

Discussion Paper Series – CRC TR 224

Discussion Paper No. 336

Project B 03, B 04, B 05

Inflated Recommendations

Martin Peitz ¹
Anton Sobolev ²

February 2022

¹Department of Economics and MaCCI, University of Mannheim, E-Mail: Martin.Peitz@gmail.com

²Department of Economics and MaCCI, University of Mannheim, E-Mail: anton.sobolev@unimannheim.de

Funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation)
through CRC TR 224 is gratefully acknowledged.

Inflated Recommendations*

Martin Peitz[†]

Anton Sobolev[‡]

first version: October 26, 2021

this version: February 14, 2022

Abstract

Biased recommendations arise naturally in a market with heterogeneous consumers: A seller offers a product to a mix of “picky” and “flexible” consumers who can purchase through an intermediary or directly from the seller. A picky consumer either encounters a good or a bad match, while a “flexible” consumer is indifferent about the product design. Consumers know whether they are picky or flexible, but picky consumers observe match quality only after purchase and, therefore, rely on the intermediary’s recommendation. We provide conditions under which the intermediary decides to recommend a welfare-reducing bad match with positive probability, resulting in inflated recommendations. A regulatory intervention that prohibits recommending bad matches may backfire. The optimal regulation that limits the rate at which the product can be recommended performs better than the *laissez-faire*.

Keywords: intermediation, digital platforms, recommendation bias, recommender system, asymmetric information, experience good, e-commerce

JEL-classification: L12, L15, D21, D42, M37

*We are grateful to Paul Heihues, Nicolas Schutz, Julian Wright, and various seminar audiences for useful comments. The authors gratefully acknowledge support by the Deutsche Forschungsgemeinschaft (DFG) through CRC TR 224 (projects B03, B04 and B05).

[†]Department of Economics and MaCCI, University of Mannheim, 68131 Mannheim, Germany; E-Mail: Martin.Peitz@gmail.com; also affiliated with CEPR, CESifo, and ZEW.

[‡]Department of Economics and MaCCI, University of Mannheim, 68131 Mannheim, Germany; E-Mail: anton.sobolev@uni-mannheim.de.

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 recommender systems, it has become possible for a firm to pick its target audience carefully and to provide recommendations that fit a consumer’s taste. People may differ in how they intend to use a product. Intermediaries may then help people in identifying products that fit a specific purpose.¹

We develop a parsimonious model in which a profit-maximizing intermediary decides whether or not to give a personalized recommendation of a new product to a consumer. While many e-commerce sites provide more than one recommendation, several such as Amazon assign the buy button to a single seller. Also, consumers will typically receive only a single recommendation if recommendations are provided by voice (as in the case of virtual assistants Alexa, Cortana, Google’s Assistant, and Siri). Relatedly, if an algorithm decides whether or not a particular consumer is shown an “editor’s pick” of a new product such recommendation is in line with our model. More broadly, an intermediary may increase the visibility of certain offers while reducing the visibility of others. Recommendation algorithms can be expected to ultimately serve the interests of the platform.²

What could make the intermediary recommend products in the interest of consumers? If consumers find out that they received a recommendation for a product that does not fit their taste they may take note when responding to receiving further recommendations in the future. This suggests that, in its own interest, the intermediary will not recommend a product that does not fit a consumer’s taste and a subsequent transaction would reduce social surplus. When the intermediary can publicly commit to its recommendation policy this would imply that it does not recommend a bad match. As we show in this paper, this reasoning is incomplete because a profit-maximizing intermediary catering to a diverse set of consumers may recommend bad matches to a certain extent and, thus, inflate recommendations.

Our starting point is the following: The intermediary’s policy to inflate recommendations may reduce the heterogeneity of expected consumer valuations conditional on receiving a recommendation. Then, inflated recommendations allow for better extract surplus extraction on the consumer side. We formalize this basic idea in a setting in which an intermediary offers its recommendation service to a monopoly seller who lacks information about consumer characteristics but can also sell directly to consumers and thus bypass the intermediary, albeit inefficiently.

¹For example, the intermediary Wirecutter hires people to test different products to figure out which one does best for a specific purpose. It then provides the affiliate link and takes a cut 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.” Quote is taken from Dana Mattioli, “Amazon Changed Search Algorithm in Ways That Boost Its Own Products,” *Wall Street Journal*, 16 September 2019.

To set the scene, the intermediary carries a base product in competitive supply and adds a new product to its portfolio – the product in competitive supply constitutes a consumer’s outside option. The seller of the new product can also sell the new product directly to consumers, but lacks the information to make informed purchase recommendations (or lacks the credibility to do so). The seller sets the retail prices – there is one price for each channel, but this price applies to all consumers. We consider an experience good setting in which some consumers are sensitive to the product design of the new product or, in other words, *picky*, whereas others are *flexible* and do not mind the particular features of the product. More specifically, consumers have unit demand and valuations that depend on whether they are picky or flexible and, when picky, whether the match is good or bad. In our setting, there are positive gains from trade whenever the match is not bad, while production costs exceed consumer valuations if the match turns out to be bad for a picky consumer irrespective of the sales channel used. Consumers find it more convenient to buy via the intermediary. On top, the intermediary may provide informative recommendations. Since we assume that the intermediary is fully informed about the consumer type and the match quality, it can give personalized purchase recommendations (based on the information it has on each consumer) that are conditioned on the retail prices set by the seller – this describes the intermediary’s recommendation policy. In addition, the intermediary decides about the percent fee it takes of the seller’s profit.

A profit-maximizing intermediary may want to “rotate” the demand of the picky consumers by not only recommending the product if it provides a good fit, but also sometimes when it does not. “Inflated recommendations” is the feature of a recommendation policy according to which the product is recommended to a fraction $\beta > 0$ of picky consumers with a bad match. The optimal recommendation policy with inflated recommendation is such that a picky consumer’s expected valuation after receiving a recommendation is equal to a flexible consumer’s valuation. If the fraction of picky consumers is below some critical level, this maximizes the intermediary’s profit and makes recommendations only partially informative. Compared to the first best, recommendations are inflated, as a purchase with a bad match is total surplus decreasing. If the fraction of picky consumers is above this critical level, the equilibrium features inefficient bypass: The seller sells to flexible consumers in the direct channel and to picky consumers with a good match in the indirect channel.

As we show, the intermediary either chooses inflated recommendations or allows for inefficient bypass, and the equilibrium implements the vertically integrated solution. If the intermediary cannot commit to its recommendation policy – that is, the seller sets prices before the intermediary decides on whether to recommend the product – the intermediary cannot “punish” the seller as severely for deviations from the equilibrium under commitment. Hence, the intermediary adjusts its recommendation policy without commitment and recommendation inflation will occur on a smaller set of parameters than under vertical integration.

Inspired by the debate on regulatory interventions in the case of biased recommendations,³ we consider

³Current policy proposals in Europe and the U.S. focus on the practice of self-preferencing – see the Digital Markets Act

several regulatory policies that restrict the intermediary’s choice of recommendation policy. First, as a benchmark, we allow the regulator to set the welfare-maximizing recommendation policy that conditions on retail prices, while the intermediary continues to decide on the profit share (and the seller continues to set retail prices). We show that, in our setting, this regulatory policy implements the first best. Second, we assume that the only feasible regulatory intervention is to mandate fully informative recommendations. Such regulation may improve on the laissez-faire and sometimes even implements the first-best when the laissez-faire does not. However, it runs the risk to backfire since in some environments this policy delivers lower welfare than the laissez-faire. Third, the optimal policy that imposes a recommendation cap – that is, an upper bound on the recommendation inflation – improves on the policy to mandate fully informative recommendation and always performs at least as good as the laissez-faire from a welfare perspective.

Related literature. Our paper is most closely related to the work on biased recommendations, both in the economics and the computer science literature. We focus our discussion on the former. A profit-maximizing intermediary may have incentives to provide biased recommendations for a variety of reasons. In Lee (2021), the intermediary is a mechanism designer who has to persuade consumers to buy the recommended product of oneself. 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).⁴ 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).

Our paper contributes to this literature. Inflated recommendations means that a product is recommended more often than is socially optimal — we develop our argument in the context of experience goods.⁵ By contrast, most of the industrial organization literature on the topic considers search goods. In such a context, biased recommendations describe a situation in which the number of recommended products differs from

proposed by the European Commission in the EU and the proposed U.S. Senate Bill “American Innovation and Choice Online Act.” 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 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.

⁴If consumers suffer from limited cognition an intermediary may exploit such consumers by providing biased recommendations (Heidhues et al., 2020).

⁵Empirical and theoretical work has looked at biased financial advice. In our paper, as in Inderst and Ottaviani (2012) and Teh and Wright (forthcoming), neither the seller nor the buyer have 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. While in Inderst and Ottaviani (2012) and Teh and Wright (forthcoming) sellers compete in those kickbacks, in our setting, the intermediary decides on those fees.

what is socially optimal or inferior (or higher-priced) products are recommended earlier under sequential search.

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 (see, e.g. de Cornière and Taylor, 2019). Such a firm may have an incentive to steer consumers to its own products. Self-preferencing as an allegedly anti-competitive practice is under investigation by competition authorities and raise the question as to which regulatory interventions increase consumer or total surplus (for formal investigations, see, e.g., Anderson and Bedre-Defolie, 2021; Aridor and Gonçalves, 2022; Etro, 2021; Hagiü, Teh, and Wright, forthcoming; Hervas-Drane and Shelegia, 2021; Kang and Muir, 2022; Zennyö, forthcoming).⁶ In our base model inflated recommendations constitute self-preferencing if the intermediary is vertically integrated with the seller of the new product. Some competition experts in the U.S. take issue with the hybrid mode and even consider prohibiting it (e.g., Khan, 2017). However, according to our model, vertical disintegration (of the new product) is either ineffective (in our base model) or leads to a welfare-inferior outcome (if the intermediary lacks commitment power with respect to its recommendation policy).⁷ In a modified version of our model, the intermediary is vertically integrated with the base product and has to decide about which consumers to recommend an innovative third-party product. We show that, under some conditions, the intermediary inflates the recommendation of this third-party product – this is the opposite of self-preferencing.

Absent vertical integration, with uniform seller fees,⁸ the intermediary has an incentive to steer consumers to those sellers that lead to a higher conversion rate – that is a higher probability that the transaction is concluded via the intermediary. For example, some sellers may be able to divert some consumers to a direct sales channel and, thus, have a lower conversion rate in the indirect channel. If the intermediary cannot monetize transactions that were initiated by the intermediary but were diverted, it has an incentive to bias its recommendation against such sellers.⁹ In an extension, we introduce platform leakage (Hagiü and

⁶As analyzed by de Cornière and Taylor (2014), recommendation biases may also arise in the context of ad-financed search engines and ad-finance website when consumers experience advertising as a nuisance. Vertical integration between search engine and one of the websites sometimes increases and sometimes reduces the recommendation bias. For a related model, see Burguet et al. (2015) and, for an overview, Peitz and Reisinger (2016).

⁷In Anderson and Bedre-Defolie (2021) welfare results are ambiguous.

⁸If sellers pay different fees, it is clear that the intermediary tends to prioritize sellers who pays the higher fee. This moves the transaction-based monetization model closer to an advertising-based model in which sellers are granted prominence in return for a payment – different from traditional ad funding, the seller’s payments to the intermediary depends on the level of sales. Whether such differential pricing leads to worse recommendations from a welfare perspective is a priori unclear.

⁹Hunold et al. (2020) empirically analyze recommendations by hotel booking platforms Booking and Expedia and find that hotels with a lower price outside the platform (on a rival platform or on a direct sales channel) receive a less prominent recommendation. This is compatible with a recommendation algorithm of a profit-maximizing intermediary that punishes hotels with a lower conversion rate, where, after conditioning on hotel characteristics, this lower conversion rate is due to better offers on alternative sales channels. Relatedly, biased recommendations may arise on the zero-revenue part of the search engine: A search engine may bias its organic search results because of profit incentives regarding the sponsored search results – see Xu, Chen, and Winston (2012), Taylor (2013), and White (2013).

Wright, 2021) – that is, a fraction of those consumers who started in the indirect channel and received a recommendation from the intermediary can opt for the direct channel – and show that the phenomenon of inflated recommendations becomes more pronounced.

Another rationale for biased recommendations is to affect competition between sellers of different versions of the product. A particular instance is to include unattractive versions rather early in the search process to relax competition between sellers (Eliaz and Spiegler, 2011; Chen and He, 2011; de Cornière and Taylor, 2014).¹⁰ This may be in the interest of the intermediary if it receives a fraction of industry profit.¹¹ This work is orthogonal to the economic mechanism that we develop in this paper.

An important ingredient for our model is that sellers have two sales channels available, a direct and an indirect sales channel. We assume that it is more efficient to use the indirect sales channel; thus, bypassing the intermediary is inefficient. However, rent extraction by the intermediary is constrained by the availability of a direct sales channel. The possibility of bypass also figures prominently in the literature on price parity clauses (e.g., Edelman and Wright, 2015; Mariotto and Verdier, 2020; Wang and Wright, 2020) — that is, contractual obligations according to which sellers are not allowed to offer lower prices on alternative sales channels. While in our model with commitment, price parity clauses do not affect the outcome under *laissez-faire*, in the model in which the intermediary cannot commit to its recommendation policy but best-responds with its recommendation strategy to the seller’s retail prices, best-price clauses do matter. Here, price parity clauses serve as a substitute for the lack of commitment with respect to the platform’s recommendation policy and are welfare improving.

Contributing to the literature on e-commerce intermediaries initiated by Baye and Morgan (2001), Ronayne and Taylor (forthcoming) consider seller competition in a model in which consumers search for the best price. They assume that consumers are heterogeneous with respect to their choice of sales channel; in particular, some consumers exclusively use the indirect channel, others the direct channel and yet others decide which channel to use depending on the prevailing prices. They treat this difference as exogenous, whereas in our model differences across consumer types in valuations derived on the sales channels is determined by the intermediary’s recommendation policy: More-informative recommendations make the indirect channel particularly attractive for picky consumers (under our assumptions that, for prices above marginal costs, they never have an incentive to use the direct channel).

More broadly, our paper contributes to the economic analysis of recommender systems.¹² While non-

¹⁰Yet another reason for inflated recommendations is provided by Karle and Peitz (2017) in a context with loss-aversion consumers, where including product in the recommendation set of consumer that they will never purchase affects the price elasticity of demand of the product that they end up buying.

¹¹More generally, an intermediary may have various instruments available to affect competition among sellers on its platform; see Belleflamme and Peitz (2019) and Teh (forthcoming).

¹²For an overview, see Belleflamme and Peitz (2021, chapters 2 and 6). Theoretical work on recommender systems has looked at how the intermediary generates information from observed user behavior, which in turn depends on past recommendations (e.g., Che and Hörner, 2018, in the context of a welfare-maximizing intermediary). In our paper, we simply postulate that

personalized recommendations would simply be based on the popularity of a product, personalized recommendations rely on personal characteristics of the user — these may be consumer demographics, consumer purchase and click behavior, as well as the user’s position in the social network. In our paper, we provide a formal model of personalized recommendations. Tucker and Zhang (2011) analyze a consumer choice model with quality uncertainty and uncertainty about the horizontal match value; they show in a setting with sequential consumer choice that popularity information may be useful when consumers receive noisy signals about product quality even if consumers have heterogeneous tastes regarding horizontal product characteristics. Personalized recommendations may lead to a longer tail in the sales ratio of different products meaning that such recommendations increase the sales ratio of niche products (Hervas-Drane, 2015). Our model speaks to the long-tail phenomenon since the new product can be considered a niche product (in the sense that some consumers have different match values and, adding to the model, only a fraction of consumers is potentially interested in the new product at all), while the base product can be considered a mass product as it delivers the same value to all consumers: Then, in the inflated recommendation outcome, the niche product is bought more often than would be socially optimal.

Our paper differs from recent contributions on information design in which an intermediary provides information about consumer characteristics to sellers that can use this information for price-discrimination purposes (Bergemann, Brooks, and Morris, 2015; Ali, Lewis, and Vasserman, 2019). In our model, 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). Recommendations provide partial information to consumers about the match value. With inflated recommendations, there is partial pooling in the sense that a good match is always recommended, while a bad match is recommended with some positive probability less than 1.

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. This rotates the demand curve of picky consumers (Johnson and Myatt, 2006). Recommending the new product to a picky consumer if this constitutes a bad match can be considered to be “false advertising” and relates to Rhodes and Wilson (2018) who consider advertising in the presence of quality uncertainty – see also Drugov and Troya-Martinez (2019).¹³

the intermediary already has the information on consumer types that enables it to recommend only good matches to picky consumers. Thus, our paper is complementary to the work that explores the tradeoff between exploitation and exploration.

¹³More specifically, Rhodes and Wilson (2018) consider a monopolist that privately learns its type – that is, whether its quality is high or low – 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 positive probability less than 1.

Our paper 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 intermediary’s recommendation of a product to certain subset of consumers can be seen as targeted advertising. Our result of inflated recommendations constitutes noisy targeting. As we show in our model with channel-specific retail prices but absent personalized pricing, such noisy targeting options are provided when there are not too many picky consumers. Otherwise, with a sufficiently large fraction of picky consumers, perfect targeting will be the outcome.

A different literature considers certification intermediaries who certify quality (Biglaiser, 1993; Lizzeri, 1999). In particular, an intermediary may certify a minimum quality. While we consider horizontal match value instead of quality as a product characteristics, we show that the intermediary may “certify” match value rather imperfectly, as it mixes good and bad matches. Inflated recommendations feature the recommendation of a product that reduces gains from trade, in contrast to what the certifying intermediary in Lizzeri (1999) would do.¹⁴ In our setting, the inclusion of bad matches arises because of the presence of flexible and picky consumers.

At a conceptual level, our paper also connects to work on pure bundling and refund contracts. 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 recommendation a consumer is offered a bundle of the product in the good state and, with 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, Stinchcombe, and Whinston, 2005; Haghpanah and Hartline, 2020). Our result contains a different though related message: With bundling under inflated recommendations – that is, including the product in the good state with probability 1 and including it in the bad state with some positive probability – the distribution of valuations can be made less dispersed allowing the seller to better extract the gains from trade (at the social cost of reducing the gains from trade); this applies to our setting with significant marginal costs in which a strategy to sell large bundles would not be profitable.

Recommendation policies in our setting also relate to refund contracts. In our paper, from a consumer’s perspective, the contract offer looks as follows: Flexible consumers do not face uncertainty and always receive a contract offer for each sales channel, while picky consumers with a recommendation know that – in case recommendations are inflated – with some probability less than one, they buy the product in the bad state. Such a situation resembles a contract with partial refund (e.g., Courty and Li, 2000). Instead of offering a partial refund in the bad state, the contract here has the feature that picky consumers can return the product

¹⁴Lizzeri (1999) has shown that when minimal quality generates positive gains from trade, a certifying intermediary may extract without providing any information at all. When minimal quality generates negative gains from trade, the certifying intermediary at least excludes such quality – it only certifies product with valuation above cost.

and receive a full refund with a probability less than one in the bad state. The contract that corresponds to our recommendation policy β is to allow consumers to return the product with some probability $(1 - \beta)$. Picky consumers who make the experience that the product is a bad match exercise this option. If the same refund contract has to be offered to all consumers the refund contract that corresponds to the laissez-faire outcome maximizes the expected industry profit in our setting with picky and flexible consumers. In the real world, implementing refund contracts leads to transaction costs, as the seller has to deal with product returns; such transaction costs can be avoided by using a recommendation policy instead.

The paper proceeds as follows. In Section 2, we spell out the model. In Section 3, we characterize the laissez-faire outcome when the intermediary can commit to a recommendation policy that conditions the recommendation level on the retail prices set by the seller. Here, we also comment on a number of model extensions and how this affects our result of inflated recommendations. In Section 4, we consider several classes of regulatory policies, derive the optimal policy within each class, and characterize their welfare properties. We compare the welfare properties of these different types of regulatory policies with each other, with the first best and with the laissez-faire. In Section 5, we consider the alternative version of an intermediary who cannot commit to recommendation levels upfront, but in which it responds to the retail prices set by the seller; we characterize the laissez-faire outcome and compare it to outcomes derived in Sections 3 and 4; we also show that price parity clauses resolve the intermediary’s commitment problem and are total surplus increasing. Section 6 concludes. Several appendices complement the analysis in the main text. Appendix A contains the relegated proofs. Appendix B characterizes the laissez-faire outcome under alternative parameter constellations. Appendix C considers regulation of the uniform recommendations level and compares the outcome with the optimal uniform recommendation level to the one with the optimal recommendation cap. Appendix D considers regulation of an intermediary that faces fixed costs and has to break even.

2 The model

We consider a seller offering a newly introduced product. Two sales channels are available: an indirect channel I controlled by an intermediary and a direct channel D . The newcomer competes against a base product in competitive supply that provides some base utility v_0 to all consumers irrespective of the sales channel and costs c_0 per unit to be put on the market.

The new product is an improved product compared to the base product that also has design features that turn out to be liked by some consumers, but not by others. In particular, there are two possible product designs X and Y and a fraction α of consumers cares about the design – these are the design-sensitive or “picky” consumers – whereas the remaining fraction is indifferent and simply appreciates the overall appearance of the product relative to the base product – these are the design-insensitive or “flexible”

consumers. Half of the picky consumers prefer design X and the other half prefer Y . All consumers have unit demand. The differential willingness to pay – that is, the willingness to pay on top of the one for the base product – of flexible consumers is v_m ; the differential willingness to pay of picky consumers is v_h with $v_h > v_m$ for the – from their point of view – superior design and v_l with $v_l < v_m$ for the – from their point of view – inferior design. We assume that $(v_l + v_h)/2 < v_m$. This is a shortcut for a slightly different model in which the expected valuation of picky consumers is the same as those of flexible consumer, but in which consumers are risk-averse. Consumers know whether they are picky or flexible; however, before purchase they do not know whether they like the design in case they are picky.

We assume that the intermediary affects the purchasing decision of consumers in two ways. First, all consumers obtain convenience benefits $b > 0$ from buying the new product in the indirect channel – this is motivated by difficulties of the newcomer to provide the same level of service compared to the intermediary that would allow consumers to fully enjoy the benefit from the general improvement of the new product relative to the base product. Second, the intermediary makes a purchase recommendation to each consumer. The intermediary decides whether to recommend the new product to a consumer. It is always in the intermediary’s interest to recommend the product to a picky consumer with a good match. However, it may also recommend the product to some fraction $\beta \in [0, 1]$ of picky consumers with a bad match. We note that it does not matter whether the intermediary makes recommendations to flexible consumers as long as they are aware of the option to buy through the intermediary.

The intermediary’s revenue model is to charge sellers on its platform. It sets the rate λ as the fraction of the seller’s profits it extracts (profit sharing).¹⁵ Since the base product is in competitive supply, it does not generate any revenue to the intermediary irrespective of λ .

The per-unit production cost increment relative to the base product is c . The key assumption is that $v_l + b < c < v_m$.¹⁶ The first inequality says that selling the inferior design (even through the platform) reduces total surplus compared to selling the base product. Recommending this design is therefore total surplus reducing and the first-best welfare maximum does not feature any such recommendation. The second inequality says that even in the direct channel there are gains from trade with flexible consumers. If $c > v_m + b$ the problem is not interesting, as, in equilibrium, only picky consumers with good matches will buy. We also exclude constellations with $v_m < c < v_m + b$ because otherwise the direct channel would not constrain the intermediary and the full seller’s surplus would be extracted by the intermediary.

Since the base product is in competitive supply and generates the same benefit irrespective of the sales channel, it will always be sold at price c_0 . The intermediary can recommend the base product or the new improved product. By not buying the new product, consumers choose the base product as the outside option which gives them $v_0 - c_0$. Thus, when everybody buys the base product, consumer surplus is $v_0 - c_0$ and

¹⁵We allow for alternative price instruments of the intermediary in Section 3.3.

¹⁶At the end of Section 3.2, we discuss what happens when $c < v_l + b$.

so is total surplus. For convenience, in what follows we renormalize consumer and total surplus and report them only in excess of this level $v_0 - c_0$. We also report prices for the new product as price increments on the base product, c_0 .

With these adjustments in place, our model becomes one with a monopoly seller competing against an outside option with value 0. Inefficiencies arise when some consumers buy in the direct instead of the indirect channel – this is a situation of *inefficient bypass* – or if some picky consumers with a bad match buy. The latter can only occur in the indirect channel and requires *inflated recommendations* ($\beta > 0$). Then the new product is recommended too often, whereas the base product is recommended too little. Since we assume that $c > v_l + b$ recommendations are excessive from a total surplus perspective as well as a consumer surplus perspective (as the price will not be below marginal cost).

