AER: Insights 2025, 7(1): 19–38
https://doi.org/10.1257/aeri.20230589

When Big Data Enables Behavioral Manipulation[†]

By Daron Acemoglu, Ali Makhdoumi,
Azarakhsh Malekian, and Asuman Ozdaglar*

We build a model of online behavioral manipulation driven by AI advances. A platform dynamically offers one of n products to a user who slowly learns product quality. User learning depends on a product's "glossiness," which captures attributes that make products appear more attractive than they are. AI tools enable platforms to learn glossiness and engage in behavioral manipulation. We establish that AI benefits consumers when glossiness is short-lived. In contrast, when glossiness is long-lived, behavioral manipulation reduces user welfare. Finally, as the number of products increases, the platform can intensify behavioral manipulation by presenting more low-quality, glossy products. (*JEL* C55, D83, D91, L86)

The data of billions of individuals are currently being utilized for personalized advertising, product offerings, and pricing, and these practices are set to grow in the years to come. This exponentially increasing amount of data in the hands of tech companies such as Google, Facebook, and others is raising a host of concerns, including those related to privacy, predatory price discrimination, surveillance, and manipulation. While some experts are optimistic that more data will mean more informative advertising, better product targeting, and finer service differentiation, benefiting users (see, e.g., Varian 2018; McMorrow 2023; Perrigo 2024), others voice various worries. Zuboff (2019, 9), for example, argues that the economic model of AI "claims human experience as free raw material for hidden commercial practices of extraction, prediction, and sales." She maintains that the 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

*Acemoglu: Department of Economics, Massachusetts Institute of Technology (email: daron@mit.edu); Makhdoumi: Fuqua School of Business, Duke University (email: ali.makhdoumi@duke.edu); Malekian: Rotman School of Management, University of Toronto (email: azarakhsh.malekian@rotman.utoronto.ca); Ozdaglar: Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology (email: asuman@mit.edu). Dirk Bergemann was the coeditor for this article. 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.

[†]Go to https://doi.org/10.1257/aeri.20230589 to visit the article page for additional materials and author disclosure statement(s).

(1999, 635), 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."

Yet these arguments are insufficient since the modification of behavior may be for the user's good (directing her toward better products). In this paper, we explore the conditions under which behavioral modification harms consumers. We develop a behavioral model to theoretically explore the "bad" type of modification, which we refer to as *behavioral manipulation*.

Our Model and Results.—We suppose that consumers are sufficiently used to and can thus act rationally in the status quo environment. But they are not "hyperrational" enough to understand how the environment changes after technology platforms can collect vast amounts of data both on them and other consumers with similar characteristics, acquiring greater behavioral predictive power (meaning prediction of future behavior conditional on current inducements).

We consider a platform that offers one of n products to a user together with 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. But some products may be in a state that masks bad news (e.g., people tend to be favorably disposed toward products that have salient attractive characteristics or hidden or delayed costs). We refer to this as *glossiness* and denote it by $\alpha_{i,t}$. 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. This state is temporary, indicating that a low-quality product may initially appear good, but the user will eventually recognize its low quality. In particular, $\alpha_{i,t}$ follows a continuous-time Markov chain and transitions from a state of 1 to an absorbing state of $\alpha = 0$ at the rate ρ. Finally, if a low-quality product is not glossy, the user receives bad news about its quality at the rate γ. In this dynamic setup, the platform cares not just about the current offered product but also about the user's learning trajectory.

To motivate our model and assumptions, consider mobile gaming companies like Zynga or King. Their well-known games such as *Candy Crush* and *FarmVille* provide one example of our notion of glossiness, since they entice users with sophisticated graphics and trailers and also often enable early success for gamers. The companies then combine data analytics, machine learning techniques, and information on user experiences to raise revenues by means of in-game purchases, in-game currency, and other microtransactions (see, e.g., Costes and Bonnaire 2022). Platforms that sell products to users, such as Amazon, and those that market services and digital products, such as app markets, also engage in similar strategies and sometimes prioritize glossy products that are more likely to lead to quick purchases (see, e.g., Lindecrantz, Gi, and Zerbi 2022; Bradbury 2024; Wang et al. 2018). Another example is provided by social media platforms including Facebook, Instagram, and YouTube. Although these platforms do not charge users, they set prices to advertisers as a way of monetizing user data. By offering catchy content early on, these platforms are often able to intensify user engagement, enabling them to generate higher revenues from advertisers (see, e.g., Ryan-Mosley 2023).

We study two informational environments. One is a *pre-AI environment* (or status quo), in which the platform has no informational advantage. By observing the signals—specifically, whether bad news has arrived or not—the platform and the user update their beliefs about both the product's underlying quality and the glossiness state. The other is a *post-AI* (*or post–big data*) *environment*, in which the platform can leverage the large amount of data it collects from users and estimate the product's glossiness for a specific user from the behavior of others with similar characteristics from whom it has collected data.[1] The behavioral bias in our model is that users do not understand that the platform knows and can exploit this knowledge, which also implies that the user does not learn from the platform's decisions.

Our main results are as follows. First, when ρ is sufficiently large—which means that the initially glossy products remain so for a short period—(ex ante) expected user utility, welfare, and platform profits in the post-AI environment are larger than in the pre-AI environment. This confirms the view that data collected by online platforms can enable better product offerings and higher-quality matches for users. The reason is that, with large ρ, offering initially glossy products is not profitable, and it is better for the platform to guide users toward higher-quality products.

Second, when the transition rate ρ is small—the initially glossy products remain glossy for a long period—expected user utility and welfare in the post-AI environment will be smaller than in the pre-AI environment. While this happens, platform profits increase because platforms benefit from behavioral manipulation. Intuitively, in this case, the platform has an incentive to encourage users to consume low-quality glossy products, which will remain in this glossy state for a while.[2]

Third, for a sufficiently small transition rate ρ, as the number of products increases, 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 platform's ability to collect more information about users and widen their offerings, there may be a double whammy from the viewpoint of the consumers.

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 potentially harmful power over users. Similarly, Calo (2014); Lin (2017); and Zarsky (2019) argue that digital platforms can manipulate consumer behavior and suggest regulatory measures (see also Tirole 2023 for a discussion of policy implications).

More closely related are several works investigating whether information harms consumers, either because of surveillance (Tirole 2021) or because of price

[1] Specifically, using AI or conventional machine learning methods, the platform can harness the purchase, web page visit, and review data of users with similar characteristics to estimate which products have strong short-term appeal to the relevant demographic group. See Ekmekci (2011); Che and Hörner (2018); and Acemoglu et al. (2022a) for other analyses of the information content of online recommendation systems.

