

ARTICLE OPEN ACCESS

Inflated Recommendations

Martin Peitz^{1,2} | Anton Sobolev¹

¹Department of Economics and MaCCI, University of Mannheim, Mannheim, Germany | ²Also affiliated with CEPR, London, United Kingdom and CESifo, Munich, Germany

Correspondence: Martin Peitz (Martin.Peitz@gmail.com)

Received: 6 November 2025 | **Accepted:** 6 November 2025

Keywords: asymmetric information | experience good | intermediation | recommendation bias | recommender system

ABSTRACT

Biased recommendations arise naturally in markets with heterogeneous consumers. We study a model in which a monopolist offers an experience good to a population of consumers with heterogeneous tastes and makes personalized purchase recommendations. We provide conditions under which a firm makes welfare-reducing purchase recommendations with positive probability, resulting in inflated recommendations. We extend this insight to a setting in which an intermediary makes the recommendations, whereas a seller sets the retail price. Regulatory interventions that forbid inflated recommendations may lead to higher social welfare or may backfire.

JEL Classification: L12, L15, D21, D42, M37

1 | Introduction

A quote attributed to the 15th century monk and poet John Lydgate says, “You can’t please all of the people all of the time.” With the advance of consumer tracking and recommendation algorithms, it is now possible for a firm to carefully and deliberately select its target audience and to provide recommendations that fit an individual consumer’s taste. People may differ in how they intend to use a product; firms (including intermediaries) may therefore help people to identify the products that fit a specific purpose.¹

We develop a parsimonious model in which a profit-maximizing firm decides whether or not to recommend a new product to a consumer. For example, an e-commerce marketplace such as Amazon decides whether to assign the “buy” button to a seller or a website decides whether or not a particular consumer is shown an “editor’s pick” for a new product. Similarly, consumers typically receive a single recommendation if recommendations are provided by voice (as with virtual assistants, such as Alexa, Cortana, Google Assistant, and Siri). More broadly, a firm may

increase the visibility of certain offers while reducing the visibility of others.

Recommendation algorithms ultimately serve the interests of the firm providing the recommendation service.² We show that a firm catering to a diverse set of consumers may decide to recommend the product to some consumers with a “bad” match and, thus, inflate recommendations. This reduces the heterogeneity of expected consumer valuations conditional on receiving a recommendation and allows the firm to increase its profit.

In our model, some consumers are sensitive to the product design of the product (“picky” consumers), whereas others do not mind the particular features of the product (“flexible” consumers). Consumers have unit demand and valuations that depend on whether they are picky or flexible (the ex ante type) and, when they are picky, whether the match is good or bad (the match value). The firm can use two sales channels: a superior (indirect) channel, in which it can also provide personalized recommendations and an inferior (direct) channel, in which it can not. Marginal costs are higher than the valuation for a bad

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2025 The Author(s). *The RAND Journal of Economics* published by Wiley Periodicals LLC on behalf of The RAND Corporation.

match. Consumers decide which sales channel to use and receive a recommendation if they have chosen the indirect channel.

A recommendation policy with “inflated recommendations” has the feature that the product is recommended to a fraction $\beta > 0$ of picky consumers with a bad match. The profit-maximizing recommendation policy with inflated recommendations is such that a picky consumer’s expected valuation conditional on receiving a recommendation is equal to a flexible consumer’s valuation. The firm then sells to picky consumers with a recommendation and all flexible consumers in the superior channel. If the fraction of picky consumers is below some critical level, this maximizes the firm’s profit and makes recommendations only partially informative. Otherwise, the firm sells to flexible consumers through the inferior channel (if at all) and to picky consumers with a good match through the superior channel. In this case, the firm may use the two different channels as a self-selection device for consumers.

To contribute to the policy debate on biased recommendations by digital platforms,³ we enrich our base model to distinguish between an intermediary that provides recommendations and a seller that sets retail prices and has to pay the intermediary for its intermediation services. We show that an intermediary who commits to its recommendation policy and the intermediation fee implements the vertically integrated solution. If consumers can buy in the direct channel after first going to the intermediary to benefit from its recommendation service (leading to “platform leakage”), the intermediary has to adjust its recommendation policy. Inflated recommendations may still be the equilibrium outcome, but they occur less frequently.

Adding a regulator who commits to its regulatory policy, we show that the welfare-maximizing regulator implements the first best if it can mandate any recommendation policy that conditions recommendations on retail prices (but does not engage in price regulation). If the regulator mandates fully informative recommendations, such regulation implements the first best in some environments in which this would not be achieved without regulation. However, in other environments, the policy backfires and delivers lower welfare than the *laissez-faire*. Furthermore, with the added optimal regulation of the fee charged by the intermediary, the regulation may still backfire compared with the *laissez-faire*.

1.1 | Related Literature

Our article contributes to the literature on information design, as the firm manages consumer expectations. Lewis and Sapington (1994) consider a firm that provides informative signals to consumers about their valuation, where the signal is fully informative with some probability and drawn from the initial distribution with the remaining probability. They show that the firm either perfectly informs consumers of their match value or does not provide any information at all. Their and our article is an instance of Bayesian persuasion (Kamenica and Gentzkow 2011), as the firm commits to a recommendation policy that allows consumers to update their beliefs about the match value. Our article contributes to the literature on information design by showing that a firm may optimally provide limited information to

consumers such that some consumers will buy even when gains from trade are negative.⁴

Our extended model with an intermediary relates to the work on biased recommendations. A profit-maximizing intermediary may want to provide biased recommendations for a variety of reasons. In Lee (2021), the intermediary is a mechanism designer who must persuade consumers to buy the recommended product. Monetizing only on the seller side, the intermediary may provide biased recommendations when seller profits are not aligned with consumer benefits (seller prices are treated as exogenous in their setting).⁵ Consumers may be exposed to biased recommendations in the presence of price effects, as recent theoretical contributions have pointed out (e.g., Armstrong and Zhou 2011; Hagiu and Jullien 2011; de Cornière and Taylor 2019). We develop our argument in the context of experience goods and endogenous prices.⁶

A particular instance of biased recommendations is “self-preferencing,” which may arise if an intermediary is also a seller and, thus, operates in a hybrid mode (e.g., de Cornière and Taylor 2019). Such a firm may have an incentive to steer consumers toward its own products. Self-preferencing as an allegedly anti-competitive practice is investigated by competition authorities and is prohibited for gatekeeper platforms under the Digital Markets Act in the EU. This raises the question of which regulatory interventions increase consumer or total surplus (for formal investigations, see, e.g., Anderson and Bedre-Defolie 2024; Aridor and Gonçalves 2022; Etro 2021; Hagiu et al. 2022; Hervas-Drane and Shelegia 2025; Kang and Muir 2022; Zenny 2022).⁷

Product recommendations may provide information on product quality. This can also be achieved through advertising and certification. Our setting connects to work on content advertising (Anderson and Renault 2006) in which advertising contains match-relevant information. Advertising in our context contains “real information” (Johnson and Myatt 2006) as picky consumers update their beliefs depending on whether or not they receive a recommendation. Recommending the new product to a picky consumer with a bad match can be considered “false advertising” and relates to Rhodes and Wilson (2018), who considers advertising in the presence of quality uncertainty (see also Drugov and Troya-Martinez 2019; Aköz et al. 2020).⁸ Our article also speaks to the literature on targeted advertising (e.g., Anand and Shachar 2009; Johnson 2013), whereby advertisers can address a group of consumers with particular characteristics. The recommendation of a product to a certain subset of consumers can be seen as targeted advertising and the inflated recommendations outcome constitutes noisy targeting.

1.2 | Roadmap

The article proceeds as follows. In Section 2, we introduce our base model. In Section 3, we analyze the firm’s profit-maximizing strategy and establish our main result of inflated recommendations. In Section 4, we distinguish between a seller and an intermediary where the seller sets the retail price and, before that, the intermediary makes purchase recommendations and charges the seller for its service. First, we show that the

vertically integrated solution is implemented and discuss several extensions of the model. Second, we characterize the outcome under three different regulatory policies. Section 5 contains further discussions and a conclusion.

In Appendix A, we collect the relegated proofs and provide a concise treatment of the extensions discussed in Section 4. In Appendix B (whose content is summarized in Section 5), we allow for alternative versions of consumers' distribution of willingness to pay and consider independent draws of the ex ante type and the match realization.

2 | The Base Model

We consider a monopoly firm offering an experience good at a marginal cost $c > 0$. Two sales channels are available: an indirect channel I that allows for personalized recommendations and a direct channel D that does not.⁹ Moreover, it is less convenient for consumers to use the direct channel; we denote by $d > 0$ the utility loss from less convenience in the direct channel. The firm sets prices p^I and p^D in these channels to maximize its profit. We introduce the two sales channels to show a trade-off between second-degree price discrimination and uniform pricing with inflated recommendations.

There is a unit mass of risk-neutral consumers, and each consumer has unit demand. A consumer may buy the experience good through either sales channel or choose an outside option normalized to zero.

A picky consumer with a good match has a valuation of v_h in the indirect channel, and a picky consumer with a bad match has v_l . The remaining consumers are "flexible" in the sense that they do not face any uncertainty about the realized match value and always have a valuation of v_m . We assume that $(v_l + v_h)/2 < v_m < v_h$; this implies that flexible consumers have an expected valuation above picky consumers and a lower variance.¹⁰ A fraction α of consumers are picky and the remaining fraction $1 - \alpha$ are flexible.

We will focus our analysis on $v_l < c < v_m$.¹¹ The first inequality says that selling the bad match to picky consumers in the indirect channel reduces total surplus compared to the outside option. Recommending the product, therefore, reduces total surplus, and the first-best welfare maximum does not feature any such recommendation. The second inequality says that there are gains from trade with flexible consumers in the indirect channel.

A firm's recommendation policy is a message sent to a consumer that depends on the consumer's type and match value. By construction, it is always in the firm's interest to recommend the product to a picky consumer with a good match. However, the firm may also recommend the product with a probability of $\beta \in [0, 1]$ to picky consumers with a bad match; it does not matter whether the firm makes recommendations to flexible consumers as long as they are aware of the option to buy in either channel.¹²

Inefficiencies arise when some consumers buy in the direct channel instead of the indirect channel (*inefficient bypass*), when some consumers with strictly positive gains from trade do not




FIGURE 1 | Consumer demand for different recommendation policies β .

buy at all (*limited sales*), or when some picky consumers with a bad match buy. The last can only occur in the indirect channel and requires a recommendation policy with $\beta > 0$. As $c > v_l$, recommendations with $\beta > 0$ are excessive from a total surplus perspective; in such a case we speak of *inflated recommendations*. As $p \geq c$, if consumers follow a recommendation that reduces total surplus, this recommendation necessarily reduces consumer surplus (and leads to ex post regret from purchase).

The *timing* is as follows:

1. The firm commits to a recommendation policy β and sets its prices in the direct channel, p^D , and in the indirect channel, p^I .
2. After observing prices and β , consumers decide which sales channel to choose.
3. Picky consumers in the indirect channel receive personalized recommendations and all consumers make their purchasing decisions.

We abstract from price opacity and, thus, assume that consumers observe prices before deciding which channel to use. In addition, consumers observe the recommendation policy.¹³

We characterize the profit-maximizing solution. We note that consumers engage in Bayesian updating. However, this belief updating is pinned down by the recommendation policy the firm has committed to. If the implemented recommendation policy is β , a picky consumer who receives a recommendation updates the belief that the match is good from $1/2$ to $1/(1 + \beta)$. A picky consumer who does not receive a recommendation updates the belief that the match is good to 0 and, thus, is convinced that the match is bad.

This leads to a demand curve that depends on the chosen recommendation policy. Figure 1 illustrates different demand curves when consumers can buy in the indirect channel or not at all (i.e., the direct channel is not available or sufficiently inconvenient to use). For $\beta = 0$, picky consumers who receive a recommendation have a valuation of v_h , flexible consumers have v_m , and picky consumers without a recommendation have v_l . For fully uninformative recommendations $\beta = 1$, flexible consumers have a valuation of v_m and picky consumers with a

recommendation have $(v_h + v_l)/2$. For recommendation policy β^* given by $(v_h + \beta^*v_l)/(1 + \beta^*) = v_m$, picky consumers with a recommendation and flexible consumers have a valuation of v_m and picky consumers without a recommendation have v_l .

Before turning to the analysis of the monopoly problem, we characterize the first best. As the first-best allocation does not involve inflated recommendations and the indirect channel is more attractive, the socially optimal allocation is that all flexible and all picky consumers with a good match buy in the indirect channel. Welfare in the first best, $(1 - \alpha)(v_m - c) + (\alpha/2)(v_h - c)$ is linear in α . It is increasing in α if and only if $(v_h - c)/2 > (v_m - c)$.

To see that inflated recommendations can only occur with ex ante heterogeneity, we take a look at the two edge cases $\alpha = 0$ and $\alpha = 1$. If all consumers are flexible ($\alpha = 0$), the recommendation policy does not affect the consumers' choice of sales channel. If all consumers are picky ($\alpha = 1$), the firm maximizes its profits by fully expropriating the expected consumer surplus conditional on β . As $c > v_l$, total surplus decreases in β and a higher β implies a lower profit when the firm sets its profit-maximizing price. Thus, it is not in the interest of the firm to set β larger than 0. To summarize, in the edge cases of our model, the monopoly outcome is efficient and does not involve any inflated recommendations.

Note that a monopolist who can engage in third-degree price-discrimination would also implement the first best, as it would choose $\beta = 0$, sell in the indirect channel at price v_h to flexible consumers with a good match and at price v_m to flexible consumers. This would generate a profit of $(1 - \alpha)(v_m - c) + (\alpha/2)(v_h - c)$, which is equal to the total gains from trade.