Since 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, $(1 - \alpha)(v_m + b - c) + (\alpha/2)(v_h + b - c)$ is linear in α . It is increasing in α if and only if $(v_h + b - c)/2 > (v_m + b - c)$.

The *timing of the game* is as follows. First, the intermediary sets λ and publicly commits to a recommendation policy $\beta(\cdot, \cdot)$ that conditions on the seller's prices (p^I, p^D) . Second, the seller sets its prices on the direct channel p^D and at the intermediary p^I . Third, after observing prices and $\beta(p^I, p^D)$ consumers decide which sales channel to choose. Fourth, picky consumers in the indirect channel receive personalized recommendations and all consumers make their purchasing decisions.

We characterize subgame perfect Nash equilibria of this game. We note that consumers engage in Bayesian updating. However, this belief updating is pinned down by the recommendation policy the intermediary 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.

The timing is motivated by the following considerations. We are abstracting from price opacity and, thus, assume that consumers observe prices before deciding which channel to use. Furthermore, we assume that the intermediary has full commitment over its recommendation policy. This means that the seller understands how a change of prices affects the recommendation policy by the intermediary. To be able to do so, the seller must observe the function $\beta(\cdot, \cdot)$. One way to interpret this is that the seller can test the intermediary's recommendation algorithm and, thus, understands how the recommendation policy responds to price changes. Consumers only need to observe the recommendation policy at prevailing prices, $\beta(p^I, p^D)$, and not the function at different prices. One way to motivate this is to consider consumers arriving in two batches. 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 $\beta(p^I, p^D)$

before deciding what to do. In the limit as ϵ turns to zero, this implies that consumers of measure 1 observe $\beta(p^I, p^D)$.¹⁷

3 The intermediary's recommendation and pricing policy

In this section, we derive the intermediary's profit-maximizing recommendation policy and characterize the equilibrium.

3.1 Preliminaries: Polar cases

Suppose that all consumers are flexible ($\alpha = 0$) and thus the recommendation policy does not affect the consumers' choice of sales channel. The seller can make $v_m - c$ selling directly or $(1 - \lambda)(v_m + b - c)$ by selling through the intermediary. From the intermediary's point of view the optimal profit sharing contract makes the seller indifferent, thus

$$\lambda^* = \frac{b}{v_m + b - c}.$$

The seller and the intermediary earn $v_m - c$ and b respectively. Total surplus is maximal and equal to $v_m + b - c$.

Suppose next that all consumers are picky ($\alpha = 1$). Consider the intermediary committing to a recommendation policy $\beta = \beta(p^I, p^D)$. The seller cannot make positive profits by selling to the picky consumers directly. Thus, the intermediary will find it optimal to fully expropriate the profits of the seller in the indirect channel by setting $\lambda = 1$. Therefore, along the equilibrium path, the intermediary does not have an incentive to inflate recommendations and sets $\beta = 0$ and the seller sets $p^I = v_h + b$ and $p^D = c$.

An equilibrium with fully specified strategies is that the intermediary sets $\lambda = 1$ and $\beta(p^I, p^D) = 0$ for all (p^I, p^D) ; the seller sets $p^I = v_h + b$, $p^D = c$; and picky consumers go to the intermediary and buy if the product is recommended to them.¹⁸

The seller and consumers obtain a net surplus of 0 and the intermediary earns $(v_h - b - c)/2$. Thus, the intermediary extracts the entire surplus. The outcome is efficient and does not involve any inflated

¹⁷Another construction is consider an infinite repetition of the multi-stage game from above with discounting in which the recommendation policy is not observable but in which $\beta(p^I, p^D)$ becomes public at the end of the stage (e.g., because of consumer feedback). If the future figures sufficiently prominently (discount factor sufficiently close to 1), then the equilibrium allocation of the one-shot game with public commitment can be supported as an equilibrium allocation of the infinitely repeated game in which consumers "punish" the intermediary for any deviation from the equilibrium policy of the one-shot game by believing that from that point on $\beta = 1$, which implies that flexible consumers buy in the indirect channel at $p^I = v_m + b$ and picky consumers do not buy at all, and $\lambda = b/(v_m + b - c)$.

¹⁸As a sanity check, the reader may want go back to the model without normalization. Then, the outside option is sold at c_0 and gives a net benefit to consumers of $v_0 - c_0$. Prices for the new product are $p^D = c + c_0$ and $p^I = v_h + b + c_0$. Hence, a picky consumer who receives a recommendation and buys the new product receives net benefit $(v_h + b + v_0) - (v_h + b + c_0) = v_0 - c_0$ and indeed does not have an incentive to revise her decision.

recommendations. This efficiency result does not longer hold with a mix of picky and flexible consumers.

3.2 A mix of picky and flexible consumers

In the previous section we have seen that either recommendations do not matter or they are as precise as possible (i.e., $\beta = 0$). In the latter case, picky consumers receive a recommendation if and only if the match is good. The key ingredient of our model is that there are both flexible and picky consumers – that is, $\alpha \in (0, 1)$. In our setting, flexible consumers buying from the direct channel constitutes inefficient bypass. This bypass possibility ensures that, in any equilibrium, the seller must make a profit at least equal to $(1 - \alpha)(v_m - c)$ – that is, the profit from selling to all flexible consumers in the direct channel at price $p^I = v_m$.

We proceed throughout the remaining analysis in the main text under the assumption that there are no gains from trade with picky consumers even in the indirect channel if these consumers do not receive any informative recommendations:

Assumption 1. *The unconditional expected total surplus when buying from the intermediary is negative for picky consumers, $(v_h + v_l)/2 + b - c < 0$.*

In other words, the intermediary must choose an informative recommendation policy – that is $\beta < 1$. Since $b \geq 0$, it also implies that in equilibrium flexible consumers will not buy in the direct channel. Picky consumers, by contrast, may want to buy directly from the new seller since $v_m > c$.

With Assumption 1 in place, we now turn to the possible consumer choices.¹⁹

Given (p^D, p^I) , what are the incentives of consumers to buy at the intermediary instead of buying in the direct channel? Flexible consumers prefer to buy in the indirect channel if and only if

$$v_m + b - p^I \geq v_m - p^D \text{ and} \tag{1}$$

$$v_m + b - p^I \geq 0. \tag{2}$$

Picky consumers decide based on seller's prices and the intermediary's recommendation policy $\beta(p^D, p^I)$. They buy in the indirect channel (if at all) if and only if

$$\frac{1}{2}(1 + \beta) \max \left\{ \frac{v_h + \beta v_l}{1 + \beta} + b - p^I, 0 \right\} + \frac{1}{2}(1 - \beta) \max \{ v_l + b - p^I, 0 \} \geq \max \left\{ \frac{v_h + v_l}{2} - p^D, 0 \right\}. \tag{3}$$

The vertically integrated solution Before characterizing the solution in our setting in which the intermediary and the seller are different economic agents, we look at the vertically integrated solution which is the outcome of the problem in which stages 1 and 2 of the 4-stage game are collapsed and a vertically

¹⁹In Appendix B, we analyze the case in which Assumption 1 does not hold – that is $(v_h + v_l)/2 + b - c > 0$. As we discuss at the end of this section (and show in that appendix), our main insights are robust in the sense that we identify parameter constellations that give rise to an inflated recommendation equilibrium.

integrated firm acts as recommender and price setter. The next lemma partially characterizes the two types of outcomes that can arise.

Lemma 1. *Under vertical integration there are two types of possible outcomes. In the first, all flexible, all picky consumers with a good match, and a fraction $\beta \geq 0$ of picky consumers with a bad match buy in the indirect channel – we call this the inflated recommendations outcome. In the second, all picky consumers with a good match buy in the indirect channel and flexible consumers buy in the direct channel – we call this the inefficient bypass outcome.*

Proof. Suppose that only flexible consumers buy in the indirect channel and picky consumers either buy directly or opt out. To sustain the former as an equilibrium outcome, the firm must set (p^I, p^D) such that $p^I \leq p^D + b$. It is straightforward to see that the firm has no incentive to set $p^D < (v_h + v_l)/2$, as this price is strictly lower than the marginal cost and picky consumers buy in the direct channel resulting in losses.

Instead, suppose that the picky consumers opt out, which implies that $p^D > (v_h + v_l)/2$ and flexible consumers buy in the indirect channel. With such consumer choices, the firm's profit is maximized at $p^I = v^m + b$. In this case the firm makes profit $(1 - \alpha)(v^m + b - c)$, which is less than what it obtains by an informative policy $\beta > 0$ sufficiently small such that also picky consumers with a recommendation buy.

It is straightforward to exclude all other outcomes but the two outcomes stated in the lemma. \square

In the profit-maximizing inflated recommendation outcome, the firm sets β such that $v_m = (v_h + \beta v_l)/(1 + \beta)$ or, equivalently $\beta = (v_h - v_m)/(v_m - v_l)$ and $p^I = v_m + b$ (and $p^D \geq v_m$). Its profit is

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

In the profit-maximizing inefficient bypass outcome, the firm sets $\beta = 0$, serves picky consumers with a good match at $p^I = v_h + b$ in the indirect channel and serves flexible consumers in the direct channel at $p^D = v_m$. Its profit is

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

Comparing the two cases above we find that the maximal profit is given by

$$(1 - \alpha)(v_m - c) + \max \left\{ \frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) + (1 - \alpha)b, \frac{\alpha}{2} (v_h + b - c) \right\}.$$

In the proof of Proposition 1, there is a uniquely defined $\bar{\alpha} \in (0, 1)$ such that the two cases give the same profit, which is given by $\frac{\bar{\alpha}}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) + (1 - \bar{\alpha})b = \frac{\bar{\alpha}}{2} (v_h + b - c)$. For $\alpha < \bar{\alpha}$, the firm maximizes its profit with inflated recommendations, while, for $\alpha > \bar{\alpha}$, it does so by inducing inefficient bypass.

Proposition 1. *The vertically integrated solution is characterized as follows:*

- for $\alpha < \bar{\alpha}$, where $\bar{\alpha} \in (0, 1)$ the firm sets $\beta^* = \frac{v_h - v_m}{v_m - v_l}$ and $(p^I, p^D) = (v_m + b, v_m)$. All consumers go to the indirect channel. All flexible consumers and all picky consumers with a recommendation buy. Welfare losses compared to the first best are given by $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$.

- for $\alpha \geq \bar{\alpha}$, the firm sets $\beta = 0$ and $(p^I, p^D) = (v_h + b, v_m)$. 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. Welfare losses compared to the first best are given by $(1 - \alpha)b$.

Proof. Comparing the maximal profit in the two possible solutions, the critical $\bar{\alpha}$ is implicitly determined by

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

which can be rewritten as

$$\frac{\bar{\alpha}}{2} (v_h + b - c - (1 + \beta)(v_m + b - c)) = (1 - \bar{\alpha})b$$

and, after further manipulation gives an explicit expression for $\bar{\alpha}$:

$$\bar{\alpha} = \frac{b}{b + \frac{1}{2}(v_h + b - c - (1 + \beta)(v_m + b - c))}.$$

To show that $\bar{\alpha} \in (0, 1)$, it is sufficient show that the expression $v_h + b - c - (1 + \beta)(v_m + b - c)$ is positive, which is equivalent to

$$\frac{v_h + b - c}{v_m + b - c} > \frac{v_h - v_l}{v_m - v_l}. \quad (4)$$

As we assumed that $v_l + b - c < 0$, it must be that $v_l < c - b$. Since $\frac{v_h - x}{v_m - x}$ is increasing in x , we must have that indeed inequality (4) is satisfied. \square

The inefficient bypass outcome has the same flavor as a simple screening contract between picky consumers and flexible consumers. Picky consumers with a good match have a valuation higher than flexible consumers and, thus constitute the high type. Since 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 then buy low quality – that is, they buy in the direct channel and do not obtain benefit b . The inflated recommendation outcome can be seen as the outcome of a pooling contract. The key difference compared to standard price discrimination problems is the possibility of partially informative recommendations, which equalizes the willingness to pay of picky consumers (with a recommendation) and flexible consumers. Some picky consumers with a bad match are made believe that the product is likely to be a good match. A picky consumer not knowing the match quality thus has a lower expected valuation after receiving such a noisy recommendation. In effect, demand of the picky type is rotated and, in the profit-maximizing solution, consumer net surplus is zero also in the inflated recommendation outcome.

As already mentioned in the literature review, the inflated recommendations outcome can be implemented with a partial refund contract: The product is offered at price $v_m + b$ and consumers have the option to return the product with probability $1 - \beta$. Picky consumers with a bad match will make use of this option.²⁰

²⁰For the equivalence to hold true, there must be zero transaction costs involved with the return and the firm must be able to resell returned items at no cost. Whenever return policies are costly to run, the recommendation policy generates higher profits.

The intermediary’s profit-maximizing strategy We now turn to the equilibrium analysis of the game in which first the intermediary chooses its strategy and then the seller sets retail prices. We will show that the two possibly profit-maximizing vertically integrated solutions can be decentralized – that is, the intermediary can choose 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. 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.

First, we consider inflated recommendations as an equilibrium outcome in which all consumers visit the intermediary. Clearly, the intermediary will recommend the product to all picky consumers with a good match. It also recommends the product to a fraction $\beta > 0$ of picky consumers with a bad match.

No matter what the intermediary does, the seller can secure some minimal profit for itself simply by selling only in the direct channel: It can set $p^D = v_m$ (and $p^I > v_m + b$) and make the profit $(1 - \alpha)(v_m - c)$. Thus, the intermediary has to provide at least such a profit level to the seller. In the vertically integrated solution with inflated recommendation, $p^I = v_m + b$, $p^D \geq v_m$ and all trade takes place in the indirect channel. Let us set the intermediary’s recommendation policy as $\beta(p^I, p^D) = \beta^*$ for $p^I = v_m + b$ and $p^D \geq v_m$ and $\beta(p^I, p^D) = 1$ for all other prices. By setting $\beta = 1$ the intermediary ensures that the indirect channel will not be used and, thus, the seller will have to comfort itself by setting only to flexible consumers. It constitutes the maximal “punishment” the intermediary can inflict on the seller. Clearly, the intermediary does not do well itself and obtains only zero profit, but it may want to commit to such a policy, as long as it does not become part of the equilibrium play.²¹ Other less severe punishments may also be used, but it is sufficient to restrict the analysis to the particular recommendation policy to show that the vertically integrated solution can be implemented.

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 such that the seller just obtains profit $(1 - \alpha)(v_m - c)$ by λ^* :

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

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

This shows that the vertically integrated solution with inflated recommendations can be implemented by the intermediary. The surplus distribution is that the seller obtains $(1 - \alpha)(v_m - c)$, the intermediary

²¹In Section 5 we consider the model in which the intermediary cannot commit to the recommendation policy ex ante and adjusts its recommendation policy in response to the seller’s prices.

obtains

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

and consumers obtain a net surplus of zero.

Second, we consider the outcome with inefficient bypass. Suppose that the intermediary takes the (almost) entire profit in the indirect channel – that is, $\lambda^* = 1$. In this case the seller can only make profit in the direct channel. It will sell to flexible consumers at $p^D = v_m$ and it will set the price $p^I = v_h + b$ (for any infinitesimally small profit fraction it maintains). This implements the vertically integrated solution with profits $(1 - \alpha)(v_m - c)$ going to the seller and $\frac{\alpha}{2}(v_h + b - c)$ going to the intermediary. Consumers obtain a net surplus of zero.

The intermediary can choose between inducing inflated recommendations or inefficient bypass. In either case, the seller makes the profit $(1 - \alpha)(v_m - c)$. Thus, the comparison of the intermediary's profits yields the same critical α as the comparison of the vertically integrated firm's profit. We summarize our findings in the following proposition:

Proposition 2. *The equilibrium when the intermediary commits to its recommendation policy is characterized as follows:*

- for $\alpha \geq \bar{\alpha}$, the intermediary sets $\lambda^* = 1$, $\beta = 0$ for all (p^I, p^D) . Equilibrium prices are given by $(p^I, p^D) = (v_h + b, v_m)$. Picky consumers go to the indirect channel and buy if they receive the recommendation to buy, whereas all the flexible consumers buy in the direct channel. Welfare losses are given by $(1 - \alpha)b$.
- for $\alpha < \bar{\alpha}$, the intermediary sets

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

$\beta(p^I = v_m + b, p^D \geq v_m) = \frac{v_h - v_m}{v_m - v_l}$ and $\beta = 1$ otherwise. Equilibrium prices are given by $(p^I, p^D) = (v_m + b, v_m)$. All consumers go to the indirect channel. The fraction of $\frac{1 + \beta}{2} = \frac{1}{2} \frac{v_h - v_l}{v_m - v_l}$ of picky consumers buy. Welfare losses are given by $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$.

With a large fraction of picky consumers in the population, the intermediary prefers to induce the outcome with inefficient bypass. By making it very expensive to sell through the intermediary, the intermediary induces the seller to sell to flexible consumers directly. The welfare loss due to inefficient bypass is the foregone benefit of flexible consumers that equals to $(1 - \alpha)b$.

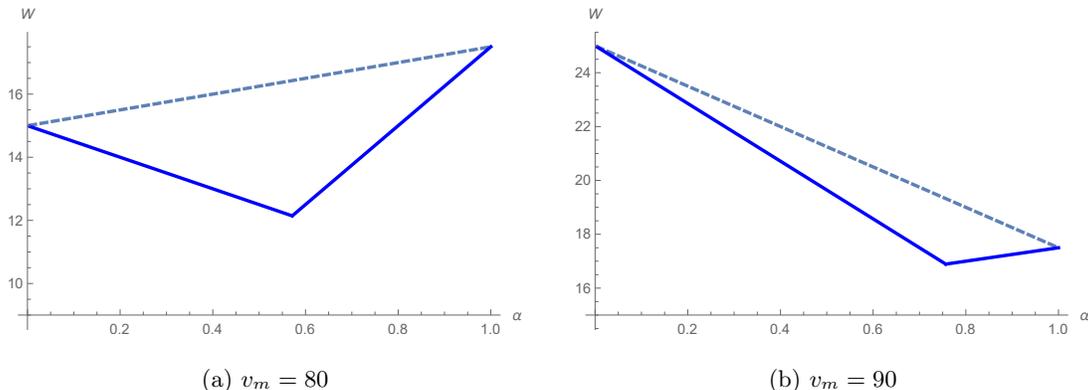


Figure 1: Total welfare in α ; $v_h = 100, v_l = 20, b = 10, c = 75$. First best outcome (dashed), private solution (solid).

By contrast, with a small fraction of picky consumers in the population, it is better to inflate 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 commits to stop providing informative recommendations to the picky consumers which renders the deviation unprofitable. The welfare loss from inflated recommendations is equal to $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$. This loss becomes more pronounced as α , c , and v_h increase and as b , v_m , and v_l decrease.

To summarize, in the equilibrium of Proposition 2, the intermediary decentralizes the vertically integrated solution.

Discussion of outcomes under alternative parameter constellations Before considering various extensions of the model, we discuss what happens in our setting with different parameter constellations – a detailed analysis is relegated to Appendix B. We continue to assume that $(v_h + v_l)/2 + b < v_m$, which is equivalent to $b < v_m - (v_h + v_l)/2$, as we want to include cases in which the convenience benefit of the indirect channel is arbitrarily small. This implies that $b < (v_h + v_l)/2 - v_l$, which is equivalent to $v_l + b < (v_h + v_l)/2$. So far, we assumed that $(v_h + v_l)/2 + b < c < v_m$. If marginal costs were higher ($c > v_m$), the intermediary would only recommend good matches to picky consumers (that is, $\beta = 0$) and sell at $v_h + b$. Thus, the equilibrium outcome would implement the first best.

What happens if marginal costs are lower than $(v_h + v_l)/2 + b$? Consider the case that there are never gains from trade for a picky consumer with a bad match – that is, $c > v_l + b$. We distinguish between two alternative parameter constellations: $(v_h + v_l)/2 < c < (v_h + v_l)/2 + b$ and $v_l + b < c < (v_h + v_l)/2$. If $(v_h + v_l)/2 < c < (v_h + v_l)/2 + b$, our results remain unchanged (some analysis needs to be added though). If $v_l + b < c < (v_h + v_l)/2$ the tradeoff between inflated recommendations and inefficient bypass continues to

apply. The novel feature is that for high α the outside option for the seller is to sell to all consumers in the direct channel and not only to flexible consumers. This implies that the intermediary has to leave a higher profit share to the seller and, in particular, the seller keeps some of the profit in the indirect channel also with inefficient bypass.

Qualitatively different results hold when there are possible gains from trade selling a bad match to picky consumers – that is, $v_l + b > c$. The first best is then to sell to all consumers in the indirect channel. So far we spoke of inflated recommendations when total-surplus-decreasing recommendations are made. This is clearly not the case for the low marginal costs considered here, as the first best outcome is that everybody buys in the indirect channel. We observed in the analysis above that the inflated recommendations outcome ($\beta > 0$) always features recommendations that reduce the surplus of picky consumers with a bad match at prevailing prices. As we show in Appendix B.3, for the share of picky consumers, α , sufficiently small, the intermediary will choose a recommendation policy with $\beta = \beta^* \in (0, 1)$ – that is, recommendations are inflated from a consumer surplus perspective because picky consumers would choose not to buy in the indirect channel if they were fully informed about match quality. For larger α , the intermediary recommends the product to everybody (which, from a consumer surplus perspective, also constitutes inflated recommendations at given prices) and, thus, implements the first best. Here, flexible consumers obtain a positive net surplus, as $p^I = (v_h + v_l)/2 + b$. In either case, there will be no sales in the direct channel. Thus, there is no counterpart to the inefficient bypass outcome that can be obtained for higher marginal costs. From a total surplus perspective, there are too few recommendations for low α .

3.3 Extensions

In this section, we explore several variations of our setting. The upfront takeaway of all these extensions is that unless the seller can set personalized prices or has a sufficiently rich portfolio of product versions, the equilibrium outcome features either inefficient bypass or inflated recommendations.

The intermediary’s price instrument So far we 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 to be paid by the seller to be carried by the intermediary, a per-unit transaction fee t or a percent transaction fee τ . As we will see, the equilibrium outcome is invariant to the particular type of pricing instrument available to the intermediary.²²

When the equilibrium features inflated recommendations, we have shown in Proposition 2 that the intermediary sets λ to guarantee the profit $(1 - \alpha)(v_m - c)$ to the seller and $\beta^* = \frac{v_h - v_m}{v_m - v_l}$. This gives the intermediary’s equilibrium profit $\frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b$. If, instead, only a listing fee is available,

²²We return to the role of price instruments in modified settings below, where the equilibrium outcome will depend on the particular type of instrument used by the intermediary.

this fee is set equal to this profit and, together with $\beta(p^I = v_m + b, p^D \geq v_m) = \beta^*$ implements the same outcome. If the intermediary uses a per-unit fee, it sets this fee

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

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

The reason that the intermediary achieves the same equilibrium profit under different classes of price instruments is that a seller deviating from equilibrium prices 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 $p^D = v_m$.

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, $\frac{\alpha}{2}(v_h + b - c)$ and the seller makes profit $(1 - \alpha)(v_m - c)$ 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, this fee is set to $T = \frac{\alpha}{2}(v_h + b - c)$; if only a per-unit transaction fee is available, this fee is set to $t = v_h + b - c$; and if only a percent transaction fee is available, that fee is set to $\tau = \frac{v_h + b - c}{v_h + b}$.

Price-parity clause What happens if the intermediary can impose the price-parity clause on the seller? We will see that such a clause does not have any bite in our setting with commitment. Under the price parity clause (PPC) the intermediary requires the seller to set $p^I \leq p^D$. We assume that the seller first decides whether or not to sell through the intermediary and then sets prices. It is easy to see that the PPC does not affect the analysis of the extreme cases as the seller's deviations to selling only directly satisfy the PPC.

For intermediate $\alpha \in (0, 1)$, we continue working under the assumption 1 – that is, $(v_h + v_l)/2 + b - c < 0$. The proof of Lemma 1 continues to apply under the PPC and, thus, it is sufficient to consider the two cases in which the intermediary induces to sell to either all or only picky consumers with a good match. It is straightforward to see that commitment power is sufficient to make the deviation as unattractive for the seller as possible. In the first case (all consumers visit the intermediary), the deviation to sell directly to flexible consumers involves no recommendation to the picky consumers and no profit from sales through the intermediary. The PPC does not add to it as in the worst case the seller has to forego selling to the picky consumers and only sells directly to the flexible consumers. In the second case (in which only picky

consumers use the indirect channel), the intermediary can set $\lambda^* = 1$ and preclude any attempts of the seller to sell to flexible consumers through the intermediary. We have seen that the intermediary's commitment power inflicts the worst payoff for the seller if it decides to use the direct channel; therefore, price parity clauses do not confer any further advantages to the intermediary.