[2] This result also implies that our model derives "endogenous privacy costs" (due to manipulation) when ρ is small, while there are no such privacy costs when ρ is large.

discrimination (Taylor 2004; Acquisti and Varian 2005; Hidir and Vellodi 2019; Bergemann, Brooks, and Morris 2015; Bonatti and Cisternas 2020). See Acquisti, Taylor, and Wagman (2016); Bergemann and Bonatti (2019); and Agrawal, Gans, and Goldfarb (2018) for excellent surveys of this literature. We depart from this literature by focusing on behavioral manipulation—distortion of product choice and experimentation—rather than price discrimination.[3] Within this literature, our paper is most closely related to independent work by Liu, Sockin, and Xiong (2023), which examines the impact of data sharing on consumer welfare and algorithmic inequality. They argue 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 from but is complementary to our emphasis on platforms identifying sources of mislearning due to glossiness and other extraneous factors.

Our paper also relates to the literature on information markets such as Bergemann and Bonatti (2015); Goldfarb and Tucker (2011); and Montes, Sand-Zantman, and Valletti (2019), who focus on the use of personal data for improved allocation of online resources, and Admati and Pfleiderer (1986); Anton and Yao (2002); Hörner and Skrzypacz (2016); Bergemann, Bonatti, and Smolin (2018), and Begenau, Farboodi, and Veldkamp (2018), who investigate how information can be monetized either by dynamic sales or optimal mechanisms.

Our work more directly builds on the literature on experimentation and multiarmed bandits such as Rothschild (1974); Bolton and Harris (1999); Keller and Rady (2010); Felli and Harris (1996); and Keller, Rady, and Cripps (2005). The mathematical arguments in our paper are distinct from those in these past works because we have to keep track of the beliefs of the platform and the user, which follow different laws of motion. Even more closely related is Bergemann and Välimäki (2000), who develop an experimentation game where a seller sets prices that guide the learning trajectory of buyers, but crucially without behavioral manipulation.

Finally, our work is inspired by a growing body of evidence that platforms are acquiring a large amount of information about users and influencing their behavior. In this context, Susser and Grimaldi (2021); Acquisti, Brandimarte, and Loewenstein (2015); and Martin and Nissenbaum (2017) have empirically studied individuals' ability to control their personal information, while Köbis, Bonnefon, and Rahwan (2021) investigate broader social implications of extensive data collection and use of machine learning techniques and Calvo et al. (2020) focus on how platform algorithms can enable price manipulation.

The rest of the paper proceeds as follows. Section I presents our model. Section II provides several results useful for the rest of the analysis. Sections III and IV characterize the equilibrium in the pre-AI and post-AI environments, respectively. Section V establishes that manipulation occurs in equilibrium—it is possible for the user to be worse off in the post-AI environment. Section VI concludes, and the Supplemental Appendix presents the omitted proofs.

[3] Our paper also relates to the works that study privacy concerns and externalities in data markets, such as Choi, Jeon, and Kim (2019); Bergemann and Bonatti (2019); Acemoglu et al. (2022b); and Gradwohl (2017); to Fainmesser, Galeotti, and Momot (2023) and Jullien, Lefouili, and Riordan (2020), who consider the negative effects of leaking users' (private) personalized data; and to Ichihashi (2020, 2023), who explores the role of information intermediaries and dynamic data collection by platforms.

I. 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 $\theta_i \in \{0, 1\}$ 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 $p_{i,t}$ denote 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:

- If $\theta_i = 1$, then there is no bad news about product quality.
- If $\theta_i = 0$ and $\alpha_{i,t} = 0$, then bad news (fully revealing signal about product quality) arrives at the rate $\gamma > 0$.
- 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 $S_{i,t} = B$ to designate the arrival of bad news for product i at time t and $S_{i,t} = NB$ to designate no bad news for product i at time t. When purchasing product i, the user receives no signal about other products $j \in \mathcal{N} \backslash \{i\}$.

Glossiness $\alpha_{i,t}$ follows a homogeneous continuous-time Markov chain with two states $\{0, 1\}$ with *transition rate matrix*

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

This means that a low-quality product appears high quality for a while, but this type of glossiness is temporary and comes to an end at the rate ρ.

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

$$(1)\ \Pr(\alpha_{i,0} = 0 \,|\, \theta_i = 1) = 1, \Pr(\alpha_{i,0} = 1 \,|\, \theta_i = 0) = \lambda, \text{ and } \mu_{i,0} = \Pr(\theta_i = 1)$$

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

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

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

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

be the indicator that product i's quality at time t is not 0 from the user's perspective.

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:

- *pre-AI environment*, in which the platform does not have an information advantage and observes the same signals as the user
- *post-AI environment*, in which the platform has an information advantage and observes the initial glossiness state $\alpha_{i,0}$ for all $i \in \mathcal{N}$

We consider the informational assumptions in the post-AI environment as capturing the fact that platforms have access to a wealth of data about the behavior and preferences of users with similar characteristics and can use these to estimate which products will (temporarily) appear to a user as more attractive than their quality warrants. For simplicity, these data are assumed not to be directly informative about the true quality of the product.

A few points about the model are worth mentioning. First, for simplicity, we assume that in the post-AI environment, the platform can perfectly estimate the initial glossiness state $\alpha_{i,0}$. Second, the platform's information advantage is about the initial glossiness state and not the product's true quality, which could be idiosyncratic across users. Third, the platform's only information advantage is about the initial glossiness state and not the subsequent glossiness trajectory (which is likely to be idiosyncratic across users within the same demographic group). We establish that even with this limited form of informational advantage, the platform can have a profitable behavioral manipulation strategy. Our main results continue to hold even if we relax these three assumptions. Fourth, we assume that a glossy product is low quality to simplify the analysis and note that our results continue to hold more generally so long as low-quality products are more likely to be glossy. Finally, the user remains 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 preferred interpretation is that most users do not understand the informational capabilities of digital platforms in today's economy.

B. *Utilities*

If the platform offers item i at time t at price $p_{i,t}$, the user's instantaneous utility from its purchase is

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

and the platform's instantaneous utility, 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 the rate r for some $r \in \mathbb{R}_+$ and update their beliefs using Bayes' rule. The platform's expected profit (utility) at time 0 is given by

$$E_{P,0}\left[\int_0^\infty re^{-rt}\sum_{i=1}^{n} x_{i,t}p_{i,t}b_{i,t}\,dt\right],$$

and the user's expected utility at time 0 is

$$E_{U,0}\left[\int_0^\infty re^{-rt}\sum_{i=1}^{n} x_{i,t}b_{i,t}\big(\theta_i - p_{i,t}\big)\,dt\right],$$

where $x_{i,t}$ and $p_{i,t}$ are adapted to the natural filtration of the platform at time t, $b_{i,t}$ is adapted to the natural filtration of the user's information set at time t, and subscripts P and U denote expectations from the platform's and user's perspectives. Throughout, we assume that the user and the platform do not observe the realized utility (and thus cannot back out θ_i from $\theta_i - p_{i,t}$).