3 | The Firm's Recommendation and Pricing Policy

The key ingredient of our model is that there is ex ante taste heterogeneity among consumers—that is, a fraction $\alpha \in (0, 1)$ of consumers are picky and the remaining are flexible—and the firm can not engage in third-degree price discrimination between picky and flexible consumers. The monopolist maximizes profits by setting $p^D \geq 0$, $p^I \geq 0$, and $\beta \in [0, 1]$.

The profit-maximizing outcome has the property that picky consumers never buy in the direct channel. By contradiction, suppose that the picky consumers buy in the direct channel. This implies that $p^D + d \leq p^I$. If this inequality is strict, all flexible consumers buy in the direct channel as well. Then, the seller can reach higher profits by shutting down the direct channel, setting $\hat{p}^I = p^D + d$ and serving all consumers in the indirect channel. Otherwise, if $p^D + d = p^I$, the seller can reach higher profits by slightly raising its price in the direct channel and inducing all consumers to buy in the indirect channel. Consequently, we obtain a contradiction.

Furthermore, some picky consumers buy in the indirect channel. By contradiction, suppose that the intermediary only sells to flexible consumers. Then, the seller sets $p^I = v_m$ and extracts the full surplus from the flexible consumers. However, the seller can further increase its profits by keeping the same prices, setting $\beta =$

0, and serving all the picky consumers with a good match, which will lead to higher profits. Hence, the profit-maximizing outcome has the property that no picky consumer buys in the direct channel and at least some of them buy in the indirect channel.

Flexible consumers then buy in the indirect channel, in the direct channel, or not at all. In the first case, the firm sets $p^I = v^m$ to sell to the flexible consumers in the indirect channel; it sets $\beta > 0$ such that picky consumers have expected valuation v_m conditional on receiving a recommendation. Note that the firm cannot do better by further increasing β and adjusting prices accordingly ($p^I = (v_h + \beta v_l)/(1 + \beta)$) because profits from each consumer type would go down. Profits from picky consumers are reduced due to the observation made in the edge case with picky consumers only that maximal profits for given β are decreasing in β . The reason is that the firm fully extracts the expected surplus of picky consumers with a recommendation and this surplus is decreasing in β . Profits made from flexible consumers are also reduced (as an increase in β implies a price p^I less than v_m). In the second case, the firm maximizes its profit by extracting the full surplus from flexible consumers. This is achieved by setting $p^D = v_m - d$ and selling to all picky consumers with a good match by setting $p^I = v_h$ and $\beta = 0$. In the third case, the firm maximizes its profit from selling to picky consumers only, which implies $\beta = 0$ and $p^I = v_h$.

Thus, potentially profit-maximizing outcomes are:

- The inflated recommendations outcome in which the firm sets β such that, $v_m = (v_h + \beta v_l)/(1 + \beta)$ or, equivalently, $\beta = (v_h - v_m)/(v_m - v_l) \equiv \beta^*$, $p^I = v_m$, and $p^D \geq \max\{v_m - d, 0\}$, and makes a profit of

$$\left[\frac{\alpha}{2} \left(1 + \frac{v_h - v_m}{v_m - v_l} \right) + (1 - \alpha) \right] (v_m - c);$$

- The inefficient bypass outcome in which the firm sets $\beta = 0$, serves picky consumers with a good match in the indirect channel at $p^I = v_h$ and serves flexible consumers in the direct channel at $p^D = v_m - d$, and makes a profit of

$$\frac{\alpha}{2} (v_h - c) + (1 - \alpha)(v_m - d - c);$$

- The limited sales outcome in which the firm sets $\beta = 0$, serves picky consumers with a good match at $p^I = v_h$ in the indirect channel and does not sell to flexible consumers, and makes a profit of $\frac{\alpha}{2}(v_h - c)$.

The maximal profit from implementing the inefficient bypass or limited sales outcome is $\frac{\alpha}{2}(v_h - c) + \max\{(1 - \alpha)(v_m - d - c), 0\}$, where the firm implements the former if $v_m - d - c > 0$.

We define critical values of α as a function of c that separates the inflated recommendations outcome from the alternative outcome as follows,

$$\bar{\alpha}(c) = \frac{\min\{v_m - c, d\}}{\min\{v_m - c, d\} + \frac{1}{2}(c - v_l) \frac{v_h - v_m}{v_m - v_l}},$$

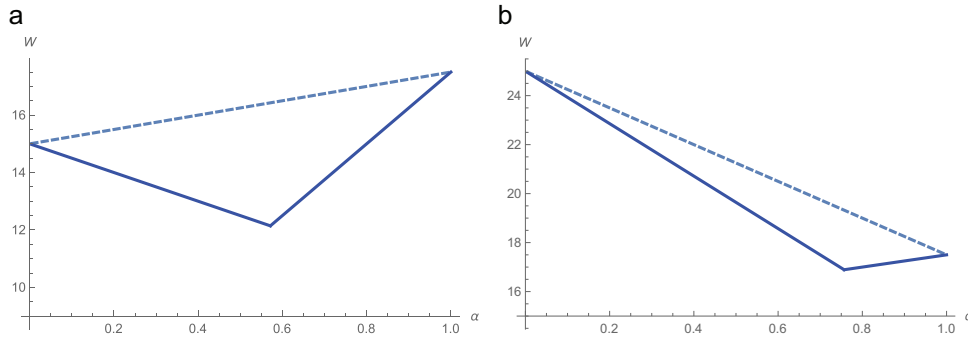


FIGURE 2 | Total welfare in α ; $v_h = 110, v_l = 30, d = 10, c = 75$. First-best outcome (dashed), private solution (solid).

where $\tilde{\alpha}(c) \in (0, 1)$. As we state in the following proposition, for $\alpha < \tilde{\alpha}(c)$, the firm maximizes its profit with inflated recommendations, whereas, for $\alpha \geq \tilde{\alpha}(c)$, it does so by inducing limited sales for $c \in (v_m - d, v_m)$ and inefficient bypass for $c \in (v_l, v_m - d)$.

Proposition 1. *Suppose that $c \in (v_l, v_m)$. The monopoly solution is characterized as follows:*

- For $\alpha < \tilde{\alpha}(c)$, the firm implements the inflated recommendations outcome: it sets $\beta = \beta^*$, $p^I = v_m$ and $p^D \geq v_m - d$. All flexible consumers and all picky consumers with a recommendation buy in the indirect channel;
- For $\alpha \geq \tilde{\alpha}(c)$, the firm implements the limited sales or inefficient bypass outcome: it sets $\beta = 0$, $p^I = v_h$, $p^D = \max\{v_m - d, c\}$. Picky consumers go to the indirect channel and buy if they receive the recommendation to buy, whereas all flexible consumers buy in the direct channel for $c < v_m - d$.

Total surplus implications are illustrated in Figure 2, where the kink in the total surplus function occurs at $\tilde{\alpha}(c)$. Compared to the first best, for $\alpha < \tilde{\alpha}(c)$, welfare losses are given by $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - v_l)$. For $\alpha \geq \tilde{\alpha}(c)$, they are given by $(1 - \alpha) \min\{d, v_m - c\}$.

Inefficient bypass and inflated recommendations: Some interpretations: The inefficient bypass outcome can be seen as the result of a simple screening contract using a damaged good strategy (Deneckere and McAfee 1996). Picky consumers with a good match have a valuation higher than flexible consumers. Because recommendations can only be given in the indirect channel and picky consumers are more quality-sensitive than flexible consumers, the single-crossing property is satisfied. Flexible consumers buy the new product in the direct channel and incur the inconvenience cost d —this inferior sales channel resembles a damaged good. In the limited sales outcome, the firm foregoes any sales to flexible consumers.

The inflated recommendations outcome can be seen as the outcome of a partial pooling contract. The key difference compared to standard price discrimination problems is the possibility of partially informative recommendations, which equalizes the valuation of picky consumers (with a recommendation) and flexible consumers. Some picky consumers with a bad match are made to believe that the product tends to be a good match. A picky consumer who does not know the match quality thus has a lower expected valuation after receiving the recommendation to buy.

From a consumer's perspective, the contract offer with inflated recommendations looks as follows: flexible consumers do not face uncertainty and always receive a contract offer for each sales channel, whereas picky consumers with a recommendation know that with a probability less than one, they buy the product with a low valuation.¹⁴

The firm's decision for marginal costs $c \notin (v_l, v_m)$: So far, we have considered $c \in (v_l, v_m)$. For $c \geq v_m$, the firm makes a profit from only selling to picky consumers with a good match in the indirect channel: it chooses the recommendation policy $\beta = 0$, sets a price of $p^I = v_h$, and makes a profit of $\frac{\alpha}{2}(v_h - c)$. Thus, inflated recommendations do not arise for $c \geq v_m$.

For $c \leq v_l$, it becomes socially efficient to serve all consumers, which implies that recommendations can no longer be inflated from a social welfare perspective. To keep the exposition short, we restrict attention to the case $c \in (v_l - d, v_l]$. Consider the firm setting recommendation policy $\beta = \beta^*$ and selling only through the indirect channel,¹⁵ that is, the firm sets $p^D + d \geq p^I$ and $p^I = \frac{v_h + \beta v_l}{1 + \beta}$. Any $\beta < \beta^*$ is strictly dominated by $\beta = \beta^*$ because flexible consumers and more picky consumers would buy at price v_m . Consider the case $\beta \geq \beta^*$. The firm's profit is given by¹⁶

$$\left(\frac{\alpha}{2}(1 + \beta) + (1 - \alpha) \right) \left(\frac{v_h + \beta v_l}{1 + \beta} - c \right).$$

The profit of the firm inducing the outcome with socially insufficient recommendations (that is, $\beta = \beta^*$) is equal to $\left(\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha) \right) (v_m - c)$. We note that for $\beta < 1$, recommendations are socially insufficient but from a consumer surplus perspective excessive, as picky consumers with a bad match would be better off not buying. The other potentially profit-maximizing strategy is to set $\beta = 1$, $p^I = \frac{v_h + v_l}{2}$, and $p^D > \frac{v_h + v_l}{2}$. This induces the first-best outcome and generates a profit of $\frac{v_h + v_l}{2} - c$.

The firm prefers the first-best outcome if the fraction of picky consumers is sufficiently high such that $\frac{v_h + v_l}{2} - c \geq \left(\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha) \right) (v_m - c)$. This holds with equality for,

$$\tilde{\alpha}(c) \equiv \frac{v_m - v_l}{v_m - c},$$

which increases in c and converges to 1 as c goes to v_l . If $\alpha \geq \hat{\alpha}(c)$, the firm will sell to all consumers in the indirect channel at $p^I = \frac{v_h+v_l}{2}$ and induce the first-best outcome. Otherwise, if $\alpha < \hat{\alpha}(c)$, the firm maximizes its profit by inducing the outcome with socially insufficient recommendations ($\beta^* = \frac{v_h-v_m}{v_m-v_l}$). It sells at $p^I = v_m$ in the indirect channel and does not sell to the fraction $1 - \beta^*$ of picky consumers with a bad match. This leads to welfare loss equal to $\frac{\alpha}{2}(1 - \beta^*)(v_l - c)$.

4 | Intermediaries as Recommenders

In this section, we introduce an intermediary as an economic actor that plays the role of a recommender and is different from the seller that sets retail prices. We assume that an intermediary has collected consumer data that allows it to identify consumer type and match value, whereas the seller lacks this information. The intermediary can then provide personalized recommendations to consumers in the indirect channel. Whether a consumer receives a recommendation may depend on the retail prices set by the seller in both channels, $\beta(p^I, p^D)$. In this context, we may think of both channels as digital sales channels, with the direct channel being the seller's website and the indirect channel being the intermediary's marketplace.

4.1 | The Profit-Maximizing Intermediary

The intermediary's revenue model is to charge sellers on its platform. It sets the rate $\lambda \in [0, 1]$ as the fraction of the seller's profits it extracts (profit sharing).¹⁷ We analyze the game in which, first, the intermediary chooses its recommendation policy and sets λ and then the seller sets the retail prices. The timing of the subsequent decisions is the same as in the base model. We will show that, in subgame-perfect equilibrium, the vertically integrated solution can be decentralized; that is, the intermediary chooses its strategy such that the seller will optimally respond by setting the same retail prices along the equilibrium path as in the vertically integrated solution that we analyzed in the previous section. We will then see that the seller always obtains the same profit. Therefore, the trade-off for the intermediary between the two possibly profit-maximizing strategies is the same as for the vertically integrated firm.

In this section, we focus on the parameter range with $c \in \left(\frac{v_h+v_l}{2}, v_m - d\right)$ in which either inflated recommendations or inefficient bypass is the outcome.¹⁸ For this interval to be non-empty, d cannot be too large.

The seller can secure some minimal profit for itself by selling only in the direct channel: it can set $p^D = v_m - d$ (and $p^I > v_m$) and make the profit $(1 - \alpha)(v_m - d - c)$. Thus, the intermediary has to provide such a profit level to the seller at the very least.

First, consider the inflated recommendations outcome. In the vertically integrated solution with inflated recommendations, $p^I = v_m$, $p^D \geq v_m - d$ and all sales occur in the indirect channel. Consider the intermediary's recommendation policy,

$$\beta(p^I, p^D) = \begin{cases} \beta^* & \text{for } p^I = v_m \text{ and } p^D \geq v_m - d \\ 1 & \text{otherwise.} \end{cases}$$

To induce the inflated recommendations outcome, the intermediary has to afford a sufficient fraction of profit, $1 - \lambda$, to the seller. We denote the λ for which the seller obtains profit $(1 - \alpha)(v_m - d - c)$ by λ^* , which solves,

$$(1 - \lambda^*) \left[\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha) \right] (v_m - c) = (1 - \alpha)(v_m - d - c).$$

Given λ^* , the seller has no incentive to deviate from $p^I = v_m$, $p^D \geq v_m - d$. If the seller sets a price p^I different from v_m , it will make profit $\max\{(1 - \alpha)(v_m - d - c), (1 - \lambda^*)(1 - \alpha)(v_m - c)\}$. From the definition of λ^* , it follows that the second expression is less than the first one and the best deviation is to serve flexible consumers in the direct channel. A deviation to $p^D < v_m$ is not profitable, as this gives profits less than $(1 - \alpha)(v_m - d - c)$.

