimate a common component from the behavior of several users with shared characteristics. The new interesting questions in such an environment relate to whether the platform would introduce additional distortions in order to raise the speed of its own learning while slowing down that of the users. Third, in practice, there are users with different levels of sophistication. An interesting direction is to analyze a game in which the platform simultaneously interacts with both sophisticated and naïve users, and may or may not know which users are which. Finally, there is much need for new empirical research to document and quantify the presence of the type of behavioral manipulation we have highlighted in this paper.

A1 Proofs

This Appendix includes the omitted proofs from the text.

Proof of Lemma 1

If $x_{i,t}b_{i,t} = 0$, the user belief about θ_i does not change. Next, we consider $x_{i,t}b_{i,t} = 1$ and, for notational convenience, let us suppress subscript i. Then, by using Bayes' rule, the probability of $\theta_i = 1$ is

$$\begin{aligned}
\mu_{t+dt} &= \frac{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 1]\mathbb{P}[\theta = 1]}{\mathbb{P}[\mathrm{NBN}_{t+dt}]} \\
&= \frac{\mathbb{P}[\mathrm{NBN}_t \mid \theta = 1]\mathbb{P}[\mathrm{NBN}_{t,t+dt} \mid \mathrm{NBN}_t, \theta = 1]\mathbb{P}[\theta = 1]}{\mathbb{P}[\mathrm{NBN}_t]\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t]} \\
&= \mu_t \frac{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t, \theta = 1]}{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 1, \mathrm{NBN}_t]\mathbb{P}[\theta = +1 \mid \mathrm{NBN}_t] + \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t]\mathbb{P}[\theta = 0 \mid \mathrm{NBN}_t]} \\
&\overset{(a)}{=} \frac{\mu_t}{\mu_t + (1-\mu_t)\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t]} \\
&= \frac{\mu_t}{\mu_t + (1-\mu_t)\left(\lambda_t + (1-\lambda_t)(1-\gamma dt)\right)} \\
&= \frac{\mu_t}{1 - (1-\mu_t)(1-\lambda_t)\gamma dt} \\
&\overset{(b)}{=} \mu_t \left(1 + (1-\mu_t)(1-\lambda_t)\gamma dt\right)
\end{aligned} \tag{A1}$$

where (a) follows from

$$\begin{aligned}
\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t] =& \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_t = 1]\mathbb{P}[\alpha_t = 1 \mid \theta = 0, \mathrm{NBN}_t] \\
&+ \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_t = 0]\mathbb{P}[\alpha_t = 0 \mid \theta = 0, \mathrm{NBN}_t]
\end{aligned}$$

and (b) follows by using Taylor expansion of $1/(1-x)$ around 0 and dropping the terms of the order $(dt)^2$. From (A1), we then have

$$d\mu_t = \mu_{t+dt} - \mu_t = \mu_t(1-\mu_t)(1-\lambda_t)\gamma dt.$$

Again, by using Bayes' rule,

$$\begin{aligned}
\lambda_{t+dt} &= \mathbb{P}[\alpha_{t+dt} = 1 \mid \theta = 0, \mathrm{NBN}_{t+dt}] \\
&= \frac{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \alpha_{t+dt} = 1]\mathbb{P}[\alpha_{t+dt} = 1 \mid \theta = 0]}{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0]} \\
&= \mathbb{P}[\alpha_t = 1 \mid \theta = 0]\mathbb{P}[\alpha_{t+dt} = 1 \mid \alpha_t = 1, \theta = 0] \\
&\times \frac{\mathbb{P}[\mathrm{NBN}_t \mid \theta = 0, \alpha_{t+dt} = 1]\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_{t+dt} = 1]}{\mathbb{P}[\mathrm{NBN}_t \mid \theta = 0]\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t, \theta = 0]} \\
&= \frac{\lambda_t \mathbb{P}[\alpha_{t+dt} = 1 \mid \alpha_t = 1, \theta = 0]\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_{t+dt} = 1]}{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t]}
\end{aligned}$$

$$
\begin{aligned}
&\overset{(a)}{=} \frac{\lambda_t(1-\rho dt)}{\lambda_t + (1-\lambda_t)(1-\gamma dt)} \\
&= \frac{\lambda_t(1-\rho dt)}{1-(1-\lambda_t)\gamma dt} \\
&\overset{(b)}{=} \lambda_t - \lambda_t \rho dt + \lambda_t(1-\lambda_t)\gamma dt
\end{aligned} \tag{A2}
$$

where (a) follows from

$$
\begin{aligned}
\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t] = &\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \alpha_t = 1, \theta = 0, \mathrm{NBN}_t]\mathbb{P}[\alpha_t = +1 \mid \theta = 0, \mathrm{NBN}_t] \\
&+ \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \alpha_t = 0, \theta = 0, \mathrm{NBN}_t]\mathbb{P}[\alpha_t = 0 \mid \theta =, \mathrm{NBN}_t]
\end{aligned}
$$

and (b) follows by dropping the terms of the order $(dt)^2$. From (A2), we now have

$$
d\lambda_t = \lambda_t((1-\lambda_t)\gamma - \rho)dt.
$$

Finally, given $x_{i,t}b_{i,t} = 1$, we have

$$
\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1] = (1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt.
$$

This completes the proof. ■

Proof of Lemma 2

Similar to the proof of Lemma 1, if $x_{i,t}b_{i,t} = 0$, the user belief about θ_i does not change, and when $x_{i,t}b_{i,t} = 1$, we let

$$
\mu_{i,t}^{(P)} = \mathbb{P}\left[\theta = +1 \mid \alpha_{i,0} = 0, \mathrm{NBN}_t\right]
$$

be the probability of a product being high quality if the initial glossiness state is zero and no bad news has arrived by time t, and

$$
\lambda_{i,t}^{(P)} = \mathbb{P}\left[\alpha_{i,t} = +1 \mid \alpha_{i,0} = 1, \mathrm{NBN}_t\right]
$$

be the probability of the glossiness state being 1 if the initial glossiness state is 1 and no bad news has arrived by time t. Again, to make the notation easier, in this proof, we drop the subscript i. Notice that conditioning on $\alpha_0 = 1$, the platform knows the product is of low quality and therefore $\mu_t^{(P)} = 0$ for all t. Also, conditioning on $\alpha_0 = 0$, the glossiness state remains at zero and therefore $\lambda_t^{(P)} = 0$. We next evaluate $\mu_t^{(P)}$ conditioning on $\alpha_0 = 0$ and $\lambda_t^{(P)}$ conditioning on $\alpha_0 = 1$.

By using Bayes' rule, for the probability of $\theta_i = 1$ when $\alpha_0 = 0$ we have

$$
\begin{aligned}
\mu_{t+dt}^{(P)} &= \frac{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 1, \alpha_0 = 0]\mathbb{P}[\theta = 1 \mid \alpha_0 = 0]}{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \alpha_0 = 0]} \\
&= \frac{\mathbb{P}[\mathrm{NBN}_t \mid \theta = 1, \alpha_0 = 0]\mathbb{P}[\mathrm{NBN}_{t,t+dt} \mid \mathrm{NBN}_t, \theta = 1, \alpha_0 = 0]\mathbb{P}[\theta = 1 \mid \alpha_0 = 0]}{\mathbb{P}[\mathrm{NBN}_t \mid \alpha_0 = 0]\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t, \alpha_0 = 0]}
\end{aligned}
$$

$$
\begin{aligned}
&= \frac{\mu_t^{(P)}\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t, \theta = 1, \alpha_0 = 0]}{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t, \alpha_0 = 0]} \\
&\overset{(a)}{=} \frac{\mu_t^{(P)}}{\mu_t^{(P)} + (1-\mu_t^{(P)})(1-\gamma dt)} \\
&= \frac{\mu_t^{(P)}}{1-(1-\mu_t^{(P)})\gamma dt} \\
&\overset{(b)}{=} \mu_t^{(P)}\left(1+(1-\mu_t^{(P)})\gamma dt\right)
\end{aligned} \tag{A3}
$$

where (a) follows from

$$
\begin{aligned}
\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t, \alpha_0 = 0] =& \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 1, \mathrm{NBN}_t, \alpha_0 = 0]\mathbb{P}[\theta = 1 \mid \mathrm{NBN}_t, \alpha_0 = 0] \\
&+ \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_0 = 0]\mathbb{P}[\theta = 0 \mid \mathrm{NBN}_t, \alpha_0 = 0] \\
&= \mu_t^{(P)} + \left(1-\mu_t^{(P)}\right)(1-\gamma dt)
\end{aligned}
$$

and (b) follows by dropping the terms of the order $(dt)^2$. By using (A3),

$$
d\mu_t^{(P)} = \mu_{t+dt}^{(P)} - \mu_t^{(P)} = \mu_t^{(P)}(1-\mu_t^{(P)})\gamma dt.
$$