C. *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 state $\alpha_{i,0}$ for $i \in \mathcal{N}$. Therefore, the platform's decision at time t, $\big\{\big(x_{i,t}, p_{i,t}\big)\big\}_{i\in\mathcal{N}}$, 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 have the same belief about the products' quality and their glossiness state. We adopt (Bayesian) Markov perfect equilibrium (MPE) as the solution concept, with the payoff-relevant states being beliefs. This means that the platform's and the user's strategies at each time are functions of the state at that time and are optimal given the strategy of the other player and the evolution of beliefs.

In the post-AI era, the platform possesses superior knowledge about the initial glossiness state $\alpha_{i,0}$, which enables a more accurate estimation of $\alpha_{i,t}$ than the user for any t. In this environment, we again use the notion of MPE, except that, in this case, each player evaluates payoffs according to their beliefs, which differ because they have access to different information. This makes our equilibrium notion the analogue of the Berk–Nash equilibrium of Esponda and Pouzo (2016) in this setting (see also Eyster and Rabin 2005 and Esponda 2008). In particular, the platform's decision at time t, $\big\{\big(x_{i,t}, p_{i,t}\big)\big\}_{i\in\mathcal{N}}$, contains additional information about the quality

of the product and its glossiness state, but the user does not recognize this superior information and does not update her beliefs based on the platform's offers. We view this behavioral aspect as plausible since users do not fully understand the superior information of platforms stemming from advances in big data and AI.[4]

II. Belief Dynamics and Equilibrium Characterization

A. *Belief Dynamics*

Pre-AI Environment.—At any 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 low quality. Otherwise, their belief is

$$\mu_{i,t} = \Pr\big(\theta_i = 1 \,|\, NBN_{i,t}\big).$$

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:

$$\lambda_{i,t} = \Pr\big(\alpha_{i,t} = 1 \,|\, \theta_i = 0, NBN_{i,t}\big).$$

Standard arguments imply that for any $\{x_{i,t}, p_{i,t}, b_{i,t}\}_{i\in\mathcal{N},t}$, in the pre-AI environment, common beliefs evolve according to[5]

$$(2) \qquad d\mu_{i,t} = x_{i,t} b_{i,t} \mu_{i,t}\big(1 - \mu_{i,t}\big)\big(1 - \lambda_{i,t}\big)\gamma dt \quad \text{with initialization} \quad \mu_{i,0}$$

$$d\lambda_{i,t} = x_{i,t} b_{i,t} \lambda_{i,t}\Big[\big(1 - \lambda_{i,t}\big)\gamma - \rho\Big]dt \quad \text{with initialization} \quad \lambda_{i,0} = \lambda.$$

These beliefs highlight that the user is Bayesian and knows about the (existence of) glossiness state $\alpha_{i,t}$. The probability of receiving bad news for product i at time t, given $NBN_{i,t}$, is

$$\Pr\big(I_{i,t+dt} = 0 \,|\, I_{i,t} = 1\big) = x_{i,t} b_{i,t}\big(1 - \mu_{i,t}\big)\big(1 - \lambda_{i,t}\big)\gamma dt.$$

Post-AI Environment.—User belief without bad news is

$$\mu_{i,t} = \Pr\big(\theta_i = 1 \,|\, NBN_{i,t}\big).$$

[4] (Nonbehavioral) 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.

[5] For completeness, we derive these standard equations in the Supplemental Appendix.

The main difference from the pre-AI case is that the platform now knows the initial glossiness state $\alpha_{i,0}$ of all products and therefore has different beliefs about product qualities:

$$\mu_{i,t}^{(P)} = \Pr\big(\theta_i = 1 \,|\, NBN_{i,t}, \alpha_{i,0}\big),$$

where the superscript P signifies that this is the platform's belief. Likewise,

$$\lambda_{i,t}^{(P)} = \Pr\big(\alpha_{i,t} = 1 \,|\, NBN_{i,t}, \alpha_{i,0}\big).$$

The user's beliefs are the same as in (2), and the platform's beliefs now evolve:

$$\begin{aligned}(3)\quad &\text{When } \alpha_{i,0} = 1,\ \mu_{i,t}^{(P)} = 0 \text{ and } d\lambda_{i,t}^{(P)} = x_{i,t} b_{i,t} \lambda_{i,t}^{(P)} \left[(1-\lambda_{i,t}^{(P)})\gamma - \rho\right] dt \\ &\quad \text{with } \lambda_{i,0}^{(P)} = 1 \\ &\quad \text{when } \alpha_{i,0} = 0,\ d\mu_{i,t}^{(P)} = x_{i,t} b_{i,t} \mu_{i,t}^{(P)} (1-\mu_{i,t}^{(P)})\,\gamma dt \\ &\quad \text{with } \mu_{i,0}^{(P)} = \frac{\mu_{i,0}}{\mu_{i,0} + (1-\mu_{i,0})(1-\lambda)} \text{ and } \lambda_{i,t}^{(P)} = 0.\end{aligned}$$

From the platform's perspective, the probability of receiving bad news for product i at time t, given $NBN_{i,t}$, is

$$\Pr\big(I_{i,t+dt} = 0 \,|\, I_{i,t} = 1, \alpha_{i,0} = 0\big) = x_{i,t} b_{i,t} \Big(1 - \mu_{i,t}^{(P)}\Big)\gamma dt \quad \text{for } \alpha_{i,0} = 0$$

$$\Pr\big(I_{i,t+dt} = 0 \,|\, I_{i,t} = 1, \alpha_{i,0} = 1\big) = x_{i,t} b_{i,t} \Big(1 - \lambda_{i,t}^{(P)}\Big)\gamma dt \quad \text{for } \alpha_{i,0} = 1.$$

B. *Preliminary Equilibrium Characterization*

We first prove a simple lemma that characterizes the platform's equilibrium pricing and the user's equilibrium purchasing strategies.