Second, we consider the inefficient bypass outcome. Suppose that the intermediary takes (almost) the entire profit in the indirect channel (i.e., $\lambda = 1$) and sets $\beta = 0$ for all prices. In this case, the seller can only make a profit in the direct channel. It will sell to flexible consumers at $p^D = v_m - d$ and it will set the price $p^I = v_h$ (for any infinitesimally small profit fraction it maintains). This implements the vertically integrated solution, with profits $(1 - \alpha)(v_m - d - c)$ going to the seller and $\frac{\alpha}{2}(v_h - c)$ going to the intermediary.

The intermediary can choose between inducing inflated recommendations or inefficient bypass. In either case, the seller makes the profit $(1 - \alpha)(v_m - d - c)$, which is the minimal profit that the seller can always obtain. Thus, the comparison of the intermediary's profits yields the same critical $\bar{\alpha}(c)$ as the comparison of the vertically integrated firm's profit and we have proved the following proposition.

Proposition 2. *Suppose that $c \in \left(\frac{v_h+v_l}{2}, v_m - d\right)$. The intermediary achieves the integrated monopoly solution as follows:*

- For $\alpha < \bar{\alpha}(c)$, the intermediary implements the inflated recommendations outcome by setting,

$$\lambda = \lambda^* \equiv 1 - \frac{(1 - \alpha)(v_m - c - d)}{\left[\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} + (1 - \alpha)\right](v_m - c)},$$

$\beta(p^I = v_m, p^D \geq v_m - d) = \beta^*$ and $\beta = 1$ otherwise. Equilibrium prices are given by $(p^I, p^D) = (v_m, v_m - d)$. All flexible consumers and the fraction $\frac{1+\beta^*}{2}$ of picky consumers buy in the indirect channel;

- For $\alpha \geq \bar{\alpha}(c)$, the intermediary implements the inefficient bypass outcome by setting $\lambda = 1$, $\beta = 0$ for all (p^I, p^D) . Equilibrium prices are given by $(p^I, p^D) = (v_h, v_m - d)$. Half of the picky consumers buy in the indirect channel and all the flexible consumers buy in the direct channel.

With a small fraction of picky consumers in the population, the intermediary does best by inflating recommendations. The intermediary induces the seller to sell to flexible consumers in the indirect channel by leaving a sufficient fraction of profits to the seller. If the seller deviates and sells to flexible consumers in the direct channel (by setting low p^D), then the intermediary is committed to stop providing informative recommendations

to picky consumers, which renders the deviation unprofitable. By contrast, with a large fraction of picky consumers in the population, the intermediary prefers to induce the outcome with inefficient bypass. By making it very unattractive to sell through the intermediary, the intermediary induces the seller to sell to flexible consumers directly.

The intermediary's decision for $c \in (v_l, (v_h + v_l)/2 - d)$: In this parameter range, the tradeoff between inflated recommendations and inefficient bypass continues to apply. The novel feature is that, for high α , the most attractive outside option for the seller may be to sell to all consumers in the direct channel at $p^D = (v_h + v_l)/2 - d$. The seller then makes a profit of $(v_h + v_l)/2 - d - c$. Alternatively, it may sell to flexible consumers only at $p^D = v_m - d$ in which case it makes $(1 - \alpha)(v_m - d - c)$. If α is sufficiently small – that is, $\alpha < \frac{v_m - (v_h + v_l)/2}{v_m - d - c} \equiv \tilde{\alpha}(c)$ – the latter dominates the former and the intermediary has to make sure that it offers a contract to the seller that allows the seller to make at least $(1 - \alpha)(v_m - d - c)$. Here, the analysis from above applies.

Otherwise—that is, $\alpha > \tilde{\alpha}(c)$ —the intermediary has to come to terms with a smaller share of industry profits, compared to the case in which the seller can only profitably cater to informed consumers in the direct channel (i.e. $c \geq (v_h + v_l)/2 - d$). The intermediary has to compensate the seller for not making $(v_h + v_l)/2 - d - c$ and thus $\lambda < 1$ even under inefficient bypass, in contrast to what happens for marginal costs $c \geq \frac{v_h + v_l}{2} - d$.

This is seen as follows. Consider the difference in seller profits $(1 - \tilde{\alpha}(c))(v_m - d - c) - \left(\frac{v_h + v_l}{2} - d - c\right)$, which can be rewritten as $\left(v_m - \frac{v_h + v_l}{2}\right) - \tilde{\alpha}(c)(v_m - d - c)$. The derivative with respect to c is given by $\tilde{\alpha}(c) + \frac{\tilde{\alpha}(c)\beta^*}{d + \frac{1}{2}(v_h - c - (1 + \beta^*)(v_m - c))}(v_m - d - c)$, which must be positive. Furthermore, the value of the difference at $c = v_l$ is negative (as $\tilde{\alpha}(v_l) = 1$) and the value at $c = \frac{v_h + v_l}{2} - d$ is positive. Therefore, for any $d \in \left(0, \frac{v_h - v_l}{2}\right)$ there exists $\bar{c}(d)$ that solves

$$\frac{v_m - \frac{v_h - v_l}{2}}{v_m - d - c} = \frac{d}{d + \frac{1}{2}(c - v_l)(v_h - v_m)}$$

such that the considered difference is positive for $c \in (\bar{c}(d), (v_h + v_l)/2 - d)$ and is negative for $c \in (v_l, \bar{c}(d))$.

This shows that if $c \in (v_l, \bar{c}(d))$ then $(1 - \tilde{\alpha}(c))(v_m - d - c) < \left(\frac{v_h + v_l}{2} - d - c\right)$, which implies that $\tilde{\alpha}(c) > \tilde{\alpha}(c)$. Therefore, we have that the outcome features inflated recommendations with λ from Proposition 2 for $\alpha < \tilde{\alpha}(c)$, inflated recommendations with $\lambda = 1 - ((v_h + v_l)/2 - d - c)/(1 - \alpha + (\alpha/2)(1 + \beta^*))$ for $[\tilde{\alpha}(c), \tilde{\alpha}(c))$ and inefficient bypass with $\lambda = 1 - ((v_h + v_l)/2 - d - c - (1 - \alpha)(v_m - d - c))/((\alpha/2)(v_h - c))$ for $\alpha \geq \tilde{\alpha}(c)$. If $c \in [\bar{c}(d), (v_h + v_l)/2 - d)$, then $\tilde{\alpha}(c) \leq \tilde{\alpha}(c)$. Thus, the outcome features inflated recommendations with λ from Proposition 2 for $\alpha < \tilde{\alpha}(c)$ and inefficient bypass with $\lambda = 1 - ((v_h + v_l)/2 - d - c - (1 - \alpha)(v_m - d - c))/((\alpha/2)(v_h - c))$ for $\alpha \geq \tilde{\alpha}(c)$.

Discussion of the result with intermediation: Inflated recommendations can also occur in the case of “platform leakage” (Hagiu and Wright 2024); that is, a fraction of consumers can first visit the indirect channel to obtain a recommendation from the intermediary and can then buy in the direct channel. In other words, some consumers can use the intermediary for showrooming and the seller may free-ride on the intermediary's recommendation service (Wang and Wright 2020). For simplicity, consider an exogenous fraction ν of picky consumers who can switch without cost. We show in Appendix A that the critical α until which there are inflated recommendations is larger than $\tilde{\alpha}$ and, in this sense, inflated recommendations become more prevalent under platform leakage. Showrooming undermines the incentives for fully informative recommendations, which makes the outcome with inefficient bypass less profitable. Thus, the intermediary is more inclined to induce the inflated recommendations outcome.

In our model, consumers can buy the product or choose an outside option and the intermediary does not make any profit when consumers choose the outside option. However, the intermediary may produce its own well-known base product at marginal cost c_0 and sell it at price p_0 to consumers with valuation $v_0 > c_0$ in the indirect channel and extract the surplus by setting $p_0 = v_0$. Then, whenever a consumer chooses the outside option in the indirect channel, the intermediary makes a profit of $v_0 - c_0$. We can show that, when $v_0 - c_0$ is small enough, our finding that the inflated recommendations outcome prevails over the inefficient bypass outcome continues to hold for a range of α .¹⁹

So far, we have assumed that the intermediary's price instrument is the fraction of industry profits λ in the indirect channel that it asks from the seller. Alternative price instruments, for example, are a listing fee T that has to be paid by the seller to be listed by the intermediary, a per-unit transaction fee of t (or, equivalently, per-click fee if there is a linear relationship between clicks and transactions), or an ad valorem transaction fee of τ . We note that, in our model, the equilibrium outcome is invariant to the particular type of pricing instrument available to the intermediary. Details on platform leakage, the intermediary with a profitable outside option, and the intermediary's price instrument are provided in Appendix A.

Inflated recommendations even occur when the intermediary cannot condition the recommendation on retail prices; that is, the intermediary chooses (β, λ) such that β can not be conditioned on (p^I, p^D) . In an inflated recommendations outcome with $\beta = \beta^*$, the seller obtains a profit of $(1 - \lambda)((1 - \alpha) + \alpha(1 + \beta^*)/2)(v_m - c)$. Instead, the seller could set a price below $v_m - d$ in the direct channel and induce flexible consumers to choose the direct channel. Then the seller would still sell to picky consumers with a recommendation at price v_m in the indirect channel and, thus, make a profit of $(1 - \alpha)(v_m - d - c) + (1 - \lambda)\alpha(1 + \beta^*)/2(v_m - c)$. Hence, the intermediary has to set λ such that it respects the seller's incentive constraint $(1 - \lambda)(v_m - c) \geq v_m - d - c$. This ties down λ : the intermediary sets λ equal to $1 - \frac{v_m - d - c}{v_m - c}$. This expression is smaller than λ^* given in Proposition 2 that applies for $\alpha < \tilde{\alpha}(c)$. This means that the intermediary has to leave a larger fraction of profits to the seller under inflated recommendations. Hence, the intermediary is more inclined to

serve picky consumers only and induce the inefficient bypass outcome. This implies that the critical α below which the intermediary inflates recommendations is less than in the case in which the intermediary can condition β on the seller's prices.

4.2 | Regulating the Intermediary

We explore whether (and if so, how) a regulator could improve welfare if it were able to restrict the intermediary's choices. Whenever the regulator assumes control, we postulate that it operates under commitment. First, we consider the problem in which the planner fully controls the recommendation policy—that is, the planner mandates β as a function of retail prices—and show that the first best can be implemented even though price setting is decentralized. Second, we consider what happens when the planner mandates fully informative recommendations and show that in some situations this policy improves on the outcome under laissez-faire, and in others it is strictly worse. Third, even if, in addition to imposing fully informative recommendations, the regulator can impose a limit on the profit share that the intermediary can ask from the seller, the regulatory policy may backfire. In these problems, the regulator moves first, then the intermediary sets λ , followed by the seller setting prices (p^I, p^D) . To keep the exposition short, we restrict attention to the case that marginal cost c is in the interval $((v_h + v_l)/2, v_m - d)$.

4.2.1 | Full Regulation of the Recommendation Policy

We consider a regulator that conditions the mandated recommendation policy on prices, $\beta(p^I, p^D)$, but cannot directly affect the intermediary's fee λ . We will show that full control over the recommendation policy allows the regulator to reach the first-best total surplus.

To reach the first-best outcome, the regulator has to: first, induce the intermediary to make the seller sell to flexible consumers through the intermediary; second, minimize recommendations to picky consumers with a bad match under the constraint that the profit of the intermediary is non-negative. As the tie-breaking rule, we assume that the regulator who is indifferent between inducing different prices picks the ones that maximize consumer surplus.

We will show that it is sufficient to restrict attention to the regulator imposing recommendation policies that reveal some information if and only if the seller sets a price p_0^I in the indirect channel and reveal no information otherwise; that is,

$$\beta(p^I, p^D) = \begin{cases} \beta_0 & \text{for some } p_0^I \text{ and any } p^D, \\ 1 & \text{otherwise.} \end{cases}$$

Such a recommendation policy makes deviations for the seller costly: Any deviation in the price in the indirect channel by the seller leads to the loss of all profit from picky consumers.

As, by assumption, for a given total surplus, the regulator prefers an outcome with a higher consumer surplus, the regulator mandates fully informative recommendations for

the lowest possible price p_0^I that implements the first best. We can restrict attention to $p_0^I < v_m$. Given λ , if the seller sets $p^I = p_0^I$, the intermediary obtains fraction λ of industry profit $(1 - \alpha/2)(p_0^I - c)$. If the seller sets a price different from p_0^I , the intermediary and the seller either would share the maximal profit with uninformative recommendations on the indirect channel $(1 - \alpha)(v_m - c)$ or the seller sells directly and obtains a profit of $(1 - \alpha)(v_m - d - c)$, whereas the intermediary obtains a profit of zero. For given λ and p_0^I , the seller sets p_0^I and induces the inflated recommendations outcome if $(1 - \lambda)(1 - \alpha/2)(p_0^I - c) \geq \max\{(1 - \lambda)(1 - \alpha)(v_m - c), (1 - \alpha)(v_m - d - c)\}$. To make sure that the seller prefers the inflated recommendations outcome over selling to flexible consumers only in the indirect channel, the regulator must pick p_0^I in its recommendation policy such that $(1 - \alpha/2)(p_0^I - c) \geq (1 - \alpha)(v_m - c)$. With the lowest price satisfying this inequality (which is less than v_m), the intermediary then sets λ to extract $\lambda(1 - \alpha/2)(p_0^I - c)$. The largest λ that is compatible with the seller's incentive constraint solves $(1 - \lambda)(1 - \alpha/2)(p_0^I - c) = (1 - \alpha)(v_m - d - c)$. As $(1 - \alpha/2)(p_0^I - c) = (1 - \alpha)(v_m - c)$, the binding constraint can be rewritten as $(1 - \lambda)(1 - \alpha)(v_m - c) = (1 - \alpha)(v_m - d - c)$, which simplifies to $\lambda = d/(v_m - c)$. Therefore, we obtain the following proposition.