By using Bayes' rule, when $\alpha_0 = 1$ we obtain

$$
\begin{aligned}
\lambda_{t+dt}^{(P)} &= \mathbb{P}[\alpha_{t+dt} = 1 \mid \theta = 0,\ \mathrm{NBN}_{t+dt}, \alpha_0 = 1] \\
&= \frac{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \alpha_{t+dt} = 1, \alpha_0 = 1]\mathbb{P}[\alpha_{t+dt} = 1 \mid \theta = 0, \alpha_0 = 1]}{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \alpha_0 = 1]} \\
&\overset{(a)}{=} \frac{\mathbb{P}[\alpha_t = 1 \mid \theta = 0, \alpha_0 = 1]\mathbb{P}[\alpha_{t+dt} = 1 \mid \alpha_t = 1, \theta = 0, \alpha_0 = 1]\mathbb{P}[\mathrm{NBN}_t \mid \theta = 0, \alpha_{t+dt} = 1, \alpha_0 = 1]}{\mathbb{P}[\mathrm{NBN}_t \mid \theta = 0, \alpha_0 = 1]\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \mathrm{NBN}_t, \theta = 0, \alpha_0 = 1]} \\
&= \frac{\lambda_t^{(P)}\mathbb{P}[\alpha_{t+dt} = 1 \mid \alpha_t = 1, \theta = 0, \alpha_0 = 1]}{\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_0 = 1]} \\
&\overset{(b)}{=} \frac{\lambda_t^{(P)}(1-\rho dt)}{\lambda_t^{(P)} + (1-\lambda_t^{(P)})(1-\gamma dt)} \\
&= \frac{\lambda_t^{(P)}(1-\rho dt)}{1-(1-\lambda_t^{(P)})\gamma dt} \\
&\overset{(c)}{=} \lambda_t^{(P)} - \lambda_t^{(P)}\rho dt + \lambda_t^{(P)}(1-\lambda_t)\gamma dt,
\end{aligned} \tag{A4}
$$

where (a) follows from $\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_{t+dt} = 1, \alpha_0 = 1] = 1$, (b) follows from

$$
\begin{aligned}
&\mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_0 = 1] \\
&= \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_0 = 1, \alpha_t = 1]\mathbb{P}[\alpha_t = 1 \mid \theta = 0, \mathrm{NBN}_t, \alpha_0 = 1] \\
&+ \mathbb{P}[\mathrm{NBN}_{t+dt} \mid \theta = 0, \mathrm{NBN}_t, \alpha_0 = 1, \alpha_t = 0]\mathbb{P}[\alpha_t = 0 \mid \theta = 0, \mathrm{NBN}_t, \alpha_0 = 1],
\end{aligned}
$$

and (c) follows by dropping the terms of the order $(dt)^2$. By using (A4), we obtain

$$d\lambda_t^{(P)} = \lambda_t^{(P)}((1-\lambda_t^{(P)})\gamma - \rho)dt.$$

Finally, we have

$$\begin{aligned}\mathbb{P}[\text{ bad news at } t \mid \alpha_0 = 0, \mathrm{NBN}_t] =& \mathbb{P}[\text{ bad news at } t \mid \alpha_0 = 0, \theta = 0, \mathrm{NBN}_t]\mathbb{P}[\theta = 0 \mid \alpha_0 = 0, \mathrm{NBN}_t] \\ &+ \mathbb{P}[\text{ bad news at } t \mid \alpha_0 = 0, \theta = 1, \mathrm{NBN}_t]\mathbb{P}[\theta = 1 \mid \alpha_0 = 0, \mathrm{NBN}_t] \\ =& \gamma dt \left(1 - \mu_t^{(P)}\right) + 0\end{aligned}$$

and

$$\begin{aligned}&\mathbb{P}[\text{ bad news at } t \mid \alpha_0 = 1, \mathrm{NBN}_t] = \mathbb{P}[\text{ bad news at } t \mid \alpha_0 = 1, \mathrm{NBN}_t, \theta = 0] \\ &= \mathbb{P}[\text{ bad news at } t \mid \alpha_0 = 1, \theta = 0, \mathrm{NBN}_t, \alpha_t = 1]\mathbb{P}[\alpha_t = 1 \mid \alpha_0 = 1, \theta = 0, \mathrm{NBN}_t] \\ &+ \mathbb{P}[\text{ bad news at } t \mid \alpha_0 = 1, \theta = 0, \mathrm{NBN}_t, \alpha_t = 0]\mathbb{P}[\alpha_t = 0 \mid \alpha_0 = 1, \theta = 0, \mathrm{NBN}_t] \\ &= 0 + \gamma dt \left(1 - \lambda_t^{(P)}\right).\end{aligned}$$

The initializations follow by using Bayes' rule, completing the proof. ■

Proof of Lemma 3

For any $\{x_{i,t}\}_{i \in \mathcal{N}, t \in \mathbb{R}_+}$ the equilibrium pricing strategy of the platform is to offer price $p_{i,t} = \mu_{i,t}$ when $x_{i,t} = 1$ and the equilibrium purchasing strategy of the user is to purchase that product. This can be seen by using the one-stage deviation principal: if the platform deviates and offers a lower price, then the continuation game does not change and the platform would make strictly lower profits from this lower current price. If the platform offers a higher price, then if the user purchases, the continuation game would be the same and her continuation payoff (from her perspective) remains zero. But, her current and discounted utility would then be negative, and thus not purchasing is a better response, leading to lower payoff for the platform. This establishes that $p_{i,t} = \mu_{i,t}$ with $x_{i,t} = 1$ is optimal for the platform. ■

Proof of Lemma 4

The differential equation of the evolution of $\lambda_{i,t}$, as we established in Lemma 1, is given by

$$\frac{d}{dt}\lambda_{i,t} = \lambda_{i,t}\left((1-\lambda_{i,t})\gamma - \rho\right)$$

whose solution is

$$\frac{\gamma - \rho}{\gamma + ce^{(\rho-\gamma)t}}$$

for some constant c. Using the initial condition $\lambda_{i,0} = \lambda$, we then have

$$\lambda_{i,t} = \frac{(\gamma - \rho)\lambda}{\lambda\gamma + e^{(\rho-\gamma)t}((1-\lambda)\gamma - \rho))}.$$

The differential equation of the evolution of $\mu_{i,t}$, as in Lemma 1, is given by

$$\frac{d}{dt}\mu_{i,t} = \mu_{i,t}\left(1 - \mu_{i,t}\right)\left(1 - \lambda_{i,t}\right)\gamma$$

whose solution is

$$\mu_{i,t} = \frac{e^{(\gamma+\rho)t}}{e^{(\gamma+\rho)t} + ce^{\gamma t}\lambda\gamma + ce^{\rho t}((1-\lambda)\gamma - \rho)}$$

for some constant c. Using the initial condition $\mu_{i,0}$, we obtain

$$\mu_{i,t} = \frac{e^{(\gamma+\rho)t}\mu_{i,0}(\gamma - \rho)}{e^{(\gamma+\rho)t}\mu_{i,0}(\gamma-\rho) + (1-\mu_{i,0})\left(e^{\gamma t}\lambda\gamma + e^{\rho t}\left((1-\lambda)\gamma - \rho\right)\right)}.$$

The rest of the proof follows by evaluating the monotonicity of $\lambda_{i,t}$ and $\mu_{i,t}$. In particular, to see that $\mu_{i,t}$ is increasing in the initial belief $\mu_{i,0}$, let us take the derivative of $\mu_{i,t}$ with respect to $\mu_{i,0}$ which is

$$\frac{(\gamma-\rho)\left(e^{-\rho t}\lambda\gamma + e^{-\gamma t}\left((1-\lambda)\gamma - \rho\right)\right)}{\left(\mu_{i,0}(\gamma-\rho) + (1-\mu_{i,0})\left(e^{-\rho t}\lambda\gamma + e^{-\gamma t}\left((1-\lambda)\gamma-\rho\right)\right)\right)^2}.$$

This expression is nonnegative. To see this, notice that the denominator is always positive, and as we show next, the numerator is also nonnegative. Let us consider the following (exhaustive) three possibilities:

1. $\gamma(1-\lambda) - \rho \geq 0$. In this case, we have $\gamma - \rho \geq 0$ and therefore both terms in the numerator are nonnegative.