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

$$E\big[\theta_i \,|\, NBN_{i,t}\big] - p_{i,t} = \mu_{i,t} - p_{i,t}.$$

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

$$p_{i,t} = E\big[\theta_i \,|\, NBN_{i,t}\big] = \mu_{i,t}.$$

For a product i at time t for which bad news has arrived ($I_{i,t} = 0$), *the equilibrium price is zero.*[6]

[6] All equalities should be read as "almost everywhere" since two stochastic processes for, say, prices that differ on measure zero sets are equivalent. We do not specify this additional qualification to simplify the exposition.

PROOF OF LEMMA 1:

The proof of the lemma follows from three claims.

CLAIM 1: *The user accepts any price* $p_{i,t} \leq \mu_{i,t}$.

To see this, consider a deviation by the user of rejecting prices $p_{i,t} \leq \mu_{i,t}$ for $t \in [T, T+\Delta t]$, where $\Delta t > 0$. Without a purchase, beliefs at $T+\Delta t$ are the same as in T, and from the one-shot deviation principle, the continuation utility from $T+\Delta t$ onward remains unchanged, so the deviation cannot improve the user utility for any Δt.

CLAIM 2: *The platform offers* $p_{i,t} \geq \mu_{i,t}$ *for all t.*

Suppose, to obtain a contradiction, that the platform offers $p_{i,t} < \mu_{i,t}$ for some interval $[T, T+\Delta t]$. Consider a deviation such that the platform raises prices to $p_{i,t} + \varepsilon_{i,t}$ for $t \in [T, T+\Delta t]$, with $p_{i,t} + \varepsilon_{i,t} \leq \mu_{i,t}$, and this offer would be accepted. Once again, this deviation would lead to the same beliefs and continuation utility, and because it has greater revenues during the interval $t \in [T, T+\Delta t]$, it is profitable, yielding a contradiction.

CLAIM 3: *The platform offers* $p_{i,t} = \mu_{i,t}$ *for all t.*

Fix some time T. From Claim 2, we have $p_{i,t} \geq \mu_{i,t}$ for all $t \geq T+\Delta t$, implying a nonpositive continuation utility for the user from $T+\Delta t$ onward, and thus the user anticipates this and she will reject any $p_{i,t} > \mu_{i,t}$ for $t \in [T, T+\Delta t]$. Therefore, the platform's equilibrium strategy is $p_{i,t} = \mu_{i,t}$ for $t \in [T, T+\Delta t]$. ∎

We next provide expressions for the platform and the user utilities.

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

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

and the platform's pricing strategy is $p_{i,t}^{pre\text{-}AI} = \mu_{i,t}$ (see Lemma 1). Here, the subscript $P,0$ denotes that the expectation is with respect to the platform's information

at time 0, and $\Pi^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right)$ denotes the platform's payoff starting from the user's belief $\{\mu_i\}_{i=1}^n$. The user's purchasing decision $b_{i,t}$ does not appear in the above expression because the platform always offers a price equal to the user's belief, and the user always accepts (Lemma 1).

We evaluate the user's utility according to the expectations of "true events." Given the platform's equilibrium strategy, the expectation of the user's utility is

$$U^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right) = E_0\left[\int_0^\infty re^{-rt}\sum_{i=1}^n x_{i,t}^{pre\text{-}AI}(\theta_i - \mu_{i,t})dt\right]. \tag{5}$$

The expected utilitarian welfare is the sum of the user's utility given in (5) and the platform's payoff given in (4):

$$W^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right) = E_0\left[\int_0^\infty re^{-rt}\sum_{i=1}^n x_{i,t}^{pre\text{-}AI}\theta_i dt\right].$$

Post-AI Environment.—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 $\left\{x_{i,t}^{post\text{-}AI}\right\}_{i\in\mathcal{N},t\in\mathbb{R}_+}$ solves

$$\begin{aligned}
\Pi^{post\text{-}AI}\left(\{\mu_i,\alpha_i\}_{i=1}^n\right) &= \max_{\{x_{i,t}\}_{i\in\mathcal{N},t\in\mathbb{R}_+}} E_{P,0}\left[\int_0^\infty re^{-rt}\sum_{i=1}^n x_{i,t}\mu_{i,t}dt\right] \qquad (6)\\
d\mu_{i,t} &= I_{i,t}x_{i,t}\mu_{i,t}(1-\mu_{i,t})(1-\lambda_{i,t})\gamma dt \quad \mu_{i,0} = \mu_i\\
d\lambda_{i,t} &= I_{i,t}x_{i,t}\lambda_{i,t}\left[(1-\lambda_{i,t})\gamma - \rho\right]dt \quad \lambda_{i,0} = \lambda\\
d\mu_{i,t}^{(P)} &= I_{i,t}x_{i,t}\mu_{i,t}^{(P)}\left(1-\mu_{i,t}^{(P)}\right)\gamma dt \text{ and } \lambda_{i,t}^{(P)} = 0 \text{ if } \alpha_i = 0\\
\mu_{i,t}^{(P)} &= 0 \text{ and } d\lambda_{i,t}^{(P)} = I_{i,t}x_{i,t}\lambda_{i,t}^{(P)}\left[\left(1-\lambda_{i,t}^{(P)}\right)\gamma - \rho\right]dt\\
&\text{if } \alpha_i = 1\\
\mu_{i,t} &= \mu_{i,t}^{(P)} = 0 \text{ if } I_{i,t} = 0\\
\Pr\left(I_{i,t+dt} = 0 \,|\, I_{i,t} = 1, \alpha_i = 0\right) &= x_{i,t}b_{i,t}\left(1-\mu_{i,t}^{(P)}\right)\gamma dt\\
\Pr\left(I_{i,t+dt} = 0 \,|\, I_{i,t} = 1, \alpha_i = 1\right) &= x_{i,t}b_{i,t}\left(1-\lambda_{i,t}^{(P)}\right)\gamma dt \text{ with } I_{i,0} = 1,
\end{aligned}$$

and the platform's pricing strategy in equilibrium is $p_{i,t}^{post\text{-}AI} = \mu_{i,t}$ (see Lemma 1), and moreover,

$$U^{post\text{-}AI}\left(\{\mu_i,\alpha_i\}_{i=1}^n\right) = E_0\left[\int_0^\infty re^{-rt}\sum_{i=1}^n x_{i,t}^{post\text{-}AI}(\theta_i - \mu_{i,t})dt\right]. \tag{7}$$

Similar to the pre-AI environment, the expected utilitarian welfare is the sum of user utility in (7) and platform profits in (6).

III. Equilibrium Characterization in the Pre-AI Environment

The pre-AI equilibrium is characterized in the following theorem.