Proposition 3. *Suppose that $c \in (\frac{v_h + v_l}{2}, v_m - d)$. If the regulator can set any recommendation policy $\beta(p^I, p^D)$, it achieves the first-best outcome by setting $\beta(p_0^I, p^D) = \beta_0 = 0$ for $p_0^I = \frac{\alpha/2}{1 - \alpha/2}c + \frac{1 - \alpha}{1 - \alpha/2}v_m$ and any p^D and $\beta(p^I, p^D) = 1$ for all other (p^I, p^D) . In equilibrium, the intermediary sets $\lambda = \frac{d}{v_m - c}$ and the seller sets price $p^I = p_0^I$. Flexible consumers and picky consumers with a good match buy the product through the indirect channel.*

The proposition shows that the regulator's recommendation policy fully determines the joint profit of the intermediary and the seller when all sales occur in the indirect channel. Therefore, the regulator does not need control over the intermediary's revenue policy λ to achieve the first best. Moreover, this recommendation policy also maximizes consumer surplus because, with either objective function, the optimal regulation makes the seller's and the intermediary's incentive constraint binding. Flexible and picky consumers are always strictly better off with regulation.

4.2.2 | Mandated Fully Informative Recommendations

Consider a policy intervention of the regulator in which the regulator requires recommendations to be fully informative (i.e., $\beta(p^I, p^D) = 0$ for all p^I, p^D). This policy is motivated by an interpretation of what may constitute the prohibition of a more-favorable treatment of the product under consideration compared to the outside option. The following proposition shows that for small α the efficient outcome is implemented, but that, for intermediate values of α , the policy backfires.

Proposition 4. *Suppose that $c \in (\frac{v_h + v_l}{2}, v_m - d)$. When the regulator mandates that recommendations must be fully informative (i.e., $\beta(p^I, p^D) = 0$ for all p^I, p^D), the equilibrium is characterized as follows: There is a critical value $\alpha_{F^I}(c) \in (0, \bar{\alpha}(c))$ such that,*




FIGURE 3 | Total welfare in α ; $v_h = 110, v_l = 30, d = 10, c = 75$. First-best outcome, private solution (solid), fully informative recommendations (dot-dashed).

- If $\alpha < \alpha_{FI}(c)$, then the intermediary sets $\lambda = \hat{\lambda} \equiv 1 - \frac{(1-\alpha)(v_m-d-c)}{(1-\frac{\alpha}{2})(v_m-c) - \frac{\alpha}{2}(v_h-c)}$ and the first-best allocation is implemented;
- If $\alpha \geq \alpha_{FI}(c)$, then the intermediary sets $\lambda = 1$; equilibrium prices are given by $(p^I, p^D) = (v_h, v_m - d)$ and the inefficient bypass outcome is implemented.

The critical value $\alpha_{FI}(c)$ below which imposing fully informative recommendations does not violate the incentive compatibility constraints of intermediary and seller, is given by the solution to,

$$\frac{(1-\alpha)(v_m-d-c)}{(1-\frac{\alpha}{2})(v_m-c) - \frac{\alpha}{2}(v_h-c)} = \frac{(1-\frac{\alpha}{2})(v_m-c) - \frac{\alpha}{2}(v_h-c)}{(1-\frac{\alpha}{2})(v_m-c)}$$

(as derived in the proof of Proposition 4) and satisfies that $\alpha_{FI}(c) < \bar{\alpha}(c)$.

As shown in Proposition 4, for $\alpha < \alpha_{FI}(c)$, the planner's policy implements the first best and strictly increases total surplus compared to the laissez-faire. For $\alpha \geq \bar{\alpha}(c)$, the regulator does not improve on the laissez-faire, as mandating $\beta = 0$ implies that flexible consumers buy in the direct channel, which also happens under laissez-faire. The regulation $\beta = 0$ performs worse than laissez-faire for $\alpha \in [\alpha_{FI}(c), \bar{\alpha}(c)]$. Mandating fully informative recommendations does not allow the intermediary to inflate recommendations. Instead, the intermediary has two potentially profit-maximizing options: first, it can set λ such that it collects profits from all flexible consumers and those picky consumers with a good match (respecting the seller's incentive compatibility constraint); or second, it can extract all surplus from picky consumers with a good match. For $\alpha \geq \alpha_{FI}(c)$, it prefers the latter and the regulation leads to inefficient bypass by flexible consumers. By contrast, under laissez-faire, by inflating recommendations, the intermediary can drive the expected gross surplus of picky consumers who receive a recommendation in the indirect channel down to the one of flexible consumers. This makes the former strategy more attractive to the intermediary simply because more consumers buy. Hence, with $\beta = 0$, the regulator gives up on the welfare generated from flexible consumers buying in the indirect instead of the direct channel. This welfare loss is larger than the welfare gain among picky consumers (under laissez-faire, some picky consumers with a bad match buy).

We illustrate these findings in Figure 3. The upper gray line depicts welfare in the first best and the lower gray line depicts welfare under laissez-faire. Welfare under mandated fully informative recommendations is depicted by the solid line.

Without regulatory intervention, consumer surplus is fully extracted by the seller. This also holds in the case in which a regulator mandates fully informative recommendations and, in the equilibrium outcome, only picky consumers with a good match are served in the indirect channel. Hence, for $\alpha \geq \alpha_{FI}(c)$, regulation is neutral to consumer surplus. However, if, under the regulation, picky consumers with a good match and flexible consumers buy in equilibrium, all consumers buy at a price of v_m and, thus, picky consumers with a good match obtain a strictly positive surplus. Hence, regulation either increases consumer surplus or leaves it unchanged and, thus, regulation cannot backfire under the consumer surplus criterion.

Although, as we showed in Section 3, without regulation, the outcome is the same regardless of whether the intermediary and seller are integrated, this is not true when the regulator mandates fully informative recommendations. The reason is as follows. The vertically integrated firm maximizes the total profit, whereas the intermediary maximizes the total profit minus the seller's profit. Under inefficient bypass, the seller obtains $(1-\alpha)(v_m-d-c)$. However, when all trade takes place in the indirect channel, the seller can always deviate and sell to picky consumers with a good match in the indirect channel at a price of v_h and to flexible consumers in the direct channel at a price of v_m . This means that the intermediary must guarantee the seller a profit strictly greater than $(1-\alpha)(v_m-d-c)$. Therefore, the intermediary has a more favorable view of inefficient bypass than the vertically integrated firm and, thus, $\alpha_{FI}(c)$ is less than the critical α under the regulation $\beta = 0$ of a vertically integrated firm (which solves $[\alpha/2 + (1-\alpha)](v_m-c) = (\alpha/2)(v_h-c) + (1-\alpha)(v_m-d-c)$ and, thus, is $d/[d + (v_h - v_m)/2]$). This shows that the welfare loss is reduced if the intermediary and the seller vertically integrate. Nevertheless, even under vertical integration, the regulatory intervention to mandate fully informative recommendations can backfire.

4.2.3 | Regulating the Intermediary's Rent Extraction

Until now, we have considered regulations that impose restrictions on the intermediary's recommendation policy. Alterna-

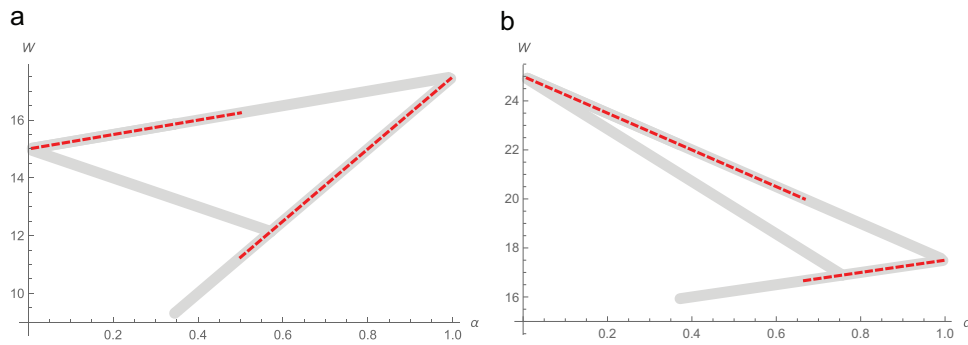


FIGURE 4 | Total welfare in α ; $v_h = 110, v_l = 30, c = 75$. First-best outcome, private solution, fully informative recommendations (solid), rent extraction regulation, $\beta = 0$ and $\bar{\lambda} = 0$ (dashed).

tively, the regulator may consider intervening by limiting the intermediary's rent extraction possibilities. We recall that the first best involves all flexible consumers and all picky consumers with a good match buying in the indirect channel and the picky consumers with a bad match not buying at all. Consider a regulator who only imposes a cap on the fraction of profits λ that the intermediary extracts from the seller. The regulator may want to choose this cap strictly less than 1 to encourage the seller to also serve flexible consumers in the indirect channel and, thus, inefficient bypass is avoided. However, if the recommendation policy remains unregulated, this encourages the intermediary to inflate recommendations. As we have shown, the laissez-faire equilibrium features $\lambda < 1$ with inflated recommendations, whereas inefficient bypass has $\lambda = 1$. In other words, a cap $\bar{\lambda}$ of slightly less than 1 has no repercussions for the intermediary's profit with inflated recommendations but reduces its profit under inefficient bypass. This makes inefficient bypass less attractive and implies that the critical α under a uniform regulated cap $\bar{\lambda}$ is larger than in the laissez-faire equilibrium. Hence, such price regulation more often leads to inflated recommendations than without regulation. Although consumers are not affected in our setting, the total surplus is reduced and rents are partly redistributed from the intermediary to the seller whenever the intermediary decides to continue to induce inefficient bypass in equilibrium.

If the regulator, in response to more-inflated recommendations, decides to mandate fully informative recommendations (i.e., $\beta(p^I, p^D) = 0$ for all p^I, p^D), either the outcome is efficient and, thus, picky and flexible consumers with a good match are served in the indirect channel yielding profit $(1 - \min\{\hat{\lambda}, \bar{\lambda}\})(1 - \alpha/2)(v_m - c)$ for the seller or there is inefficient bypass yielding profit $(1 - \alpha)(v_m - d - c) + (1 - \bar{\lambda})\frac{\alpha}{2}(v_h - c)$ for the seller.

For any given $\bar{\lambda} < 1$, there exists a uniquely defined $\alpha(\bar{\lambda}) \in (\alpha_{FI}(c), \bar{\alpha}(c))$ such that the equilibrium is characterized by the efficient outcome for $\alpha < \alpha(\bar{\lambda})$ and by the inefficient bypass outcome, otherwise (for details, see Appendix A). The regulator sets $\bar{\lambda}$ to maximize total surplus. The following remark (proved in Appendix A) summarizes the welfare implications of the regulation.

Remark 1. Regulating rent extraction on top of mandating fully informative recommendations weakly increases total surplus. For an intermediate range of α it will set $\bar{\lambda} \leq \bar{\lambda}(\alpha)$ and is indifferent

between all $\bar{\lambda} \in [0, \bar{\lambda}(\alpha)]$, where $\bar{\lambda}(\alpha) < 1$ is an upper bound that is a decreasing function of α (defined in Appendix A). In this range, such regulation may still backfire compared to the laissez-faire. Outside this range, the intermediary is indifferent between any $\bar{\lambda} \in [0, 1]$ and the regulation does not backfire.

Figure 4 illustrates our findings. For an intermediate range of α , the regulation leads to inefficient bypass instead of inflated recommendations under laissez-faire. In this case, the regulation backfires from a total surplus perspective.

5 | Discussion and Conclusion

A monopolist that sets the same price for all consumers may also make personalized recommendations to buy the product to consumers who lack information on match quality and, thus, may rely on recommendations. When all consumers are ex ante identical, the firm will never recommend the product to a consumer whose valuation is below marginal cost. In our terminology this means that the firm will not inflate recommendations. However, as we show in this article, with ex ante taste heterogeneity, the firm commits to a recommendation policy according to which a product is sometimes recommended even when the valuation is below marginal costs. As the profit-maximizing price is above marginal costs, this implies that some consumers ex post regret that they have bought the recommended product.

In our base model, the firm offers its product in a direct and an indirect sales channel; recommendations can only be made in the indirect channel. There are two ex ante types of consumers: picky consumers who experience a high or a low gross valuation and flexible consumers with an intermediate valuation. Selling to picky consumers with a low gross valuation is socially inefficient when the gross valuation is below marginal costs. The indirect channel offers two advantages: first, its use increases the benefit for all consumers; and second, picky consumers appreciate informative recommendations. By recommending the product to picky consumers with a bad match, which is known to the firm ex ante, but not to anybody else, the firm can bring the expected valuation of picky consumers with a recommendation down to that of flexible consumers and increase the sales volume without lowering the price it would set to serve flexible consumers.

5.1 | Inflated Recommendations in Other Settings

The main takeaway that recommendations may be inflated—that is, products are recommended and consumers follow this recommendation even though ex post consumers end up buying a product that generated negative gains from trade—does not rely on our modeling choice that the distribution of match values depends on the ex ante type of the consumer.