Results with price-parity clauses are markedly different if the intermediary cannot commit to its recommendation policy, as we will show in Section 5.3.

Unconditional recommendation strategy In our model, the intermediary commits to a recommendation policy that conditions recommendations on seller prices. Here, we consider the alternative environment in which the recommendation level is price-independent. We show that the critical α above which flexible consumers buy directly is lower than in the case in which the intermediary can condition β on the seller's prices.

Suppose that the intermediary chooses (β, λ) . To make the flexible consumers choose the indirect channel, the intermediary has to set λ such that it respects the seller's incentive constraint $(1 - \lambda)(v_m + b - c) \geq v_m - c$. This ties down λ : the intermediary sets $\lambda = 1 - \frac{v_m - c}{v_m + b - c}$. This expression is smaller than λ^* given in Proposition 2 that applies for $\alpha < \bar{\alpha}$. This means that the intermediary has to leave a larger fraction of profits to the seller. Therefore, the intermediary is more inclined to serve picky consumers only, in which case recommendations are fully informative. Therefore, the critical α below which the intermediary inflates recommendations is less than $\bar{\alpha}$.

Intermediary with a profitable outside option Our setting features a base product in competitive supply, which implies that the intermediary does not make any profit from recommending the base product. Here, we consider a situation in which selling the base product generates positive profits for the intermediary. In particular, suppose that the intermediary provides the base product itself – this is a different type of vertical integration compared to the vertically integrated solution above. It implies that the consumer's outside option gives a net surplus of zero to consumers but confers a profit of $v_0 - c_0$ per unit to the intermediary if the consumer chooses the outside option at price v_0 . Since the intermediary also sets the price for the base product, there is an additional pricing decision that needs to be included in the game. For simplicity, we postulate that the intermediary can respond to the prices set by the seller; this price is set before consumers decide which channel to go to. While the intermediary continues to commit to a recommendation policy, the pricing of the intermediary after the seller's pricing implies only partial commitment regarding the full set of the intermediary's instruments.

Compared to the analysis in Section 3.2, the seller's prices are increased by $v_0 - c_0$. We modify Assumption 1 to $\frac{v_h + v_l}{2} + b + v_0 - c - c_0 < 0$, which implies that picky consumers do not buy the new product when recommendations are completely uninformative ($\beta = 1$).

We will show that for small α the equilibrium features inflated recommendations. Suppose that the seller diverts flexible consumers to the direct channel at price p^D . The intermediary profitably undercuts by setting the price of the base product $p_0 = p^D - v_m$ if $p^D - v_m - c_0 \geq \alpha(v_0 - c_0)$, where on the right-hand side we have the profit of the intermediary to sell the base product to picky consumers only. This implies that the seller's best deviation is to set $p^D = v_m + c_0 + \alpha(v_0 - c_0)$, which gives a deviation profit of $(1 - \alpha)(v_m - c + \alpha(v_0 - c_0))$. Thus, the intermediary has to allow the seller to absorb part of the surplus generated from selling the base product.²³

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

We now turn to the outcome with inefficient bypass. Here, $\lambda = 1$ and $\beta(\cdot, \cdot) = 1$ 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 $p^I = v_0 + v_h + b$ and p^D such that the intermediary does not have an incentive to set the price for the base product such that flexible consumers buy from the intermediary. With $\lambda = 1$, the seller makes profit from selling to the flexible consumers, $(1 - \alpha)(p^D - (c + c_0))$. In the equilibrium candidate the intermediary makes profit $\frac{\alpha}{2}(v_0 - c_0) + \frac{\alpha}{2}(v_0 + v_h + b - (c + c_0))$. If the intermediary sets $p^D - v_m$, all consumers will buy the base product and the intermediary will make $p^D - v_m - c_0$. The maximal price flexible consumers are willing to pay in the direct channel is $v_0 + v_m$. Thus, 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 $p^D = v_0 + v_m$, if $\frac{\alpha}{2}(v_0 - c_0) + \frac{\alpha}{2}(v_0 + v_h + b - (c + c_0)) \geq v_0 - c_0$, which is equivalent to $\frac{\alpha}{2}(v_h + b - c) \geq (1 - \alpha)(v_0 - c_0)$. This defines the critical $\tilde{\alpha} = \frac{v_0 - c_0}{v_0 - c_0 + \frac{1}{2}(v_h + b - c)}$.

When does the intermediary prefer the inflated recommendation outcome? This is the case if

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

For $v_0 - c_0$ not too large, we can show that, at $\tilde{\alpha}$ given above, this inequality is satisfied. Thus, there is a particular value α above this value such that the intermediary is indifferent between the two outcomes. For larger values of α , we are in the inefficient bypass equilibrium and for lower α in the inflated recommendations equilibrium. Thus, our main result is robust to introducing an outside option that generates profits for

²³In the analysis in Section 3.2 consumers obtained the surplus $v_0 - c_0$ (we referred to the increment to this surplus when reporting consumers surplus) since the base product was in competitive supply. Here, the intermediary controls the base product. In effect, consumers do not obtain any surplus and the gains from trade $v_0 - c_0$ that would be generated from selling the base product is split between intermediary and seller.

the intermediary and means that the opposite of self-preferencing can be an equilibrium outcome if the intermediary owns the base product but not the new product.

We also note that the equilibrium outcome is different from the outcome under full vertical integration in which the intermediary owns base product and new product. As we have shown above, the intermediary has to give a larger profit to the seller under inefficient bypass than under inflated recommendation. This implies that the critical α that separates inflated recommendation from inefficient bypass is higher than the critical α under full vertical integration. In other words, the intermediary is less inclined to opt for inefficient bypass instead of inflated recommendation when the intermediary faces an independent seller of the new product compared to the case of full vertical integration.

Platform leakage In the base model we assumed that a consumer receives a recommendation only if using the indirect channel and that it does not switch to the direct channel after receiving the recommendation. Consider now the case of “platform leakage” – that is, a fraction of consumers in the indirect channel can use the direct channel after having received a recommendation from the intermediary. Here, the intermediary offers a show-rooming service and the seller may free-ride on the the intermediary’s service (Wang and Wright (2020)).

For simplicity, consider an exogenous fraction ν of picky consumers who can switch without cost. Consider 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 destabilizes the outcome that all trade occurs in the indirect channel by setting a price p^D below v_m , its profit will be bounded by $(1 - \alpha + \nu \frac{\alpha}{2}(1 + \beta^*))(v_m - c)$. Therefore, the intermediary inducing the seller to serve flexible consumers in the indirect channel has to respect the seller’s incentive constraint that is given by

$$(1 - \lambda) \left(\frac{\alpha}{2}(1 + \beta^*) + 1 - \alpha \right) (v_m + b - c) \geq \left(1 - \alpha + \nu \frac{\alpha}{2}(1 + \beta^*) \right) (v_m - c).$$

It is straightforward to see that the intermediary has to settle for a lower fraction of industry profit λ to satisfy the seller’s incentive constraint than in the absence of platform leakage. We note that platform leakage happens only off the equilibrium path in an inflated recommendation equilibrium.

In the case of inefficient bypass, the seller sells directly to flexible consumers at $p^D = v_m$ and aims to serve picky consumers with a good match at $p^I = v_h + b$. Given these prices, with platform leakage, the fraction ν of picky consumers with a good match buys directly and each consumers obtains a net surplus of $v_h - v_m$. The seller’s profit is $(1 - \alpha + \nu \frac{\alpha}{2})(v_m - c)$ and platform leakage occurs along the equilibrium path with inefficient bypass.

How does the intermediary choose between these two options? With inefficient bypass the intermediary

obtains $(1 - \nu)\frac{\alpha}{2}(v_h + b - c)$ while with inflated recommendations it obtains

$$\begin{aligned}\lambda\left(\frac{\alpha}{2}(1 + \beta^*) + 1 - \alpha\right)(v_m + b - c) &= \left(\frac{\alpha}{2}(1 + \beta^*) + 1 - \alpha\right)(v_m + b - c) - \left(1 - \alpha + \nu\frac{\alpha}{2}(1 + \beta^*)\right)(v_m - c) \\ &= \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b - \nu\frac{\alpha}{2}(1 + \beta^*)(v_m - c).\end{aligned}$$

For $\alpha = \bar{\alpha}$ defined in Proposition 1, the profit under inflated recommendation is equal to $\frac{\bar{\alpha}}{2}(v_h + b - c) - \nu\frac{\bar{\alpha}}{2}(1 + \beta^*)(v_m - c)$. This profit is strictly higher than the profit under inefficient bypass since $(1 + \beta^*)(v_m - c) < v_h + b - c$. 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 and approaches 1 when ν increases and turns to some critical value less than 1.

Intuitively, there are two opposing effects of how platform leakage changes the intermediary's tradeoff. First, with inflated recommendations, if more consumers can show-room, then the intermediary faces tougher competition from the direct channel and has to leave more surplus to the seller. Second, show-rooming undermines the incentives for fully informative recommendations and this makes the outcome with inefficient bypass less profitable. We showed that for all parameters the second effect dominates and the intermediary is more inclined to induce the inflated recommendations outcome.

Network effects between consumers in the indirect channel For applications to digital platforms it is important to include network effects into the analysis. More buyers in the indirect channel may allow the intermediary to improve its services, for example, by reducing delivery time or by providing add-on services that improve with a larger number of buyers (e.g., if users provide feedback on how to use the product, which is beneficial to all other users in the indirect channel). The benefit in the indirect channel b becomes a function that is increasing in the number of users. Denote $b_0 = b(0)$, $b_1 = b(\alpha/2)$, and $b_2 = b((1 - \alpha) + (1 + \beta^*)\alpha/2)$. There is no change in the number of participating users under inflated recommendations if $b'((1 - \alpha) + (1 + \beta^*)\alpha/2)$ is sufficiently small (assuming that b is continuously differentiable).

With inefficient bypass, using the indirect channel gives consumption benefit b_1 . With inflated recommendations, this benefit is b_2 . As follows from the analysis in Section 3.2, the intermediary's profit is $\frac{\alpha}{2}(1 + \beta^*)(v_m + b_2 - c) + (1 - \alpha)b_2 = \frac{\alpha}{2}(1 + \beta^*)(v_m + b_0 - c) + (1 - \alpha)b_0 + (\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha))(b_2 - b_0)$ with inflated recommendations, while it is $\frac{\alpha}{2}(v_h + b_1 - c) = \frac{\alpha}{2}(v_h + b_0 - c) + \frac{\alpha}{2}(b_1 - b_0)$ with inefficient bypass. With network effects, the intermediary's option to induce the inflated recommendations outcome rather than the one with inefficient bypass becomes more attractive for two reasons. First, network effects increase b and this applies to more consumers in the inflated recommendation outcome. Second, since there will be more users in the indirect channel under inflated recommendations than under inefficient bypass, the increase in benefit becomes stronger with inflated recommendations. As a result, there is a larger range of the parameter α such that the intermediary will go for inflated recommendations.

Multiple product categories In our model, the intermediary interacts with a single seller and chooses its price instrument optimally. In reality, the intermediary faces sellers in many different product categories, and these categories differ in market characteristics (including α); yet, the intermediary chooses the fee applicable to this product category from a small number of different fees. The intermediary then picks the best available $\lambda < 1$ from the available set that induces an inflated recommendations outcome (i.e., the largest available λ that induces the seller not to use the direct channel) or $\lambda = 1$, which induces the inefficient bypass outcome. With a limited set of fees, the inflated recommendations outcome continues to occur but becomes less prevalent.

Personalized pricing Our next two extensions clarify that two model ingredients are essential to obtain inflated recommendations. The first concerns pricing by the seller: if personalized prices were feasible, the intermediary would not have any incentive to inflate recommendations. Thus, the inability of the seller to engage in personalized prices is essential to obtain inflated recommendations. According to our theory, sellers rely on the intermediary's inflated recommendations when personalized pricing is not available.²⁴

If the seller could offer personalized prices, it would be able to extract the entire expected willingness to pay from consumers. Since selling the product to picky consumers with the wrong taste reduces surplus, the seller's profit would be reduced with inflated recommendations. Thus, for given fee λ such that $(1 - \lambda)(v_m + b - c)$ is weakly less than b , the seller's profit is maximized by selling to flexible consumers at price $v_m + b$ and to picky consumers with a good match at $v_h + b$. The seller's profit is $(1 - \lambda) \left[(1 - \alpha)(v_m + b - c) + \frac{\alpha}{2}(v_h + b - c) \right]$. The intermediary sets λ such that $\lambda(v_m + b - c) = b$ when it is optimal to serve flexible consumers. Hence, $\lambda = b/(v_m + b - c)$. Thus, when recommending the product to fraction β of picky consumers with the wrong taste, it obtains

$$\begin{aligned} & \lambda \left[(1 - \alpha)(v_m + b - c) + \frac{\alpha}{2}(1 + \beta) \left[\beta(v_l + b - c)/(1 + \beta) + (v_h + b - c)/(1 + \beta) \right] \right] \\ &= \lambda \left[(1 - \alpha)(v_m + b - c) + \frac{\alpha}{2}(1 + \beta) \left[(v_h + \beta v_l)/(1 + \beta) + b - c \right] \right] \\ &= \lambda \left[(1 - \alpha)(v_m + b - c) + \frac{\alpha}{2}(v_h + b - c - \beta(c - b - v_l)) \right]. \end{aligned}$$

The intermediary maximizes its profit with respect to beta and this gives $\beta = 0$ – that is, the intermediary does not have an incentive to inflate recommendations.

This shows that it is the inability to set personalized prices (e.g., because of consumer backlash or the risk of fines imposed by the competition or consumer protection agency) that can give rise to inflated recommendations.

²⁴While a number of theoretical works have looked at personalized pricing (or first-degree price discrimination), the empirical evidence on the use of such pricing is at best mixed (OECD (2018)). Hannak et al. (2014) and Seidenschwarz et al. (2021) do not find clear evidence of the use of personalized pricing in e-commerce.

Multiple new products A key feature of our model is that picky consumers may encounter a good or a bad match and that only one version of the new product is available implying that for some consumers there is no new product available that delivers a good match. In other words, the seller can not provide a menu of different product versions.

If the intermediary could offer a menu of product versions, it may offer specific designs that fit the different tastes of picky consumers and a generic version that works well for flexible consumers but never satisfies picky consumers. With such a portfolio of different product versions the firm could make a monopoly profit on each consumer segment – in our model with two possible taste realizations of picky consumers, there would be three consumer segments, one for each taste realization of picky consumers and one for flexible consumers. This suggests that with a sufficiently large portfolio of product versions, the intermediary would not have an incentive to inflate recommendations.

If the intermediary can offer product versions that work well for picky consumers, but does not have a separate version for flexible consumers, the analysis needs modification, but recommendations will not be inflated either. In our model, this is seen when introducing a second new product that is a good match to picky consumers who would experience a bad match with the first product. Then either one of the following configurations is an equilibrium. Inefficient bypass may arise in which the seller charges $p^I = v_h + b$ and $p^D = v_m$, the intermediary makes personalized recommendations such that each picky consumer receives a recommendation of the new product with a good match. Picky consumers buy in the indirect channel, while flexible consumers bypass the intermediary. The welfare loss compared to the first best is $(1 - \alpha)b$. Alternatively, pooling may arise. Here, the seller charges $p^I = v_m + b$ and $p^D \geq v_m$ and the intermediary recommends the product with a good match to each picky consumer. This implements the first best and consumers obtain a positive net surplus of $\alpha(v_h - v_m)$. This will occur in equilibrium if $\alpha(v_h - v_m) < (1 - \alpha)b$ or, equivalently, α less than $b/(v_h - v_m + b)$. Otherwise, inefficient bypass will occur in equilibrium. Clearly, inflated recommendations are not part of the picture.²⁵ Thus, to obtain inflated recommendations, it is essential to assume that the new product does not provide a good match to all picky consumers – more broadly speaking, there are more different taste realizations among picky consumers than different product designs.

4 Regulatory policy

In this section we explore whether and if so how a regulator could improve welfare if it were able to restrict the intermediaries choices. Whenever the planners assumes control, we postulate that it operates under

²⁵For α below the threshold, there are also equilibria with inflated recommendations, as picky consumers continue to buy indirectly as long as the recommendation is sufficiently informative. However, the intermediary does not improve by introducing a bias and recommending one product more often than others. This would be different if the intermediary were integrated with one of the two new products. Here, the intermediary would engage in self-preferencing for α sufficiently low.

commitment. We are particularly interested in restrictions imposed by the regulator on the recommender system implemented by the intermediary. Let us preview the different types of policy and their welfare properties. 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 this policy in some situations improves on the laissez-faire and in others does strictly worse. Third, we analyze the problem in which the planner mandates a uniform recommendation level and show that this policy welfare-dominates the fully informative recommendation policy. However, in some situations, it does strictly worse than the laissez-faire. Fourth, we consider the planner’s policy to mandate a cap on inflated recommendations, noting that the two polar cases are the obligation of recommending only perfect matches ($\beta = 0$) and of allowing any recommender policy by the intermediary up to some cap $\bar{\beta}$. We show that the policy with the optimal recommendation cap welfare-dominates the fully informative recommendation policy and the laissez-faire. While both the optimal uniform recommendation policy and the optimal recommendation cap policy dominate the fully informative recommendation policy, the welfare ranking among the two policies is ambiguous.

4.1 Full control over the recommendation policy

We consider a regulator that sets recommendation policy $\beta = \beta(\cdot, \cdot)$ but cannot directly affect the intermediary’s price λ . To do so, we assume that the setup cost of the intermediary is 0.²⁶ We will show that the full control over the recommendation policy allows the regulator to reach the first-best total welfare.

To reach the first-best outcome the regulator has to *i*) induce the intermediary to make the seller to sell to flexible consumers through the intermediary as well as to *ii*) minimize excessive recommendations 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 benefits.

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 some predetermined prices (p^I, p^D) and to reveal no information (by recommending to always buy the product) if some different prices are set, so

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

This recommendation policy makes deviations for the seller maximally costly: any deviation in prices by the

²⁶In Appendix D we fully characterize the solution of the regulator that has to ensure some minimal profit of the intermediary to cover the setup cost of the intermediary, $K > 0$. We show that for high enough K , the regulator cannot reach the first best outcome and, under some parameter restrictions optimally induces inflated recommendations (compared to the first best) that allow the intermediary to generate sufficiently high profits to cover its setup cost.

seller leads to the loss of all profit from the picky consumers. To reach the first best, the regulator has to choose minimal β_0 for some prices (p^I, p^D) such that the intermediary and the seller prefer the outcome in which all sales take place in the indirect channel.

The incentive compatibility constraint of the seller is given by

$$(1 - \lambda) \left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha \right) (p^I - c) \geq \max\{(1 - \alpha)(v_m - c), (1 - \lambda)(1 - \alpha)(v_m + b - c)\},$$

where the right-hand represents the maximum of the profits from serving the flexible consumers in the direct and the indirect channels respectively. The intermediary can always ensure the profits of $(1 - \alpha)b$ by setting $\lambda = b/(v_m + b - c)$ and making it more profitable for the seller to serve the flexible consumers in the indirect channel rather than serving them in the direct channel.

We begin by establishing the optimal price p^I that the regulator would set for a given recommendation policy β_0 . The following lemma characterizes the relationship between recommendation policy and induced price in the indirect channel that maximizes social welfare for the case in which the seller operates only through the indirect channel.

Lemma 2. *If sales take place exclusively in the indirect channel, then the regulator induces $p^I = v^m + b \leq \frac{v_h + \beta_0 v_l}{1 + \beta_0} + b$ if $\beta_0 > 0$ and $p^I = \frac{\alpha/2}{\alpha/2 + 1 - \alpha}c + \frac{1 - \alpha}{\alpha/2 + 1 - \alpha}(v_m + b)$ if $\beta_0 = 0$. The welfare loss is given by $\frac{\alpha}{2}\beta_0(c - b - v_l)$.*

The proof is relegated to Appendix A. Intuitively, if $\beta_0 > 0$ and $p^I < v_m + b$, then the regulator could slightly decrease β_0 and increase p^I in its recommendation policy without changing the incentive compatibility constraint of the seller. If $\beta_0 = 0$, then according to the tie-breaking rule the regulator sets fully informative recommendation policy for price p^I that solves $(1 - \lambda)(\alpha/2 + 1 - \alpha)(p^I - c) = (1 - \lambda)(1 - \alpha)(v_m + b - c)$. It is easy to see that p^I maximizes consumer surplus keeping the incentive constraint of the seller satisfied. To see this, note that for $\lambda \leq b/(b + v_m - c)$ we have that

$$(1 - \lambda)(\alpha/2 + 1 - \alpha)(p^I - c) = (1 - \lambda)(1 - \alpha)(v_m + b - c) \geq (1 - \alpha)(v_m - c).$$

Therefore, we obtain that the regulator can reach the first best by setting $\beta_0 = 0$ since the total generated profits in the indirect channel are large enough to induce the first-best outcome. We summarize the preceding analysis by the following proposition.

Proposition 3. *Suppose that the regulator has full control over the recommendation policy. Then it can achieve the first-best outcome by setting $\beta \left(p^I = \frac{\alpha/2}{\alpha/2 + 1 - \alpha}c + \frac{1 - \alpha}{\alpha/2 + 1 - \alpha}(v_m + b), p^D = v_m \right) = 0$ and $\beta = 1$ otherwise. In equilibrium, the intermediary sets $\lambda = \frac{b}{v_m + b - c}$ and the seller sets price $p^I = \frac{\alpha/2}{\alpha/2 + 1 - \alpha}c + \frac{1 - \alpha}{\alpha/2 + 1 - \alpha}(v_m + b)$. All flexible consumers and all picky consumers with a good match buy the product through the indirect channel.*

Since we assumed that the regulator selects the solution that is best for consumers and there is some leeway in the final retail prices, flexible consumers are better off than under laissez-faire. Picky consumers are necessarily better off as they benefit from fully informative product recommendations, whereas under laissez-faire recommendations would be inflated.

The proposition shows that the regulator's recommendation policy fully determines the total profit of the intermediary and the seller when selling only in the indirect channel. We showed that the total profit with an efficient allocation can be made large enough to induce the intermediary to share profits such that the seller finds it optimal to sell only through the indirect channel. Therefore, the regulator does not need control over the intermediary's revenue policy λ to achieve the first best

The regulator's optimal recommendation policy does not change if the regulator maximizes consumer surplus, as the following remark shows.

Remark 1. *Suppose that the regulator's objective is to maximize consumer surplus. If all sales happen in the indirect channel, consumer surplus is equal to*

$$(1 - \alpha)(v_m + b) + \frac{\alpha}{2}(v_h + b) + \frac{\alpha}{2}\beta_0(v_l + b) - \left(\frac{\alpha}{2}(1 + \beta_0) + (1 - \alpha)\right)p^I.$$

Clearly, the regulator has the incentive to induce the lowest possible price p^I and recommendation policy β_0 such that the incentive compatibility constraint of the seller is satisfied

$$(1 - \lambda) \left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha\right) (p^I - c) \geq \max\{(1 - \alpha)(v_m - c), (1 - \lambda)(1 - \alpha)(v_m + b - c)\}.$$

We showed that if $\left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha\right) (p^I - c) \geq (1 - \alpha)(v_m + b - c)$, then the intermediary will induce sales in the indirect channel at p^I by setting λ such that the seller does not find it profitable to serve the flexible consumers directly. Therefore, consumers surplus is equal to

$$\begin{aligned} (1 - \alpha)(v_m + b) + \frac{\alpha}{2}(v_h + b) + \frac{\alpha}{2}\beta_0(v_l + b) - \left(\frac{\alpha}{2}(1 + \beta_0) + (1 - \alpha)\right) c - (1 - \alpha)(v_m + b - c) \\ = \frac{\alpha}{2}\beta_0(v_l + b - c) + \frac{\alpha}{2}(v_h + b - c) \end{aligned}$$

and is maximized at $\beta_0 = 0$. This is because the consumer surplus-maximizing regulator has to induce intermediary and seller to sell only through the indirect channel and, thus, has to permit high enough profits. The regulator maximizing consumer welfare chooses the same recommendation policy as the one specified in Proposition 3.