THEOREM 1: *Without loss of generality, let* $\mu_{1,0} \geq \ldots \geq \mu_{n,0}$ *be the initial beliefs for the n products. Then, in the pre-AI environment, the platform's equilibrium decision is to offer product 1 until bad news occurs for this product (i.e.,* $x_{1,t}^{\text{pre-AI}} = 1$ *for all* $t \in [0, \tau_1)$ *where* $\tau_1 \in \mathbb{R}_+ \cup \{\infty\}$ *is the stochastic time at which bad news for product 1 occurs) and then to offer product 2 and so on.*

PROOF:

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 being glossy, given that no bad news has occurred. For any product $i \in \mathcal{N}$, we introduce the Gittins index:

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

with the same evolution as the ones in (4) and where the supremum is over all stopping times. Using the Gittins index theorem (see, e.g., Gittins, Glazebrook, and Weber 2011; Chapter 2), the platform always chooses the product with the highest Gittins index. Next, note that $\mu_{i,t}$ is increasing in t, and thus the optimal stopping time is to offer product i until bad news occurs for it and that, absent bad news, the Gittins index increases. Moreover, products with higher beliefs have initially higher Gittins indices. This follows because if $\mu_{i,0} \geq \mu_{i',0}$, then $\mu_{i,t} \geq \mu_{i',t}$ for all t and the first time for which bad news occurs for i first-order stochastically dominates that of i'. Therefore, the platform offers the product with the highest belief until bad news occurs for this product, at which point it switches to the product with the second-highest belief and so on. ∎

IV. Equilibrium Characterization in the Post-AI Environment

The next theorem characterizes the post-AI equilibrium.

THEOREM 2: *Without loss of any generality, let* $\mu_{i_1,0} \geq \ldots \geq \mu_{i_{n_0},0}$ *be the initial beliefs for initially nonglossy products and* $\mu_{j_1,0} \geq \ldots \geq \mu_{j_{n_1},0}$ *be the initial beliefs for the initially glossy products. In the post-AI environment, the platform's equilibrium strategy involves first offering either* i_1 *or* j_1*. If the platform first offers* i_1 (j_1)*, the platform continues to offer this product until bad news occurs. The platform then offers either* i_2 *or* j_1 $(i_1$ *or* $j_2)$ *and so on.*

The proof of this theorem is similar to that of Theorem 1 and is omitted. The key aspect of the platform's decision that concerns the choice of i_1 and j_1 is studied in our next theorem.

THEOREM 3: *Consider two products* $i,j \in \mathcal{N}$ *in the post-AI environment with* $\alpha_{j,0} = 1$ *and* $\alpha_{i,0} = 0$*.*

(*i*) (*Manipulation Effect*) *There exists a constant* $\underline{\rho}$ *and a function* $\mu_{i,0} \mapsto f(\mu_{i,0}; \lambda, \gamma, r)$ *parameterized by* λ, γ, r *that is everywhere below* $\mu_{i,0}$ *such that for* $\rho \leq \underline{\rho}$ *and*

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

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

(*ii*) (*Helpfulness Effect*) *There exists a constant* $\bar{\rho}$ *and a function* $\mu_{j,0} \mapsto g(\mu_{j,0}; \lambda, \gamma, r)$ *parametrized by* λ, γ, r *that is everywhere below* $\mu_{j,0}$ *such that for* $\rho \geq \bar{\rho}$ *and*

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

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

PROOF:

Let $\mu_1 = \mu_{j,0}$ and $\mu_0 = \mu_{i,0}$, and $U_{0,1}(\rho, \mu_0, \mu_1)$ and $U_{1,0}(\rho, \mu_0, \mu_1)$ be the platform payoffs from offering these two products. Then

$$\lim_{\rho \to 0} U_{1,0}(\rho, \mu, \mu) - U_{0,1}(\rho, \mu, \mu) = \left(\int_0^\infty re^{-rt} a(t) dt\right)\left(\int_0^\infty re^{-rt} b(t)\, dt\right) - \int_0^\infty re^{-rt} a(t) b(t)\, dt,$$

where $a(t) = \frac{\mu\gamma}{\mu\gamma + (1-\mu)(\lambda\gamma + e^{-\gamma t}(1-\lambda)\gamma)}$ and $b(t) = \frac{\mu + e^{-\gamma t}(1-\lambda)(1-\mu)}{1 - \lambda(1-\mu)}$. The above expression is strictly positive from the FKG–Harris inequality (see, e.g., Alon and Spencer 2016; Chapter 6.2), which states that if $a{:}\mathbb{R}_+ \to \mathbb{R}_+$ is an increasing and $b{:}\mathbb{R}_+ \to \mathbb{R}_+$ is a decreasing function, then $E[a(t)b(t)] \leq E[a(t)]E[b(t)]$ for random variable t distributed over $\mathbb{R}_+$ (and in our setting the density is $\Pr(t = s) = re^{-rs}$). Moreover, this inequality is strict if both functions are strictly monotone and X has a strictly positive mass everywhere. This, together with the continuity of $U_{0,1}(\cdot)$ and $U_{1,0}(\cdot)$, proves the first part. The second part is proved analogously by considering the same functions for the limit $\rho \to \infty$ and again invoking the FKG–Harris inequality. ∎

The platform offers a low-quality product with a lower user belief because of a *manipulation effect*: when $\alpha_{j,0} = 1$, 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 the lower-quality product j. This is countered by a *helpfulness effect*: when $\alpha_{j,0} = 1$, the platform knows that the quality is low, and eventually, the user receives bad news. This force incentivizes the platform not to offer product j. For sufficiently large ρ, the helpfulness effect dominates, and the platform guides users away from low-quality products. When ρ is sufficiently small, then behavioral manipulation can go on for a long time, and the platform favors offering the low-quality product j.[7]

[7] Note that even if $\mu_{j,0} > \mu_{i,0}$, so long as the belief $\mu_{i,0}$ is not too small (i.e., $\mu_{i,0} > g(\mu_{i',0}; \lambda, \gamma, r)$), the platform prefers offering product i instead of product j. A similar point applies in the case where the manipulation effect dominates. We also note that both the f and g functions can be written explicitly in terms of the primitives $(\mu, \lambda, \gamma, r)$.

V. Who Benefits from the Platform's Superior Information?

A. *When Users Benefit from the Platform's Superior Information*

Our next theorem confirms what might be viewed as the conventional wisdom in the literature: more data enables better product allocation and benefits users.