2. $\gamma(1-\lambda) - \rho \leq 0$ and $\gamma - \rho \geq 0$. In this case, $e^{-\rho t} \geq e^{-\gamma t}$ and we have

$$\begin{aligned}(\gamma-\rho)\left(e^{-\rho t}\lambda\gamma + e^{-\gamma t}\left((1-\lambda)\gamma-\rho\right)\right) \geq& (\gamma-\rho)\left(e^{-\gamma t}\lambda\gamma + e^{-\gamma t}\left((1-\lambda)\gamma-\rho\right)\right)\\ =&(\gamma-\rho)e^{-\gamma t}\left(\lambda\gamma + ((1-\lambda)\gamma-\rho)\right)\\ =&(\gamma-\rho)^2 e^{-\gamma t} \geq 0.\end{aligned}$$

3. $\gamma - \rho \leq 0$. In this case, $(\gamma-\rho)\left((1-\lambda)\gamma-\rho\right) \geq 0$ and $e^{-\gamma t} \geq e^{-\rho t}$ and we have

$$\begin{aligned}(\gamma-\rho)\left(e^{-\rho t}\lambda\gamma + e^{-\gamma t}\left((1-\lambda)\gamma-\rho\right)\right) \geq& (\gamma-\rho)e^{-\rho t}\lambda\gamma + (\gamma-\rho)\left((1-\lambda)\gamma-\rho\right)e^{-\rho t}\\ =&(\gamma-\rho)e^{-\rho t}\left(\lambda\gamma + ((1-\lambda)\gamma-\rho)\right)\\ =&(\gamma-\rho)^2 e^{-\rho t} \geq 0.\end{aligned}$$

This completes the proof. ■

Proof of Theorem 1

In the pre-AI environment, the state of each product is a pair (μ_i, λ_i), where μ_i is the user and platform belief about the product's quality and λ_i is the probability of having a glossiness $\alpha = 1$, given no bad news has occurred. For any product $i \in \mathcal{N}$, we introduce the Gittins index:

$$
\begin{aligned}
M(\mu_i, \lambda_i) = \sup_{\tau \geq 0} & \frac{\mathbb{E}_0\left[\int_0^\tau e^{-rt}\mu_{i,t}dt\right]}{\mathbb{E}_0\left[\int_0^\tau e^{-rt}dt\right]} \\
& d\mu_{i,t} = I_{i,t}\mu_{i,t}(1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt \quad \text{with } \mu_{i,0} = \mu_i \\
& d\lambda_{i,t} = I_{i,t}\lambda_{i,t}\left((1-\lambda_{i,t})\gamma - \rho\right)dt \quad \text{with } \lambda_{i,0} = \lambda_i \\
& \mu_{i,t} = 0 \quad \text{if } I_{i,t} = 0 \\
& \mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1] = (1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt \quad \text{with } I_{i,0} = 1,
\end{aligned}
\tag{A5}
$$

where the supremum is over all stopping times. In what follows, for analogy with the Gittens index literature, we refer to product i as arm i, and we refer to offering product i as pulling arm i.

Lemma A1. *In the pre-AI environment, the platform's equilibrium decision at any time t is to offer i^* where*

$$i^* \in \arg\max_{i \in \mathcal{N}} M(\mu_i, \lambda_i).$$

This lemma is an immediate implication of the optimality of the Gittins index established in, e.g., [Gittins et al., 2011, Chapter 2]. We make use of the following three lemmas in this proof.

Lemma A2. *Consider product i with the state (μ_i, λ_i). The optimal stopping time in the definition of Gittins index is to stop pulling arm i until bad news occurs for it. That is*

$$\tau_i \triangleq \inf\{t \in \mathbb{R}_+ \ : \ I_t = 0\}$$

Proof of Lemma A2: This lemma follows because $\mu_{i,t}$ is increasing in t (Lemma 4). ■

Lemma A3. *For two products i and i' with $\lambda_i = \lambda_{i'}$, we have*

$$M(\mu_i, \lambda_i) \geq M(\mu_{i'}, \lambda_{i'}) \text{ for } \mu_i \geq \mu_{i'}.$$

Proof of Lemma A3: First note that by using Lemma 4, we know that $\mu_{i,t}$ is increasing in $\mu_{i,0}$. Therefore, $\mu_{i,t} \geq \mu_{i',t}$ for all t. Now consider the first time for which bad news occurs for both arms i and i' and let us denote them by τ_i and $\tau_{i'}$. We claim that τ_i first-order stochastically dominates $\tau_{i'}$, as

$$
\begin{aligned}
\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1] &= (1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt \\
&\overset{(a)}{=} (1-\mu_{i,t})\left(1-\lambda_{i',t}\right)\gamma dt \\
&\leq \left(1-\mu_{i',t}\right)\left(1-\lambda_{i',t}\right)\gamma dt = \mathbb{P}[I_{i',t+dt} = 0 \mid I_{i',t} = 1],
\end{aligned}
$$

where (a) follows from the fact that, as established in Lemma 4, given $\lambda_{i,0} = \lambda_{i',0} = \lambda$, we have $\lambda_{i,t} = \lambda_{i',t}$.

Using Lemma A2 and given the above two properties, we have that

$$M(\mu_i, \lambda_i) = \frac{\mathbb{E}_{t\sim\tau_i}[\int_0^t e^{-rs}\mu_{i,s}ds]}{\mathbb{E}_{t\sim\tau_i}[\int_0^t e^{-rs}ds]} \geq \frac{\mathbb{E}_{t\sim\tau_{i'}}[\int_0^t e^{-rs}\mu_{i',s}ds]}{\mathbb{E}_{t\sim\tau_{i'}}[\int_0^t e^{-rs}ds]} = M(\mu_{i'}, \lambda_{i'}).$$

This completes the proof. ■

Lemma A4. *Consider arm i with initial state (μ_i, λ_i). If bad news does not occur, the Gittin's index of this arm increases.*

Proof of Lemma A4: Letting τ_i be the first time for which bad news occurs and using Lemma A2, we can write

$$\begin{aligned} M(\mu_i, \lambda_i) = \frac{\mathbb{E}_{t\sim\tau_i}[\int_0^t e^{-rs}\mu_{i,s}ds]}{\mathbb{E}_{t\sim\tau_i}[\int_0^t e^{-rs}ds]} &= \frac{\mathbb{P}[I_{i,0}=0]\mu_{i,0} + \mathbb{P}[I_{i,0}=1]\mathbb{E}_{t\sim\tau_i}\left[\int_0^t e^{-rs}\mu_{i,s}ds \mid \mathrm{NBN}_{i,0}\right]}{\mathbb{P}[I_{i,0}=0] + \mathbb{P}[I_{i,0}=1]\mathbb{E}_{t\sim\tau_i}\left[\int_0^t e^{-rs}ds \mid \mathrm{NBN}_{i,0}\right]} \\ &\overset{(a)}{\leq} \frac{\mathbb{E}_{t\sim\tau_i}\left[\int_0^t e^{-rs}\mu_{i,s}ds \mid \mathrm{NBN}_{i,0}\right]}{\mathbb{E}_{t\sim\tau_i}\left[\int_0^t e^{-rs}ds \mid \mathrm{NBN}_{i,0}\right]}, \end{aligned}$$

where (a) follows from the fact that $\mu_{i,t}$ is increasing in t, as established in Lemma 4. ■

We now continue with the proof of Theorem 1. Without loss of generality, let us suppose $\mu_{n,0} \geq \cdots \geq \mu_{1,0}$. Using Lemma A3, and given $\lambda_{i,0} = \lambda$ for all $i \in \mathcal{N}$, we have $M(\mu_n, \lambda_n) \geq M(\mu_i, \lambda_i)$. Therefore, the platform starts by offering product n. If bad news does not occur for this product, using Lemma A4, the Gittins index of arm n weakly increases while the other Gittins indices remain unchanged. Therefore, using Lemma A1, the platform finds it optimal to keep pulling arm n until bad news occurs. Whenever bad news occurs for product n (if at all), the platform's problem becomes offering the best arm among $n-1$ products, and by induction, again, the platform starts offering product $n-1$ and so on. ■