4.2 Mandated fully informative recommendations

Consider a policy intervention of the regulator in which β is chosen by the regulator and does not depend on the seller's prices, so $\beta(p^I, p^D) = \beta_{\text{UNI}}$ for all (p^I, p^D) . In this section we suppose that the regulator may impose the policy that recommendations must be fully informative ($\beta_{\text{UNI}} = 0$). In the context of the intermediary being vertically integrated with the seller of the new product, this can be interpreted to correspond

to the prohibition of self-preferencing since this implies that no consumer receives a recommendation for a bad match, which would not be in the consumer's interest.

As we will see, for sufficiently small α mandating fully informative recommendations implements the first best. As derived in the proof of Proposition 2, the critical value α_{FI} below which imposing fully informative recommendations does not violate the incentive compatibility constraints of intermediary and seller is given by the solution to

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

For high α , the regulator implements the laissez-faire ($\alpha \geq \bar{\alpha}$). However, for intermediate values of α – that is, $\alpha \in (\alpha_{FI}, \bar{\alpha})$ the naive policy performs even worse than the laissez-faire, as the unconstrained intermediary inflates its recommendation and all sales occur in the indirect channel.

Proposition 4. *When the regulator mandates that recommendations must be fully informative (i.e., $\beta = 0$), the equilibrium is characterized as follows:*

- if $\alpha < \alpha_{FI}$, then the first-best allocation is implemented.
- if $\alpha \geq \alpha_{FI}$ where $\alpha_{FI} \in (0, \bar{\alpha})$, then the intermediary sets $\lambda^* = 1$, for all (p^I, p^D) ; equilibrium prices are given by $(p^I, p^D) = (v_h + b, v_m)$. Picky consumers go to the indirect channel and buy if they receive recommendations to buy, whereas all the flexible consumers buy in the direct channel. For $\alpha > \bar{\alpha}$ this is the same outcome as under laissez-faire. Welfare losses are given by $(1 - \alpha)b$.

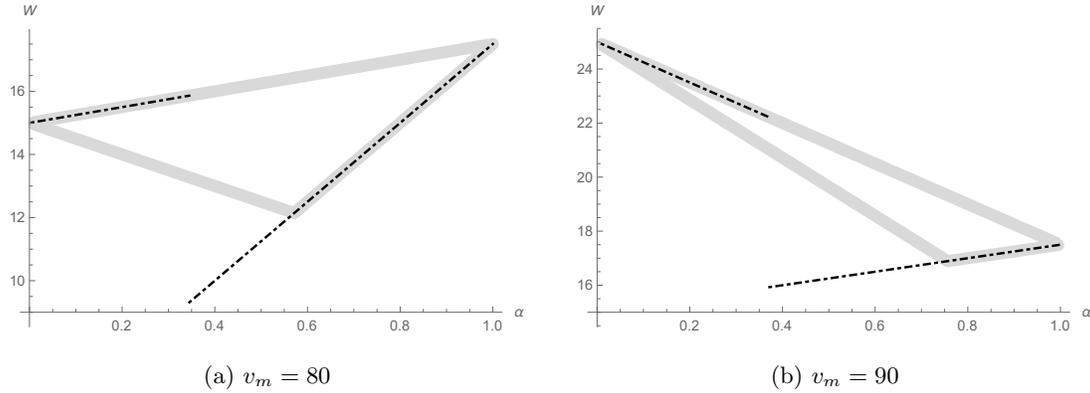


Figure 2: Total welfare in α ; $v_h = 100, v_l = 20, b = 10, c = 75$. First best outcome, private solution (solid), fully informative recommendations (dot-dashed).

The proof is relegated to Appendix A. We illustrate our findings in Figure 2. The upper grey line depicts welfare in the first best, while the lower grey line depicts welfare under laissez-faire. Welfare under mandated

fully informative recommendations is depicted by the solid line. As shown in Proposition 4, for $\alpha \leq \alpha_{\text{FI}}$, the planner's policy implements the first best. For $\alpha \geq \bar{\alpha}$, the planner does not improve on the laissez-faire outcome, as $\beta_{\text{UNI}} = 0$ implies that flexible consumers buy in the direct channel, which is also happening under laissez-faire. Remarkably, the planner's policy $\beta_{\text{UNI}} = 0$ performs worse than the laissez-faire for $\alpha \in (\alpha_{\text{FI}}, \bar{\alpha})$. Mandating fully informative recommendations in this range does not allow the intermediary to inflate recommendations. This implies that any picky consumers who end up buying generate a higher surplus than any flexible consumer in the direct channel ($v_h + b > v_m + b$). The intermediary has two potentially profit-maximizing options: Either 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 it can extract all surplus from picky consumers with a good match. For $\alpha > \alpha_{\text{FI}}$ it prefers the latter. Therefore, the policy $\beta_{\text{UNI}} = 0$ 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 $\beta_{\text{UNI}} = 0$, the planner 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, which leads to a welfare loss.

While, as we have shown in Section 3.2, under laissez-faire the outcome is the same independent of whether intermediary and seller are integrated, this is not true when the regulator mandates fully informative recommendations. The reason is as follows. Under inefficient bypass, the seller obtains $(1 - \alpha)(v_m - c)$, whereas, when all trade takes place in the indirect channel, the seller obtains $(1 - \alpha)(v_m - c) + (1 - \lambda)\frac{\alpha}{2}(v_h + b - c)$, which is larger than under inefficient bypass. The vertically integrated firm maximizes total profit, whereas the intermediary maximizes total profit minus the seller's profit. Therefore, the intermediary has a relatively more favorable views of inefficient bypass than the vertically integrated firm and, thus, α_{FI} is less than the critical α under the current regulation of a vertically integrated firm, which solves $[\alpha/2 + (1 - \alpha)](v_m + b - c) = (\alpha/2)(v_h + b - c) + (1 - \alpha)(v_m - c)$ and, thus, is $b/[b + (v_h - v_m)/2]$. 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.

4.3 Recommendation cap

We consider a regulator who can set the upper bound on the recommendation level $\bar{\beta}$ and does not control λ .

Whatever the planner does, the intermediary can always induce the seller to sell to picky consumers in the indirect channel by setting $\beta = 0$ and $\lambda = 1$. Thus, the lower bound on the intermediary's profit is

$$\frac{\alpha}{2}(v_h + b - c).$$

We explore how the optimal recommendation cap affects the intermediary's strategy. A sufficiently high recommendation cap strictly less than one is welfare-dominated by laissez-faire because the intermediary is limited in its ability to punish a seller who deviates to selling directly to flexible consumers. This implies that the intermediary has to offer a larger fraction of the total profit to the seller. If the recommendation cap is too high, this, however, incentivizes the intermediary to not induce the seller to use the indirect channel only. Instead, it recommends the product only to picky consumers with a good match and extracts the full surplus from those consumers. Such an outcome is welfare-inferior than the laissez-faire because the welfare gain from fully informative recommendations is less than the welfare loss arising from flexible consumers using the inefficient direct channel.

A lower recommendation cap may improve on the laissez-faire. While this does not happen for high values of α , for lower values of α we distinguish between two regimes. For very low values of α the optimal recommendation cap is $\bar{\beta} = 0$, which coincides with the fully informative recommendation policy. Here, the concern for flexible consumers is overwhelming and the intermediary is keen on inducing sales in the indirect channel only. For intermediate values of α , the regulator has to allow for some inflated recommendation (compared to the first best) to satisfy the incentive constraints of the seller and the intermediary. This policy improves on the laissez-faire and the fully informative recommendation policy. The comparison to the optimal uniform recommendation policy is more intricate and will be discussed in the following section.

The optimal recommendation cap policy is characterized in the following proposition; the proof is relegated to Appendix A.

Proposition 5. *Suppose that the regulator is restricted to set a recommendation cap $\bar{\beta}$. Then*

- *if $\alpha \leq \alpha_{FI}$, then the optimal policy is $\bar{\beta} = 0$. The outcome is first best: All flexible consumers and picky consumers with good matches buy through the indirect channel.*
- *if $\alpha \in (\alpha_{FI}, \alpha_{CAP})$, then the optimal policy has the property $\bar{\beta} \in (0, \frac{v_h - v_m}{v_m - v_l})$. The intermediary sets $\beta = \bar{\beta}$ and $\lambda \in (0, 1)$. Equilibrium prices are given by $(p^I, p^D) = (v_m + b, v_m)$. All flexible consumers, picky consumers with good matches, and a fraction $\bar{\beta}$ of picky consumers with bad matches buy through the indirect channel. The welfare loss compared to the first best is equal to $\frac{\alpha}{2}\bar{\beta}(c - b - v_l)$.*
- *$\alpha \geq \alpha_{CAP}$, then the SP sets $\bar{\beta} = 1$. The equilibrium coincides with the laissez-faire.*

We observe that the optimal policy $\bar{\beta}(\alpha)$ is weakly increasing in α – see Figure 3. In the interval $[0, \alpha_{FI}]$ the regulator imposes the most stringent cap ($\bar{\beta} = 0$) and still achieves that flexible consumers use the indirect channel. For higher values of $\alpha \in (\alpha_{FI}, \alpha_{CAP})$ it has to be more accommodating to the intermediary and allow for some inflated recommendations albeit less than under laissez-faire. This is no longer feasible for higher values of α in which case the optimal policy is to be completely unrestrictive. This means that at

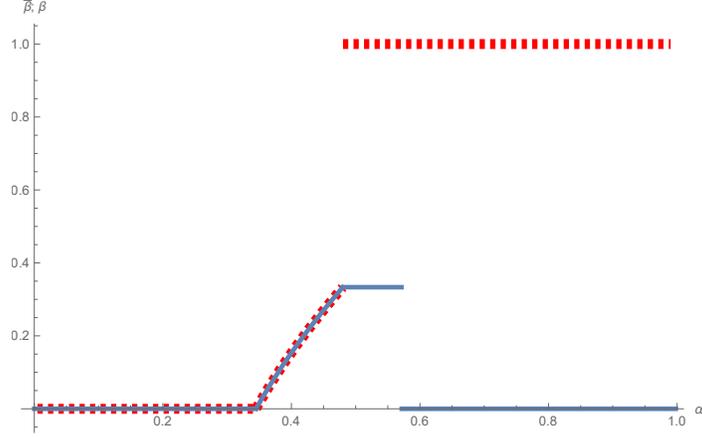


Figure 3: Optimal recommendation cap in α (dotted), equilibrium recommendation level (solid); $v_h = 100$, $v_m = 80$, $v_l = 20$, $b = 10$, $c = 75$

α_{CAP} there is an upward jump of the optimal policy to $\bar{\beta} = 1$. The policy is then no longer binding along the equilibrium path. However, such an unrestricted policy is strictly preferred by the regulator to more restrictive policies because the latter limit the intermediary's ability to punish a seller deviating by selling to flexible consumers in the direct channel. Thus, the laissez-faire outcome in which flexible consumers use the indirect channel is implemented for $[\alpha_{CAP}, \bar{\alpha}]$ – we note that the regulator prefers this outcome over the alternative that flexible consumers buy directly and only picky consumers with good matches buy in the indirect channel. For yet higher values $\alpha > \bar{\alpha}$ flexible consumers use the direct channel, which, in this range of α , is also preferred by the regulator.

Thus, we have three parameter regions of α . For low α , the optimal recommendation cap regulation implements the first best; for intermediate values of α , it implements an allocation that is strictly better than the laissez-faire but cannot implement the first best; and, for high values of α , it does not impose any restriction on the intermediary and, therefore, implements the laissez-faire. The associated welfare is illustrated in Figure 4 with the two parameter constellations considered in the previous subsections.

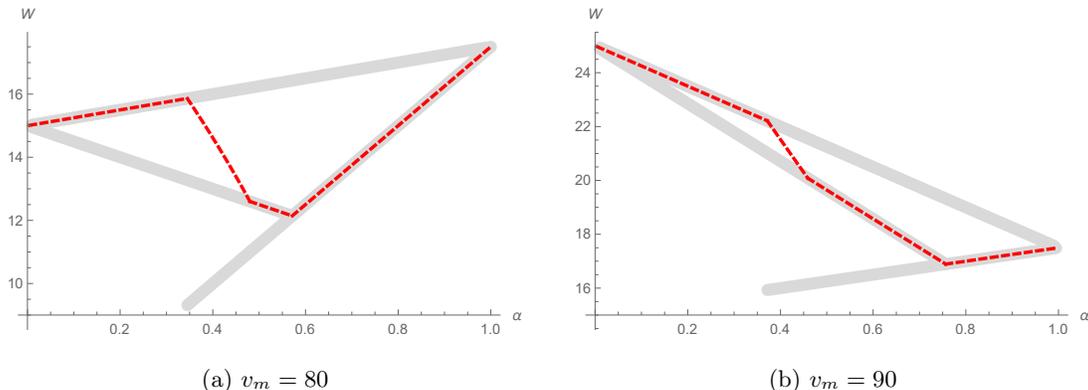


Figure 4: Total welfare in α ; $v_h = 100, v_l = 20, b = 10, c = 75$. First best outcome, private solution, fully informative recommendations (solid), recommendation cap (dashed).

In Appendix C we consider a regulator who imposes a uniform recommendation level. The regulator may optimally choose inflated recommendations, but a lower level than what would prevail under *laissez-faire*. While the optimal uniform policy weakly improves on the policy that mandates fully informative recommendations, it may still backfire and deliver lower welfare than the *laissez-faire*. The uniform recommendation level regulation is sometimes superior and at other times inferior to the recommendation cap regulation. However, for practical considerations, recommendation cap regulation appears to be more relevant than uniform recommendation level regulation that imposes inflated recommendations, as it seems to be hard to implement a policy in which a firm may be fined for providing too informative recommendations.

4.4 Regulating the intermediary's rent extraction

Up to now we considered regulations that impose restrictions on the intermediary's recommendation policy. Alternatively, 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 to buy in the indirect channel and the picky consumers with a bad match not to buy. 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 that the seller serves also 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* features $\lambda < 1$ with inflated recommendations, while inefficient bypass has $\lambda = 1$ under *laissez-faire*. In other words, a cap on λ 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 on λ is larger than under

laissez-faire. Hence, such price regulation leads more often to inflated recommendations than the laissez-faire. While consumers do not benefit from such a policy, rents are 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 severe recommendation inflation, decides to impose a fully informative recommendation policy $\beta = 0$, the intermediary's hands are completely tied and it is only the seller which determines the outcome. The seller decides whether to sell to all flexible and all picky consumers with a good match in the indirect channel at price $v_m + b$ or to sell to flexible consumers directly at price v_m and to picky consumers with a good match at price $v_h + b$. In the former, the seller obtains $(1 - \lambda)(\alpha/2 + (1 - \alpha))(v_m + b - c)$ and, in the latter, $(1 - \alpha)(v_m - c) + (1 - \lambda)(\alpha/2)(v_h + b - c)$. From the viewpoint of the seller the best regulation would be $\lambda = 0$ in which case we obtain the critical α equal to $b/(b + (v_h - v_m)/2)$, which is strictly less than 1. For larger α , the seller decides to inefficiently sell to flexible consumers in the direct channel. For positive λ , inefficient bypass becomes more attractive. Thus, the regulator fails to implement the first best for α above the threshold.

Furthermore, this regulatory policy ($\beta = 0$ and $\lambda < 1$) may backfire and give lower welfare than under laissez-faire because the critical α under this regulation is strictly lower than under laissez-faire. Formally, this is seen as follows. The critical α under laissez-faire is given by the solution to $\frac{\alpha}{2}(1 + \beta)(v_m + b - c) + (1 - \alpha)b = \frac{\alpha}{2}(v_h + b - c)$, while the critical α under the current regulation satisfies $(\alpha/2 + (1 - \alpha))(v_m + b - c) = \frac{1 - \alpha}{1 - \lambda}(v_m - c) + \frac{\alpha}{2}(v_h + b - c)$, which can be rewritten as $(\alpha/2 + (1 - \alpha))(v_m + b - c) - \frac{1 - \alpha}{1 - \lambda}(v_m - c) = \frac{\alpha}{2}(v_h + b - c)$. Since, by simple manipulation, one can show that $\frac{\alpha}{2}(1 + \beta)(v_m + b - c) + (1 - \alpha)b$ is larger than $(\alpha/2 + (1 - \alpha))(v_m + b - c) - \frac{1 - \alpha}{1 - \lambda}(v_m - c)$ for any λ , it must be that $\bar{\alpha}$ is necessarily larger than the critical α under the regulation considered here.

To summarize, even if, in addition to imposing fully informative recommendations, the regulator forces the intermediary not to absorb any profit (or up to a fraction thereof), the regulatory policy may backfire. It is now the seller who prefers to set prices such that it extracts all consumer surplus with inefficient bypass instead of implementing the efficient allocation in which case it has to leave a positive net surplus on the table for picky consumers with a good match.

5 The intermediary's policy without commitment

In this section we consider the alternative timing in which the intermediary can adjust its recommendations in response to the prices set by the seller.²⁷

The timing of the game is as follows. First, the intermediary sets λ . Second, the seller sets its prices in

²⁷The dynamic interpretation is similar to the one given for the model with commitment. The important difference is that consumers expect the intermediary to best-respond to the seller's prices p^D and p^I instead of expecting the intermediary to stick to some policy $\beta(p^D, p^I)$. In other words, consumers think that the intermediary is accommodating to the seller and does not punish it for behaving this way.

the direct channel p^D and in the indirect channel p^I . Third, the intermediary observes prices and chooses its recommendation policy β . Fourth, consumers observe prices p^D and p^I and the recommendation policy β and then decide which sales channel to choose. Fifth, picky consumers in the indirect channel receive personalized recommendations and all consumers make their purchasing decisions.

With this alternative timing, the intermediary best-responds to the seller's prices with its recommendation level β .

Before considering a mix of flexible and picky consumers let us point out that the outcomes in the two polar cases do not differ from the ones in the commitment case. If all consumers are flexible ($\alpha = 0$), then $p^D = v_m$ and $p^I = v_m + b$ and the seller makes profit that is equal to $\max\{v_m - c, (1 - \lambda)(v_m + b - c)\}$. The intermediary chooses λ such that the seller is indifferent between serving flexible consumers in the indirect and the direct channels. Thus, the optimal sharing contract is characterized by

$$\lambda^* = \frac{b}{v_m + b - c}.$$

The seller and the intermediary earn $v_m - c$ and b respectively. Total surplus is maximal and equals to $v_m + b - c$.

If all consumers are picky ($\alpha = 1$), the seller cannot sell to consumers in the direct channel and, thus, it will accept any $\lambda \leq 1$. In every subgame with $\lambda \in [0, 1]$, the seller sets $p^I = v_h + b$ and the intermediary makes fully informative recommendations – this maximizes the intermediary's and the seller's profit. In the equilibrium of the full game, total surplus is maximal and equal to $\frac{1}{2}(v_h + b - c)$. It is fully extracted by the intermediary (since $\lambda = 1$).

5.1 Equilibrium characterization under laissez-faire

We recall consumer choices after observing prices (p^I, p^D) , recommendation policy β , and, for picky consumers, personalized recommendations. The flexible consumers decide to buy in the indirect channel if and only if $v_m + b - p^I \geq \max\{v_m - p^D, 0\}$. Picky consumers buy in the indirect channel, after receiving a personalized recommendation, if and only if $\frac{v_h + \beta v_l}{1 + \beta} + b - p^I \geq 0$. Picky consumers who did not receive a recommendation do not buy.

Next, we investigate how the intermediary chooses its recommendation policy in response to the sharing contract $\lambda \in [0, 1]$ and the seller's prices (p^I, p^D) . If $p^D \leq p^I - b$, we can restrict attention to $p^I \in (c, v_h + b]$, and $p^D \leq v_m$. Then, the seller will induce the inefficient bypass outcome – that is, the flexible consumers will buy in the direct channel and the picky consumers will visit the indirect channel. The profit of intermediary is maximized when the expected match value of picky consumers in the indirect channel is equal to $p^I - b$ – that is, $\frac{v_h + \beta v_l}{1 + \beta} + b = p^I$. This implies that $\beta = \frac{v_h - (p^I - b)}{p^I - b - v_l}$. Then, the profit of the seller is

$$(1 - \alpha)(p^D - c) + (1 - \lambda) \frac{\alpha}{2} (v_h - v_l) \frac{p^I - c}{p^I - b - v_l}.$$

It is easy to see that, for any $\lambda > 0$, the profit of the seller is maximized at $p^D = v_m$ and $p^I = v_h + b$, which results in $(1 - \alpha)(v_m - c) + (1 - \lambda)\frac{\alpha}{2}(v_h + b - c)$.

If $p^D > p^I - b$, we can restrict attention to $p^I \leq v_m + b$ because otherwise, the inefficient bypass outcome from above is profit-increasing. The seller induces the outcome with inflated recommendations and all consumers will visit the indirect channel. The profit of the intermediary is maximized when all flexible and picky consumers with the recommendation to buy end up with a net surplus of zero – that is, $\beta = \frac{v_h - (p^I - b)}{p^I - b - v_l}$. The profit of the seller is equal to

$$(1 - \lambda) \left(\frac{\alpha}{2}(1 + \beta) + (1 - \alpha) \right) (p^I - c) = (1 - \lambda) \left(\frac{\alpha}{2} \frac{v_h - v_l}{p^I - b - v_l} + (1 - \alpha) \right) (p^I - c).$$

The seller maximizes its profit by setting $p^D = v_m$ and $p^I = v_m + b$ and obtains

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

Thus, the intermediary obtains the remaining fraction λ of total surplus $\left(\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} + (1 - \alpha) \right) (v_m + b - c)$.

The seller finds it optimal to induce the outcome with inflated recommendations if and only if

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

Note that if $\alpha \geq \bar{\alpha}$ (as defined in section 3.2), then, for all λ , the seller will find it optimal to induce inefficient bypass because

$$\begin{aligned} & (1 - \lambda) \left(\left(\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} + (1 - \alpha) \right) (v_m + b - c) - \frac{\alpha}{2}(v_h + b - c) \right) \\ & \leq \underbrace{\left(\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) + (1 - \alpha)b - \frac{\alpha}{2}(v_h + b - c) \right)}_{\leq 0} + (1 - \alpha)(v_m - c) \\ & \leq (1 - \alpha)(v_m - c). \end{aligned}$$

The profit of the intermediary is $\frac{\alpha}{2}(v_h + b - c)$.

Otherwise, if $\alpha < \bar{\alpha}$, the intermediary can induce the outcome with inflated recommendations by setting

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

It is easy to establish that λ^* decreases from $b/(v_m + b - c)$ to 0 when α goes from 0 to $\bar{\alpha}$. The profit of the intermediary is equal to

$$\begin{aligned} & \lambda^* \left[\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} + (1 - \alpha) \right] (v_m + b - c) \\ & = \frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) + (1 - \alpha)b - (1 - \lambda^*)\frac{\alpha}{2}(v_h + b - c). \end{aligned}$$

The intermediary will find it optimal to induce the outcome with inflated recommendations if and only if the profits under inflated recommendation are higher than the profits under inefficient bypass – that is,

$$-\frac{\alpha}{2} \left[\underbrace{(v_h + b - c) - \frac{v_h - v_l}{v_m - v_l} (v_m + b - c)}_{>0} \right] + (1 - \alpha)b - (1 - \lambda^*(\alpha)) \frac{\alpha}{2} (v_h + b - c) \geq 0.$$

We observe that all three terms on the left-hand side above are continuously decreasing in α . The left-hand side is positive at $\alpha = 0$ and negative at $\alpha = \bar{\alpha}$. Thus, there exists a unique $\alpha_3 \in (0, \bar{\alpha})$ such that for all $\alpha \leq \alpha_3$ the equilibrium features inflated recommendations. Otherwise, if $\alpha > \alpha_3$, then the equilibrium features inefficient bypass. We summarize our analysis with the following proposition.

Proposition 6. *The equilibrium when the intermediary cannot commit to its recommendation policy is characterized as follows:*

- for $\alpha \geq \alpha_3$, the intermediary sets $\lambda^* = 1$ in the first stage, and $\beta = 0$ along the equilibrium path in the third stage. The seller sets $(p^I, p^D) = (v_h + b, v_m)$. Picky consumers visit the intermediary and buy if and only if they receive the recommendation to buy; all flexible consumers buy in the direct channel. Welfare losses are equal to $(1 - \alpha)b$.
- for $\alpha < \alpha_3$, the intermediary sets

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

and $\beta = \frac{v_h - v_m}{v_m - v_l}$ along the equilibrium path. The seller sets $(p^I, p^D) = (v_m + b, v_m)$. All consumers visit the intermediary. All flexible consumers and picky consumers with a recommendation buy in the indirect channel. Welfare losses are equal to $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$.