THEOREM 4: *Suppose that 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*

$$E\left[\Pi^{post\text{-}AI}\left(\{\mu_i, \alpha_i\}_{i=1}^n\right)\right] > E\left[\Pi^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right)\right]$$

$$E\left[U^{post\text{-}AI}\left(\{\mu_i, \alpha_i\}_{i=1}^n\right)\right] > E\left[U^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right)\right]$$

$$E\left[W^{post\text{-}AI}\left(\{\mu_i, \alpha_i\}_{i=1}^n\right)\right] > E\left[W^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right)\right],$$

where the expectations are over $\mu_i \sim \text{Unif}[0,1]$, $\theta_i \sim \text{Bern}(\mu_i)$, *and* α_i *drawn according to (1) for all* $i \in \mathcal{N}$.

PROOF:

The platform's information advantage in the post-AI implies the first inequality. We next prove the second inequality, from which the third follows immediately. Without loss of generality, consider $\mu_{1,0} \leq \ldots \leq \mu_{n,0}$ as a realization of the products' beliefs. Theorem 1 implies that in the pre-AI environment, the platform first offers product n, and subsequently, after bad news for this product, it offers product $n-1$ and so on.

In contrast to the pre-AI environment, Theorem 3 shows that the platform will no longer offer products in sequence in the post-AI environment. Nevertheless, we can characterize the platform's optimal strategy and its welfare consequences by considering a series of *helpful swaps*: Suppose that i is the highest index for which $\alpha_{i,0} = 0$, and 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. We prove in the Supplemental Appendix that expected user welfare increases from helpful swaps provided that ρ is large enough. Specifically, let product n have the highest belief at time 0 and suppose that $\alpha_{n,0} = 1$. Let i be the largest index for which $\alpha_{i,0} = 0$. Then, with positive probability we have $\mu_{i,0} > g(\mu_{n,0}, \lambda, \gamma, r)$ and the platform can profitably offer product i in the post-AI, which improves the user's expected utility. ∎

The informational advantage of the platform always increases its own profits, as it enables the platform to modify the user's behavior. Theorem 4 establishes that this informational advantage also increases the expected user's utility and welfare for large enough ρ. This is again because, for large enough ρ, the helpfulness effect dominates the manipulation effect (characterized in Theorem 3).

B. *When Behavioral Manipulation Harms Users*

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

THEOREM 5: *Suppose that 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*

$$E\left[\Pi^{post\text{-}AI}\left(\{\mu_i, \alpha_i\}_{i=1}^n\right)\right] > E\left[\Pi^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right)\right]$$

$$E\left[U^{post\text{-}AI}\left(\{\mu_i, \alpha_i\}_{i=1}^n\right)\right] < E\left[U^{pre\text{-}AI}\left(\{\mu_i\}_{i=1}^n\right)\right]$$

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

PROOF:

The first inequality is immediate. We prove the third inequality, which, together with the first inequality, also implies the second inequality. Given the continuity of the utilities, it suffices to prove that in the limit when $\rho \to 0$, the difference between its right-hand side and left-hand side is bounded by some strictly positive quantity (that does not depend on ρ). Consider a realization of the products' beliefs and, without loss of generality, assume that

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

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

(i) $\alpha_{n,0} = 1$: In both environments, the platform starts offering product n. Since $\rho = 0$, bad news does not arrive for this product, and the expected welfare in both environments is the same.

(ii) $\alpha_{n,0} = 0$ and for all products $j \in \mathcal{N}\backslash\{n\}$ for which $\alpha_{j,0} = 1$, we have $\mu_{j,0} \leq f(\mu_{n,0}, \lambda, \gamma, r)$: In both environments, 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.

(iii) $\alpha_{n,0} = 0$ and there exists product $j \in \mathcal{N}\backslash\{n\}$ for which $\alpha_{j,0} = 1$ and $\mu_{j,0} > f(\mu_{n,0}, \lambda, \gamma, r)$: Letting 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, the platform starts by offering product n, and therefore, with a positive probability, the welfare is strictly positive.

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

With a low ρ, the user is pushed toward initially glossy products. 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 profits. Because glossy products are of low quality, such behavioral manipulation is bad for the user utility and utilitarian welfare.

We emphasize that the uniform assumption for the initial beliefs in Theorems 4 and 5 is without loss of generality: the comparisons hold with weak inequalities for any fixed initial beliefs, and the strict inequalities hold whenever the platform's equilibrium strategies in the two environments are different.

C. *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 increases as well.

THEOREM 6: *Suppose that the initial beliefs* $\{\mu_i\}_{i=1}^{n+1}$ *are i.i.d. and uniform over* $\left[\frac{1}{2}-\Delta,\frac{1}{2}+\Delta\right]$. *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* $n \geq \left\lceil 1-\frac{\log\lambda}{\log(2-\lambda)/(1-\lambda)}\right\rceil$ *in the post-AI environment we have*:

$$E\left[\Pi^{\textit{post-AI}}\left(\{\mu_i,\alpha_i\}_{i=1}^{n+1}\right)\right] > E\left[\Pi^{\textit{post-AI}}\left(\{\mu_i,\alpha_i\}_{i=1}^{n}\right)\right]$$

$$E\left[U^{\textit{post-AI}}\left(\{\mu_i,\alpha_i\}_{i=1}^{n+1}\right)\right] < E\left[U^{\textit{post-AI}}\left(\{\mu_i,\alpha_i\}_{i=1}^{n}\right)\right]$$

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

PROOF:

Similar to the argument of Theorem 5, it suffices to prove the third inequality. We let $\Delta \leq \Delta_1$, where Δ_1 is small enough so that

$$f\left(\frac{1}{2}+\Delta_1;\lambda,\gamma,r\right) \leq \frac{1}{2}-\Delta_1. \tag{8}$$

Such Δ_1 exists because $f(\mu;\lambda,\gamma,r)$ is everywhere below μ (and continuous). Conditioned on the belief realizations, we have the following cases:

(i) There exists $i \in \mathcal{N}$ with $\alpha_{i,0} = 1$: Given $\Delta \leq \Delta_1$ and (8), the platform's equilibrium strategy with and without the $n+1$th product is to offer one of the initially glossy products. Therefore, the expected welfare with and without the extra product is 0.

(ii) $\alpha_{i,0} = 0$ for all $i \in \mathcal{N}$ and $\alpha_{n+1,0} = 1$: Given $\Delta \leq \Delta_1$ and (8), the platform's equilibrium strategy with the extra product is to offer that product

whose expected welfare is 0. The platform's equilibrium strategy with n products, however, is to offer them in decreasing order of their belief.

(iii) $\alpha_{i,0} = 0$ for all $i \in \{1, \ldots, n\}$ and $\alpha_{n+1,0} = 0$: The platform's equilibrium strategy with and without the extra product is to offer them in decreasing order of their belief.