Proof of Lemma 5

The differential equation of the evolution of $\lambda_{i,t}^{(P)}$ when $\alpha_{i,0} = 1$, as we established in Lemma 2, is given by

$$\frac{d}{dt}\lambda_{i,t}^{(P)} = \lambda_{i,t}^{(P)}\left((1-\lambda_{i,t}^{(P)})\gamma - \rho\right)$$

whose solution is

$$\frac{e^{\gamma t+\rho c}(\gamma-\rho)}{e^{\rho t+\gamma c} - e^{\gamma t+\rho c}\gamma}$$

for some constant c. Using the initial condition $\lambda_{i,0}^{(P)} = 1$, we obtain

$$\lambda_{i,t}^{(P)} = \frac{\gamma-\rho}{\gamma - e^{(\rho-\gamma)t}\rho}$$

The differential equation of the evolution of $\mu_{i,t}^{(P)}$ when $\alpha_{i,0} = 0$, as shown in Lemma 2, is given by

$$\frac{d}{dt}\mu_{i,t}^{(P)} = \mu_{i,t}^{(P)}\left(1 - \mu_{i,t}^{(P)}\right)\gamma$$

whose solution is

$$\mu_{i,t}^{(P)} = \frac{1}{1 + ce^{-\gamma t}}$$

for some constant c. Using the initial condition $\mu_{i,0}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0}+(1-\mu_{i,0})(1-\lambda)}$, we obtain

$$\mu_{i,t}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0} + e^{-\gamma t}(1-\lambda)(1-\mu_{i,0})}.$$

For $\alpha_{i,0} = 1$, we have

$$\mathbb{P}[\mathrm{NBN}_t \mid \alpha_{i,0} = 1] = e^{-\int_0^t \left(1 - \frac{\gamma-\rho}{\gamma - e^{(\rho-\gamma)s}\rho}\right)\gamma ds} = \frac{e^{-\rho t}\gamma - e^{-\gamma t}\rho}{\gamma - \rho}.$$

Also, the probability of bad news arriving for the first time in $(t, t+dt]$ is

$$\begin{aligned}\mathbb{P}[I_{t+dt} = 0, I_t = 1 \mid \alpha_{i,0} = 1] =& \mathbb{P}[\mathrm{NBN}_t \mid \alpha_{i,0} = 1]\left(1 - \frac{\gamma-\rho}{\gamma - e^{(\rho-\gamma)t}\rho}\right)\gamma dt \\ =& \frac{\left(e^{-\rho t} - e^{-\gamma t}\right)\rho}{\gamma - \rho}\gamma dt.\end{aligned}$$

For $\alpha_{i,0} = 0$, we have

$$\mathbb{P}[\mathrm{NBN}_t \mid \alpha_{i,0} = 0] = e^{-\int_0^t \left(1 - \frac{\mu_{i,0}}{\mu_{i,0} + e^{-\gamma s}(1-\lambda)(1-\mu_{i,0})}\right)\gamma ds} = \frac{\mu_{i,0} + e^{-\gamma t}(1-\lambda)(1-\mu_{i,0})}{1 - \lambda(1-\mu_{i,0})}.$$

Also, the probability of bad news arriving for the first time in $(t, t+dt]$ is

$$\begin{aligned}\mathbb{P}[I_{t+dt} = 0, I_t = 1 \mid \alpha_{i,0} = 0] &= \mathbb{P}[\mathrm{NBN}_t \mid \alpha_{i,0} = 0]\left(1 - \frac{\mu_{i,0}}{\mu_{i,0} + e^{-\gamma t}(1-\lambda)(1-\mu_{i,0})}\right)\gamma dt \\ &= \frac{e^{-\gamma t}(1-\lambda)(1-\mu_{i,0})}{1 - \lambda(1-\mu_{i,0})}\gamma dt.\end{aligned}$$

This completes the proof. ■

Proof of Theorem 2

In the post-AI environment, the state of each product is given by the tuple $(\mu_i, \lambda_i, \mu_i^{(P)}, \lambda_i^{(P)})$, where μ_i is the user belief about product's quality, λ_i is the probability of having a glossiness state $\alpha = 1$ given no bad news has occurred and $(\mu_i^{(P)}, \lambda_i^{(P)})$ are the same quantities but from the platform's perspective.

Similar to the analysis of pre-AI, for any product $i \in \mathcal{N}$, we introduce the Gittins index as

$$M(\mu_i, \lambda_i, \mu_i^{(P)}, \lambda_i^{(P)}) = \sup_{\tau \geq 0} \frac{\mathbb{E}_0^P \left[\int_0^\tau e^{-rt} \mu_{i,t} dt\right]}{\mathbb{E}_0^P \left[\int_0^\tau e^{-rt} dt\right]} \tag{A6}$$

$$\begin{aligned}
&d\mu_{i,t} = I_{i,t}\mu_{i,t}(1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt \quad \text{ with } \mu_{i,0} = \mu_i \\
&d\lambda_{i,t} = I_{i,t}\lambda_{i,t}\left((1-\lambda_{i,t})\gamma - \rho\right) dt \quad \text{ with } \lambda_{i,0} = \lambda \\
&d\mu_{i,t}^{(P)} = I_{i,t}\mu_{i,t}^{(P)}(1-\mu_{i,t}^{(P)})\gamma dt \quad \text{with } \mu_{i,0}^{(P)} = \mu_i^{(P)} \\
&d\lambda_{i,t}^{(P)} = I_{i,t}\lambda_{i,t}^{(P)}\left((1-\lambda_{i,t}^{(P)})\gamma - \rho\right) dt \quad \text{with } \lambda_{i,0}^{(P)} = \lambda_i^{(P)} \\
&\mu_{i,t} = \mu_{i,t}^{(P)} = 0 \quad \text{ if } I_{i,t} = 0 \\
&\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1, \alpha_{i,0} = 0] = \left(1 - \mu_{i,t}^{(P)}\right)\gamma dt \\
&\mathbb{P}[I_{i,t+dt} = 0 \mid I_{i,t} = 1, \alpha_{i,0} = 1] = \left(1 - \lambda_{i,t}^{(P)}\right)\gamma dt,
\end{aligned}$$

where the supremum is over all stopping times. Notice that the platform, in the post-AI environment, observes the initial $\alpha_{i,0}$ which we can compactly write as some initialization for $\mu_{i,0}^{(P)}$ and $\lambda_{i,0}^{(P)}$:

If the initial glossiness state is $\alpha_{i,0} = 1$, then

$$\mu_{i,0}^{(P)} = 0, \quad \lambda_{i,0}^{(P)} = 1,$$

and if the initial glossiness state is $\alpha_{i,0} = 0$, then

$$\mu_{i,0}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0} + (1-\mu_{i,0})(1-\lambda)}, \quad \lambda_{i,0}^{(P)} = 0.$$

Lemma A5. *In the post-AI environment, the platform's equilibrium decision at any time t is to offer i^* where*

$$i^* \in \arg\max_{i \in \mathcal{N}} M(\mu_i, \lambda_i, \mu_i^{(P)}, \lambda_i^{(P)}).$$

Using a similar argument to that of Lemma A3, and given $\lambda_{i,0} = \lambda$ for all $i \in \mathcal{N}$, the platform starts by offering the product with the highest user belief either among those with the initial glossiness state $\alpha = 1$ or among those with the initial glossiness state of 0. If bad news does not occur for this product, again using a similar argument to Lemma A4, establishes that the Gittins index of this arm weakly increases while the other Gittins indices remain unchanged. Therefore, the platform finds it optimal to keep offering this product until bad news occurs. Whenever bad news occurs for this product (if at all), the platform's problem becomes offering the best arm among the $n-1$ remaining products, and the proof follows by induction on the number of products. ■

Proof of Proposition 1

By using Lemma A5, it suffices to consider the platform's problem with two products with initial beliefs $\mu_{i',0} = \mu_1, \alpha_{i',0} = 1$ and $\mu_{i,0} = \mu_0, \alpha_{i,0} = 0$ and characterize the platform's equilibrium decision. Also, notice that by using Lemma A4, the platform's equilibrium strategy is to either offer product i until bad news occurs and then switch to the product i' or to offer product i' until bad news occurs and then switch to the product i. We next evaluate the platform's utility with these two strategies and compare them. We find it useful to write the platform's payoff with these two strategies as a function of ρ.