We recall that in the inflated recommendations outcome, a fraction β of picky consumers with a bad match buy. We note that whenever there are inflated recommendations, the level of recommendations is the same as under commitment. However, the critical α below which inflated recommendations are the equilibrium outcome is lower than with commitment ($\alpha_3 < \bar{\alpha}$). This is due to the fact that it becomes harder for the intermediary to induce the seller to not sell directly to flexible consumers. Without commitment the intermediary has to leave some extra rents – to be precise, $(1 - \lambda^*) \frac{\alpha}{2} (v_h + b - c)$ – to the seller as the severe “punishments” under commitment are no longer credible.

The proposition shows that, for $\alpha \notin (\alpha_3, \bar{\alpha})$, commitment power does not change the outcome of the game and the intermediary induces the outcome of the vertically-integrated firm. By contrast, for $(\alpha_3, \bar{\alpha})$, the vertically integrated outcome cannot be obtained when the intermediary cannot commit to its recommendation policy.

We report the welfare findings in Figure 5. For $\alpha \in (\alpha_3, \bar{\alpha})$, total welfare is even less without commitment than with commitment. If we interpret the move from commitment to no commitment as weakening the intermediary, then weakening the intermediary is bad for social welfare.

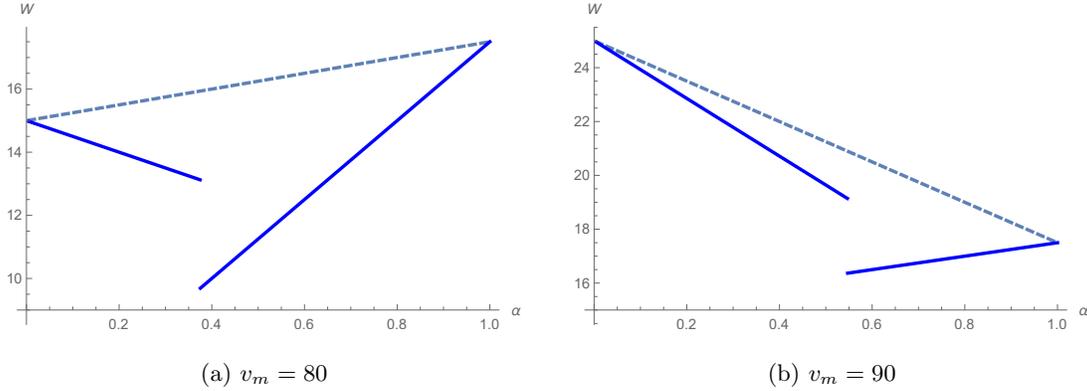


Figure 5: Total welfare in α ; $v_h = 100, v_l = 20, b = 10, c = 75$. First best outcome (dashed), private solution under no commitment (solid).

5.2 Mandated fully informative recommendations v. laissez-faire without commitment

In this section, we compare the equilibrium without commitment to the outcome when a regulator mandates fully informative recommendations. The latter has been characterized in Section 4.2 and is a meaningful comparison because the regulator fixes $\beta = 0$ and, thus, does not leave any discretion to the intermediary regarding its recommendation policy. The question for intermediary and seller is whether to sell to flexible consumers directly or indirectly. If flexible consumers buy in the direct channel, the inefficient bypass outcome is the same under laissez-faire and regulation. What are the seller's incentives to set prices such that flexible consumers buy in the indirect channel? Under regulation with $\beta = 0$ the seller obtains a fraction $1 - \lambda$ of the joint profit $(\frac{\alpha}{2} + 1 - \alpha)(v_m + b - c)$, while, under laissez-faire, it obtains a fraction of $(\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} + (1 - \alpha))(v_m + b - c)$. Since the former is less than the latter, the intermediary has to give a larger fraction of joint profits to the seller under regulation. Since the joint profit from selling to flexible consumers in the indirect channel is strictly smaller under regulation and the intermediary has to confine itself with a strictly smaller share λ , the critical α below which flexible consumers buy in the indirect channel must be strictly lower under regulation than under laissez-faire without commitment – that is, $\alpha_{FI} < \alpha_3$.

Remark 2. *Mandated fully informative recommendations lead to higher welfare than the laissez-faire without commitment for $\alpha < \alpha_{FI}$, but lower welfare for $\alpha \in (\alpha_{FI}, \alpha_3)$. Welfare is the same for $\alpha \geq \alpha_3$.*

While the welfare comparison can go in either way, the range of values for α such that regulation backfires

is smaller than in an environment in which the intermediary can commit to its recommendation policy. Figure 6 illustrates our finding.

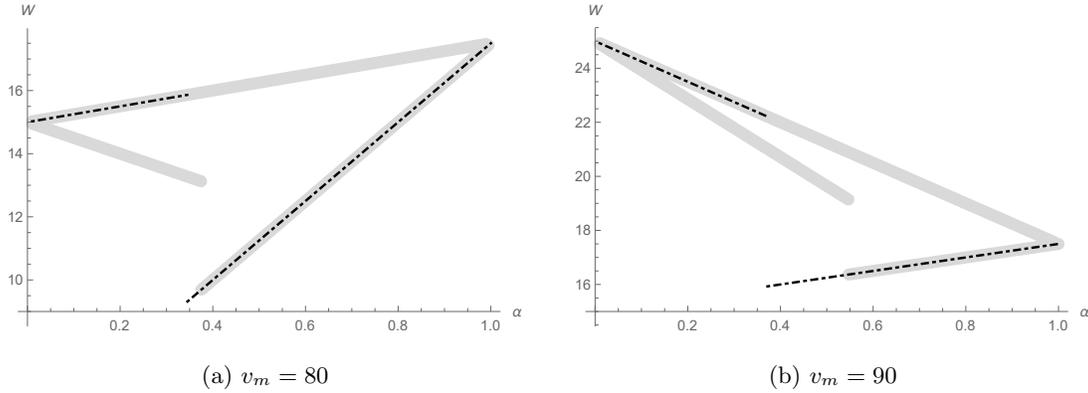


Figure 6: Total welfare in α ; $v_h = 100, v_l = 20, b = 10, c = 75$. First best outcome, private solution (solid), fully informative recommendations (dot-dashed).

5.3 Price Parity Clauses (PPC)

An important policy concern in recent years has been the use of price parity clauses imposed by intermediaries that restrict the pricing behavior of sellers. In our context, we consider price parity clauses that restrict the seller in its pricing in the direct channel – that is, the seller is not allowed to set a lower price when selling directly, $p^D \geq p^I$. This implies that price parity clauses do not allow the seller to sell to the flexible consumers in the direct channel and set a high price to the picky consumers in the indirect channel, which makes the inefficient bypass outcome in which all flexible consumers buy directly and some picky consumer buy indirectly unattainable. In particular, we assume that together with λ the intermediary decides whether to impose a price parity clause in the first stage.

By imposing price parity the intermediary can make sure that by not selling to flexible consumers in the indirect channel the seller will make profit $(1 - \alpha)(v_m - c)$. This is the same profit the seller can always achieve under commitment. Thus, to implement inflated recommendations, the intermediary will impose price parity and the same outcome as under commitment will be implemented whenever the intermediary prefers to implement this outcome. Alternatively, the intermediary may want to refrain from imposing price parity. This is in its best interest if it wants the inefficient bypass outcome to be implemented. The intermediary's tradeoff is thus the same as under commitment and the critical α is given by $\bar{\alpha}$.

Remark 3. *If the intermediary can impose a price parity clause at the first stage, it uses this option for $\alpha < \bar{\alpha}$ and the same inflated recommendations outcome as in Proposition 2 will prevail, while it does not use the option for $\alpha > \bar{\alpha}$ and the same inefficient bypass outcome as in Proposition 2 will prevail.*

This implies that, whenever price parity clauses make a difference, they are welfare-increasing as they solve the intermediary’s commitment problem and limit the seller’s profit.

6 Conclusion

An intermediary may recommend a new product to consumers in the indirect sales channel; alternatively the product may be purchased directly from the seller. The indirect channel offers two advantages: its use increases the benefit of all consumers and consumers may appreciate if the intermediary recommends to buy the product. The main result is that the intermediary may inflate recommendations – this is, it recommends the new product to more consumers than is socially optimal. In our setting, this happens because the seller cannot set different prices to different consumers using the same sales channel depending on whether or not they care about the specific design of the product. By recommending the product to picky consumers with a bad match, which is known to the intermediary *ex ante*, but not to anybody else, the intermediary can bring the expected valuation of picky consumers down to the one of flexible consumers and increase the sales volume of the seller without lowering the price. When the intermediary commits to extract a certain profit share for the seller and a recommendation policy, the intermediary either induces inflated recommendations or inefficient bypass. This insight is robust to a number of extensions which, depending on the particular extension, makes inflated recommendations more or less pervasive. Essential for the result is that the seller cannot use personalized prices and that, in a more general context, there are fewer product versions than there are taste realizations.

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 recommendation policy as a function of retail prices, it would be able to implement the first best. However, if the regulator can only impose a cap, the first best cannot always be implemented. A particularly simple and tempting regulatory intervention is 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. When there are not too many picky consumers in the population, this regulation can implement the first best. However, above a critical threshold, inefficient bypass will happen and the regulation backfires, as it delivers lower welfare than the *laissez-faire*. The optimal regulatory recommendation cap improves on the mandated fully informative recommendation policy and cannot backfire relative to the *laissez-faire*.

In our analysis, we assumed that the intermediary has full commitment power. As a result, the equilibrium implements the vertically integrated solution in which intermediary and seller act jointly as a single decision-maker. If the intermediary decides on whom to recommend the product after the seller has set retail prices, certain recommendation policies are no longer credible. This implies that the intermediary has to give larger rents to the seller to implement the outcome that all trade takes place in the indirect channel.

Consequently, inefficient bypass will be the equilibrium outcome for a larger range of parameter values. The intermediary is no longer able to implement the vertically integrated solution and welfare is lower than with commitment. However, even without the intermediary's ability to commit to a recommendation policy upfront, the regulation that recommendations must be fully informative may still backfire, albeit for a smaller parameter range than with commitment. The ability to impose price parity compensates for the inability to commit to the recommendation policy upfront. Thus, the same outcome as under laissez-faire with commitment (in which case price parity clauses are irrelevant for the outcome) is obtained, which implies that price parity clauses lead to higher welfare, as the laissez-faire with commitment has better welfare properties than the laissez-faire without commitment. In other words, prohibiting price parity clauses leads to lower welfare.

Appendix

A Relegated Proofs

Proof of Lemma 2. First, consider the case of $\beta_0 > 0$.

Suppose that $p^I > v_m + b$ in the equilibrium. Then, the flexible consumers do not buy in the indirect channel and it is optimal for the regulator to ensure sales in the direct channel inducing $p^D = v^m$ in the equilibrium. This contradicts the assumption that sales take place only in the indirect channel and implies that $p^I < v_m + b$.

Suppose that $\frac{v_h + \beta_0 v_l}{1 + \beta_0} + b < p^I \leq v_m + b$ and the picky consumers do not buy in the indirect channel. The welfare loss is given by $\frac{\alpha}{2}(v_h + b - c)$. Note that the regulator can always reach welfare loss of $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$ by mimicking the strategy of the intermediary that induces the outcome with inflated recommendations (see Proposition 2). Since $v_m > c - b$ we have that

$$\frac{v_h - v_m}{v_m - v_l} < \frac{v_h - (c - b)}{c - b - v_l},$$

which implies that the outcome with inflated recommendations described in Proposition 2 leads to lower welfare losses. Therefore, $p^I \leq \min\{\frac{v_h + \beta_0 v_l}{1 + \beta_0} + b, v_m + b\}$.

Assume for a contradiction that $p^I < \frac{v_h + \beta_0 v_l}{1 + \beta_0} + b < v_m + b$. Since the total welfare loss depends primarily on β_0 but not on p^I , the regulator can slightly decrease β_0 and induce a slightly higher price in the indirect channel p^I that would satisfy the seller's and the intermediary's incentive constraints. This leads to strictly lower welfare loss, a contradiction.

Next, we show that $p^I = \frac{v_h + \beta_0 v_l}{1 + \beta_0} + b < v_m + b$ cannot be in the equilibrium as the regulator will always have incentives to decrease β_0 and reduce the number of inflated recommendations. In this case the total profit of the seller and the intermediary is given by

$$\left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha\right) \left(\frac{v_h + \beta_0 v_l}{1 + \beta_0} + b - c\right).$$

The sign of the derivative with respect to β_0 , for $\beta_0 > \frac{v_h - v_m}{v_m - v_l}$ is determined by

$$\begin{aligned} & \frac{\alpha}{2} \left(\frac{v_h - v_l}{1 + \beta_0} + v_l + b - c\right) - \left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha\right) \frac{v_h - v_l}{(1 + \beta_0)^2} \\ & - \frac{\alpha}{2}(c - b - v_l) - (1 - \alpha) \frac{v_h - v_l}{(1 + \beta_0)^2} < 0, \end{aligned}$$

which implies that $\beta_0 = \frac{v_h - v_m}{v_m - v_l}$ results in a higher total profit. This recommendation policy would allow the regulator to reach higher social welfare keeping the incentive constraint of the intermediary satisfied. Therefore, $\beta_0 \leq \frac{v_h - v_m}{v_m - v_l}$.

Finally, assume for a contradiction that $p^I < v_m + b \leq \frac{v_h + \beta_0 v_l}{1 + \beta_0} + b$. If $\beta_0 > 0$, then the regulator can slightly decrease β_0 and induce a slightly higher price p^I such that the seller's and the intermediary's incentive constraints are satisfied. This would lead to strictly higher social welfare, a contradiction.

The remaining case of $p^I = v^m + b \leq \frac{v_h + \beta_0 v_i}{1 + \beta_0} + b$ establishes the result of the lemma for $\beta_0 > 0$.

Second, suppose that $\beta_0 = 0$ and all sales take place in the indirect channel. Then, according to the tie-breaking rule, the regulator will select the lowest price p^I such that the seller will not find it optimal to deviate. Suppose that $p^I < v_m + b$. The seller can always sell to the flexible consumers in the indirect channel by setting price equal to $v_m + b$ and earn $(1 - \lambda)(1 - \alpha)(v_m + b - c)$. The incentive compatibility constraint of the seller implies that $(\alpha/2 + 1 - \alpha)(p^I - c) \geq (1 - \alpha)(v_m + b - c)$ or equivalently $p^I \geq \frac{\alpha/2}{\alpha/2 + 1 - \alpha}c + \frac{1 - \alpha}{\alpha/2 + 1 - \alpha}(v_m + b)$. At this price the intermediary will find sufficiently low λ such that the seller does not deviate to sell to the flexible consumers directly. To see this note that for $\lambda \leq b/(v_m + b - c)$ the seller's profit from the first best outcome is weakly higher than the profit from the deviation

$$(1 - \lambda)(\alpha/2 + (1 - \alpha))(p^I - c) = (1 - \lambda)(1 - \alpha)(v_m + b - c) \geq (1 - \alpha)(v_m - c).$$

□

Proof of Proposition 4. With fully informative recommendations, picky consumers value the good at $v_h + b$ if they buy through the indirect channel. With this value, we derive the incentive compatibility constraints for seller and intermediary, which differ from those under laissez-faire. The intermediary will set λ to make the seller weakly prefer catering to flexible consumer via the direct and the indirect channel than using the indirect channel for picky consumers only:

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

The intermediary also has to respect the incentive compatibility constraint that it prefers to sell also to flexible consumers instead of selling only to picky consumers with a good match:

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

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

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

Clearly, if there exists $\lambda \in [0, 1]$ satisfying these two inequalities, then the naive policy will implement the highest λ that makes the seller indifferent. This implements the first-best outcome.

We see that, for $\alpha \approx 0$, there is a non-empty set of λ satisfying the two inequalities. The maximal λ in this set then implements the first best and the naive policy is optimal. If $\alpha \approx 1$, then the seller always prefers to sell directly to the flexible consumers because $(1 - \frac{\alpha}{2})(v_m + b - c) - \frac{\alpha}{2}(v_h + b - c) < 0$. 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 $\alpha_{\text{FI}} \in (0, 1)$ such that for all $\alpha \leq \alpha_{\text{FI}}$, the naive policy β_{UNI} is optimal and results in zero losses. Otherwise, if $\alpha > \alpha_{\text{FI}}$ the seller will serve the flexible consumers through the direct channel D

leading to welfare loss equal to $(1 - \alpha)b$. We note that $\alpha_{\text{FI}} < \bar{\alpha}$. This means that for $\alpha \in (\alpha_{\text{FI}}, \bar{\alpha})$ the laissez-faire is welfare-superior. As we have shown in Section 3.2, for $\alpha < \bar{\alpha}$, the welfare loss under laissez-faire is $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$. Thus, we have to show that, for $\alpha < \bar{\alpha}$, the inequality $(1 - \alpha)b > \frac{\alpha}{2} \beta (c - b - v_l)$ must be satisfied. For $\alpha < \bar{\alpha}$ we obtain

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

□

Proof of Proposition 5. First, we will show that the regulator prefers the cap $\bar{\beta} = 1$ over any cap $\bar{\beta} \in \left(\frac{v_h - v_m}{v_m - v_l}, 1\right)$. Suppose that the regulator has chosen a cap $\bar{\beta} \in \left(\frac{v_h - v_m}{v_m - v_l}, 1\right)$. To induce the seller to use the indirect channel, the best the intermediary can do is to minimize the profit of the seller when deviating to a price vector different from $(p^I, p^D) = (v_m + b, v_m)$. The only possibly profitable deviation by the seller is to sell directly to flexible consumers. The profit contribution of flexible consumers is at most $(1 - \alpha)(v_m - c)$. The seller's profit from picky consumers served through the indirect channel is minimized at $\beta = \bar{\beta}$ for $\beta \in [0, \bar{\beta}]$. Hence, the intermediary will set $\beta = \bar{\beta}$ for any $(p^I, p^D) \neq (v_m + b, v_m)$. For prices $(v_m + b, v_m)$, it will set the same recommendation level as under laissez-faire, $\beta = \frac{v_h - v_m}{v_m - v_l}$ and the welfare loss compared to the first best is equal to $\frac{\alpha}{2} \beta (c - b - v_l)$.

Regarding the choice of λ , the intermediary has to respect the seller's incentive compatibility constraint to not use the direct channel:

$$(1 - \lambda) \left[\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} + (1 - \alpha) \right] (v_m + b - c) \geq (1 - \alpha)(v_m - c) + (1 - \lambda) \frac{\alpha}{2} (1 + \bar{\beta}) \max \left\{ \frac{v_h + \bar{\beta} v_l}{1 + \bar{\beta}} + b - c, 0 \right\}.$$

If $\frac{v_h - v_m}{v_m - v_l} < \bar{\beta} < \frac{v_h + b - c}{c - b - v_l}$, then the seller's deviation profit is larger than under laissez-faire because the intermediary is constrained in its choice of punishment strategy. Then, to satisfy the seller's incentive compatibility constraint, the intermediary has to leave a larger fraction of total profit to the seller. As this additional profit left to the seller is increasing in α , for α slightly less than $\bar{\alpha}$, the intermediary will find it optimal to give up flexible consumers and only induce picky consumers to buy through the indirect channel (and reveal information fully) – that is, $\beta = 0$ and $\lambda = 1$.

The welfare loss under the recommendation cap compared to laissez-faire is $(1 - \alpha)b + \alpha/2\beta(c - b - v_l)$ for $\alpha < \bar{\alpha}$ in the vicinity of $\bar{\alpha}$ such that the intermediary optimally chooses $\beta = 0$. This expression is positive, as shown in the previous section. For other values of α the welfare loss of the recommendation cap coincides with the one under laissez-faire. To summarize, the regulator cannot improve on $\bar{\beta} = 1$ by setting $\bar{\beta} \in \left(\frac{v_h - v_m}{v_m - v_l}, 1\right)$.

Second, we consider a recommendation cap that is stricter than the recommendation level resulting under laissez-faire – that is, $0 < \bar{\beta} \leq \frac{v_h - v_m}{v_m - v_l}$. In this part of the proof, we will show that the best such policy is the minimal cap such that the corresponding incentive compatibility constraints of seller and intermediary are both satisfied with equality and we will characterize this policy.

If the intermediary induces the seller to serve flexible consumers in the indirect channel, then the total profits are maximal when the expected match value of picky consumers is as close as possible to $v_m + b$. Therefore, the intermediary sets the highest permitted recommendation level $\beta = \bar{\beta}$. To minimize the seller's deviation profits (which result from serving flexible consumers in the direct channel), the intermediary sets the recommendation level (for out-of-equilibrium prices of the seller) that minimizes the expected valuation of picky consumers – that is, $\beta = \bar{\beta}$. Furthermore, the intermediary sets λ as high as possible, respecting the seller's incentive compatibility constraint that is given by

$$(1 - \lambda) \left[\frac{\alpha}{2}(1 + \bar{\beta}) + (1 - \alpha) \right] (v_m + b - c) \geq (1 - \alpha)(v_m - c) + (1 - \lambda) \frac{\alpha}{2}(1 + \bar{\beta}) \left[\frac{v_h + \bar{\beta}v_l}{1 + \bar{\beta}} + b - c \right]. \quad (5)$$

The resulting intermediary's profit from inducing flexible consumers to buy in the indirect channel must be higher than what it would make by extracting all surplus from the picky consumers with good matches (by setting $\beta = 0$ and $\lambda = 1$) – that is,

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

We will show that if – with some recommendation cap $\bar{\beta} \leq \frac{v_h - v_m}{v_m - v_l}$ – the regulator cannot induce the outcome in which flexible consumers are served in the indirect channel, then it will not be able to do so with an even stricter cap. To show this, it is useful to rewrite the seller's incentive compatibility constraint (5) as follows:

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

Thus, as $\bar{\beta}$ is reduced, it becomes harder to satisfy the inequality. This, in turn, forces the intermediary to reduce λ . This implies that the minimal recommendation cap that satisfies the incentive constraints – both, of the seller and the intermediary – makes them binding.

We now turn to the characterization of the recommendation cap policy. We introduce the function W as the difference between the (maximal) total profits in the indirect channel with and without flexible consumers buying in the indirect channel, which depends on β ,

$$W(\beta) = \left[\frac{\alpha}{2}(1 + \beta) + 1 - \alpha \right] (v_m + b - c) - \frac{\alpha}{2}(1 + \beta)(p^I - c).$$

By solving both of the incentive compatibility constraints with respect to $1 - \lambda$ we find that the minimal $\bar{\beta}$ has to satisfy

$$1 - \lambda = \frac{(1 - \alpha)(v_m - c)}{W(\bar{\beta})} = 1 - \frac{\frac{\alpha}{2}(v_h + b - c)}{\left(\frac{\alpha}{2}(1 + \bar{\beta}) + 1 - \alpha \right) (v_m + b - c)} \quad (6)$$

if it belongs to $(0, \frac{v_h - v_m}{v_m - v_l}]$.

Third, we determine the overall optimal recommendation cap policy. There are three candidates: $\bar{\beta} = 0$; $\bar{\beta} = 1$; and, if it exists, $\bar{\beta}$ characterized in 6 to implement that all flexible consumers buy in the indirect channel. The associated three possible outcomes are: first best in which flexible consumers and picky consumers with a good match buy in the indirect channel; the laissez-faire outcome in which also some picky consumers with a bad match buy in the indirect channel; an intermediate outcome in which fewer picky consumers with a bad match buy in the indirect channel – this is, the intermediary chooses a more informative recommendation policy. A fourth possible outcome is that flexible consumers buy in the direct channel and only picky consumers with a good match buy in the direct channel. We start by comparing the outcomes in which no consumers uses the direct channel.

It is clear that $\bar{\beta} = 0$ is optimal for $\alpha \leq \alpha_{\text{FI}}$ as the regulator can achieve the first best by picking $\bar{\beta} = 0$. For α slightly above α_{FI} the regulator can no longer achieve the first best, but it can ensure that only a fraction of picky consumer with a bad match receive a recommendation. This policy, whenever feasible, is preferred to the laissez-faire (with the outcome that the direct channel is not used) because the only difference in the allocation is that picky consumers with a bad match receive a recommendation less often. The regulator will then pick $\bar{\beta}$ characterized in 6. Whenever this solution is feasible, the weaker cap $\bar{\beta} = \frac{v_h - v_m}{v_m - v_l}$ would make the intermediary choose β and λ such that all flexible consumers buy in the indirect channel. If this is impossible, the regulator will set $\bar{\beta} = 1$.