To make this transparent, we remove the direct sales channel (i.e., d is sufficiently large). Suppose that consumers are heterogeneous along two dimensions (θ, ε) . The indirect utility of a consumer (θ, ε) buying the product at price p is given by,

$$u^i(\theta, \varepsilon) = \theta + \varepsilon - p,$$

with $E[\varepsilon|\theta] = 0$ for all θ . We model ex ante taste heterogeneity as differences in expected valuations θ and refer to θ as the consumers' type and to ε as the match value. We consider the two-type case and the continuous-type case. Match values can be discrete or drawn from a continuous distribution. In the discrete case, we allow for two realizations and write $\varepsilon(\theta) \in \{\underline{\varepsilon}(\theta), \bar{\varepsilon}(\theta)\}$. Consumers observe their type θ but are uncertain about the match value ε . The firm sets price p and can provide recommendations that reveal some information about the match value ε to consumers. We assume that the product recommendations can be conditioned on θ .

Our base model is a special case of this setting with $\theta_l = (v_l + v_h)/2$, ε conditional on type θ_l drawn from $\{(v_l - v_h)/2, (v_h - v_l)/2\}$ with equal probability, $\theta_h = v_m$ and $\underline{\varepsilon}(\theta_h) = \bar{\varepsilon}(\theta_h) = 0$. There, the low type has the lower expected valuation and the higher variance. Our assumption that the high type has zero variance may be seen as rather special.

In Appendix B, we analyze environments in which the match value distribution is independent of the consumer type θ . When there are two ex ante types θ_l, θ_h and two match value realizations $\underline{\varepsilon}, \bar{\varepsilon}$ that are drawn independently across types, then depending on the parameter values, the firm chooses inflated recommendations, socially insufficient recommendations, or fully extracts the surplus of consumers of high-type consumers with the highest match value realization (see Proposition 6 in Appendix B.1). Inflated recommendations either involve selling to high-type consumers with a bad match or to low-type consumers with a bad match.

We consider two other settings: two ex ante types and a continuum of match value realizations for each ex ante type and a continuum of ex ante types and a continuum of match value realizations for each ex ante type (both in Appendix B.2). In each of these settings, we provide the conditions such that some consumers receive the recommendation to buy, follow that recommendation, and then learn that their valuation is not only lower than the price they paid, but also lower than the marginal cost incurred by the firm—that is, recommendations are inflated.

In our analysis, we assume that the ex ante heterogeneity is due to differences in tastes. An alternative approach would be to assume that all consumers have ex ante the same tastes but they have differential information in the sense that some consumers obtain a

more-informative signal about their ex post valuation than others. An analysis with ex ante heterogeneously informed consumers is more intricate, but inflated recommendations also arise in such alternative settings, as we explore in Peitz and Sobolev (2025).

5.2 | The Policy Debate on Biased Recommendations

The policy debate on biased recommendations is focused on intermediaries' incentives in a vertically related market in which consumers receive the recommendation by an intermediary, whereas retail prices are set by the seller of the respective product. We extend our base model by introducing an intermediary and a seller. The intermediary commits to extracting a certain profit share from the seller and a recommendation policy. Given those terms, the seller sets the price for its product in the direct and indirect channels. As we show in a certain parameter range, the intermediary either induces inflated recommendations or inefficient bypass and the vertically integrated solution is implemented.

A regulator may want to remedy the welfare loss stemming from inflated recommendations or inefficient bypass. If the regulator were able to impose a sophisticated recommendation policy that ex ante specifies the recommendation policy as a function of retail prices, it would be able to implement the first best. However, if the regulator were to require that recommendations must be fully informative—that is, the intermediary is not allowed to recommend the product to picky consumers with a bad match—welfare effects are ambiguous: when there is only a small fraction of picky consumers in the population, this regulation can implement the first best; however, above a critical threshold, inefficient bypass will occur and the regulation backfires, as welfare is lower with the regulation than without.

Acknowledgments

We are grateful to the editor, Gary Biglaiser, two reviewers, Sarah Auster, Heski Bar-Isaac, Paul Heidhues, David Ronayne, Nicolas Schutz, Radostina Shopova, Julian Wright, and various seminar audiences for their useful comments. The authors gratefully acknowledge support from the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (projects B03, B04, and B05).

Open access funding enabled and organized by Projekt DEAL.

Endnotes

¹For example, "Wirecutter" is an intermediary that hires people to test different products to assess which one performs best for a specific purpose. It then provides the affiliate link and takes a percentage fee from any sales this generates. See reporting in Amanda Mull, "There Is Too Much Stuff," *The Atlantic*, May 24, 2019.

²It has been reported that, in 2018, "Amazon optimized the secret algorithm that ranks listings so that instead of showing customers mainly the most-relevant and best-selling listings when they search – as it had for more than a decade—the site also gives a boost to items that are more profitable for the company." This quote is taken from Dana Mattioli, "Amazon Changed Search Algorithm in Ways That Boost Its Own Products," *Wall Street Journal*, September 16, 2019.

³As a legislative response to this debate, the Digital Markets Act in the EU prohibits self-preferencing by gatekeeper platforms through its core

platform services. The Tenth Amendment of the German Competition Act from 2021 also explicitly states that the competition authority may prohibit self-preferencing by digital gatekeepers. However, recommendation biases are of policy concern more broadly. For example, as part of its consumer protection mandate, the Competition Markets Authority (CMA) in the UK has formulated some principles in the hotel booking sector with one of the aims to provide transparency about hidden payments from sellers to the intermediary (see CMA, “Consumer Protection Law Compliance: Principles for Businesses Offering Online Accommodation Booking Services,” February 26, 2019).

⁴Our article differs from recent contributions on intermediaries’ information design in which an intermediary provides information about consumer characteristics to sellers that can then use this information for price-discrimination purposes (Bergemann et al. 2015; Ali et al. 2023). In our extended model with intermediation, the seller can only set channel-specific prices but does not have further information available to price discriminate among consumers (or prefers not to use this information).

⁵If consumers suffer from limited cognition, an intermediary may exploit such consumers by providing biased recommendations (Heidhues et al. 2023).

⁶Empirical and theoretical work has looked at biased financial advice, whereas most of the industrial organization literature on this topic considers the recommendation of search goods. In our article, as in Inderst and Ottaviani (2012) and Teh and Wright (2022), neither the seller nor the buyer has private information about the match value between product design and consumer tastes. Instead, it is the intermediary who possesses this information and makes recommendations in return for a fee. Whereas in Inderst and Ottaviani (2012) and Teh and Wright (2022) sellers compete for those kickbacks, in our setting, the intermediary decides on those fees.

⁷As analyzed by de Cornière and Taylor (2014), recommendation biases may also arise in the context of an ad-financed search engine and ad-financed websites when consumers experience advertising as a nuisance. Vertical integration between the search engine and one of the websites has an ambiguous effect on the size of the recommendation bias. For a related model, see Burguet et al. (2015) and, for an overview, Peitz and Reisinger (2016).

⁸More specifically, Rhodes and Wilson (2018) consider a monopolist that privately learns its type—that is, whether its product is of low or high quality—and incurs production costs independent of quality. Low quality generates a lower but positive profit in the market than high quality if truthfully revealed. “False advertising” is a situation in which the low-quality type claims to be of high quality and mimics the high-quality type. In their model, advertising claims that are proven to be false are penalized. For moderate penalties, the low-quality type pools with the high-quality type with a positive probability of less than 1.

⁹To fix ideas, the direct channel could be a physical outlet, whereas the indirect channel could be an e-commerce site in which, thanks to consumer tracking, personalized recommendations are possible.

¹⁰In Section 5, we introduce consumers (θ, ε) with indirect utility $\theta + \varepsilon - p^I$ in the indirect channel and $\theta + \varepsilon - d - p^D$ in the direct channel. Here, we consider two types of θ , where the high type has $\varepsilon = 0$ and the low type has a high or low realization of ε with probability 1/2. In Section 5 and Appendix B, we consider other distributions of consumer valuations in which all consumer types θ draw ε from the same distribution.

¹¹At the end of Section 3, we discuss what happens under alternative parameter constellations.

¹²We use the framing that the firm provides personalized recommendations. The same analysis applies to environments in which all picky consumers agree on whether a match is good or bad and thus the draws of ε are perfectly correlated over consumers. In this alternative

setting, the firm is able to assess the quality of the experience good and decides whether to recommend the product to picky consumers. For example, a product (such as outdoor equipment) may work well under normal conditions, but consumers do not know whether a product will continue to function under extreme conditions. Picky consumers are those consumers who use the product under extreme conditions and rely on the firm’s recommendation, whereas flexible consumers only use the product under normal conditions (and are of the high type because they are rich urban consumers).

¹³One way to motivate this is to consider consumers arriving in two batches facing the same prices. A fraction ε arrives early and the remaining fraction arrives late. Early arrivals publicly report their experience and whether the product was recommended to them. In this way, late arrivals learn β before deciding what to do. In the limit as ε turns to zero, this implies that consumers of measure 1 observe β .

¹⁴The offer relates to offering a pure bundle. From a picky consumer’s perspective, we can distinguish two states of the world: the state of the world is good if the product constitutes a good match and bad otherwise. With inflated recommendations, a consumer is offered a bundle of the product in the good state and, with a positive probability, in the bad state. An important insight from the bundling literature, when marginal costs are negligible is that it may pay to sell large bundles, as the distribution of valuations becomes less dispersed (Bakos and Brynjolfsson 1999; Geng et al. 2005; Haghpanah and Hartline 2021). Our result contains a different, though related, message: under inflated recommendations—that is, including the product in the good state with a probability of 1 and including it in the bad state with a positive probability—the distribution of valuations becomes less dispersed, allowing the seller to better extract the gains from trade (at the social cost of reducing the gains from trade). Indeed, this holds with significant marginal costs when a strategy to sell large bundles would not be profitable.

¹⁵We can show that it is never profit-maximizing to sell to some consumers in the direct channel.

¹⁶It is easy to see that the second derivative of the profit function is positive for all $\beta \in [\beta^*, 1]$, and thus the profit function is convex. The maximum is reached either at $\beta = \beta^*$ or at $\beta = 1$.

¹⁷As the intermediary can implement the vertically integrated solution with this instrument (which we establish in this section in Proposition 2), it cannot do better with a different price instrument. Moreover, there are many price instruments that can be used to implement this solution (which we show in the discussion below).

¹⁸If $c \in \left(v_l, \frac{v_h + v_l}{2} - d \right]$, the trade-off between inflated recommendations and inefficient bypass remains the same, but the seller would obtain higher profits—we comment on this case below. If $c \in \left(\frac{v_h + v_l}{2} - d, \frac{v_h + v_l}{2} \right]$, our results remain unchanged, but proofs would need to be adjusted. If $c \in [v_m - d, v_m)$, the seller can not make any profit in the direct channel; thus, the intermediary can ask for a profit share of 100% and the intermediary trivially implements the vertically integrated solution.

¹⁹One may want to compare this setting to one in which a base product is offered by third-party sellers and is in competitive supply. Then, the model with the base product in competitive supply implements the vertically integrated solution, whereas the model in which the intermediary controls the base product as a monopolist does not and, thus, introduces a further inefficiency.

²⁰This shows that although for $c < (v_h + v_l)/2$, the firm might also want to serve all consumers in the indirect channel at $p^I = (v_h + v_l)/2$ yielding profit $(v_h + v_l)/2 - c$, such a strategy is not profit-maximizing.

²¹The firm does not find it profitable to sell to types θ , for which $\theta + \bar{\varepsilon} < c$. Therefore, in this case, we can always relabel the lowest type that satisfies $\theta + \bar{\varepsilon}$ as $\underline{\theta}$.