Using Lemma 5, the platform's payoff when it offers product i first and when bad news occurs, switches to i' is

$$\begin{aligned}
U_{0,1}(\rho,\mu_0,\mu_1) \triangleq & \frac{\mu_0}{\mu_0+(1-\lambda)(1-\mu_0)} \int_0^\infty re^{-rt}\mu_{i,t}dt \\
&+ \frac{(1-\lambda)(1-\mu_0)}{\mu_0+(1-\lambda)(1-\mu_0)} \int_0^\infty re^{-rt}\mu_{i,t}\mathbb{P}\left[\mathrm{NBN}_{i,t} \mid \theta_i=0,\alpha_{i,0}=0\right]dt \\
&+ \frac{(1-\lambda)(1-\mu_0)}{\mu_0+(1-\lambda)(1-\mu_0)} \left(\int_0^\infty e^{-rt}\mathbb{P}\left[I_{i,t+dt}=0, I_{i,t}=1 \mid \theta_i=0,\alpha_{i,0}=0\right]\right) \\
&\times \left(\int_0^\infty re^{-rt}\mu_{i',t}\mathbb{P}\left[\mathrm{NBN}_{i',t} \mid \alpha_{i',0}=1\right]\right) \\
=& \int_0^\infty re^{-rt}\mu_{i,t}\frac{\mu_0+(1-\mu_0)(1-\lambda)e^{-\gamma t}}{\mu_0+(1-\lambda)(1-\mu_0)}dt \\
&+ \frac{(1-\lambda)(1-\mu_0)}{\mu_0+(1-\lambda)(1-\mu_0)} \left(\int_0^\infty e^{-rt}\gamma e^{-\gamma t}dt\right)\left(\int_0^\infty re^{-rt}\mu_{i',t}\frac{\gamma e^{-\rho t}-\rho e^{-\gamma t}}{\gamma-\rho}dt\right),
\end{aligned}$$

where, using Lemma 4, we have

$$\begin{aligned}
\mu_{i,t} =& \frac{\mu_0(\gamma-\rho)}{\mu_0(\gamma-\rho)+(1-\mu_0)\left(e^{-\rho t}\lambda\gamma+e^{-\gamma t}\left((1-\lambda)\gamma-\rho\right)\right)} \text{ and} \\
\mu_{i',t} =& \frac{\mu_1(\gamma-\rho)}{\mu_1(\gamma-\rho)+(1-\mu_1)\left(e^{-\rho t}\lambda\gamma+e^{-\gamma t}\left((1-\lambda)\gamma-\rho\right)\right)}.
\end{aligned}$$

Again, using Lemma 5, the platform's payoff by offering product i' first and when bad news occurs then switching to product i is given by

$$\begin{aligned}
U_{1,0}(\rho,\mu_0,\mu_1) \triangleq & \int_0^\infty re^{-rt}\mu_{i',t}\mathbb{P}\left[\mathrm{NBN}_{i',t} \mid \alpha_{i',0}=1\right]dt \\
&+ \left(\int_0^\infty e^{-rt}\mathbb{P}\left[I_{i',t+dt}=0, I_{i,t}=1 \mid \alpha_{i',0}=1\right]\right)\left(\int_0^\infty re^{-rt}\mu_{i,t}\mathbb{P}\left[\mathrm{NBN}_{i,t} \mid \alpha_{i,0}=0\right]dt\right) \\
=& \int_0^\infty re^{-rt}\mu_{i',t}\frac{\gamma e^{-\rho t}-\rho e^{-\gamma t}}{\gamma-\rho}dt \\
&+ \left(\int_0^\infty e^{-rt}\gamma\rho\frac{e^{-\rho t}-e^{-\gamma t}}{\gamma-\rho}dt\right)\left(\int_0^\infty re^{-rt}\mu_{i,t}\frac{\mu_0+e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)}dt\right).
\end{aligned}$$

Notice that both functions are (uniformly) continuous in ρ, and therefore it suffices to establish the statement in the limit of $\rho\to 0$. We can write

$$\lim_{\rho\to 0} U_{1,0}(\rho,\mu_0,\mu_1) - U_{0,1}(\rho,\mu_0,\mu_1)$$

$$
\begin{aligned}
&= \int_0^\infty re^{-rt} \frac{\mu_1\gamma}{\mu_1\gamma + (1-\mu_1)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)} dt \\
&- \int_0^\infty re^{-rt} \frac{\mu_0\gamma}{\mu_0\gamma + (1-\mu_0)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)} \frac{\mu_0 + e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)} dt \\
&- \left(\int_0^\infty e^{-rt} \frac{e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)} \gamma dt\right) \left(\int_0^\infty re^{-rt} \frac{\mu_1\gamma}{\mu_1\gamma + (1-\mu_1)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)} dt\right) \\
&= \left(\int_0^\infty re^{-rt} \frac{\mu_1\gamma}{\mu_1\gamma + (1-\mu_1)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)} dt\right) \left(1 - \frac{(1-\lambda)(1-\mu_0)\gamma}{\left(1-\lambda(1-\mu_0)\right)(r+\gamma)}\right) \\
&\quad - \int_0^\infty re^{-rt} \frac{\mu_0\gamma}{\mu_0\gamma + (1-\mu_0)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)} \frac{\mu_0 + e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)} dt \\
&\quad = \left(\int_0^\infty re^{-rt} \frac{\mu_1\gamma}{\mu_1\gamma + (1-\mu_1)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)} dt\right) \left(\int_0^\infty re^{-rt} \frac{\mu_0 + e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)} dt\right) \\
&\quad - \int_0^\infty re^{-rt} \frac{\mu_0\gamma}{\mu_0\gamma + (1-\mu_0)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)} \frac{\mu_0 + e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)} dt. \qquad \text{(A7)}
\end{aligned}
$$

We next prove that the above expression is strictly positive for $\mu_0 = \mu_1 = \mu$. This follows from FKG-Harris inequality (see, e.g., [Alon and Spencer, 2016, Chapter 6.2]) that states if $a : \mathbb{R}_+ \to \mathbb{R}_+$ is an increasing and $b : \mathbb{R}_+ \to \mathbb{R}_+$ is a decreasing function, then we have

$$
\mathbb{E}\left[a(X)b(X)\right] \leq \mathbb{E}\left[a(X)\right]\mathbb{E}\left[b(X)\right]
$$

for random variable X distributed over $\mathbb{R}_+$. Moreover, the above inequality is strict if both functions are strictly monotone and X has a strictly positive mass everywhere. In particular, the positivity of the right-hand side of (A7) follows by invoking the above inequality for

$$
a(t) = \frac{\mu\gamma}{\mu\gamma + (1-\mu)\left(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma\right)}, b(t) = \frac{\mu + e^{-\gamma t}(1-\lambda)(1-\mu)}{1-\lambda(1-\mu)},
$$

and t distributed as $\mathbb{P}[t = s] = re^{-rs}$. Therefore, there exists a function $f(\mu, \lambda, \gamma, r)$ that is everywhere below μ, such that for $\mu_1 \geq f(\mu_0, \lambda, \gamma, r)$, the right-hand side of inequality (A7) is positive. ■

Proof of Proposition 2

Similar to the proof of Proposition 1, it suffices to consider the platform's problem with two products with initial beliefs $\mu_{i',0} = \mu_1, \alpha_{i',0} = 1$ and $\mu_{i,0} = \mu_0, \alpha_{i,0} = 0$ and characterize the platform's equilibrium decision. Again, using a similar notation to that of Proposition 1, we let $U_{1,0}(\rho, \mu_0, \mu_1)$ denote the platform's payoff from offering product i' first and $U_{0,1}(\rho, \mu_0, \mu_1)$ denote the platform's payoff from offering product i first. Again, notice that both functions are (uniformly) continuous in ρ (over $\mathbb{R}_+ \cup \{\infty\}$), and therefore it suffices to establish the statement in the limit of $\rho \to \infty$. We can write

$$
\begin{aligned}
&\lim_{\rho\to\infty} U_{0,1}(\rho, \mu_0, \mu_1) - U_{1,0}(\rho, \mu_0, \mu_1) \\
&= \int_0^\infty re^{-rt} \frac{\mu_0}{\mu_0 + e^{-\gamma t}(1-\mu_0)} \frac{\mu_0 + e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)} dt
\end{aligned}
$$

$$
\begin{aligned}
&+\left(\int_0^\infty e^{-rt}\frac{e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)}\gamma dt\right)\left(\int_0^\infty re^{-rt}\frac{\mu_1}{\mu_1+e^{-\gamma t}(1-\mu_1)}e^{-\gamma t}dt\right)\\
&-\int_0^\infty re^{-rt}\frac{\mu_1}{\mu_1+e^{-\gamma t}(1-\mu_1)}e^{-\gamma t}dt\\
&-\left(\int_0^\infty e^{-rt}\gamma e^{-\gamma t}dt\right)\left(\int_0^\infty re^{-rt}\frac{\mu_0}{\mu_0+e^{-\gamma t}(1-\mu_0)}\frac{\mu_0+e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)}dt\right)\\
&=\left(\int_0^\infty re^{-rt}\frac{\mu_1}{\mu_1+e^{-\gamma t}(1-\mu_1)}e^{-\gamma t}dt\right)\left(\frac{\gamma(1-\lambda)(1-\mu_0)}{(1-\lambda(1-\mu_0))(\gamma+r)}-1\right)\\
&+\int_0^\infty re^{-rt}\frac{\mu_0}{\mu_0+e^{-\gamma t}(1-\mu_0)}\frac{\mu_0+e^{-\gamma t}(1-\lambda)(1-\mu_0)}{1-\lambda(1-\mu_0)}dt\left(1-\frac{\gamma}{r+\gamma}\right).
\end{aligned}
\tag{A8}
$$