To be feasible, the seller must find it optimal to serve flexible consumer through the indirect channel. This is the case if and only if $(1-\alpha)(1-\lambda)(v_m+b-c) \geq (1-\alpha)(v_m-c)$ or, equivalently, $(1-\lambda)(v_m+b-c) \geq (v_m-c)$. Solving for the maximal λ and plugging it into the intermediary's incentive compatibility constraint we obtain that the condition on α is

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

Note that the left-hand side of this inequality is higher than the right-hand side for $\alpha \approx 0$. Conversely, the right-hand side is higher for $\alpha \approx 1$. To see this we use the fact that

$$\frac{v_h + b - c}{v_h - v_l} - \frac{b}{v_m - v_l} = \left(\frac{v_h + b - c}{v_h - v_l} - \frac{v_m + b - c}{v_m - v_l} \right) + \frac{v_m - c}{v_m - v_l} > 0,$$

such that $\frac{x+b-c}{x-v_l}$ increases in x as $v_l + b < c$. Therefore, there exists a cutoff level of the fraction of picky consumers

$$\alpha_{\text{CAP}} = \frac{b}{b + \frac{1}{2} \left(v_h + b - c - \frac{v_h - v_l}{v_m - v_l} b \right)}$$

such that for all $\alpha < \alpha_{\text{CAP}}$ the solution is feasible. Otherwise, the regulator does not have a (strict) incentive to impose a recommendation cap. A cap less than 1 is strictly worse if the laissez-faire outcome features that no trade takes place in the direct channel; it is immaterial if flexible consumers use the direct channel under laissez-faire, which is the fourth possible outcome mentioned above.

To complete the picture and to rule out that the above characterization is upset by the regulator's preference for the fourth possible outcome, we return to the characterization of the laissez-faire outcome. As we have shown in Section 3.2, under laissez-faire, flexible consumers buy in the direct channel if and only if $\alpha > \bar{\alpha}$ and this is what the regulator would do as well. It is straightforward to see that $\bar{\alpha} > \alpha_{\text{CAP}}$ as the recommendation cap $\bar{\beta} = \frac{v_h - v_m}{v_m - v_l}$ makes it more costly for the intermediary to induce sales to flexible consumer in the indirect channel. This implies that the fourth possible outcome prevails only for values $\alpha > \bar{\alpha}$. \square

B Laissez-faire under alternative assumptions on the model parameters

We continue to work under the assumption that $v_m \geq \frac{v_h + v_l}{2} + b$ – that is, the convenience benefit b is sufficiently small such that, absent any information about the match quality, a flexible consumer obtains a higher expected evaluation in the direct channel than does a picky consumer in the indirect channel. This assumption implies that $v_l + b < v_l + v_m - \frac{v_h + v_l}{2} < \frac{v_h + v_l}{2}$. In the main text, we considered marginal costs c that satisfy $(v_h + v_l)/2 + b < c < v_m$. In this appendix we consider alternative values of c in the following ranges: first, $(v_h + v_l)/2 < c < (v_h + v_l)/2 + b$; second, $v_l + b < c < (v_h + v_l)/2$; and, third, $v_l < c < v_l + b$.

B.1 Marginal costs $c \in ((v_h + v_l)/2, (v_h + v_l)/2 + b)$

Since $(v_h + v_l)/2 > v_l + b$, picky consumers do not buy in the indirect channel if they learn that they have a bad match, but they would buy if no information is revealed and they base their decision on their prior.

We begin by characterizing the solution of the vertically integrated firm. Since $v_l + b < c$, we have that if the firm induces the outcome with inefficient bypass, then it will set the fully informative recommendation policy and set prices $p^I = v_h + b$ and $p^D = v_m$. The resulting profit is equal to $\frac{\alpha}{2}(v_h + b - c) + (1 - \alpha)(v_m - c)$. If the firm induces the outcome with inflated recommendation with $\beta \geq \beta^*$, defined in 3.2, then the profit is equal to

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

It can be shown that since $v_l + b < c$, the profit from inflated recommendation decreases in β and is maximal at $\beta = \beta^*$. The resulting maximal profit in the outcome with inflated recommendations is $\left(\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} + (1 - \alpha)\right)(v_m + b - c)$. Since these profits for both outcomes are exactly the same as in Section 3.2, we obtain that the critical level of α for the vertically integrated firm is equal to $\bar{\alpha}$, defined in Section 3.2. We note that the firm could also serve all consumers in the indirect channel at $p^I = (v_h + v_l)/2 + b$ yielding profit $(v_h + v_l)/2 + b - c$. However, this outcome gives a lower profit than the inefficient bypass outcome because

$\frac{\alpha}{2}(v_h+b-c)+(1-\alpha)(v_m-c) > (v_h+v_l)/2+b-c$ simplifies to $(\alpha/2)(c-v_l-b)+(1-\alpha)(v_m-(v_h+v_l)/2-b) > 0$ which holds under our assumption on c .

We now turn to the characterization of the optimal strategy of the intermediary that sets λ and commits to some recommendation policy $\beta(p^i, p^d)$. Suppose that the intermediary induces the outcome with inflated recommendations and sets $\beta = \beta^*$ for prices $p^I = v_m + b$ and $p^D \geq v_m$. To minimize the seller's deviation profit (from diverting flexible consumers to the direct channel), the intermediary sets $\beta = 1$ for all other prices. Consider the seller's pricing problem if it decides to serve flexible consumers in the direct channel. It can always ensure the profit from serving flexible consumers in the direct channel $(1-\alpha)(v_m-c)$ by setting $p^D = v_m$ and $p^I = +\infty$. We will show that the seller cannot make higher profits by selling to the picky consumers in the indirect channel. If it sets $p^I \leq \frac{v_h+v_l}{2} + b$, then in order to keep flexible consumers in the direct channel it has to set p^D that satisfies

$$v_m - p^D \geq v_m + b - p^I \geq v_m + b - \left(\frac{v_h + v_l}{2} + b \right),$$

or equivalently $p^D \leq \frac{v_h+v_l}{2}$, which results in non-positive profits from the flexible consumers and a decrease in profits from the picky consumers. Therefore, the incentive constraint of the seller is given by

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

If the intermediary induces the outcome with inefficient bypass, then it sets $\lambda = 1$ and $\beta = 0$ for all prices. In this case the profit of the intermediary is given by $\frac{\alpha}{2}(v_h + b - c)$.

Since the intermediary can achieve the outcome of the vertically integrated firm and has to always leave profits of $(1-\alpha)(v_m - c)$ to the seller, the optimal strategy of the intermediary is characterized by the critical level $\bar{\alpha}$ and therefore, the equilibrium coincides with the one in Proposition 2. Also, the first-best allocation remains unchanged compared to the one in the main text. We summarize the analysis by the following proposition.

Proposition 7. *Suppose that $v_l + b < (v_h + v_l)/2 < c < (v_h + v_l)/2 + b$. Then, the equilibrium when the intermediary commits to its recommendation policy coincides with the equilibrium characterized by Proposition 2.*

B.2 Marginal costs $c \in (v_l + b, (v_h + v_l)/2)$

The vertically integrated firm has three potentially optimal strategies. The first is to sell to picky consumers with a good match in the indirect channel at $p^I = v_h + b$ and to flexible consumers in the direct channel at $p^D = v_m$. This yields a profit of $(\alpha/2)(v_h + b - c) + (1-\alpha)(v_m - c)$. The second is to sell only in the indirect channel at $p^I = v_m + b$ and to set $\beta = \beta^*$. This yields a profit of $(1-\alpha + (\alpha/2)(1+\beta^*))(v_m + b - c)$. The third is to serve all consumers in the indirect channel at $p^I = (v_h + v_l)/2 + b$ yielding profit $(v_h + v_l)/2 + b - c$.

If $c > v_l + b$, as has been assumed here, the profit with the third strategy cannot be larger than with the first. Thus, the vertically integrated outcome does not change compared to the main text. Also, the first best is the same: sell to all flexible consumers and to all picky consumers with a good match in the indirect channel.

Moving to disintegration, compared to the previous parameter constellations, the novel feature is that, absent the intermediary, the seller may want to sell to all consumers in the direct channel. The best way of doing so is to set $p^D = (v_h + v_l)/2$. The seller then makes profit $(v_h + v_l)/2 - c$. Alternatively, it may sell to flexible consumers only at $p^D = v_m$ in which case it makes $(1 - \alpha)(v_m - c)$. If α is sufficiently small – that is, $\alpha < \hat{\alpha} = \frac{v_m - (v_h + v_l)/2}{v_m - c}$ the latter dominates the former, the intermediary has to make sure that it offers a contract to the seller that allows the seller to make at least $(1 - \alpha)(v_m - c)$. Here, the analysis in the main text applies. In the opposite case $\alpha > \hat{\alpha}$, the incentive constraint in the inflated recommendation outcome reads $(1 - \lambda)(1 - \alpha + (\alpha/2)(1 + \beta^*))(v_m + b - c) \geq (v_h + v_l)/2 - c$. In this range of α the intermediary has to do with a smaller share of industry profits compared to the case in which the seller can only cater to informed consumers in the direct channel. This applies to the inflated recommendation outcome and the inefficient bypass outcome alike. Regarding the latter, the intermediary selling to picky consumers with a good match in the indirect channel has to compensate the seller for not making $\alpha[(v_h + v_l)/2 - c]$ and thus λ must be less than one under inefficient bypass.

To characterize the equilibrium, we have to compare $\hat{\alpha}$ and $\bar{\alpha}$ defined in Section 3.2 and given by

$$\bar{\alpha} = \frac{b}{b + \frac{1}{2}(v_h + b - c - (1 + \beta^*)(v_m + b - c))}.$$

It is easy to see that $\bar{\alpha}$ increases in b on $(0, c - v_l)$. If b is sufficiently close to 0, then $\bar{\alpha} \approx 0 < \hat{\alpha}$. If b is sufficiently close to $c - v_l$, then $\bar{\alpha} \approx 1 > \hat{\alpha}$. Thus, there exists a unique $\hat{b} \in (0, c - v_l)$ that solves

$$\frac{\hat{b}}{\hat{b} + \frac{1}{2} \left(v_h + \hat{b} - c - \frac{v_h - v_l}{v_m - v_l} (v_m + \hat{b} - c) \right)} = \frac{v_m - (v_h + v_l)/2}{v_m - c}$$

such that we have $\hat{\alpha} \geq \bar{\alpha}$ for all $b \leq \hat{b}$ and $\hat{\alpha} < \bar{\alpha}$ for $b > \hat{b}$.

If $b \in (\hat{b}, c - v_l)$, then $\hat{\alpha} < \bar{\alpha}$. Recommendations are inflated with λ as given in Proposition 1 for $\alpha \leq \hat{\alpha}$, while they are inflated with a lower λ than the one in Proposition 1 for $\alpha \in (\hat{\alpha}, \bar{\alpha}]$. More specifically, λ solves $(1 - \lambda)((1 - \alpha) + (\alpha/2)(1 + \beta^*))(v_m + b - c) = (v_h + v_l)/2 - c$. Note that since $\alpha \in (\hat{\alpha}, \bar{\alpha}]$, we have that

$$\begin{aligned} \left(1 - \alpha + \frac{\alpha}{2}(1 + \beta^*)\right) (v_m + b - c) &> \alpha \frac{v_h + b - c}{2} + (1 - \alpha)(v_m - c) \\ &> \alpha \frac{v_h - (c - v_l) - c}{2} + (1 - \alpha) \left(\frac{v_h + v_l}{2} - c \right) \\ &= \frac{v_h + v_l}{2} - c, \end{aligned}$$

which implies that $\lambda \in (0, 1)$.

If $\alpha > \bar{\alpha}$, then the intermediary induces the outcome with inefficient bypass such that flexible consumers buy in the direct instead of the indirect channel with $\lambda < 1$. In particular, if $\alpha > \bar{\alpha}$, then λ solves $(1 - \lambda)(\alpha/2)(v_h + b - c) + (1 - \alpha)(v_m - c) = (v_h + v_l)/2 - c$ and is equal to

$$\lambda = 1 - \frac{(v_h + v_l)/2 - c - (1 - \alpha)(v_m - c)}{(\alpha/2)(v_h + b - c)}.$$

Since

$$(v_h + v_l)/2 - c - (1 - \alpha)(v_m - c) = (\alpha/2)(v_h - (c - v_l) - c) - (1 - \alpha)(v_m - (v_h + v_l)/2) < (\alpha/2)(v_h + b - c),$$

we obtain that $\lambda \in (0, 1)$.

In the opposite case when $0 < b \leq \hat{b}$, then $\bar{\alpha} \leq \hat{\alpha}$ and there is inflated recommendation with λ as given in Proposition 2 for $\alpha \leq \hat{\alpha}$ and inefficient bypass such that flexible consumers buy in the direct channel for $\alpha > \bar{\alpha}$. The intermediary asks for $\lambda = 1$ if $\alpha \in (\bar{\alpha}, \hat{\alpha})$ and a profit share less than 1 for larger α . If $\alpha > \bar{\alpha}$, then $\lambda = 1 - [(v_h + v_l)/2 - c - (1 - \alpha)(v_m - c)] / [(\alpha/2)(v_h + b - c)]$.

We summarize the analysis by the following proposition.

Proposition 8. *Suppose that $v_l + b < c < (v_h + v_l)/2$. Then, the equilibrium when the intermediary commits to its recommendation policy is characterized by the outcomes with inflated recommendations and inefficient bypass as follows:*

- For $b \in [0, \hat{b})$, the outcome features inflated recommendations with λ as given in Proposition 2 for $\alpha \leq \hat{\alpha}$ and inefficient bypass with $\lambda = 1 - [(v_h + v_l)/2 - c - (1 - \alpha)(v_m - c)] / [(\alpha/2)(v_h + b - c)]$ for $\alpha > \hat{\alpha}$.
- For $b \in [\hat{b}, c - v_l)$, the outcome features inflated recommendations with λ as given in Proposition 2 for $\alpha \leq \hat{\alpha}$, inflated recommendation with $\lambda = 1 - [(v_h + v_l)/2 - c] / [(1 - \alpha) + (\alpha/2)(1 + \beta^*)(v_m + b - c)]$ for $\alpha \in (\hat{\alpha}, \bar{\alpha}]$, and inefficient bypass with $\lambda = 1 - [(v_h + v_l)/2 - c - (1 - \alpha)(v_m - c)] / [(\alpha/2)(v_h + b - c)]$ for $\alpha > \bar{\alpha}$.

B.3 Marginal costs $c \in (v_l, v_l + b)$

We begin by characterizing the optimal solution of the vertically integrated firm. We will show that the firm finds it optimal to sell only in the indirect channel.

First, we show that $p^I \leq v_m + b$ in the equilibrium. Suppose, by contradiction, that $p^I > v_m + b$. This implies that flexible consumers do not buy in the indirect channel. Any price $p^I > v_h + b$ cannot be profit-maximizing since no consumers buy in the indirect channel and the firm can strictly increase its profit by diverting all sales from the direct channel to the indirect channel. If $p^I \in (v_m + b, v_h + b]$, then the corresponding recommendation policy β must make picky consumers who received recommendations indifferent – that is, $p^I = \frac{v_h + \beta v_l}{1 + \beta} + b$. The firm's profit from the indirect channel is $\frac{\alpha}{2}(1 + \beta) \left(\frac{v_h + \beta v_l}{1 + \beta} + b - c \right)$.

Since this profit is increasing in β , the firm always has an incentive to lower its price and increase β . This implies that $p^I > v_m + b$ cannot be profit-maximizing.

Second, $p^I \leq v_m + b$ implies that it is profit-maximizing to sell only in the indirect channel. Suppose, by contradiction, that picky consumers with recommendations buy in the indirect channel, while flexible consumers buy in the direct channel at price $p^D < p^I - b$. Then it is optimal to shut down the direct channel and divert flexible consumers to the indirect channel. If the picky consumers are served in the direct channel, then it is also optimal to shut down the direct channel by setting $p^D > p^I - b$, adjust β and sell only in the indirect channel.

Consider the integrated firm setting recommendation policy $\beta = \beta^*$ and selling only through the indirect channel – that is, it sets $p^D \geq p^I - b$ and $p^I = \frac{v_h + \beta v_l}{1 + \beta} + b$. Any $\beta < \beta^*$ is strictly dominated by $\beta = \beta^*$ since flexible consumers and more picky consumers would buy at price $v_m + b$. Consider the case $\beta \geq \beta^*$. The profit of the integrated firm is given by

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

It is easy to see that the second derivative of the profit function is positive for all $\beta \in [\beta^*, 1]$ and therefore the profit function is convex. The maximum is reached either at $\beta = \beta^*$ or at $\beta = 1$. The profit of the firm inducing the outcome with “inflated” recommendations – that is, $\beta = \beta^*$, is equal to $\left(\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha)\right) (v_m + b - c)$. The profit of the firm setting $\beta = 1$ and inducing the first best outcome is $\frac{v_h + v_l}{2} + b - c$. We continue to use the term “inflated recommendations” because from a consumer surplus perspective recommendations are also inflated in this case as picky consumers with a bad match would be better off not buying. However, from a total surplus perspective, recommendations are deflated.

The vertically integrated firm will find it optimal to induce the first best outcome if and only if $\alpha > \alpha_4$, where

$$\alpha_4 = \frac{v_m - v_l}{v_m + b - c}. \quad (7)$$

For any $\alpha \geq \alpha_4$, we have that

$$\begin{aligned} & \left(\frac{\alpha}{2}(1 + \beta^*) + 1 - \alpha\right) (v_m + b - c) = \left(1 - \frac{\alpha}{2}(1 - \beta^*)\right) (v_m + b - c) \\ & \leq \left(1 - \frac{\alpha_4}{2}(1 - \beta^*)\right) (v_m + b - c) = v_m + b - c - \frac{1}{2}(2v_m - (v_h + v_l)) \\ & = \frac{v_h + v_l}{2} + b - c. \end{aligned}$$

Otherwise, if $\alpha < \alpha_4$, then the firm finds it optimal to induce the outcome with inflated recommendations ($\beta^* = \frac{v_h - v_m}{v_m - v_l}$) which leads to welfare loss equal to $\frac{\alpha}{2}(1 - \beta^*)(v_l + b - c)$.

We move to the case of disintegration and show that in the game in which the intermediary that commits to a recommendation policy and a profit-sharing rule λ implements the vertically integrated solution. The seller can always guarantee the profits from serving flexible consumers or all consumers in the direct channel – that is, $\max\{(1 - \alpha)(v_m - c), (v_h + v_l)/2 - c\}$.

First, suppose that the intermediary induces the first-best outcome by setting $\beta = 1$ for all prices. The seller sets $p^I = \frac{v_h + v_l}{2} + b$ and $p^D \geq \frac{v_h + v_l}{2}$ in the equilibrium. The seller cannot deviate and serve some consumers directly and other consumers in the indirect channel, since any $p^D < \frac{v_h + v_l}{2}$ would attract all consumers to the direct channel. This implies that the seller can either serve all consumers in the direct channel and earn $\frac{v_h + v_l}{2} - c$ or serve only flexible consumers directly and earn $(1 - \alpha)(v_m - c)$.

Thus, the incentive constraint of the seller is given by

$$(1 - \lambda) \left(\frac{v_h + v_l}{2} + b - c \right) \geq \max \left\{ (1 - \alpha)(v_m - c), \frac{v_h + v_l}{2} - c \right\}$$

The intermediary will induce the first best outcome if and only if $\alpha > \hat{\alpha}$ or $\alpha \leq \hat{\alpha}$ and $\frac{v_h + v_l}{2} + b - c - (1 - \alpha)(v_m - c) > 0$. Thus, the intermediary can induce the first best for all

$$\alpha \geq \alpha_5 = 1 - \frac{(v_h + v_l)/2 + b - c}{v_m - c}.$$

Since the profit under inflated recommendations $\left(\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha)\right)(v_m + b - c)$ increases in α and

$$\begin{aligned} & \left(\frac{\alpha_4}{2}(1 + \beta^*) + (1 - \alpha_4)\right)(v_m + b - c) \\ &= \frac{v_h + v_l}{2} + b - c \\ &= (1 - \alpha_5)(v_m - c) \\ &< \left(\frac{\alpha_5}{2}(1 + \beta^*) + (1 - \alpha_5)\right)(v_m + b - c) \end{aligned}$$

we have that $\alpha_4 > \alpha_5$. Therefore, for all $\alpha \geq \alpha_4$ the intermediary can induce the first best outcome and obtain the profits of the vertically integrated firm net of the minimal profits $\max\{(1 - \alpha)(v_m - c), (v_h + v_l)/2 - c\}$ that have to be left to the seller.

Second, suppose that the intermediary induces the outcome with inflated recommendations by setting $\beta = \beta^*$ for $(p^I, p^D) = (v_m + b, v_m)$ and $\beta = 1$ for all other prices. If the seller decides to serve flexible consumers in the direct channel, then the seller can not make positive profits from the picky consumers in the indirect channel – that is, the picky consumers will either switch to the direct channel as well ($p^D = (v_h + v_l)/2$) or will not buy at all ($p^D = v_m, p^I > v_m + b$). This implies that the deviating seller cannot earn more than $\max\{(1 - \alpha)(v_m - c), (v_h + v_l)/2 - c\}$.

The incentive constraint of the seller is given by

$$(1 - \lambda) \left(\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha) \right) (v_m + b - c) \geq \max \left\{ (1 - \alpha)(v_m - c), \frac{v_h + v_l}{2} - c \right\}$$

We will show that the intermediary can induce the outcome with inflated recommendations for all α . If $\alpha \geq \hat{\alpha}$, then it is straightforward to see that there exists $\lambda \in (0, 1)$ that induces the outcome with inflated recommendations. Otherwise, if $\alpha < \hat{\alpha}$, then the IC constraint is given by

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

We will show that there exists $\lambda \in (0, 1)$ for which the IC constraint is satisfied. Note that

$$\begin{aligned} & \left(\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha) \right) (v_m + b - c) - \left(\frac{v_h + v_l}{2} - c \right) \\ &= (1 - \alpha) \left(v_m + b - \frac{v_h + v_l}{2} \right) + \alpha \left(\frac{1}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) - \frac{v_h + v_l}{2} + c \right) \end{aligned}$$

Since $v_m > (v_h + v_l)/2$, the first term of this expression is positive. The sign of the second expression is also positive since

$$\begin{aligned} & \alpha \left(\frac{1}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) - \frac{v_h + v_l}{2} + c \right) \\ & > \alpha \left(\frac{1}{2} \frac{v_h - v_l}{v_m - v_l} (v_m - c) - \frac{v_h + v_l}{2} + c \right) \\ &= \alpha (v_m - c) \left(\frac{\frac{v_h + v_l}{2} - v_l}{v_m - v_l} - \frac{\frac{v_h + v_l}{2} - c}{v_m - c} \right) > 0, \end{aligned}$$

where we used the fact that $c > v_l$ and $\frac{(v_h + v_l)/2 - x}{v_m - x}$ decreases in x .

Therefore, for all α the intermediary can induce the outcome with inflated recommendations. This implies that for all $\alpha < \alpha_4$ the intermediary can reach the profit of the vertically integrated firm and leave the profits of $\max\{(1 - \alpha)(v_m - c), (v_h + v_l)/2 - c\}$ to the seller.

To fully characterize the equilibrium, it remains to compare α_4 and $\hat{\alpha}$. We will show that $\alpha_4 > \hat{\alpha}$ for all $b \in (0, v_m - (v_h + v_l)/2]$. When b is sufficiently close to 0, then

$$\alpha_4 - \hat{\alpha} \approx \frac{v_m - v_l}{v_m - c} - \frac{v_m - \frac{v_h + v_l}{2}}{v_m - c} = \frac{\frac{v_h + v_l}{2} - v_l}{v_m - c} > 0.$$

When b is sufficiently close to $v_m - \frac{v_h + v_l}{2}$, then

$$\alpha_4 - \hat{\alpha} = \frac{v_m - v_l}{v_m - c + v_m - \frac{v_h + v_l}{2}} - \frac{v_m - \frac{v_h + v_l}{2}}{v_m - c}.$$

The sign of this expression is determined by the sign of

$$\begin{aligned} & \left(v_m - \frac{v_h + v_l}{2} + \frac{v_h + v_l}{2} - v_l \right) (v_m - c) - \left(v_m - c + v_m - \frac{v_h + v_l}{2} \right) \left(v_m - \frac{v_h + v_l}{2} \right) \\ &= \left(\frac{v_h + v_l}{2} - v_l \right) (v_m - c) - \left(v_m - \frac{v_h + v_l}{2} \right)^2 \\ &> \left(v_m - \frac{v_h + v_l}{2} \right) \left(\frac{v_h + v_l}{2} - v_l - v_m + \frac{v_h + v_l}{2} \right) = \left(v_m - \frac{v_h + v_l}{2} \right) (v_h - v_m) > 0, \end{aligned}$$

where we used the fact that $(v_h + v_l)/2 > c$. Since α_4 as a function of b is monotone for all $b \in (0, v_m - (v_h + v_l)/2]$ and is higher than $\hat{\alpha}$ at the end points of this interval, we obtain that $\alpha_4 > \hat{\alpha}$ for all $b \in (0, v_m - (v_h + v_l)/2]$.