References

- Aköz, K. K., C. E. Arbatli, and L. Celik. 2020. "Manipulation Through Biased Product Reviews." *Journal of Industrial Economics* 68: 591–639.
- Ali, S. N., G. Lewis, and S. Vasserman. 2023. "Voluntary Disclosure and Personalized Pricing." *Review of Economic Studies* 90: 538–571.
- Anand, B. N., and R. Shachar. 2009. "Targeted Advertising as a Signal." *Quantitative Marketing and Economics* 7: 237–266.
- Anderson, S. P., and Bedre-Defolie, O. 2024. "Hybrid Platform Model: Monopolistic Competition and a Dominant Firm." *RAND Journal of Economics* 55: 684–718.
- Anderson, S. P., and R. Renault. 2006. "Advertising Content." *American Economic Review* 96: 93–113.
- Aridor, G., and D. Gonçalves. 2022. "Recommenders' Originals: The Welfare Effects of the Dual Role of Platforms as Producers and Recommender Systems." *International Journal of Industrial Organization* 83: 102845.
- Armstrong, M., and J. Zhou. 2011. "Paying for Prominence." *Economic Journal* 121: 368–395.
- Bakos, Y., and E. Brynjolfsson. 1999. "Bundling Information Goods: Pricing, Profits, and Efficiency." *Management Science* 45: 1613–1630.
- Bergemann, D., B. Brooks, and S. Morris. 2015. "The Limits of Price Discrimination." *American Economic Review* 105: 921–957.
- Burguet, R., R. Caminal, and M. Ellman. 2015. "In Google We Trust?" *International Journal of Industrial Organization* 39: 44–55.
- de Cornière, A., and G. Taylor. 2015. "Integration and Search Engine Bias." *RAND Journal of Economics* 45: 576–597.
- de Cornière, A., and G. Taylor. 2019. "A Model of Biased Intermediation." *RAND Journal of Economics* 50: 854–882.
- Deneckere, R. J., and R. P. McAfee. 1996. "Damaged Goods." *Journal of Economics & Management Strategy* 5: 149–174.
- Drugov, M., and M. Troya-Martinez. 2019. "Vague Lies and Lax Standards of Proof: On the Law and Economics of Advice." *Journal of Economics & Management Strategy* 28: 298–315.
- Etro, F. 2021. "Product Selection in Online Marketplaces." *Journal of Economics & Management Strategy* 30: 614–637.
- Geng, X., M. B. Stinchcombe, and A. B. Whinston. 2005. "Bundling Information Goods of Decreasing Value." *Management Science* 51: 662–667.
- Haghpanah, N., and J. Hartline. 2021. "When is Pure Bundling Optimal?" *Review of Economic Studies* 88: 1127–1156.
- Hagiu, A., and B. Jullien. 2011. "Why do Intermediaries Divert Search?" *RAND Journal of Economics* 42: 337–362.
- Hagiu, A., T.-H. Teh, and J. Wright. 2022. "Should Platforms be Allowed to Sell on Their Own Marketplaces?" *RAND Journal of Economics* 53: 297–327.
- Hagiu, A., and J. Wright. 2024. "Marketplace Leakage." *Management Science* 70: 1529–1553.
- Heidhues, P., M. Köster, and B. Kőszegi. 2023. "Steering Fallible Consumers." *Economic Journal* 133: 1430–1465.
- Hervas-Drane, A., and S. Shelegia. 2025. "Retailer-Led Marketplaces." *Management Science* 71: 9650–9669.
- Inderst, R., and M. Ottaviani. 2012. "Competition Through Commissions and Kickbacks." *American Economic Review* 102: 780–809.
- Johnson, J. P. 2013. "Targeted Advertising and Advertising Avoidance." *RAND Journal of Economics* 44: 128–144.
- Johnson, J. P., and D. P. Myatt. 2006. "On the Simple Economics of Advertising, Marketing, and Product Design." *American Economic Review* 96: 756–784.
- Kamenica, E., and M. Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101: 2590–2615.
- Kang, Z. Y., and E. V. Muir. 2022. "Contracting and Vertical Control by a Dominant Platform." Working Paper (presented at ACM EC 22), Stanford University.
- Lee, C. 2021. "Unpublished manuscript. Optimal Recommender System Design." Last modified December 7, 2023. University of Pennsylvania.
- Lewis, T. R., and D. E. M. Sappington. 1994. "Supplying Information to Facilitate Price Discrimination." *International Economic Review* 35: 214–231.
- Peitz, M., and M. Reisinger. 2016. "Media Economics of the Internet." In *Handbook of Media Economics*, edited by S. Anderson, D. Stromberg, and J. Waldfogel. Elsevier.
- Peitz, M., and A. Sobolev. 2025. "Biased Recommendations and Differentially Informed Consumers." *Economic Theory Bulletin* 13: 245–261.
- Rhodes, A., and C. M. Wilson. 2018. "False Advertising." *RAND Journal of Economics* 49: 348–369.
- Teh, T. H., and J. Wright. 2022. "Intermediation and Steering: Competition in Prices and Commissions." *American Economic Journal: Microeconomics* 14: 281–321.
- Wang, C., and J. Wright. 2020. "Search Platforms: Showrooming and Price Parity Clauses." *RAND Journal of Economics* 51: 32–58.
- Zenny, Y. 2022. "Platform Encroachment and Own-Content Bias." *Journal of Industrial Economics* 70: 684–710.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.

Appendix S1: Internet Appendix.

Appendix

Appendix A: Relegated Proofs and Further Details on the Extensions in Section 4.1

A.1 | Relegated Proofs

Proof of Proposition 1. We start by proving that the profit-maximizing seller induces one of three outcomes: either the inflated recommendations outcome, the inefficient bypass outcome, or the limited sales outcome.

First, we show that in a profit-maximizing outcome, the picky consumers are never served in the direct channel. By contradiction, suppose that some picky consumers choose the direct channel. A necessary condition is that $p^D \leq p^I - d$. If this is held with equality, then picky and flexible consumers are indifferent between the two channels. By slightly increasing p^D the firm can divert all consumers to the indirect channel earning an extra profit of d for each diverted consumer. If this is held with strict inequality, all picky and flexible consumers must buy in the direct channel. The firm could then set instead $\hat{p}^I = p^D + d$ and $\hat{p}^D > \hat{p}^I - d$, which would increase the firm's profit because the same number of consumers will buy and consumers who bought in the direct channel are diverted to the indirect channel generating an extra profit of d per consumer.

Second, we show that some picky consumers buy in the indirect channel. By contradiction, suppose that the intermediary only sells to flexible consumers. The profit-maximizing strategy to do so is to set $p^I = v_m$ and $p^D > v_m - d$ and, thereby, extract the full surplus from flexible consumers. For $\beta = 0$, also picky consumers with a high valuation will buy, which will lead to higher profits. Hence, the profit-maximizing outcome has the property that no picky consumer buys in the direct channel and at least some of them buy in the indirect channel.

Flexible consumers can be served in the indirect channel, in the direct channel, or not served at all. If the flexible consumers buy in the indirect channel, then $p^I \in [(v_l + v_h)/2, v_m]$. The number of picky consumers buying at p^I is maximized when the recommendation policy β solves $\frac{v_h + \beta v_l}{1 + \beta} = p^I$. The resulting profit from the picky consumers at price $p^I \in [(v_l + v_h)/2, v_m]$ is given by,

$$\frac{\alpha}{2}(1 + \beta)(p^I - c) = \frac{\alpha}{2}(v_h - c - \beta(c - v_l)),$$

and is decreasing in β .²⁰ Thus, if the flexible consumers buy in the indirect channel, the firm sets $p^I = v_m$, $p^D \geq \max\{v_m - d, 0\}$ and $\beta = \beta^*$ that solves $\frac{v_h + \beta v_l}{1 + \beta} = v_m$. This is the inflated recommendations outcome resulting in profits of

$$\left[\frac{\alpha}{2} \left(1 + \frac{v_h - v_m}{v_m - v_l} \right) + (1 - \alpha) \right] (v_m - c).$$

If the flexible consumers are served in the direct channel, then the profit is maximized when the recommendations are efficient, $\beta = 0$, and prices are set to $p^I = v_h$ and $p^D = v_m - d$. This is the inefficient bypass outcome. The firm makes a profit of

$$\frac{\alpha}{2}(v_h - c) + (1 - \alpha)(v_m - d - c).$$

If the flexibles do not buy at all, then the firm sets $\beta = 0$ and sells to the picky consumers with a good match at price $p^I = v_h$. The price in the direct channel $p^D > v_m$. This is the limited sales outcome.

Next, we compare the profits of the three potentially profit-maximizing outcomes. We derive the critical α as a function of c . In the case that $v_m - c - d < 0$, the inflated recommendations outcome leads to higher profits than the limited sales outcome if and only if,

$$\left[\frac{\alpha}{2}(1 + \beta) + (1 - \alpha) \right] (v_m - c) \geq \frac{\alpha}{2}(v_h - c),$$

which after substituting for $\beta = \frac{v_h - v_m}{v_m - v_l}$, can be rewritten as,

$$(1 - \alpha)(v_m - c) \geq \frac{\alpha}{2} \frac{1}{v_m - v_l} (c - v_l)(v_h - v_m).$$

This defines the critical $\bar{\alpha}(c)$, which can be rewritten as,

$$\bar{\alpha}(c) = \frac{v_m - c}{v_m - c + \frac{1}{2}(c - v_l) \frac{v_h - v_m}{v_m - v_l}}, \tag{1}$$

for $c \in (v_m - d, v_m)$.

In the other case that $v_m - d - c > 0$, the inflated recommendations outcome leads to higher profits than the inefficient bypass outcome if and only if,

$$(1 - \alpha)d \geq \frac{\alpha}{2}(v_h - c - (1 + \beta)(v_m - c)) = \frac{\alpha}{2}(c - v_l) \frac{v_h - v_m}{v_m - v_l}.$$

When satisfied with equality, this defines the critical $\bar{\alpha}$, which can be rewritten as,

$$\bar{\alpha}(c) = \frac{d}{d + \frac{1}{2}(c - v_l) \frac{v_h - v_m}{v_m - v_l}}, \tag{2}$$

for $c \in (\frac{v_l + v_h}{2}, v_m - d)$.

As $v_l - c < 0$, it must be that $\bar{\alpha}(c) \in (0, 1)$ in Equations (1) and (2). \square

Proof of Proposition 4. The seller weakly prefers catering to flexible consumers and picky consumers with a good match via the indirect channel rather than using the indirect channel for picky consumers only if,

$$(1 - \lambda) \left(\frac{\alpha}{2} + 1 - \alpha \right) (v_m - c) \geq (1 - \alpha)(v_m - d - c) + (1 - \lambda) \frac{\alpha}{2} (v_h - c).$$

The seller prefers to serve flexible consumers in the indirect channel instead of serving only picky consumers with a good match if,

$$\lambda \left(\frac{\alpha}{2} + 1 - \alpha \right) (v_m - c) \geq \frac{\alpha}{2} (v_h - c).$$

Solving these two inequalities for $1 - \lambda$, we obtain that,

$$\frac{(1 - \alpha)(v_m - d - c)}{(1 - \frac{\alpha}{2})(v_m - c) - \frac{\alpha}{2}(v_h - c)} \leq 1 - \lambda \leq \frac{(1 - \frac{\alpha}{2})(v_m - c) - \frac{\alpha}{2}(v_h - c)}{(1 - \frac{\alpha}{2})(v_m - c)}.$$

If there exists $\lambda \in [0, 1]$ satisfying these two inequalities, the intermediary picks the highest such λ ; that is, $\lambda = \hat{\lambda} = 1 - \frac{(1 - \alpha)(v_m - d - c)}{(1 - \frac{\alpha}{2})(v_m - c) - \frac{\alpha}{2}(v_h - c)}$. This implements the first-best outcome.

We see that, for positive $\alpha \approx 0$, there is a non-empty set of λ satisfying the two inequalities. The maximal λ in this set is then chosen by the intermediary and implements the first best. If $\alpha \approx 1$, then the seller always prefers to sell directly to the flexible consumers because $(1 - \frac{\alpha}{2})(v_m - c) - \frac{\alpha}{2}(v_h - c) < 0$. It is easy to see that the upper bound for $1 - \lambda$ decreases in α whereas the lower bound for $1 - \lambda$ increases in α . Therefore, there exists a unique $\alpha_{FI}(c) \in (0, 1)$ such that for all $\alpha \leq \alpha_{FI}(c)$, the regulation $\beta = 0$ is optimal and implements the first best. Otherwise, if $\alpha > \alpha_{FI}(c)$, the seller will serve the flexible consumers through the direct channel, leading to welfare loss equal to $(1 - \alpha)d$. We note that $\alpha_{FI} < \bar{\alpha}(c)$. This means that, for $\alpha \in (\alpha_{FI}(c), \bar{\alpha}(c))$, the laissez faire is welfare-superior. As we show in Section 3, for $\alpha < \bar{\alpha}(c)$, the welfare loss under laissez-faire is $\frac{\alpha}{2}\beta^*(c - v_l)$. Thus, we have to show that, for $\alpha < \bar{\alpha}(c)$, the inequality

$(1 - \alpha)d > \frac{\alpha}{2}\beta^*(c - v_l)$ must be satisfied. For $\alpha < \bar{\alpha}(c)$, we obtain,

$$\begin{aligned} (1 - \alpha)d - \frac{\alpha}{2}\beta^*(c - v_l) &> (1 - \bar{\alpha}(c))d - \frac{\bar{\alpha}(c)}{2}\beta^*(c - v_l) \\ &= \frac{\bar{\alpha}(c)}{2}(v_h - c - (1 + \beta^*)(v_m - c) - \beta^*(c - v_l)) \\ &= \frac{\bar{\alpha}(c)}{2}(v_h - v_m - \beta^*(v_m - v_l)) \\ &= 0. \end{aligned}$$

□

Proof of Remark 1. Suppose that the regulator sets $\beta = 0$ and a cap on the fraction of the rent extracted by the intermediary, $\bar{\lambda} \in [0, 1)$. The case of $\bar{\lambda} = 1$ was covered in Proposition 4. Define α_{UP} as the solution to

$$(1 - \alpha)d = \frac{\alpha}{2}(v_h - v_m).$$

Note that $\alpha_{UP} > \alpha_{FI}$ defined in Proposition 4 (for simplicity, here we no longer make the dependence on c explicit).

We show that for any $\bar{\lambda} < 1$ there exists a critical level $\alpha(\bar{\lambda}) \in (\alpha_{FI}, \alpha_{UP})$ such that for any $\alpha < \alpha(\bar{\lambda})$, the intermediary sets $\lambda = \min\{\bar{\lambda}, \hat{\lambda}\}$, where $\hat{\lambda}$ was defined in Proposition 4, and the seller induces the efficient outcome. Otherwise, the intermediary sets $\lambda = \bar{\lambda}$, and the inefficient bypass outcome is induced.

We characterize the equilibrium for $\alpha > \alpha_{UP}$ and $\alpha \leq \alpha_{UP}$ separately. Consider the case that $\alpha > \alpha_{UP}$. Note that for any $\lambda \in [0, 1]$ set by the intermediary, we have that

$$\begin{aligned} (1 - \lambda)\left(\frac{\alpha}{2} + 1 - \alpha\right)(v_m - c) &= (1 - \lambda)(1 - \alpha)(v_m - c) + (1 - \lambda)\frac{\alpha}{2}(v_m - c) \\ &< (1 - \lambda)(1 - \alpha)(v_m - c) + (1 - \lambda)\left(\frac{\alpha}{2}(v_h - c) - (1 - \alpha)d\right) \\ &< (1 - \alpha)(v_m - d - c) + (1 - \lambda)\frac{\alpha}{2}(v_h - c), \end{aligned}$$

implying that the seller's profit from serving flexible consumers in the direct channel at a price of $p^D = v_m - d$ and picky consumers with a good match in the indirect channel at a price of $p^I = v_h$ is greater than the profit from inducing the efficient outcome. Therefore, for any $\bar{\lambda}$, the intermediary optimally sets $\lambda = \bar{\lambda}$, and the seller induces the inefficient bypass outcome, earning a profit of $(1 - \alpha)(v_m - d - c) + (1 - \bar{\lambda})\frac{\alpha}{2}(v_h - c)$.