For $\mu_1=\mu_0=\mu$, the right-hand side of (A8) can be written as

$$
\left(\frac{r}{r+\gamma}\right)^2\left(\mathbb{E}\left[a(t)b(t)\right]-\mathbb{E}\left[a(t)\right]\mathbb{E}\left[b(t)\right]\right)
$$

for $a(t)=\frac{\mu}{\mu+e^{-\gamma t}(1-\mu)}$ and $b(t)=\frac{\mu e^{\gamma t}+(1-\lambda)(1-\mu)}{1-\lambda(1-\mu)}$ and t being distributed as $\mathbb{P}[t=s]=(\gamma+r)e^{-(\gamma+r)s}$. Again using FKG-Harris inequality for (strictly) increasing functions $a(\cdot)$ and $b(\cdot)$, we conclude that the right-hand side of (A8) is strictly positive. Therefore, there exists a function $g(\mu,\lambda,\gamma,r)$ that is everywhere below μ, such that for $\mu_0\geq g(\mu_1,\lambda,\gamma,r)$, the right-hand side of (A8) is positive. ■

Proof of Theorem 3

As we established in Theorems 1 and 2, the equilibrium in both pre-AI and post-AI settings start with the offering of one of the products until bad news occurs, and then following bad news, the platform offers another product and so on. Also, as in Propositions 1 and 2, both the expected utility obtained from offering any product and the probability of bad news occurring for a product are continuous functions of ρ. Therefore, it suffices to prove the theorem statement in the limit of $\rho\to\infty$ with strict inequality. This establishes the existence of large enough ρ_h such that the inequalities hold for $\rho\geq\rho_h$.

First, note that the platform's information advantage in the post-AI implies that its expected payoff is higher than in the pre-AI environment. We next prove that the expected user utility in the post-AI is higher than in the pre-AI environment. Since utilitarian welfare is the summation of the user and the platform utilities, we conclude that the expected welfare in the post-AI is higher than in the pre-AI.

Let us consider a realization of the products' beliefs and, without loss of generality, assume

$$
\mu_{1,0}\leq\cdots\leq\mu_{n,0}.
$$

As we established in Theorem 1, in the pre-AI environment, the platform first offers the product with the highest belief, that is, product n, and if bad news happens, offer the product with the second highest product, and so on. Using Proposition 2, in the post-AI environment, the platform may start offering product i where $i<n$, $\alpha_{i,0}=0$, and $\alpha_{n,0}=1$. In the post-AI environment, the platform equilibrium strategy may also differ in the order of other products. However, we can obtain the platform strategy in

the post-AI from its strategy in the pre-AI by performing a series of *helpful swap*s, defined next.

Suppose $\mu_{1,0} \leq \cdots \leq \mu_{n,0}$ and i is the highest index for which $\alpha_{i,0} = 0$. Consider the strategy of offering products in decreasing order of their beliefs. A *helpful swap* offers product i first and then offers the rest of the products in decreasing order of their beliefs.

The proof of theorem follows from the following lemma that establishes for large enough ρ, a helpful swap increases the user expected utility.

Lemma A6 (Helpful swap)**.** *The expected user utility increases for large enough ρ after performing a helpful swap.*

Proof of Lemma A6: It suffices to prove the lemma in the limit of $\rho \to \infty$ with strict inequality. This establishes the existence of large enough ρ for which the statement holds. For any $j \in \mathcal{N}$, we let τ_j be the stochastic time at which bad news occurs for product j given $\theta_j = 0$. Notice that in the limit of $\rho \to \infty$, for both a product with $\alpha_{j,0} = 1$ and a low-quality product with $\alpha_{j,0} = 0$, the platform knows that bad news arrive with rate γ. We let

$$A_j \triangleq \mathbb{E}\left[-\int_0^{\tau_j} re^{-rt}\mu_{j,t}dt\right]$$

for $j = n, n-1, i$. Using this notation, and that for $j = n, n-1, \ldots, i+1$, $\delta \triangleq \mathbb{E}_{t\sim\tau_j}[e^{-rt}] < 1$, the expected utility before the swap is

$$\begin{aligned} U_1 \triangleq & \sum_{j=n}^{i+1} A_j\delta^{n-j} + \delta^{n-i}\left(\mu_{i,0}^{(P)}\int_0^{\infty} re^{-rt}(1-\mu_{i,t})dt + (1-\mu_{i,0}^{(P)})A_i\right) \\ & + \delta^{n-i+1}(1-\mu_{i,0}^{(P)})\,(\text{Expected utility from products } i-1,\ldots,0)\,, \end{aligned} \tag{A9}$$

where $\mu_{i,0}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0}+(1-\lambda)(1-\mu_{i,0})}$ when $\alpha_{i,0} = 0$. The expected utility after the swap is

$$\begin{aligned} U_2 \triangleq & \mu_{i,0}^{(P)}\int_0^{\infty} re^{-rt}(1-\mu_{i,t})dt + (1-\mu_{i,0}^{(P)})A_i + \delta(1-\mu_{i,0}^{(P)})\sum_{j=n}^{i+1} A_j\delta^{n-j} \\ & + \delta^{n-i+1}(1-\mu_{i,0}^{(P)})\,(\text{Expected utility from products } i-1,\ldots,0) \\ & \overset{(a)}{>} \delta^{n-i}\mu_{i,0}^{(P)}\int_0^{\infty} re^{-rt}(1-\mu_{i,t})dt + (1-\mu_{i,0}^{(P)})A_i + \delta(1-\mu_{i,0}^{(P)})\sum_{j=n}^{i+1} A_j\delta^{n-j} \\ & + \delta^{n-i+1}(1-\mu_{i,0}^{(P)})\,(\text{Expected utility from products } i-1,\ldots,0) \\ & = U_1 + A_i\left((1-\mu_{i,0}^{(P)}) - \delta^{n-i}(1-\mu_{i,0}^{(P)})\right) + \sum_{j=n}^{i+1} A_j\left(\delta^{n+1-j}(1-\mu_{i,0}^{(P)}) - \delta^{n-j}\right) \\ & \overset{(b)}{\geq} U_1 + A_i\left((1-\mu_{i,0}^{(P)}) - \delta^{n-i}(1-\mu_{i,0}^{(P)})\right) + \sum_{j=n}^{i+1} A_i\left(\delta^{n+1-j}(1-\mu_{i,0}^{(P)}) - \delta^{n-j}\right) \\ & = U_1 - \mu_{i,0}^{(P)}A_i\left(\sum_{j=0}^{n-i+1}\delta^j\right) \overset{(c)}{>} U_1 \end{aligned} \tag{A10}$$

where (a) follow from $\mu_{i,0}^{(P)} \int_0^\infty re^{-rt}(1-\mu_{i,t})dt > 0$ and $1 > \delta^{n-i}$, (b) follows from $A_j < A_i$ and $\delta^{n+1-j}(1-\mu_{i,0}^{(P)}) - \delta^{n-j} < 0$ for $j = n, n-1, i+1$, and (c) follows from $A_i < 0$, completing the proof of lemma. ■

We now continue with the proof of the theorem. If $\alpha_{n,0} = 0$, then in both pre-AI and post-AI the platform starts by offering product n, and the proof follows by induction on the number of products. If $\alpha_{n,0} = 1$, then letting i be the largest index for which $\alpha_{i,0} = 0$, with positive probability we have that $\mu_{i,0} > g(\mu_{n,0}, \lambda, \gamma, r)$ and therefore Proposition 2 implies that the platform starts by offering product i in the post-AI. The above Lemma establishes the expected user utility increases by this helpful swap. The proof completes by noting that the platform strategy in the post-AI can be obtained from its strategy in the pre-AI by performing a series of helpful swaps. ■

Proof of Theorem 4

Again, as we established in Theorems 1 and 2, the equilibrium in both pre-AI and post-AI settings start offering one of the products until bad news occurs for it and then offers another product and so on. Also, as we showed in the proof of Propositions 1 and 2, both the expected utility obtained from offering any product and the probability of bad news occurring for a product are continuous functions of ρ. Therefore, it suffices to prove the theorem statement in the limit of $\rho \to 0$ with strict inequality. This establishes the existence of small enough ρ_l such that the inequalities hold for $\rho \leq \rho_l$. First, note that the platform's information advantage in the post-AI implies its expected payoff is higher than in the pre-AI environment. We next prove that the expected utilitarian welfare in the post-AI is smaller than in the pre-AI environment. Since utilitarian welfare is the summation of the user and the platform utilities, the expected user utility in the post-AI is smaller than in the pre-AI.