With the extra product, the expected welfare in case (ii) increases, while the expected welfare in case (iii) decreases. The difference between expected welfare with and without the extra product in case (ii) is upper bounded by

$$-\mu_l \frac{1 - \left[(1 - \mu_l)\delta\right]^n}{1 - (1 - \mu_l)\delta},$$

and in case (iii) is upper bounded by

$$\mu_h \frac{1 - \left[(1 - \mu_h)\delta\right]^{n+1}}{1 - (1 - \mu_h)\delta} - \mu_l \frac{1 - \left[(1 - \mu_l)\delta\right]^n}{1 - (1 - \mu_l)\delta},$$

where $\mu_l = \frac{\frac{1}{2} - \Delta_2}{\frac{1}{2} - \Delta_2 + \left(\frac{1}{2} + \Delta_2\right)(1 - \lambda)}$ and $\mu_h = \frac{\frac{1}{2} + \Delta_2}{\frac{1}{2} + \Delta_2 + \left(\frac{1}{2} - \Delta_2\right)(1 - \lambda)}$ are the smallest and largest beliefs conditional on the initial glossiness state being 0 (obtained by using Bayes' rule), respectively. Using the above expressions, it is straightforward to see that for small enough Δ, expected welfare loss in case (iii) dominates. ∎

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 that $\Delta \leq \bar{\Delta}$ and n is large are added 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 offer initially glossy items. 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 because they enable the platform to engage in more intense behavioral manipulation.

VI. Conclusion and Policy Implications

This paper has argued that the vast amounts of data collected by online platforms may also enable behavioral manipulation, harming users. We view our paper as a first step in the systematic analysis of product choice and consumer experimentation in online markets in which platforms have extensive and growing information about user preferences (and biases).

Our analysis raises the natural question of whether regulations can be used to combat behavioral manipulation. Three different sets of policies could be useful:

- *Data minimization*: Companies could be mandated to collect only the data necessary for a specific, declared purpose and prohibit the collection and use of data for other reasons unless explicit consent is provided.
- *Transparency*: Users could be provided with more information about how their data are being used. Such transparency measures could be augmented with programs that educate consumers about how to understand platform strategies and protect their data.
- *Competition*: Legal and regulatory measures could be adopted to support the entry of new platforms that could credibly commit to different data harvesting and algorithmic pricing policies.

There are several other interesting questions for future research, which we list briefly:

- Can behavioral manipulation persist in the long run even as users become accustomed to the big data environment?
- Can AI tools be deployed for correcting, rather than exploiting, behavioral biases of users?
- How does dynamic learning by platforms (rather than estimating product characteristics perfectly, as we have assumed) affect the scope for behavioral manipulation?
- How does the presence of users with different levels of sophistication affect platform behavior and welfare?

REFERENCES

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. 2022a. "Learning from Reviews: The Selection Effect and the Speed of Learning." *Econometrica* 90 (6): 2857–99.

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar. 2022b. "Too Much Data: Prices and Inefficiencies in Data Markets." *American Economic Journal: Microeconomics* 14 (4): 218–56.

Acquisti, Alessandro, Laura Brandimarte, and George Loewenstein. 2015. "Privacy and Human Behavior in the Age of Information." *Science* 347 (6221): 509–14.

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman. 2016. "The Economics of Privacy." *Journal of Economic Literature* 54 (2): 442–92.

Acquisti, Alessandro, and Hal R. Varian. 2005. "Conditioning Prices on Purchase History." *Marketing Science* 24 (3): 367–81.

Admati, Anat R., and Paul Pfleiderer. 1986. "A Monopolistic Market for Information." *Journal of Economic Theory* 39 (2): 400–438.

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston: Harvard Business Review Press.

Alon, Noga, and Joel H. Spencer. 2016. *The Probabilistic Method*. 4th ed. Hoboken, NJ: John Wiley & Sons.

Anton, James J., and Dennis A. Yao. 2002. "The Sale of Ideas: Strategic Disclosure, Property Rights, and Contracting." *Review of Economic Studies* 69 (3): 513–31.

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp. 2018. "Big Data in Finance and the Growth of Large Firms." *Journal of Monetary Economics* 97: 71–87.

Bergemann, Dirk, and Alessandro Bonatti. 2015. "Selling Cookies." *American Economic Journal: Microeconomics* 7 (3): 259–94.

Bergemann, Dirk, and Alessandro Bonatti. 2019. "Markets for Information: An Introduction." *Annual Review of Economics* 11: 85–107.

Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin. 2018. "The Design and Price of Information." *American Economic Review* 108 (1): 1–48.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris. 2015. "The Limits of Price Discrimination." *American Economic Review* 105 (3): 921–57.

Bergemann, Dirk, and Juuso Välimäki. 2000. "Experimentation in Markets." *Review of Economic Studies* 67 (2): 213–34.

Bolton, Patrick, and Christopher Harris. 1999. "Strategic Experimentation." *Econometrica* 67 (2): 349–74.

Bonatti, Alessandro, and Gonzalo Cisternas. 2020. "Consumer Scores and Price Discrimination." *Review of Economic Studies* 87 (2): 750–91.

Bradbury, John. 2024. "5 Things the Google Play Store Desperately Needs to Improve." *Android Police*, May 2, 2024. https://www.androidpolice.com/things-the-google-play-store-desperately-needs-to-improve/.

Calo, Ryan. 2014. "Digital Market Manipulation." *George Washington Law Review* 82 (4): 995–1051.

Calvo, Rafael A., Dorian Peters, Karina Vold, and Richard M. Ryan. 2020. "Supporting Human Autonomy in AI Systems: A Framework for Ethical Enquiry." In *Ethics of Digital Well-Being: A Multidisciplinary Approach*, Vol. 140, edited by Christopher Burr and Luciano Floridi, 31–54. Cham, Switzerland: Springer.

Che, Yeon-Koo, and Johannes Hörner. 2018. "Recommender Systems as Mechanisms for Social Learning." *Quarterly Journal of Economics* 133 (2): 871–925.

Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim. 2019. "Privacy and Personal Data Collection with Information Externalities." *Journal of Public Economics* 173: 113–24.

Costes, Jean-Michel, and Céline Bonnaire. 2022. "Spending Money in Free-to-Play Games: Sociodemographic Characteristics, Motives, Impulsivity and Internet Gaming Disorder Specificities." *International Journal of Environmental Research and Public Health* 19 (23): 15709.

Ekmekci, Mehmet. 2011. "Sustainable Reputations with Rating Systems." *Journal of Economic Theory* 146 (2): 479–503.