Next, consider the opposite case $\alpha \leq \alpha_{UP}$. Note that, for every $\alpha \leq \alpha_{UP}$, $\hat{\lambda}$ (as defined in Proposition 4) belongs to $[0, 1)$, and recall that, at $\lambda = \hat{\lambda}$, the seller is indifferent between inducing the efficient outcome and deviating such that it serves flexible consumers in the direct channel. Thus, the intermediary's profit when the seller weakly prefers to induce the efficient outcome is $\min\{\bar{\lambda}, \hat{\lambda}\}\left(1 - \frac{\alpha}{2}\right)(v_m - c)$. The intermediary's profit when the seller induces the inefficient bypass outcome is $\bar{\lambda}\frac{\alpha}{2}(v_h - c)$. Therefore, the intermediary sets $\lambda = \min\{\bar{\lambda}, \hat{\lambda}\}$ and the seller induces the efficient outcome if,

$$\min\{\bar{\lambda}, \hat{\lambda}\}\left(1 - \frac{\alpha}{2}\right)(v_m - c) - \bar{\lambda}\frac{\alpha}{2}(v_h - c) \geq 0. \quad (3)$$

Otherwise, the intermediary sets $\lambda = \bar{\lambda}$ and the seller induces the inefficient bypass outcome.

Note that,

$$\frac{d\hat{\lambda}}{d\alpha} = -\frac{\frac{1}{2}(v_h - v_m)(v_m - d - c)}{\left(\left(1 - \frac{\alpha}{2}\right)(v_m - c) - \frac{\alpha}{2}(v_h - c)\right)^2} < 0,$$

implying that the left-hand side of equation (3) is strictly decreasing in α . Moreover, it changes its sign from positive to negative in the interval

$(\alpha_{FI}, \alpha_{UP})$. This is seen as follows. Note that at $\alpha = \alpha_{FI}$, we have

$$\begin{aligned} \hat{\lambda}\left(1 - \frac{\alpha_{FI}}{2}\right)(v_m - c) - \bar{\lambda}\frac{\alpha_{FI}}{2}(v_h - c) &> \bar{\lambda}\left(1 - \frac{\alpha_{UP}}{2}\right)(v_m - c) - \bar{\lambda}\frac{\alpha_{UP}}{2}(v_h - c) = 0. \\ \hat{\lambda}\left(1 - \frac{\alpha_{FI}}{2}\right)(v_m - c) - \bar{\lambda}\frac{\alpha_{FI}}{2}(v_h - c) &> \hat{\lambda}\left(1 - \frac{\alpha_{FI}}{2}\right)(v_m - c) - \frac{\alpha_{FI}}{2}(v_h - c) = 0, \end{aligned}$$

implying that the left-hand side of Equation (3) is strictly positive at $\alpha = \alpha_{FI}$. Moreover, $\hat{\lambda} = 0$ at $\alpha = \alpha_{UP}$, implying the left-hand side of equation (3) is negative at $\alpha = \alpha_{UP}$. Therefore, for any $\bar{\lambda} < 1$ there exists a uniquely defined $\alpha(\bar{\lambda}) \in (\alpha_{FI}, \alpha_{UP})$ that solves $\hat{\lambda}(1 - \alpha/2)(v_m - c) = \bar{\lambda}\alpha/2(v_h - c)$, such that the equilibrium is characterized by the efficient outcome for $\alpha < \alpha(\bar{\lambda})$ and by the inefficient bypass outcome, otherwise. Note that $\alpha(\bar{\lambda})$ is strictly decreasing in $\bar{\lambda}$ and, therefore, admits an inverse function that we denote by $\bar{\lambda}(\alpha)$.

We turn to the regulator's total surplus maximization problem. If $\alpha \leq \alpha_{FI}$, the outcome is efficient for all $\bar{\lambda} \leq 1$. If $\alpha \geq \alpha_{UP}$, the equilibrium is characterized by the inefficient bypass outcome for all $\bar{\lambda} \leq 1$. Thus, for $\alpha \leq \alpha_{FI}$ and $\alpha \geq \alpha_{UP}$, the regulator is indifferent between any $\bar{\lambda}$. Otherwise, for any $\alpha \in (\alpha_{FI}, \alpha_{UP})$, the regulator strictly prefers to set $\bar{\lambda} \in [0, \bar{\lambda}(\alpha)]$ over any $\bar{\lambda} > \bar{\lambda}(\alpha)$ and is indifferent between all $\bar{\lambda}$ in $[0, \bar{\lambda}(\alpha)]$.

To compare the total surplus under regulation with the total surplus under laissez-faire, it is sufficient to compare α_{UP} with $\bar{\alpha}$, defined in Proposition 2 (again, we no longer make the dependence on c explicit). As was shown in the proof of Proposition 4, $\bar{\alpha}$ solves

$$(1 - \alpha)d - \frac{\alpha}{2}\beta^*(c - v_l) = 0.$$

We now evaluate the expression on the left-hand side at $\alpha = \alpha_{UP}$ and show that it is always positive. We obtain that,

$$\begin{aligned} (1 - \alpha_{UP})d - \frac{\alpha_{UP}}{2}\beta^*(c - v_l) &= \frac{\alpha_{UP}}{2}(v_h - v_m - \beta^*(c - v_l)) \\ &= \frac{\alpha_{UP}}{2} \frac{(v_h - v_m)(v_m - c)}{v_m - v_l} > 0, \end{aligned}$$

implying that $\alpha_{UP} < \bar{\alpha}$. Therefore, the regulatory policy $\beta = 0$ and $\bar{\lambda} = 0$ increases welfare compared to the laissez faire for $\alpha \leq \alpha_{UP}$, reduces welfare for $\alpha \in (\alpha_{UP}, \bar{\alpha})$, and has no effect for $\alpha \geq \bar{\alpha}$. □

A.2 | Details on the Extensions in Section 4.1

A.2.1 | Platform Leakage

With inflated recommendations, as in the main model, the level β will be such that picky consumers with a recommendation have the same expected valuation as flexible consumers; that is, $\beta^* = \frac{v_h - v_m}{v_m - v_l}$. If the seller were to destabilize the outcome by setting a price p^D below $v_m - d$ such that all trade occurs in the indirect channel, then the intermediary is committed to make recommendations uninformative and thereby makes sure that the seller does not have an incentive to do so. Thus, the intermediary obtains the same fraction of industry profit λ as under no leakage. With inefficient bypass, the seller sells directly to flexible consumers at $p^D = v_m - d$ and aims to serve picky consumers with a good match at $p^I = v_h$. Given these prices, with platform leakage, the fraction ν of picky consumers with a good match buys directly and each consumer obtains a net surplus of $v_h - v_m$. The seller's profit is $(1 - \alpha + \nu\frac{\alpha}{2})(v_m - d - c)$ and platform leakage occurs along the equilibrium path with inefficient bypass.

With inefficient bypass, the intermediary obtains $(1 - \nu)\frac{\alpha}{2}(v_h - c)$, which is less than in the main model and decreasing in the share ν , whereas the intermediary's profit remains unchanged when inducing the inflated recommendations outcome. Therefore, the critical α that separates inflated recommendations from inefficient bypass is higher with platform leakage. Moreover, the critical α increases in the degree of platform leakage ν .

A.2.2 | The Intermediary With a Profitable Outside Option

Our setting features an outside option that does not generate any profit for the intermediary. Suppose instead a situation in which the outside option consists of the intermediary selling its own product with gains from trade $v_0 - c_0 > 0$; that is, consumers value the intermediary's product at v_0 and the intermediary incurs a cost of c_0 per unit of providing its own product, where $0 < v_0 - c_0 < v_m - c$. The intermediary is assumed to set the price p_0 of its product after the seller has set its price.

We will show that for small α the equilibrium features inflated recommendations. In such an equilibrium, suppose that the seller diverts flexible consumers to the direct channel by setting a lower price p^D , which implies that $\beta = 1$. In response, the intermediary profitably undercuts by setting the price $p_0 = p^D - (v_m - d - v_0)$ and sells to all consumers if $p^D - (v_m - d - v_0) - c_0 \geq \alpha(v_0 - c_0)$, where on the right-hand side we have the profit of the intermediary to sell the base product to picky consumers only. To avoid the intermediary's undercutting, the seller's best deviation is to set $p^D = v_m - d - (1 - \alpha)(v_0 - c_0)$, which gives a deviation profit of $(1 - \alpha)(v_m - d - c - (1 - \alpha)(v_0 - c_0))$.

Along the equilibrium path, the seller will set $p^I = v_m$, $p^D \geq v_m - d$. The intermediary's recommendation policy is $\beta(p^I = v_m, p^D \geq v_m) = \beta^*$ and $\beta(p^I, p^D) = 1$ for all other prices. The intermediary will set λ such that $(1 - \lambda)[\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha)](v_m - c) = (1 - \alpha)(v_m - d - c - (1 - \alpha)(v_0 - c_0))$ and, along the equilibrium path, the price of the base product $p_0 \geq v_0$. This is the equilibrium characterization for a sufficiently small α .

We now turn to the outcome with inefficient bypass. Here, $\lambda = 1$ and $\beta(\cdot, \cdot) = 0$ for all prices (p^D, p^I) . We have to specify the prices set by the seller along the equilibrium path. The seller will set $p^I = v_h$ and $p^D \leq v_m - d$ such that the intermediary does not have an incentive to set the price for the base product in such a way that flexible consumers buy from the intermediary. With $\lambda = 1$, the seller makes a profit from selling to the flexible consumers, $(1 - \alpha)(p^D - c)$. In the candidate equilibrium the intermediary sells its own product to picky consumers with a bad match and makes profit from all picky consumers: $\frac{\alpha}{2}(v_0 - c_0) + \frac{\alpha}{2}(v_h - c)$. Selling the base product to all consumers gives at most $v_0 - c_0$ to the intermediary. The intermediary then does not interfere with the seller selling to flexible consumers at $p^D = v_m - d$, if $\frac{\alpha}{2}(v_0 - c_0) + \frac{\alpha}{2}(v_h - c) \geq v_0 - c_0$ or, equivalently, $\alpha \geq \check{\alpha}(v_0 - c_0) \equiv 2(v_0 - c_0)/(v_0 - c_0 + v_h - c)$. Note that $\check{\alpha}(v_0 - c_0)$ tends to 0 as $v_0 - c_0 \rightarrow 0$.

The intermediary prefers to induce the inflated recommendations outcome over the inefficient bypass outcome if,

$$\begin{aligned} & \left[\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha) \right] (v_m - c) - (1 - \alpha)(v_m - d - c - (1 - \alpha)(v_0 - c_0)) \\ & \geq \frac{\alpha}{2}(v_0 - c_0) + \frac{\alpha}{2}(v_h - c), \end{aligned}$$

which simplifies to

$$\frac{\alpha}{2}(1 + \beta^*)(v_m - c) + (1 - \alpha)d - \frac{\alpha}{2}(v_h - c) + \left((1 - \alpha)^2 - \frac{\alpha}{2} \right) (v_0 - c_0) \geq 0.$$

By Proposition 1, we have that the sum of the first three terms is positive for $\alpha < \bar{\alpha}(c)$. The fourth term is non-negative of $\alpha \leq 1/2$. Moreover, for sufficiently small $v_0 - c_0$, $\check{\alpha}(v_0 - c_0) < \bar{\alpha}(c)$. Therefore, a sufficient condition for the intermediary to induce the inflated recommendations outcome is $\alpha \in (\check{\alpha}(v_0 - c_0), \max\{\bar{\alpha}(c), 1/2\})$ for sufficiently small $v_0 - c_0$.

A.2.3 | The Intermediary's Price Instrument

When the equilibrium features inflated recommendations, we have shown in Proposition 2 that the intermediary sets λ to guarantee the profit $(1 - \alpha)(v_m - c - d)$ to the seller and $\beta^* = \frac{v_h - v_m}{v_m - v_l}$. This gives the intermediary's equilibrium profit $\frac{\alpha}{2}(1 + \beta^*)(v_m - c) - (1 - \alpha)d$. If, instead, only a listing fee is available, this fee is set equal to this profit and, together with $\beta(p^I = v_m, p^D \geq v_m - d) = \beta^*$ implements the same

outcome. If the intermediary uses a per-unit fee, it sets this fee as

$$t = \left(\frac{\alpha}{2}(1 + \beta^*)(v_m - c) - (1 - \alpha)d \right) / \left((1 - \alpha) + \frac{\alpha}{2}(1 + \beta^*) \right)$$

with the same recommendation policy and implements the vertically integrated solution. This is also achieved if the intermediary sets an ad valorem transaction fee: this fee τ is set, such that the seller is indifferent between selling only in the indirect channel at $p^I = v_m$ and selling to flexible consumers in the direct channel at $p^D = v_m - d$; that is, $[(1 - \tau)v_m - c](1 - \alpha + \frac{\alpha}{2}(1 + \beta^*)) = (1 - \alpha)(v_m - c - d)$.

The reason that the intermediary achieves the same equilibrium profit under different price instruments is that a deviating seller will be "punished" with $\beta = 1$ and, thus, no trade takes place in the indirect channel, making the type of price instrument off the equilibrium path irrelevant. Furthermore, each price instrument allows the intermediary to achieve the industry profit of the vertically integrated solution minus the profit of the seller's outside option to sell directly to flexible consumers at $p^D = v_m - d$.

When the equilibrium features inefficient bypass, the invariance of the allocation in response to the type of price instrument being used by the intermediary is straightforward: the intermediary absorbs the entire profit in the indirect channel, $\frac{\alpha}{2}(v_h - c)$ and the seller makes a profit of $(1 - \alpha)(v_m - d - c)$ in the direct channel. The seller will sell to picky consumers in the indirect channel if it avoids a loss there. If the intermediary sets a profit share, the entire profit in the indirect channel is extracted with $\lambda = 1$. If, instead, only a fixed fee is available, the fee is set to $T = \frac{\alpha}{2}(v_h - c)$; if only a per-unit transaction fee is available, the fee is set to $t = v_h - c$; and if only an ad valorem transaction fee is available, the fee is set to $\tau = \frac{v_h - c}{v_h - d - c}$.