This allows us to characterize the equilibrium for all α . If $\alpha > \alpha_4$, then the intermediary will induce the first best outcome by setting $\beta = 1$ for all prices. The intermediary sets λ that solves $(1 - \lambda)((v_h +$

$v_l)/2 + b - c) = (v_h + v_l)/2 - c$. If $\alpha < \alpha_4$, then the intermediary induces the outcome with inflated recommendations by setting $\beta = \beta^*$ for $(p^I, p^D) = (v_m + b, v_m)$ and $\beta = 1$ for all other prices. The intermediary sets λ that makes the seller indifferent between the profits from the outcome with inflated recommendations $(1 - \lambda) \left(\frac{\alpha}{2}(1 + \beta^*) + (1 - \alpha) \right) (v_m + b - c)$ and profits from serving all consumers in the direct channel $(v_h + v_l)/2 - c$ for any $\alpha \in (\hat{\alpha}, \alpha_4]$ or profits from serving flexible consumers in the direct channel $(1 - \alpha)(v_m - c)$ for any $\alpha < \hat{\alpha}$. We summarize the analysis in the following proposition.

Proposition 9. *Suppose that $v_l < c < v_l + b$. When the intermediary commits to its recommendation policy, the equilibrium is characterized by*

- for $\alpha > \alpha_4$, the intermediary sets

$$\lambda^* = \frac{b}{\frac{v_h + v_l}{2} + b - c},$$

$\beta = 1$ for all (p^I, p^D) . Equilibrium prices are given by $(p^I, p^D) = ((v_h + v_l)/2 + b, (v_h + v_l)/2)$. All consumers visit and buy through the intermediary. The first best outcome is implemented.

- for $\alpha \leq \alpha_4$, the intermediary sets

$$\lambda^* = \begin{cases} 1 - \frac{(1-\alpha)(v_m-c)}{\left[\frac{\alpha}{2}\frac{v_h-v_l}{v_m-v_l} + (1-\alpha)\right](v_m+b-c)}, & \text{if } \alpha \leq \hat{\alpha} \\ 1 - \frac{(v_h+v_l)/2-c}{\left[\frac{\alpha}{2}\frac{v_h-v_l}{v_m-v_l} + (1-\alpha)\right](v_m+b-c)}, & \text{if } \alpha \in (\hat{\alpha}, \alpha_4] \end{cases},$$

$\beta(p^I = v_m + b, p^D \geq v_m) = \frac{v_h - v_m}{v_m - v_l}$ and $\beta = 1$ otherwise. Equilibrium prices are $(p^I, p^D) = (v_m + b, v_m)$. All consumers go to the indirect channel. The fraction $\frac{1+\beta}{2} = \frac{1}{2} \frac{v_h - v_l}{v_m - v_l}$ of picky consumers buy. Welfare losses are equal to $\alpha \frac{v_m - (v_h + v_l)/2}{v_m - v_l} (c - b - v_l)$.

C Regulation of a uniform recommendation level and comparison to the regulation with a recommendation cap

C.1 Optimal uniform recommendation policy

In this Appendix, the regulator chooses the optimal uniform recommendation policy β_{UNI} . The regulator has to anticipate that such a uniform policy applies regardless of whether or not flexible consumers buy through the indirect channel. Therefore, the policy affects profits on and off the equilibrium path.

As mentioned in the main text, it may be difficult to think that the intermediary may be fined if it provides better recommendation than what is imposed by the regulator. However, a different way to think about regulating the recommender system is to also regulate data gathering and data storage activities of the intermediary.²⁸ If the intermediary's data gathering activities are restricted by the regulator and,

²⁸Some privacy advocates have complained about the ability of sellers to target narrow audiences. This may be remedied by restricting the data gathering and data storage activities of the intermediary.

at the same time, the regulator obliges the intermediary to use all those data, this may be interpreted as implementing a uniform recommendation level regulation. To see this, suppose that without any data gathering the intermediary does not have any information about match quality of a picky consumer and that with additional data gathering the intermediary increases the fraction of consumers (among picky consumers) that it is able of identify as constituting a bad match. With such information, the intermediary can choose any β above a critical level that is determined by its previous data gathering efforts. Thus, using all available information leads to a uniform recommendation level.

As we have shown in the previous section, if $\alpha \leq \alpha_{\text{FI}}$, then the regulator can achieve the first best by imposing that recommendations are fully informative, i.e. $\beta_{\text{UNI}} = 0$. For $\alpha > \alpha_{\text{FI}}$ the naive policy cannot induce the first best and will always result in the outcome in which all flexible consumers are served in the direct channel, resulting in welfare losses equal to $(1 - \alpha)b$. Can the regulator improve by imposing a recommendation level $\beta_{\text{UNI}} \neq 0$? The answer will be ‘yes’ for α not too large – we will characterize the optimal policy β_{UNI} and how it relates to α .

When the naive policy $\beta_{\text{UNI}} = 0$ does not implement the first best, the regulator may resort to imposing a uniform recommendation policy $\beta_{\text{UNI}} > 0$ – that is, the regulator prescribes a certain level of inflated recommendations – or it may resign itself such that only picky consumers are served in the direct channel (implying $\beta_{\text{UNI}} = 0$). It may want to do the former so as to keep all consumers who buy in the indirect channel (as this generates extra benefit b compared to the direct channel). The associated outcome is that all flexible consumers, all picky consumers with good matches, and fraction β_{UNI} of picky consumers with bad matches are served in the indirect channel. Yet another alternative for the regulator is to give up on flexible consumers buying in the indirect channel. Then only picky consumers with good matches buy through the intermediary. Below, we characterize parameter values for any of the three outcomes to prevail as the result of the optimal uniform recommendation policy.

Our first observation is that the optimal uniform recommendation policy never features a higher recommendation level than the outcome under laissez-faire, as we show with the next lemma.

Lemma 3. *The socially optimal uniform recommendation policy features weakly fewer recommendation than the solution under laissez-faire – that is, $\beta_{\text{UNI}} > \frac{v_h - v_m}{v_m - v_l}$ is suboptimal for the regulator.*

Proof. The maximal possible retail price in the indirect channel is equal to $p^I = \frac{v_h + \beta_{\text{UNI}} v_l}{1 + \beta_{\text{UNI}}} + b$, which is less or equal to $v_m + b$. We distinguish between two cases.

First, suppose that $p^I \leq v_m$. Then the seller will always find it optimal to sell to the flexible consumers directly by setting a lower price $p^D = p^I - \varepsilon$. Consequently, the resulting welfare loss is going to be strictly higher than $(1 - \alpha)b$. Thus, this strategy of the regulator is dominated by the naive policy $\beta_{\text{UNI}} = 0$, which, as we have shown in the previous section, gives $(1 - \alpha)b$ for all α .

Second, suppose that $v_m < p^I \leq v_m + b$. If the seller sold to flexible consumers directly, its profit-

maximizing price on the direct channel would be $p^D = v_m$. The associated profit is dominated if the seller finds it optimal to sell to the flexible consumers through the indirect channel. This holds if and only if

$$(1 - \lambda) \left[\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha \right] (p^I - c) \geq (1 - \alpha)(v_m - c) + (1 - \lambda) \frac{\alpha}{2}(1 + \beta_{\text{UNI}})(p^I - c),$$

which implies that $(1 - \lambda)(1 - \alpha)(p^I - c) \geq (1 - \alpha)(v_m - c)$ or, equivalently,

$$\lambda \leq 1 - \frac{v_m - c}{p^I - c}. \quad (8)$$

Note that the intermediary can always ensure profits of $\frac{\alpha}{2}(1 + \beta_{\text{UNI}})(p^I - c)$ by setting $\lambda = 1$. Therefore, to induce the intermediary to set a low enough λ such that sales take place only in the indirect channel, the regulator has to ensure that the resulting profit of the intermediary is higher than what the intermediary would make by setting $\lambda = 1$:

$$\lambda \left[\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha \right] (p^I - c) \geq \frac{\alpha}{2}(1 + \beta_{\text{UNI}})(p^I - c),$$

which is equivalent to

$$\lambda \geq 1 - \frac{1 - \alpha}{\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha} \quad (9)$$

Suppose that, for given β_{UNI} , the incentives constraints of the intermediary and the seller, inequalities (8) and (9) are satisfied – that is, there exist λ such that

$$1 - \frac{1 - \alpha}{\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha} \leq \lambda \leq 1 - \frac{v_m - c}{p^I - c}.$$

Note that the interval for admissible λ becomes wider as β_{UNI} decreases.²⁹ A more precise recommendation policy – that is, lower β_{UNI} – increases the total profit more than the payoffs from deviations for the intermediary and the seller. Therefore, the regulator that induces sales only in the indirect channel would always prefer to set $\beta_{\text{UNI}} = \frac{v_h - v_m}{v_m - v_l}$, which would result in welfare loss of $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$. Thus, we have shown that any uniform policy β_{UNI} with $\beta_{\text{UNI}} > \frac{v_h - v_m}{v_m - v_l}$ is dominated by the more precise uniform recommendation policy $\beta_{\text{UNI}} = \frac{v_h - v_m}{v_m - v_l}$. \square

Thanks to the lemma, we can restrict attention to uniform recommendation policies that satisfy $\beta_{\text{UNI}} \leq \frac{v_h - v_m}{v_m - v_l}$. Three possible types of outcomes will prevail under the optimal uniform recommendation policy. The critical α that separates the planner's policy of managed recommendation inflation (i.e., a policy $\beta_{\text{UNI}} > 0$) from giving up on keeping flexible consumers on the direct channel is denoted by α_{UNI} and derived in the proof below. To characterize the intermediary's strategy, we again use the function W (see Proof of Proposition 5, Appendix A) as the difference between the (maximal) total profits in the indirect channel with and without flexible consumers buying in the indirect channel, which now depends on β_{UNI} :

$$W(\beta_{\text{UNI}}) = \left[\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha \right] (v_m + b - c) - \frac{\alpha}{2}(1 + \beta_{\text{UNI}}) \left(\frac{v_h + \beta_{\text{UNI}}v_l}{1 + \beta_{\text{UNI}}} + b - c \right). \quad (10)$$

²⁹This is due to the fact that $1 - \frac{1 - \alpha}{\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha}$ increases in β_{UNI} and $1 - \frac{v_m - c}{p^I - c}$ decreases in β_{UNI} (as p^I decreases in β_{UNI}).

Proposition 10. *Suppose that the regulator is restricted to set a uniform recommendation policy.*

- If $\alpha \leq \alpha_{FI}$, the regulator sets $\beta_{UNI} = 0$ and the first best is implemented.
- If $\alpha \in (\alpha_{FI}, \alpha_{UNI}]$, the regulator sets β_{UNI} such that the intermediary's and the seller's incentive compatibility constraints are binding. The intermediary sets $\lambda = 1 - (1 - \alpha)(v_m - c)/W(\beta_{UNI})$ and prices are $(p^I, p^D) = (v_m + b, v_m)$ along the equilibrium path. All flexible consumers and $1/2(1 + \beta_{UNI})$ of picky consumers buy in the indirect channel. The welfare loss is equal to $\frac{\alpha}{2}\beta_{UNI}(c - b - v_l)$.
- If $\alpha > \alpha_{UNI}$, the regulator sets $\beta_{UNI} = 0$. The intermediary sets $\lambda = 1$ and prices are $(p^I, p^D) = (v_h + b, v_m)$ along the equilibrium path. All flexible consumers buy in the direct channel and $1/2$ of picky consumers buy in the indirect channel. The welfare loss is equal to $(1 - \alpha)b$.

Proof. The proof proceeds in three steps: (i) we characterize conditions on the intermediary's profit share λ for which the seller and the intermediary prefer to serve flexible consumers in the indirect channel; (ii) we show that if the regulator sets $0 < \beta_{UNI} < \frac{v_h - v_m}{v_m - v_l}$ and induces the outcome in which flexible consumers buy in the indirect channel, then β_{UNI} is chosen such that inequality (11) (as derived below) is binding; (iii) we derive the conditions under which implementing such policy $\beta_{UNI} > 0$ is optimal – that is, the regulator prefers not to set $\beta_{UNI} = 0$ with flexible consumer buying in the direct channel.

(i) We characterize conditions on the intermediary's profit share λ for which the seller and the intermediary prefer to serve flexible consumers in the indirect channel. The inequality $\beta_{UNI} \leq \frac{v_h - v_m}{v_m - v_l}$ implies that $\frac{v_h + \beta_{UNI}v_l}{1 + \beta_{UNI}} + b \geq v_m + b$ – that is, the maximal possible retail price that at which flexible consumer would buy in the indirect channel is $v_m + b$. If the seller deviates and serves flexible consumers in the direct channel, it will charge picky consumers their expected match value $p^I = \frac{v_h + \beta_{UNI}v_l}{1 + \beta_{UNI}} + b$. The incentive compatibility constraint of the seller is given by

$$(1 - \lambda) \left[\frac{\alpha}{2}(1 + \beta_{UNI}) + 1 - \alpha \right] (v_m + b - c) \geq (1 - \alpha)(v_m - c) + (1 - \lambda) \frac{\alpha}{2} (1 + \beta_{UNI})(p^I - c).$$

A stricter uniform recommendation policy results in less total profits that can be collected if all flexible consumers are served in the indirect channel, whereas it makes the seller's deviation profits (resulting from serving flexible consumer in the direct channel) higher.

In order to induce the intermediary to set sufficiently low λ such that the seller finds it optimal to serve flexible consumers in the indirect channel, the regulator has to ensure that the intermediary does not find it profitable to set $\lambda = 1$ and serve only picky consumers in the indirect channel

$$\lambda \left[\frac{\alpha}{2}(1 + \beta_{UNI}) + 1 - \alpha \right] (v_m + b - c) \geq \frac{\alpha}{2}(1 + \beta_{UNI})(p^I - c).$$

Solving both of the incentive compatibility constraints with respect to $1 - \lambda$, we find that the flexible consumers will be served in the indirect channel if and only if there exists λ satisfying

$$\frac{(1 - \alpha)(v_m - c)}{W(\beta_{UNI})} \leq 1 - \lambda \leq \frac{W(\beta_{UNI})}{\left[\frac{\alpha}{2}(1 + \beta_{UNI}) + 1 - \alpha \right] (v_m + b - c)}. \quad (11)$$

(ii) We show that if the regulator sets $0 < \beta_{\text{UNI}} \leq \frac{v_h - v_m}{v_m - v_l}$ and induces the outcome in which flexible consumers buy in the indirect the channel, then β_{UNI} is chosen such that inequality (11) is binding.

To show this, suppose that for a given uniform recommendation policy $\beta_{\text{UNI}} \leq \frac{v_h - v_m}{v_m - v_l}$ there exists λ that satisfies (11). One can show that the range of λ that satisfies both the incentive constraints of the intermediary and the seller becomes smaller for a stricter uniform recommendation policy – that is, when β_{UNI} decreases.³⁰ This is due to the fact that a stricter uniform recommendation policy results in less total profits that can be collected if flexible consumers are served in the indirect channel and increases the profits from serving only picky consumers in the indirect channel. If selling to flexible consumer in the indirect channel cannot be induced by the naive recommendation policy $\beta = 0$, then the regulator will choose the minimal β_{UNI} such that there exist values λ satisfying inequalities (11). This policy makes both the intermediary’s and the seller’s incentive compatibility constraints binding which implies that

$$W^2(\beta_{\text{UNI}}) = (1 - \alpha) \left[\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha \right] (v_m + b - c)(v_m - c). \quad (12)$$

(iii) As shown in the previous section, the first best is implemented with β_{UNI} when $\alpha \leq \alpha_{\text{FI}}$. What is the optimal policy for $\alpha > \alpha_{\text{FI}}$? We derive the conditions under which implementing the above policy $\beta_{\text{UNI}} > 0$ is optimal – that is, the regulator prefers not to set $\beta_{\text{UNI}} = 0$ and flexible consumers buy in the direct channel. The regulator can implement the intermediate uniform policy β_{UNI} that solves (12) if and only if there exists λ that satisfies 11 for the uniform recommendation policy that corresponds to the recommendation level of laissez-faire outcome. Solving inequality (11) for $\beta_{\text{UNI}} = \frac{v_h - v_m}{v_m - v_l}$, we obtain that if $\frac{v_m - c}{p^I - c} \leq \frac{1 - \alpha}{\frac{\alpha}{2}(1 + \beta_{\text{UNI}}) + 1 - \alpha}$ holds true, then the intermediate uniform policy β_{UNI} that solves equation (12) can be implemented. It is easy to check that this inequality is satisfied for all $\alpha \leq \alpha_1$, where α_1 is strictly higher than α_{FI} and is given by

$$\alpha_1 = \frac{b}{b + \frac{1}{2} \frac{v_h - v_l}{v_m - v_l} (v_m - c)}. \quad (13)$$

If $\alpha > \alpha_1$, it is not possible to achieve the outcome in which flexible consumers are served in the indirect channel and, thus, it is optimal for the regulator to choose the fully informative recommendation policy and set $\beta_{\text{UNI}} = 0$. This makes sure that only picky consumers with good matches are served in the indirect channel.

By contrast, if $\alpha \in (\alpha_{\text{FI}}, \alpha_1]$, then the regulator can induce the outcome in which sales take place in the indirect channel and fraction β_{UNI} of picky consumer with bad matches are served, resulting in welfare loss of $\frac{\alpha}{2} \beta_{\text{UNI}} (c - b - v_l)$. It remains to check whether the social planner does not find it optimal to set instead $\beta_{\text{UNI}} = 0$ with the ensuing welfare loss of $(1 - \alpha)b$.

³⁰To see this, we use the fact that $W'(\beta_{\text{UNI}}) = \frac{\alpha}{2}(v_m + b - c) + \frac{\alpha}{2}(c - v_l) > 0$ implies that $\frac{(1 - \alpha)(v_m - c)}{W(\beta_{\text{UNI}})}$ decreases in β_{UNI} . The sign of the derivative of the maximal possible value for $1 - \lambda$ is positive since it is determined by $W'(\beta_{\text{UNI}})(W(\beta_{\text{UNI}}) + \frac{\alpha}{2}(1 + \beta_{\text{UNI}})(p^I - c)) - \frac{\alpha}{2}(v_m + b - c)W(\beta_{\text{UNI}}) = \frac{\alpha}{2}(c - v_l)W(\beta_{\text{UNI}}) + \frac{\alpha}{2}(1 + \beta_{\text{UNI}})(p^I - c)W'(\beta_{\text{UNI}}) > 0$.

Recall from Section 3.2 that if $\alpha \leq \bar{\alpha}$, then the losses from serving $\frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l}$ of picky consumers with bad matches are lower than the losses from serving the flexible consumers in the direct channel, i.e. $(1 - \alpha)b \geq \frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l)$.

We have that either $\alpha_1 \leq \bar{\alpha}$ or $\alpha_1 > \bar{\alpha}$. If $\alpha_1 \leq \bar{\alpha}$, then we have that

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

This implies that for all $\alpha \leq \alpha_1$, the regulator finds it optimal to set the intermediate level of uniform policy β_{UNI} determined in equation (12).

Consider next the case where $\alpha_1 > \bar{\alpha}$. We define

$$\alpha_2 = \frac{b}{b + \frac{1}{2} \beta_{\text{UNI}} (c - b - v_l)},$$

where β_{UNI} solves (12) for $\alpha = \alpha_2$. For all $\alpha \leq \alpha_2$, the regulator will find it optimal to set the intermediate β_{UNI} that solves (12) as this policy will implement the outcome in which flexible consumers are served in the indirect channel. This policy will result in higher social welfare than the fully informative recommendation policy. If $\alpha > \alpha_2$, then the naive recommendation policy $\beta_{\text{UNI}} = 0$ is chosen by the regulator.

It remains to establish conditions on the primitives that determine whether or not $\bar{\alpha} > \alpha_1$. Consider $\bar{v}_m \in (c, v_h)$ defined as follows

$$\bar{v}_m = \frac{c - b - v_l}{c - b - v_l + v_h - v_l} v_h + \frac{v_h - v_l}{c - b - v_l + v_h - v_l} c.$$

Note that if $v_m < \bar{v}_m$, then $v_m(v_h - v_l + c - b - v_l) < v_h(c - b - v_l) + c(v_h - v_l)$ which implies that $(v_h - v_l)(v_m - c) < (v_h - v_m)(c - b - v_l)$ and therefore

$$\frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} (v_m - c) < \frac{\alpha}{2} \frac{v_h - v_m}{v_m - v_l} (c - b - v_l).$$

This implies that $\alpha_1 > \bar{\alpha}$ if and only if $v_m < \bar{v}_m$ and vice versa.

Finally, we define

$$\alpha_{\text{UNI}} = \begin{cases} \alpha_1 & \text{if } v_m \leq \bar{v}_m, \\ \alpha_2, & \text{otherwise.} \end{cases}$$

This completes the proof. □

If $\alpha \in (\alpha_{\text{FI}}, \alpha_{\text{UNI}})$, the regulator sets β_{UNI} such that this policy makes the intermediary's and the seller's incentive compatibility constraints binding. Here, the regulator admits some degree of recommendation inflation to make sure that intermediary and seller make decisions such that flexible consumers buy in the indirect channel. As stated in the proposition, when there are sufficiently many picky consumers, the regulator's concern is mostly about the allocation of picky consumers, and, by setting $\beta_{\text{UNI}} = 0$, it makes sure that only picky consumers with good matches buy.

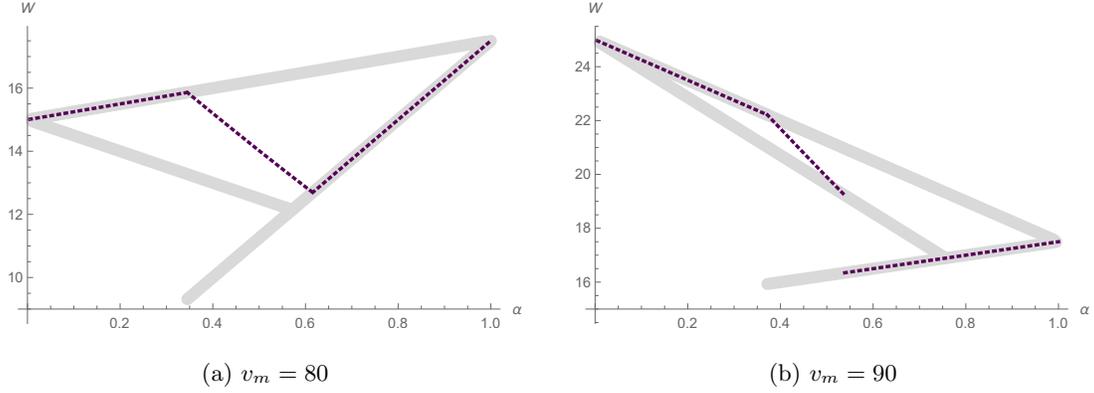


Figure 7: Total welfare in α ; $v_h = 100, v_l = 20, b = 10, c = 75$. First best outcome, private solution, fully informative recommendations (solid), uniform recommendation policy (dotted).

Figure 7 depicts the welfare associated with the optimal uniform recommendation policy compared to first best, laissez-faire, and mandated fully informative recommendations. Obviously, the optimal uniform recommendation policy performs weakly better than the mandated fully informative recommendations. It does strictly better for an intermediate range of α with $\alpha > \alpha_{FI}$. In this range, the regulator sets $\beta_{UNI} > 0$ such that all flexible consumers buy in the indirect channel. Whenever this is the case, welfare is strictly larger than under laissez-faire because the regulator permits less recommendation inflation than what would be the intermediary's choice under laissez-faire. However, since the intermediary cannot punish a deviating seller by even higher recommendation inflation, there are situations in which the intermediary chooses $\beta > 0$ along the equilibrium path under laissez-faire, whereas the regulator mandates $\beta_{UNI} = 0$ and flexible consumers buy in the indirect channel. Whenever this holds, welfare is lower under the optimal uniform recommendation policy than under laissez-faire – see the right-hand panel of Figure 7. This shows that allowing for $\beta_{UNI} > 0$ attenuates the negative welfare consequences of the uniform policy $\beta_{UNI} = 0$, but that the policy may still backfire and lead to lower welfare than under laissez-faire.