Esponda, Ignacio. 2008. "Behavioral Equilibrium in Economies with Adverse Selection." *American Economic Review* 98 (4): 1269–91.

Esponda, Ignacio, and Demian Pouzo. 2016. "Berk–Nash Equilibrium: A Framework for Modeling Agents with Misspecified Models." *Econometrica* 84 (3): 1093–130.

Eyster, Erik, and Matthew Rabin. 2005. "Cursed Equilibrium." *Econometrica* 73 (5): 1623–72.

Fainmesser, Itay P., Andrea Galeotti, and Ruslan Momot. 2023. "Digital Privacy." *Management Science* 69 (6): 3157–73.

Felli, Leonardo, and Christopher Harris. 1996. "Learning, Wage Dynamics, and Firm-Specific Human Capital." *Journal of Political Economy* 104 (4): 838–68.

Gittins, John, Kevin Glazebrook, and Richard Weber. 2011. *Multi-armed Bandit Allocation Indices*. Hoboken, NJ: John Wiley & Sons.

Goldfarb, Avi, and Catherine Tucker. 2011. "Online Display Advertising: Targeting and Obtrusiveness." *Marketing Science* 30 (3): 389–404.

Gradwohl, Ronen. 2017. "Information Sharing and Privacy in Networks." In *EC '17: Proceedings of the 2017 ACM Conference on Economics and Computation*, edited by Moshe Babaioff, Hervé Moulin, and Constantinos Daskalakis, 349–50. New York: Association for Computing Machinery.

Hanson, Jon D., and Douglas A. Kysar. 1999. "Taking Behavioralism Seriously: The Problem of Market Manipulation." *New York University Law Review* 74 (3): 630–749.

Hidir, Sinem, and Nikhil Vellodi. 2019. "Personalization, Discrimination and Information Disclosure." Unpublished.

Hörner, Johannes, and Andrzej Skrzypacz. 2016. "Selling Information." *Journal of Political Economy* 124 (6): 1515–62.

Ichihashi, Shota. 2020. "Online Privacy and Information Disclosure by Consumers." *American Economic Review* 110 (2): 569–95.

Ichihashi, Shota. 2023. "Dynamic Privacy Choices." *American Economic Journal: Microeconomics* 15 (2): 1–40.

Jullien, Bruno, Yassine Lefouili, and Michael H. Riordan. 2020. "Privacy Protection, Security, and Consumer Retention." Toulouse School of Economics Working Paper 18-947.

Keller, Godfrey, and Sven Rady. 2010. "Strategic Experimentation with Poisson Bandits." *Theoretical Economics* 5 (2): 275–311.

Keller, Godfrey, Sven Rady, and Martin Cripps. 2005. "Strategic Experimentation with Exponential Bandits." *Econometrica* 73 (1): 39–68.

Köbis, Nils, Jean-François Bonnefon, and Iyad Rahwan. 2021. "Bad Machines Corrupt Good Morals." *Nature Human Behaviour* 5 (6): 679–85.

Lin, Tom C. W. 2017. "The New Market Manipulation." *Emory Law Journal* 66 (6): 1253–314.

Lindecrantz, Erik, Madeleine Tjon Pian Gi, and Stefano Zerbi. 2022. "Personalizing the Customer Experience: Driving Differentiation in Retail." *McKinsey & Company*, April 28, 2020. https://www.mckinsey.com/industries/retail/our-insights/personalizing-the-customer-experience-driving-differentiation-in-retail.

Liu, John Zhuang, Michael Sockin, and Wei Xiong. 2023. "Data Privacy and Algorithmic Inequality." Unpublished.

Martin, Kirsten, and Helen Nissenbaum. 2017. "Measuring Privacy: An Empirical Test Using Context to Expose Confounding Variables." *Columbia Science and Technology Law Review* 18 (1): 176–218.

McMorrow, Ryan. 2023. "Andrew Ng: 'Do We Think the World Is Better Off with More or Less Intelligence?'" *Financial Times*, December 19, 2023. https://www.ft.com/content/2dc07f9e-d2a9-4d98-b746-b051f9352be3.

Montes, Rodrigo, Wilfried Sand-Zantman, and Tommaso Valletti. 2019. "The Value of Personal Information in Online Markets with Endogenous Privacy." *Management Science* 65 (3): 1342–62.

Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge, MA: Harvard University Press.

Perrigo, Billy. 2024. "Yann Lecun Is Optimistic That AI Will Lead to a Better World." *Time*, February 11, 2024. https://time.com/collection/time100-impact-awards/6692039/yann-lecun-meta-artificial-intelligence-time-award/.

Rothschild, Michael. 1974. "A Two-Armed Bandit Theory of Market Pricing." *Journal of Economic Theory* 9 (2): 185–202.

Ryan-Mosley, Tate. 2023. "Junk Websites Filled with AI-Generated Text Are Pulling In Money from Programmatic Ads." *MIT Technology Review*, June 26, 2023. https://www.technologyreview.com/2023/06/26/1075504/junk-websites-filled-with-ai-generated-text-are-pulling-in-money-from-programmatic-ads/.

Susser, Daniel, and Vincent Grimaldi. 2021. "Measuring Automated Influence: Between Empirical Evidence and Ethical Values." In *AIES '21: Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, edited by Marion Fourcade, Benjamin Kuipers, Seth Lazar, and Deirdre Mulligan, 242–53. New York: Association for Computing Machinery.

Taylor, Curtis R. 2004. "Consumer Privacy and the Market for Customer Information." *RAND Journal of Economics* 35 (4): 631–50.

Tirole, Jean. 2021. "Digital Dystopia." *American Economic Review* 111 (6): 2007–48.

Tirole, Jean. 2023. "Competition and the Industrial Challenge for the Digital Age." *Annual Review of Economics* 15: 573–605.

Varian, Hal. 2018. "Artificial Intelligence, Economics, and Industrial Organization." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, 399–419. Chicago: University of Chicago Press.

Wang, Haoyu, Zhe Liu, Jingyue Liang, Narseo Vallina-Rodriguez, Li Yao Guo, Juan Tapiador, Jingcun Cao, and Guoai Xu. 2018. "Beyond Google Play: A Large-Scale Comparative Study of Chinese Android App Markets." In *IMC '18: Proceedings of the Internet Measurement Conference 2018*, 293–307. New York: Association for Computing Machinery.

Zarsky, Tal Z. 2019. "Privacy and Manipulation in the Digital Age." *Theoretical Inquiries in Law* 20 (1): 157–88.

Zuboff, Shoshana. 2019. *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. London: Profile Books.