Appendix B: Inflated Recommendations, Match Values, and Ex Ante Consumer Types

B.1 | Two Consumer Types With the Same Binary Match Value Distribution

As in the base model, a fraction α of consumers are of low type θ_l and $1 - \alpha$ are of high type θ_h , where $\theta_h > \theta_l$. However, we assume that ε is drawn with equal probability from $\{\underline{\varepsilon}, \bar{\varepsilon}\}$, which is independent of the type θ . With $E\varepsilon = 0$, we have $\varepsilon = -\bar{\varepsilon}$. Thus, there are four ex post types, namely $\theta_l + \underline{\varepsilon}$, $\theta_l + \bar{\varepsilon}$, $\theta_h + \underline{\varepsilon}$, and $\theta_h + \bar{\varepsilon}$. In line with the base model, we assume that $\theta_l + \bar{\varepsilon} > \theta_h$ and $c > \theta_l + \underline{\varepsilon}$. Furthermore, to restrict the analysis to situations in which some of both consumer types may buy, we assume that $\theta_l + \bar{\varepsilon} > c$.

Then, the following two cases are possible: (i) $c \in (\theta_l + \underline{\varepsilon}, \theta_h + \underline{\varepsilon}]$ and (ii) $c \in (\theta_h + \underline{\varepsilon}, \theta_l + \bar{\varepsilon})$.

In case (i), the firm may want to set $p = \theta_h$ and recommend the product to a fraction β_l of low-type consumers with a bad match. The conditional expected valuation of a low-type consumer who receives a recommendation is $\frac{\theta_l + \bar{\varepsilon} + \beta_l(\theta_l + \underline{\varepsilon})}{1 + \beta_l} = \theta_l + \frac{(1 - \beta_l)}{1 + \beta_l} \bar{\varepsilon}$. The profit-maximizing β_l satisfies $\theta_l + \frac{(1 - \beta_l)}{1 + \beta_l} \bar{\varepsilon} = \theta_h$ and, therefore,

$$\beta_l = \frac{\theta_l + \bar{\varepsilon} - \theta_h}{\theta_h + \bar{\varepsilon} - \theta_l} = \frac{\bar{\varepsilon} - (\theta_h - \theta_l)}{\bar{\varepsilon} + (\theta_h - \theta_l)},$$

with the property that $\beta_l \in (0, 1)$. Here, we observe inflated recommendations to the low type.

Alternatively, the firm extracts the full surplus of low-type consumers with a good match by setting $p = \theta_l + \bar{\varepsilon}$ and recommends the product to only a fraction β_h of high-type consumers with a bad match. This fraction is chosen such that the conditional expected valuation of a high-type consumer who receives a recommendation is equal to this price $\theta_l + \bar{\varepsilon}$;

that is, $\frac{\theta_h + \bar{\varepsilon} + \beta_h(\theta_h + \underline{\varepsilon})}{1 + \beta_h} = \theta_l + \bar{\varepsilon}$, which is equivalent to

$$\beta_h = \frac{\theta_h - \theta_l}{2\bar{\varepsilon} - (\theta_h - \theta_l)},$$

with the property that $\beta_l \in (0, 1)$. Here, the firm provides socially insufficient recommendations because $\theta_h + \underline{\varepsilon} < c$.

Using the former strategy gives a profit of $(\theta_h - c)(1 - \alpha + \alpha(1 + \beta_l)/2)$, whereas the latter gives $(\theta_l + \bar{\varepsilon} - c)((1 - \alpha)(1 + \beta_h)/2 + \alpha/2)$. Therefore, the firm finds it optimal to use the former strategy and inflate recommendations for the low-type consumers if,

$$\alpha < \frac{\theta_h - c - \frac{1 + \beta_h}{2}(\theta_l + \bar{\varepsilon} - c)}{\frac{1 - \beta_l}{2}(\theta_h - c) - \frac{\beta_h}{2}(\theta_l + \bar{\varepsilon} - c)} = \frac{\frac{1 - \beta_h}{2}(\theta_h + \underline{\varepsilon} - c)}{\frac{1 - \beta_h}{2}(\theta_h + \underline{\varepsilon} - c) + \bar{\varepsilon} + \frac{\beta_l}{2}(\theta_l + \bar{\varepsilon} - c)}.$$

Otherwise, if α is greater than this threshold, the firm employs the latter strategy and provides socially insufficient recommendations to the high-type consumers.

In case (ii), the firm may want to set $p = \theta_h + \bar{\varepsilon}$ and recommend the product only to high-type consumers with a good match. Alternatively, as in case (i), the firm extracts the full surplus of low-type consumers with a good match by setting $p = \theta_l + \bar{\varepsilon}$ and recommends the product only to a fraction β_h of high-type consumers with a bad match with the same β_h as in case (i). As in case (ii), $c > \theta_h + \bar{\varepsilon}$, the firm inflates recommendations.

Using the former strategy gives a profit of $(\theta_h + \bar{\varepsilon} - c)(1 - \alpha)/2$, whereas the latter gives $(\theta_l + \bar{\varepsilon} - c)((1 - \alpha)(1 + \beta_h)/2 + \alpha/2)$. Thus, if

$$\alpha < \frac{\beta_h(c - \theta_h - \underline{\varepsilon})}{\beta_h(c - \theta_h - \underline{\varepsilon}) + (\theta_l + \bar{\varepsilon} - c)},$$

then the firm employs the first strategy and sells only to high-type consumers with a good match. Otherwise, the firm sets $p = \theta_l + \bar{\varepsilon}$ and recommends the product to some high-type consumers with a valuation less than the price. We summarize our analysis in the following proposition.

Proposition 5. *With two consumer types θ_h and θ_l with $\theta_h > \theta_l$ that face the same binary distribution of match values drawn from $\{\underline{\varepsilon}, \bar{\varepsilon}\}$ where $\theta_l + \bar{\varepsilon} > \max\{\theta_h, c\}$ and $c > \theta_l + \underline{\varepsilon}$, given a single sales channel, the monopoly solution takes the following form:*

For $c \in (\theta_l + \underline{\varepsilon}, \theta_h + \underline{\varepsilon})$,

- The firm recommends the product to a fraction of high-type consumers with a bad match, sells at $p = \theta_l + \bar{\varepsilon}$ to all high-type consumers with a recommendation and low-type consumers with a good match or
- The firm recommends the product to a fraction of low-type consumers with a bad match, sells at $p = \theta_h + \underline{\varepsilon}$ to all high-type consumers and all low-type consumers with a recommendation (inflated recommendations outcome).

For $c \in (\theta_h + \underline{\varepsilon}, \theta_l + \bar{\varepsilon})$,

- The firm recommends and sells the product at $p = \theta_h + \bar{\varepsilon}$ to high-type consumers with a good match or
- The firm recommends the product to a fraction of high-type consumers with a bad match, sells at $p = \theta_l + \bar{\varepsilon}$ to all high-type consumers with a recommendation and low-type consumers with a good match (inflated recommendations outcome).

B.2 | Continuous Distribution of Match Values

Here, we assume that the match value distribution is continuous: $\varepsilon \sim F : [\underline{\varepsilon}, \bar{\varepsilon}]$, where $\underline{\theta}, \underline{\varepsilon} \in \mathbb{R} \cup \{-\infty\}$ and $\bar{\theta}, \bar{\varepsilon} \in \mathbb{R} \cup \{+\infty\}$. Without loss of generality we assume that selling to the consumer of the lowest type and the highest match value, $(\underline{\theta}, \bar{\varepsilon})$, generates a positive surplus (i.e., $\underline{\theta} + \bar{\varepsilon} \geq c$).²¹ We also assume that selling to the consumer of the highest

type and the lowest match value, $(\bar{\theta}, \underline{\varepsilon})$, is inefficient from a total welfare perspective (i.e., $\bar{\theta} + \underline{\varepsilon} < c$).

Recall that consumers observe their type θ but are uncertain about the match value ε and that the firm sets a uniform price p and can provide product recommendations that reveal some information about the match value ε to consumers and can be conditioned on θ . In the following lemma (which is proved in the online appendix), we show that the profit-maximizing recommendation strategy has a cutoff structure for any $\theta \in [\underline{\theta}, \bar{\theta}]$.

Lemma 1. *With match values ε drawn from a continuous distribution, for any price $p \in (c, \theta + \bar{\varepsilon})$, the firm maximizes its profit by recommending a consumer of type θ to buy the good if and only if ε is greater than some cutoff level $\hat{\varepsilon} = \hat{\varepsilon}(p, \theta) \in (\underline{\varepsilon}, \bar{\varepsilon})$. The cutoff level $\hat{\varepsilon}(p, \theta)$ is chosen such that consumers are indifferent between buying and taking the outside option; it solves,*

$$p = \mathbb{E}[\theta + \varepsilon | \varepsilon \geq \hat{\varepsilon}(p, \theta)]. \tag{4}$$

To analyze the price setting, it is useful to first consider the profit-maximizing pricing and recommendation strategy of a firm that can price discriminate according to the consumer type θ . Define $\pi(p, \theta)$ as the maximal profit from consumers of type θ when price p is set, and the profit-maximizing recommendation policy is chosen. In the following lemma (which is proved in the online appendix), we show that under price discrimination, the recommendation policy is socially optimal.

Lemma 2. *Suppose that the firm can condition its price on θ . With match values ε drawn from a continuous distribution, the firm recommends consumers of type θ to buy the good if and only if $\theta + \varepsilon \geq c$ and sets price*

$$p(\theta) = \mathbb{E}[\theta + \varepsilon | \theta + \varepsilon \geq c]. \tag{5}$$

Consumers follow the recommendations. The equilibrium profit from the consumers of type θ is given by:

$$\pi(p(\theta), \theta) = (1 - F(c - \theta))\mathbb{E}[\theta + \varepsilon - c | \theta + \varepsilon \geq c]. \tag{6}$$

Note that the firm recommends buying the product if and only if $\theta + \varepsilon \geq c$, which is equivalent to the recommendation policy of the social planner.

We define the firm's recommendations to be *inflated* for type θ if the firm recommends buying the good to some consumers of type θ with $\theta + \varepsilon < c$ and consumers follow the recommendations. Similarly, the firm's recommendations are *insufficient* for type θ if the firm recommends not buying for some $\theta + \varepsilon > c$ and consumers follow the recommendations.

In the remainder of this appendix, we explore the conditions under which the firm's recommendations are inflated or insufficient from a total surplus perspective. We separately analyze the case in which θ can take only two values and the case in which θ is continuously distributed.

B.2.1 | Binary Consumer Types

Suppose that θ is distributed according to a binary distribution. We assume that $\theta = \theta_l$ with a probability of $\alpha \in (0, 1)$ and $\theta = \theta_h > \theta_l$ with a probability of $1 - \alpha$. For the model to be non-degenerate, we assume that $\theta_l + \bar{\varepsilon} > c$. The match value distribution is assumed to have a finite mean, which, without loss of generality, is set equal to zero. We assume that $1 - F$ is strictly log-concave on $[\underline{\varepsilon}, \bar{\varepsilon}]$.

As the following proposition establishes (the proof is relegated to the online appendix), under some conditions the firm maximizes its profit by inflating recommendations for type θ_h and providing insufficient recommendations for type θ_l . In the proposition, prices p_h and p_l are defined as the profit-maximizing prices of a firm that (third-degree) price discriminates between the two consumer types with p_h for the high type and p_l for the low type.

Proposition 6. *Suppose that the firm sets the same price for all consumers. Then, in the setting with binary types and a continuous match value distribution, there exists $\bar{\alpha} \in (0, 1)$ such that the firm provides insufficient*

recommendations for the low type and inflates recommendations for the high type,

- if $p_h \leq \theta_l + \bar{\varepsilon}$ or
- if $p_h > \theta_l + \bar{\varepsilon}$ and $\alpha \geq \bar{\alpha}$.

Otherwise, if $p_h > \theta_l + \bar{\varepsilon}$ and $\alpha < \bar{\alpha}$, the recommendations are efficient for the high type and consumers of the low type do not buy the product.

We note that, for $\bar{\varepsilon} = \infty$, we always have insufficient recommendations for the low type and inflated recommendations for the high type.

B.2.2 | Continuous Consumer Types

Suppose that θ is continuously distributed according to G on $[\underline{\theta}, \bar{\theta}]$, where $\underline{\theta}, \bar{\theta} \in \mathbb{R}$ and ε is continuously distributed according to F with support $(-\infty, \infty)$. We assume that $1 - F$ is log-concave.

If the firm were able to condition its price on θ , then by Lemma 2, the profit-maximizing pricing strategy is given by

$$p(\theta) = E[\theta + \varepsilon | \theta + \varepsilon \geq c],$$

and the profit-maximizing recommendation policy prescribes to recommending the product to consumers of type θ and match value ε if and only if $\theta + \varepsilon \geq c$. We note that $p(\theta)$ strictly increases in θ (implied by the strict log-concavity of $1 - F$; see the Proof of Proposition 6). The following proposition characterizes the profit-maximizing recommendation strategy when the firm has to set the same price for all consumer types—the proof is relegated to the online appendix.

Proposition 7. *Suppose that the firm sets a uniform price. Then, in the setting with a continuous type distribution and a continuous match value distribution with support $(-\infty, \infty)$, there exists a marginal type $\hat{\theta} \in (\underline{\theta}, \bar{\theta})$ such that the firm inflates recommendations for $\theta > \hat{\theta}$, provides insufficient recommendations for $\theta < \hat{\theta}$, and provides efficient recommendations for $\theta = \hat{\theta}$.*

Higher-type consumers sometimes regret the purchase ex post as they receive a recommendation too often from a consumer welfare perspective. Also from a total welfare perspective, they receive a recommendation too often, whereas lower-type consumers receive socially insufficient recommendations.