C.2 Uniform policy v. recommendation cap

In this section, we compare the welfare properties of the optimal uniform recommendation policy and the optimal recommendation cap policy with each other. Recall that the recommendation cap leaves some freedom to the intermediary to choose more informative recommendations than the cap (i.e., $\beta < \bar{\beta}$). If the only concern is excessive recommendation inflation, one may think that the two policies are equally well capable of addressing this concern and, therefore, should have the same welfare effects. Given the concern about inflated recommendations, this may suggest that the optimal recommendation cap is always better in terms of welfare than the optimal uniform policy and, thus, one may think that the recommendation cap

performs better. As we will see, the two policies differ on some range of α . To understand the difference, we have take a look at the seller’s and intermediary’s incentives.

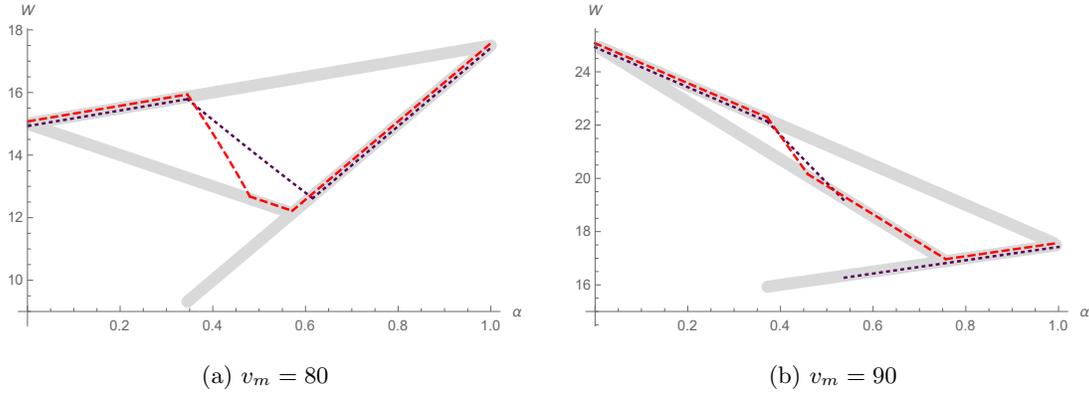


Figure 8: Total welfare in α ; $v_h = 100, v_l = 20, b = 10, c = 75$. First best outcome, private solution, fully-informative recommendations (solid), uniform recommendation policy (dotted), recommendation cap (dashed).

The key difference between the two policies is that, under strictly positive recommendation cap, the intermediary can set $\beta = 0$ and $\lambda = 0$ and extract the full surplus from picky consumers. This may constitute a profitable deviation for the intermediary. In response, the regulator has to choose a cap strictly above the one in the optimal uniform recommendation policy. This suggests that the optimal uniform recommendation policy is strictly better than the optimal recommendation cap policy. This is correct under some parameter constellations such as the one in the left-hand panel of Figure 8. However, there are other parameter constellations such that the ranking of the two policies in terms of welfare is ambiguous – see right-hand panel of Figure 8. While the optimal recommendation cap policy always weakly improves on the laissez-faire, this is not always the case with the optimal uniform recommendation policy, as explained in Section C.1. The reason that the optimal uniform recommendation policy can be welfare-inferior to the optimal recommendation cap is that the latter is never worse than the laissez-faire as it enables the intermediary to commit to a recommendation policy that conditions on the seller’s prices and, thus, to tame the seller’s deviation incentives. Hence, setting a recommendation level instead of a cap is a double-edged sword.

As illustrated in Figure 8, for $\alpha \leq \alpha_{FI}$, the optimal uniform recommendation policy and the optimal recommendation cap policy coincide and $\beta_{UNI} = \bar{\beta} = 0$; here, both policies implement the first best. For large α , the optimal uniform policy is $\beta_{UNI} = 0$ and all flexible consumers buy in the indirect channel. This is the laissez-faire outcome, which also obtains with the optimal recommendation cap, which can take any value ($\bar{\beta} \in [0, 1]$).

In the left-hand panel of Figure 8 the two policies differ on the range $(\alpha_{\text{FI}}, \alpha_2)$.³¹ The optimal uniform recommendation policy is strictly preferred. This reflects the difference between the two policies in constraining the intermediary to condition its recommendation policy on the seller's prices.

In the right-hand panel of Figure 8, the two policies differ on the range $(\alpha_{\text{FI}}, \bar{\alpha})$. As long as the optimal uniform recommendation policy can induce the intermediary and the seller to make choices such that flexible consumers buy in the indirect channel, this policy is the preferred policy.³² However, for larger α (i.e., $\alpha \in (\alpha_1, \bar{\alpha})$), the optimal uniform policy fails to deliver such an outcome and the regulator has to comfort itself with a fully informative recommendation policy that leaves flexible consumers choosing the inefficient bypass. By contrast, the recommendation cap can be set sufficiently high (for instance, $\bar{\beta} = 1$) such that the laissez-faire outcome prevails.

To summarize, while we have seen that both the optimal uniform recommendation policy and the optimal price cap regulation improve upon the fully informative recommendation policy $\beta_{\text{UNI}} = 0$ it is a priori not clear, which one delivers higher welfare.

D Regulation when the intermediary is subject to a break-even constraint

We consider the problem in which the intermediary operating the indirect channel has to incur the setup cost $K > 0$. We continue to work under Assumption 1 – that is, $(v_h + v_l)/2 + b - c < 0$.

First, consider the case in which intermediary and seller are vertically integrated. If the vertically integrated firm does not use the indirect channel, then it can maximally earn $(1 - \alpha)(v_m - c)$ by serving the flexible consumers in the direct channel. Thus, the vertically integrated firm finds it optimal to induce either the outcome with inflated recommendations or the outcome with inefficient bypass described in Section 3.2 if and only if

$$\max \left\{ \frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) + (1 - \alpha)b, \frac{\alpha}{2} (v_h + b - c) \right\} + (1 - \alpha)(v_m - c) - K \geq (1 - \alpha)(v_m - c),$$

or equivalently

$$\max \left\{ \frac{\alpha}{2} \frac{v_h - v_l}{v_m - v_l} (v_m + b - c) + (1 - \alpha)b, \frac{\alpha}{2} (v_h + b - c) \right\} \geq K. \quad (14)$$

Otherwise, it is optimal to serve only the flexible consumers in the direct channel and earn $(1 - \alpha)(v_m - c)$.

Second, consider the disintegrated case. Proposition 2 implies that the maximal profit of the intermediary is equal to the profit of the vertically integrated firm minus the minimal profit $(1 - \alpha)(v_m - c)$ that the seller can always ensure by serving the flexible consumers in the direct channel. Therefore, the intermediary will

³¹The upper bound α_2 has been defined in the proof of Proposition 10.

³²This is the case in the range $(\alpha_{\text{FI}}, \alpha_1]$, where α_1 has been defined in the proof of Proposition 10.

operate if and only if condition (14) is satisfied. If the intermediary can cover its setup cost $K > 0$, then the optimal strategy is characterized by Proposition 2.

Next, we move to the problem of the regulator who has full control over the recommendation policy of the intermediary but cannot directly affect λ . Following Section 4.1, we assume that the objective of the regulator is to maximize total welfare - we recall that the first best features that all flexible consumer and picky consumers with a good match are served in the indirect channel and picky consumers with a bad match do not buy. If the first best cannot be implemented, the regulator either ensures that all sales take place in the indirect channel and the inflation of recommendations is minimal or chooses a policy that implements inefficient bypass. When there are multiple solutions to the regulator's problem we choose those that maximize consumers surplus.

To induce the most efficient outcome in which all sales take place in the indirect channel, it is sufficient to consider the regulator's mandated recommendation policy of the following form:

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

This recommendation policy prescribes to fully inflate recommendations for the picky consumers for out-of-equilibrium prices (i.e., prices different from (p^I, p^D)), as this makes the deviations of the seller and the intermediary the least profitable. If the setup cost of the intermediary $K > 0$ is not too large, then the regulator's problem is to minimize inflated recommendations by choosing the minimal β_0 and ensure that all sales happen in the indirect channel.

The resulting total welfare for the outcome with inflated recommendations is given by

$$(1 - \alpha)(v_m + b - c) + \frac{\alpha}{2}(v_h + b - c) - \frac{\alpha}{2}\beta_0(c - b - v_l) - K.$$

If the regulator induces the outcome with inefficient bypass by setting $\beta = 0$ for all prices (p^I, p^D) , then the total welfare is given by

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

We begin by considering the regulator who induces the outcome in which all sales take place in the indirect channel. The incentive compatibility constraint of the seller is given by

$$(1 - \lambda) \left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha \right) (p^I - c) \geq \max\{(1 - \alpha)(v_m - c), (1 - \lambda)(1 - \alpha)(v_m + b - c)\}.$$

This implies that on the equilibrium path it must be that $(\alpha/2(1 + \beta_0) + 1 - \alpha)(p^I - c) \geq (1 - \alpha)(v_m + b - c)$. The intermediary will choose maximal λ ensuring that the seller does not deviate to serve the flexible consumers in the direct channel - that is, $(1 - \lambda)(\alpha/2(1 + \beta_0) + 1 - \alpha)(p^I - c) = (1 - \alpha)(v_m - c)$. The intermediary will operate if and only if its profit is higher than the setup cost K - that is,

$$\lambda \left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha \right) (p^I - c) = \left(\frac{\alpha}{2}(1 + \beta_0) + 1 - \alpha \right) (p^I - c) - (1 - \alpha)(v_m - c) \geq K.$$

If $K \leq (1 - \alpha)b$, then the first-best outcome described in Proposition 3 can be implemented.

Otherwise, if $K \in ((1 - \alpha)b, (1 - \alpha)b + \frac{\alpha}{2}(v_m + b - c)]$, the regulator can still set $\beta_0 = 0$ but has to increase the price in the indirect market p^I such that the intermediary can cover the setup cost – that is, the price must solve

$$\left(\frac{\alpha}{2} + 1 - \alpha\right)(p^I - c) = (1 - \alpha)(v_m - c) + K$$

and is equal to

$$p^I = c + \frac{1}{1 - \frac{\alpha}{2}}((1 - \alpha)(v_m - c) + K).$$

This continues to maximize total surplus, but consumer surplus is reduced. For a higher setup cost K , the regulator will have to allow for inflated recommendations by setting $\beta_0 > 0$ because the selected p^I cannot be above $v_m + b$. Lemma 2 implies that the price in the indirect channel for any β_0 has to be equal to $p^I = v_m + b$. The regulator will select β_0 such that the intermediary can cover its setup cost – that is,

$$\frac{\alpha}{2}(1 + \beta_0)(v_m + b - c) + (1 - \alpha)b = K.$$

If K is larger than the profit of the intermediary that induces the outcome with inflated recommendations under laissez-faire – that is, $\frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b$ – then the intermediary will decide not to become active and the regulator will have to induce the outcome with inefficient bypass.

If the intermediary induces the outcome with inflated recommendations, then the welfare loss is given by $K + \alpha/2\beta_0(c - b - v_l)$. The regulator will find it optimal to induce the outcome with inflated recommendations if and only if the welfare loss from the outcome with inefficient bypass $K + (1 - \alpha)b$ is larger than the loss from the outcome with inflated recommendations.

We will characterize the socially optimal outcome for $K > (1 - \alpha)b + \frac{\alpha}{2}(v_m + b - c)$ for three different ranges of α . Define α_6 as the solution of $\frac{\alpha}{2}(v_h + b - c) = \frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b$ – that is,

$$\alpha_6 = \frac{b}{b + (v_h - v_m)/2}.$$

It is easy to check that $\alpha_6 < \bar{\alpha}$. Parameter α falls in either one of the three intervals: $[0, \alpha_6)$, $[\alpha_6, \bar{\alpha})$, and $[\bar{\alpha}, 1]$.

First, if $\alpha < \alpha_6$ we have that $\frac{\alpha}{2}(v_h + b - c) < \frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b$. For any $K > \frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b$, the regulator cannot induce the outcome with inefficient bypass since the profit of the intermediary $\frac{\alpha}{2}(v_h + b - c)$ cannot cover the setup cost K . It remains to check that for any $K \leq \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b$, the intermediary will induce the outcome with inflated recommendations and this will result in total welfare

that is higher than with the intermediary being inactive. This is seen as follows:

$$\begin{aligned}
& (1 - \alpha)(v_m + b - c) + \frac{\alpha}{2}(v_h + b - c) - \frac{\alpha}{2}\beta_0(c - b - v_l) - K \\
& \geq (1 - \alpha)(v_m + b - c) + \frac{\alpha}{2}(v_h + b - c) - \frac{\alpha}{2}\beta^*(c - b - v_l) - \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) - (1 - \alpha)b \\
& = (1 - \alpha)(v_m - c) + \frac{\alpha}{2}(v_h - v_m) - \frac{\alpha}{2}\beta^*(v_m - v_l) \\
& = (1 - \alpha)(v_m - c),
\end{aligned}$$

where the last expression represents total welfare in the direct channel if the intermediary does not operate. For a setup cost K that is even higher than $\frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b$, the intermediary does not operate and total welfare is equal to $(1 - \alpha)(v_m - c)$.

Second, if $\alpha \in [\alpha_6, \bar{\alpha})$ we have that

$$\frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b \leq \frac{\alpha}{2}(v_h + b - c) < \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b.$$

For all $K \in (\frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b, \frac{\alpha}{2}(v_h + b - c)]$, the regulator will always induce the outcome with inflated recommendation since

$$\begin{aligned}
& \left(K - \left(\frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b \right) \right) \frac{c - b - v_l}{v_m + b - c} - (1 - \alpha)b \\
& \leq \left(\frac{\alpha}{2}(v_h + b - c) - \left(\frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b \right) \right) \frac{c - b - v_l}{v_m + b - c} - (1 - \alpha)b \\
& = \left(\frac{\alpha}{2}(v_h - v_m) - (1 - \alpha)b \frac{v_m - v_l}{c - b - v_l} \right) \frac{c - b - v_l}{v_m + b - c} = \left(\frac{\alpha}{2}\beta^*(c - b - v_l) - (1 - \alpha)b \right) \frac{v_m - v_l}{v_m + b - c} < 0,
\end{aligned}$$

where the last expression is negative for all $\alpha < \bar{\alpha}$. Therefore, it is optimal for the regulator to induce the outcome with inflated recommendations for all $\frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b < K \leq \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b$.

Third, suppose that $\alpha \geq \bar{\alpha}$. This implies that $\frac{\alpha}{2}(v_h + b - c) > \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b$. The regulator can induce the outcome with inflated recommendations and the outcome with inefficient bypass only for $K \in [\frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b, \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b]$. The regulator will induce the outcome with inflated recommendations if and only if

$$(1 - \alpha)b > \frac{\alpha}{2}\beta_0(c - b - v_l) = \left(K - \left(\frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b \right) \right) \frac{c - b - v_l}{v_m + b - c}.$$

If $K = \frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b$, then it is optimal for the regulator to induce the outcome with inflated recommendations. If $K = \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c) + (1 - \alpha)b$, then it is optimal to induce the outcome with inefficient bypass since for $\alpha \geq \bar{\alpha}$ the welfare loss from inflated recommendations is larger than the welfare loss from inefficient bypass. We define

$$\bar{K}(\alpha) = \frac{\alpha}{2}(v_m + b - c) + (1 - \alpha)b \frac{v_m - v_l}{c - b - v_l}$$

as the level of setup cost at which the regulator is indifferent between the two outcomes. For $K < \bar{K}(\alpha)$, the regulator will induce the outcome with inflated recommendations and for $K \in [K(\alpha), \alpha/2(v_h + b - c)]$, the regulator will induce the outcome with inefficient bypass.

We summarize this analysis by the following proposition:

Proposition 11. *Suppose that the regulator has full control over the recommendation policy.*

- *If $K \leq (1 - \alpha)b$, then the regulator implements the first best and the regulatory solution is characterized by Proposition 3.*
- *If $K \in ((1 - \alpha)b, (1 - \alpha)b + \frac{\alpha}{2}(v_m + b - c)]$, then the regulator implements the first best outcome. It can do so by setting $\beta = 0$ along the equilibrium path – that is, for prices $p^I = c + \frac{1}{1 - \alpha/2}[(1 - \alpha)(v_m - c) + K]$, $p^D = v_m$ – and $\beta = 1$ otherwise.*
- *If $K \in ((1 - \alpha)b + \frac{\alpha}{2}(v_m + b - c), (1 - \alpha)b + \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c)]$ with $\alpha < \bar{\alpha}$ or $K \in ((1 - \alpha)b + \frac{\alpha}{2}(v_m + b - c), \bar{K}(\alpha)]$ with $\alpha \geq \bar{\alpha}$, then the regulator optimally induces an outcome with inflated recommendations by setting $\beta = \frac{K - (1 - \alpha)b}{(\alpha/2)(v_m + b - c)} - 1$ along the equilibrium path – that is, for prices $p^I = v_m + b$, $p^D = v_m$ – and $\beta = 0$ otherwise. The welfare loss compared to the first best (that includes the cost K) is given by $\frac{\alpha}{2} \left(\frac{K - (1 - \alpha)b}{(\alpha/2)(v_m + b - c)} - 1 \right) (c - b - v_l)$.*
- *If $K \in (\bar{K}(\alpha), \frac{\alpha}{2}(v_h + b - c)]$ with $\alpha \geq \bar{\alpha}$, then the regulator finds it optimal to induce the outcome with inefficient bypass by setting $\beta = 0$ for all prices. The welfare loss is given by $(1 - \alpha)b$.*
- *If $K > \max \left\{ (1 - \alpha)b + \frac{\alpha}{2}(1 + \beta^*)(v_m + b - c), \frac{\alpha}{2}(v_h + b - c) \right\}$, then the indirect channel does not operate.*

As this proposition formalizes, the degree by which recommendations are inflated depends on the intermediary's set-up cost. In particular, since the intermediary has to be compensated for its costs, higher set-up costs go hand-in-hand with more-inflated recommendations.

References

- ALI, S. N., G. LEWIS, AND S. VASSERMAN (2019): “Voluntary Disclosure and Personalized Pricing,” *NBER Working Paper No. 26592*.
- ANAND, B. N. AND R. SHACHAR (2009): “Targeted Advertising as a Signal,” *Quantitative Marketing and Economics*, 7, 237–266.
- ANDERSON, S. P. AND Ö. BEDRE-DEFOLIE (2021): “Hybrid platform model,” CEPR Discussion Paper No. DP16243.
- 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,” Unpublished manuscript.
- 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.
- BAYE, M. R. AND J. MORGAN (2001): “Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets,” *American Economic Review*, 91, 454–474.
- BELLEFLAMME, P. AND M. PEITZ (2019): “Managing Competition on a Platform,” *Journal of Economics & Management Strategy*, 28, 5–22.
- (2021): *The Economics of Platforms: Concepts and Strategy*, Cambridge, UK: Cambridge University Press.
- BERGEMANN, D., B. BROOKS, AND S. MORRIS (2015): “The Limits of Price Discrimination,” *American Economic Review*, 105, 921–57.
- BIGLAISER, G. (1993): “Middlemen as Experts,” *Rand Journal of Economics*, 24, 212–223.
- BURGUET, R., R. CAMINAL, AND M. ELLMAN (2015): “In Google we trust?” *International Journal of Industrial Organization*, 39, 44–55.
- CHE, Y.-K. AND J. HÖRNER (2018): “Recommender Systems as Mechanisms for Social Learning,” *Quarterly Journal of Economics*, 133, 871–925.
- CHEN, Y. AND C. HE (2011): “Paid Placement: Advertising and Search on the Internet,” *Economic Journal*, 121, 309–328.

- COURTY, P. AND H. LI (2000): “Sequential Screening,” *Review of Economic Studies*, 67, 697–717.
- DE CORNIÈRE, A. AND G. TAYLOR (2014): “Integration and Search Engine Bias,” *Rand Journal of Economics*, 45, 576–597.
- (2019): “A Model of Biased Intermediation,” *Rand Journal of Economics*, 50, 854–882.
- 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.
- EDELMAN, B. AND J. WRIGHT (2015): “Price Coherence and Excessive Intermediation,” *Quarterly Journal of Economics*, 130, 1283–1328.
- ELIAZ, K. AND R. SPIEGLER (2011): “A Simple Model of Search Engine Pricing,” *Economic Journal*, 121, F329–F339.
- 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 (2020): “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 (forthcoming): “Should Platforms be Allowed to Sell on Their own Marketplaces?” *Rand Journal of Economics*.
- HAGIU, A. AND J. WRIGHT (2021): “Platform Leakage,” Unpublished manuscript.
- HANNAK, A., G. SOELLER, D. LAZER, A. MISLOVE, AND C. WILSON (2014): “Measuring Price Discrimination and Steering on E-Commerce Web Sites,” *Proceedings of the 2014 Conference on Internet Measurement Conference*, 305–318.
- HEIDHUES, P., M. KÖSTER, AND B. KŐSZEKI (2020): “Steering Fallible Consumers,” Unpublished manuscript.
- HERVAS-DRANE, A. (2015): “Recommended for You: The Effect of Word of Mouth on Sales Concentration,” *International Journal of Research in Marketing*, 32, 207–218.
- HERVAS-DRANE, A. AND S. SHELEGIA (2021): “Retailer-led Marketplaces,” Unpublished manuscript.

- HUNOLD, M., R. KESLER, AND U. LAITENBERGER (2020): “Rankings of Online Travel Agents, Channel Pricing, and Consumer Protection,” *Marketing Science*, 39, 92–116.
- 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.
- KANG, Z. Y. AND E. V. MUIR (2022): “Contracting and Vertical Control by a Dominant Platform,” Unpublished manuscript, Stanford University.
- KARLE, H. AND M. PEITZ (2017): “De-targeting: Advertising an Assortment of Products to Loss-Averse Consumers,” *European Economic Review*, 95, 103–124.
- KHAN, L. (2017): “Amazon’s Antitrust Paradox,” *Yale Law Journal*, 126, 710–805.
- LEE, C. (2021): “Optimal Recommender System Design,” Unpublished manuscript, University of Pennsylvania.
- LIZZERI, A. (1999): “Information Revelation and Certification Intermediaries,” *Rand Journal of Economics*, 30, 214–231.
- MARIOTTO, C. AND M. VERDIER (2020): “Platform–Merchant Competition for Sales Services,” *Journal of Economics & Management Strategy*, 29, 834–853.
- OECD (2018): “Personalised Price Discrimination,” DAF/COMP(2018)13.
- PEITZ, M. AND M. REISINGER (2016): “Media Economics of the Internet,” Elsevier, vol. 1A of *Handbook of Media Economics*, 445–530.
- RHODES, A. AND C. M. WILSON (2018): “False Advertising,” *Rand Journal of Economics*, 49, 348–369.
- RONAYNE, D. AND G. TAYLOR (forthcoming): “Competing Sales Channels with Captive Consumers,” *Economic Journal*.
- SEIDENSCHWARZ, H., J. FALTERMEIER, Q. WIERER, N. DEICHER, G. WITTMANN, A. BEER, AND J. SCHENKL (2021): “Empirie zu personalisierten Preisen im E-Commerce,” Final Report ibi research / trinnoative.

- TAYLOR, G. (2013): “Search Quality and Revenue Cannibalisation by Competing Search Engines,” *Journal of Economics & Management Strategy*, 22, 445–467.
- TEH, T. H. (forthcoming): “Platform Governance,” *American Economic Journal: Microeconomics*.
- TEH, T. H. AND J. WRIGHT (forthcoming): “Intermediation and Steering: Competition in Prices and Commissions,” *American Economic Journal: Microeconomics*.
- TUCKER, C. AND J. ZHANG (2011): “How Does Popularity Information Affect Choices? Theory and a Field Experiment,” *Management Science*, 57, 828–842.
- WANG, C. AND J. WRIGHT (2020): “Search Platforms: Showrooming and Price Parity Clauses,” *Rand Journal of Economics*, 51, 32–58.
- WHITE, A. (2013): “Search Engines: Left Side Quality Versus Right Side Profits,” *International Journal of Industrial Organization*, 31, 690–701.
- XU, L., J. CHEN, AND A. WINSTON (2012): “Effects of the Presence of Organic Listing in Search Advertising,” *Information System Research*, 23, 1284–1302.
- ZENNYO, Y. (forthcoming): “Platform Encroachment and Own-Content Bias,” *Journal of Industrial Economics*.