Let us consider a realization of the products' beliefs and, without loss of generality, assume

$$\mu_{1,0} \leq \cdots \leq \mu_{n,0}.$$

Using the characterizations of Theorems 1 and 2, we have the following cases:

1. $\alpha_{n,0} = 1$: In this case, in both pre-AI and post-AI environments, the platform starts offering product n. Since $\rho = 0$, bad news does not arrive for this product, and the expected welfare in pre-AI and post-AI environments are the same.

2. $\alpha_{n,0} = 0$, there are two cases:

   2.1. For all products $j \in \mathcal{N} \setminus \{n\}$ for which $\alpha_{j,0} = 1$, we have $\mu_{j,0} \leq f(\mu_{n,0}, \lambda, \gamma, r)$. In this case, in both the pre-AI and post-AI environment, the platform starts offering product n. Therefore, by induction on the number of products, the expected welfare in the post-AI is smaller than in the pre-AI environment.

   2.2. There exists product $j \in \mathcal{N} \setminus \{n\}$ for which $\alpha_{j,0} = 1$ and $\mu_{j,0} > f(\mu_{n,0}, \lambda, \gamma, r)$. We let i be the largest index among such products. In the post-AI environment, the platform starts offering product i, and given that bad news does not occur for this product ($\rho = 0$) and $\theta_i = 0$, the expected welfare becomes 0. In the pre-AI environment, the platform starts by offering product n, and therefore, with probability $\mu_{n,0}$ we have $\theta_n = 1$, ensuring an expected welfare of at least $\mu_{n,0}$.

Finally, notice that case 2.2 happens with a positive probability, proving that the expected welfare in the post-AI environment is strictly lower than in the pre-AI environment. ■

Proof of Theorem 5

Similar to the argument of Theorem 4, it suffices to prove the theorem statement for the expected welfare in the limit of $\rho \to 0$ with strict inequality. This establishes the existence of small enough $\tilde{\rho}_l$ such that the inequalities hold for $\rho \leq \tilde{\rho}_l$. We let Δ_1 be small enough so that

$$f\left(\frac{1}{2} + \Delta_1; \lambda, \gamma, r\right) \leq \frac{1}{2} - \Delta_1. \tag{A11}$$

Such Δ_1 exists because $f(\mu; \lambda, \gamma, r)$ is continuous in μ and is everywhere below μ and in particular, $f(\frac{1}{2}; \lambda, \gamma, r) \leq \frac{1}{2}$. The continuity of $f(\cdot)$ follows from the characterization of the function $f(\cdot)$ derived in Proposition 1 (and in particular, (A7)) and implicit function theorem. Let us compare the expected welfare conditioned on the belief realization for the first n products. Without loss of generality, we assume

$$\mu_{1,0} \leq \mu_{2,0} \leq \dots \mu_{n,0}.$$

Let us assume $\Delta \leq \Delta_1$. We have the following cases:

1. There exists $i \in \mathcal{N}$ for which $\alpha_{i,0} = 1$: In this case, given $\Delta \leq \Delta_1$ and (A11), the platform's equilibrium strategy with and without the $n+1$-th product is to offer one of the products with the initial glossiness state of 1. Therefore, the expected welfare with and without the extra product is zero.

2. $\alpha_{i,0} = 0$ for all $i \in \mathcal{N}$ and $\alpha_{n+1,0} = 1$: In this case, given $\Delta \leq \Delta_1$ and (A11), the platform's equilibrium strategy with the $n+1$-th product is to offer that product, and the expected welfare becomes 0. The platform's equilibrium strategy with n products, however, is to offer them in the decreasing order of their belief whose expected welfare is recursively defined as

$$W(\mu_{n,0}, \dots, \mu_{1,0}) = \mu_{n,0}^{(P)} + (1 - \mu_{n,0}^{(P)})\delta W(\mu_{n-1,0}, \dots, \mu_{1,0}), \tag{A12}$$

   where $\delta \triangleq \mathbb{E}_{t \sim \tau_j}\left[e^{-rt}\right] < 1$, τ_j is the arrival time of bad news for a low-quality product in the initial glossiness state of 0, and $\mu_{i,0}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0} + (1-\lambda)(1-\mu_{i,0})}$ when $\alpha_{i,0} = 0$.

3. $\alpha_{i,0} = 0$ for all $i \in \{1, \dots, n\}$ and $\alpha_{n+1,0} = 0$: In this case, again, the expected welfare with n

products is $W(\mu_{n,0},\ldots,\mu_{1,0})$ as defined in (A12). The expected welfare with $n+1$ products is

$$W(\mu_{n,0},\ldots,\mu_{i+1,0},\mu_{n+1,0},\mu_{i,0}\ldots,\mu_{1,0}) \tag{A13}$$

where i is such that $\mu_{i+1,0}\geq\mu_{n+1,0}\geq\mu_{i,0}$.

With the extra product, the expected welfare in case 2 increases while the expected welfare in case 3 decreases. We next prove that, for small enough Δ_2, when $\mu_i\in[1/2-\Delta_2,1/2+\Delta_2]$, in expectation, the expected welfare decreases. In case 2, the difference between expected welfare with and without the extra product is upper bounded by

$$-\mu_l\frac{1-((1-\mu_l)\delta)^n}{1-(1-\mu_l)\delta},\text{ where }\mu_l=\frac{\frac{1}{2}-\Delta_2}{\frac{1}{2}-\Delta_2+(\frac{1}{2}+\Delta_2)(1-\lambda)},$$

noting that μ_l is the smallest belief conditional on the initial glossiness state being 0 obtained by using Bayes' rule. In case 3, the difference between expected welfare with and without the extra product is upper bounded by

$$\mu_h\frac{1-((1-\mu_h)\delta)^{n+1}}{1-(1-\mu_h)\delta}-\mu_l\frac{1-((1-\mu_l)\delta)^n}{1-(1-\mu_l)\delta},\text{ where }\mu_h=\frac{\frac{1}{2}+\Delta_2}{\frac{1}{2}+\Delta_2+(\frac{1}{2}-\Delta_2)(1-\lambda)},$$

noting that μ_h is the largest belief conditional on the initial glossiness state being 0 obtained by using Bayes' rule. Therefore, the expected welfare with $n+1$ products minus the expected welfare with n products is upper bounded by

$$-(1-\mu_h)\lambda\mu_l\frac{1-((1-\mu_l)\delta)^n}{1-(1-\mu_l)\delta}+(1-(1-\mu_h)\lambda)\left(\mu_h\frac{1-((1-\mu_h)\delta)^{n+1}}{1-(1-\mu_h)\delta}-\mu_l\frac{1-((1-\mu_l)\delta)^n}{1-(1-\mu_l)\delta}\right).$$

We claim that the above expression is negative for small enough Δ_2. Notice this expression is continuous in Δ_2 and for $\Delta_2=0$ it becomes

$$\begin{aligned}&\frac{-(2-\lambda)(1-\lambda)\lambda+\left(\delta\frac{1-\lambda}{2-\lambda}\right)^n\left((2-\lambda)^2-\delta(1-\lambda)(2-(2-\lambda)\lambda)\right)}{(2-\delta(1-\lambda)-\lambda)(2-\lambda)^2}\\ &\overset{(a)}{<}\frac{-(2-\lambda)(1-\lambda)\lambda+\left(\frac{1-\lambda}{2-\lambda}\right)^n(2-\lambda)^2}{(2-\delta(1-\lambda)-\lambda)(2-\lambda)^2}\overset{(b)}{\leq}0\end{aligned}$$

where (a) follows from the fact that the denominator is positive and that $(2-\lambda)^2-\delta(1-\lambda)(2-(2-\lambda)\lambda)\geq(2-\lambda)^2-(1-\lambda)(2-(2-\lambda)\lambda)=2(1-\lambda)>0$, $\delta<1$, and $\delta(1-\lambda)(2-(2-\lambda)\lambda)>0$ and (b) follows from $n\geq\frac{\log\frac{2-\lambda}{(1-\lambda)\lambda}}{\log\frac{2-\lambda}{1-\lambda}}=1-\frac{\log\lambda}{\log\frac{2-\lambda}{1-\lambda}}$. This proves that for $\Delta=\min\{\Delta_1,\Delta_2\}$, the expected welfare (and therefore the expected user utility) decreases as the number of products increases. ■

References

D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Learning from reviews: The selection effect and the speed of learning. *Econometrica*, 90(6):2857–2899, 2022a.

D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4):218–256, 2022b.

A. Acquisti and H. R. Varian. Conditioning prices on purchase history. *Marketing Science*, 24(3):367–381, 2005.

A. Acquisti, L. Brandimarte, and G. Loewenstein. Privacy and human behavior in the age of information. *Science*, 347(6221):509–514, 2015.

A. Acquisti, C. Taylor, and L. Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2): 442–92, 2016.

A. R. Admati and P. Pfleiderer. A monopolistic market for information. *Journal of Economic Theory*, 39(2): 400–438, 1986.

A. Agarwal, M. Dahleh, and T. Sarkar. A marketplace for data: An algorithmic solution. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 701–726, 2019.

A. Agrawal, J. Gans, and A. Goldfarb. *Prediction machines: the simple economics of artificial intelligence*. Harvard Business Press, 2018.

N. Alon and J. H. Spencer. *The probabilistic method*. John Wiley & Sons, 2016.

J. J. Anton and D. A. Yao. The sale of ideas: Strategic disclosure, property rights, and contracting. *Review of Economic Studies*, 69(3):513–531, 2002.

M. Babaioff, R. Kleinberg, and R. Paes Leme. Optimal mechanisms for selling information. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 92–109. ACM, 2012.

A. Bardhi. Attributes: Selective learning and influence. *Available at SSRN 3468546*, 2020.

J. Begenau, M. Farboodi, and L. Veldkamp. Big data in finance and the growth of large firms. *Journal of Monetary Economics*, 97:71–87, 2018.

D. Bergemann and A. Bonatti. Selling cookies. *American Economic Journal: Microeconomics*, 7(3):259–94, 2015.

D. Bergemann and A. Bonatti. Markets for information: An introduction. *Annual Review of Economics*, 11, 2019.

D. Bergemann, B. Brooks, and S. Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–57, 2015.

D. Bergemann, A. Bonatti, and A. Smolin. The design and price of information. *American Economic Review*, 108(1):1–48, 2018.

D. Bergemann, A. Bonatti, and T. Gan-a. Markets for information. *mimeo, MIT and Yale*, 2019.

D. A. Berry and B. Fristedt. Bandit problems: sequential allocation of experiments (monographs on statistics and applied probability). *London: Chapman and Hall*, 5(71-87):7–7, 1985.

P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, 1999.

A. Bonatti and G. Cisternas. Consumer scores and price discrimination. *The Review of Economic Studies*, 87(2):750–791, 2020.

R. Calo. Digital market manipulation. *Geo. Wash. L. Rev.*, 82:995, 2013.

R. A. Calvo, D. Peters, K. Vold, and R. M. Ryan. Supporting human autonomy in ai systems: A framework for ethical enquiry. *Ethics of digital well-being: A multidisciplinary approach*, pages 31–54, 2020.

Y.-K. Che and J. Hörner. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925, 2018.

Y. Chen and S. Zheng. Prior-free data acquisition for accurate statistical estimation. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 659–677, 2019.

Y. Chen, N. Immorlica, B. Lucier, V. Syrgkanis, and J. Ziani. Optimal data acquisition for statistical estimation. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pages 27–44. ACM, 2018.

J. P. Choi, D.-S. Jeon, and B.-C. Kim. Privacy and personal data collection with information externalities. *Journal of Public Economics*, 173:113–124, 2019.

M. Ekmekci. Sustainable reputations with rating systems. *Journal of Economic Theory*, 146(2):479–503, 2011.

K. Eliaz, R. Eilat, and X. Mu. Optimal privacy-constrained mechanisms. *Available at SSRN: https://ssrn.com/abstract=3336795*, 2019.

P. Eső and B. Szentes. Optimal information disclosure in auctions and the handicap auction. *Review of Economic Studies*, 74(3):705–731, 2007.

I. Esponda. Behavioral equilibrium in economies with adverse selection. *American Economic Review*, 98 (4):1269–91, 2008.

E. Eyster and M. Rabin. Cursed equilibrium. *Econometrica*, 73(5):1623–1672, 2005.

I. P. Fainmesser, A. Galeotti, and R. Momot. Digital privacy. *Available at SSRN: https://ssrn.com/abstract=3459274*, 2019.

L. Felli and C. Harris. Learning, wage dynamics, and firm-specific human capital. *Journal of Political Economy*, 104(4):838–868, 1996.

E. Frostig and G. Weiss. Four proofs of gittins' multiarmed bandit theorem. *Annals of Operations Research*, 241(1):127–165, 2016.

D. Fudenberg and G. Lanzani. The stability of misperceptions under mutations. *Available at SSRN 3709932*, 2020.

D. Fudenberg and J. M. Villas-Boas. Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, 1:377–436, 2006.

D. Fudenberg, G. Lanzani, and P. Strack. Limits points of endogenous misspecified learning. *Available at SSRN*, 2020.

J. Gittins, K. Glazebrook, and R. Weber. *Multi-armed bandit allocation indices*. John Wiley & Sons, 2011.

A. Goldfarb and C. Tucker. Online display advertising: Targeting and obtrusiveness. *Marketing Science*, 30(3):389–404, 2011.

O. Gossner, J. Steiner, and C. Stewart. Attention please! *University of Zurich, Department of Economics, Working Paper*, (308), 2018.

R. Gradwohl. Information sharing and privacy in networks. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 349–350, 2017.

J. D. Hanson and D. A. Kysar. Taking behavioralism seriously: The problem of market manipulation. *NYUL Rev.*, 74:630, 1999.

P. Heidhues and B. Kőszegi. Behavioral industrial organization. *Handbook of Behavioral Economics: Applications and Foundations 1*, 1:517–612, 2018.

P. Heidhues, B. Kőszegi, and T. Murooka. Inferior products and profitable deception. *The Review of Economic Studies*, 84(1):323–356, 2016.

P. Heidhues, J. Johnen, and B. Koszegi. Browsing versus studying: A pro-market case for regulation. *unpublished MS*, 27, 2018.

S. Hidir and N. Vellodi. Personalization, discrimination and information disclosure. Technical report, working paper, 2019.

J. Horner and A. Skrzypacz. Selling information. *Journal of Political Economy*, 124(6):1515–1562, 2016.

S. Ichihashi. Dynamic privacy choices. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 539–540, 2020a.

S. Ichihashi. Online privacy and information disclosure by consumers. *American Economic Review*, 110 (2):569–95, 2020b.

B. Jullien, Y. Lefouili, and M. H. Riordan. Privacy protection, security, and consumer retention. *Security, and Consumer Retention (June 1, 2020)*, 2020.

E. Kamenica, S. Mullainathan, and R. Thaler. Helping consumers know themselves. *American Economic Review*, 101(3):417–22, 2011.

G. Keller and S. Rady. Strategic experimentation with poisson bandits. *Theoretical Economics*, 5(2):275–311, 2010.

G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73 (1):39–68, 2005.

N. Köbis, J.-F. Bonnefon, and I. Rahwan. Bad machines corrupt good morals. *Nature Human Behaviour*, 5 (6):679–685, 2021.

N. V. Krylov. *Controlled diffusion processes*, volume 14. Springer Science & Business Media, 2008.

A. Liang, X. Mu, and V. Syrgkanis. Dynamically aggregating diverse information. 2021.

T. C. Lin. The new market manipulation. *Emory LJ*, 66:1253, 2016.

J. Z. Liu, M. Sockin, and W. Xiong. Data privacy and algorithmic inequality. *Working Pape*, 2023.

K. Martin and H. Nissenbaum. Measuring privacy: An empirical test using context to expose confounding variables. *Colum. Sci. & Tech. L. Rev.*, 18:176, 2016.

R. Montes, W. Sand-Zantman, and T. Valletti. The value of personal information in online markets with endogenous privacy. *Management Science*, 65(3):1342–1362, 2019.

F. Pasquale. *The black box society*. Harvard University Press, 2015.

H. Robbins. Some aspects of the sequential design of experiments. *Bulletin of the American Mathematical Society*, 58(5):527–535, 1952.

A.-K. Roesler and B. Szentes. Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–80, 2017.

M. Rothschild. A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202, 1974.

D. Susser and V. Grimaldi. Measuring automated influence: Between empirical evidence and ethical values. In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, pages 242–253, 2021.

C. R. Taylor. Consumer privacy and the market for customer information. *RAND Journal of Economics*, pages 631–650, 2004.

J. Tirole. Digital dystopia. 2019.

J. Tirole. Competition and the industrial challenge for the digital age. 2020.

T. Z. Zarsky. Privacy and manipulation in the digital age. *Theoretical Inquiries in Law*, 20(1):157–188, 2019.

W. Zhong. Optimal dynamic information acquisition. *Available at SSRN 3024275*, 2017.

S. Zuboff. *The age of surveillance capitalism: The fight for a human future at the new frontier of power: Barack Obama's books of 2019*. Profile books, 2019.