



^b
**UNIVERSITÄT
BERN**

Faculty of Business, Economics
and Social Sciences

Department of Economics

A Theory of Recommendations

Jean-Michel Benkert, Armin Schmutzler

24-06

This version: October, 2024

First version: August, 2024

DISCUSSION PAPERS

Schanzeneckstrasse 1
CH-3012 Bern, Switzerland
<http://www.vwi.unibe.ch>

A Theory of Recommendations

Jean-Michel Benkert and Armin Schmutzler*

October 11, 2024

Abstract

This paper investigates the value of recommendations for disseminating economic information, with a focus on frictions resulting from preference heterogeneity. We consider Bayesian expected-payoff maximizers who receive non-strategic recommendations by other consumers. The paper provides conditions under which different consumer types accept these recommendations. Moreover, we assess the overall value of a recommendation system and the determinants of that value. Our analysis highlights the importance of disentangling objective information from subjective preferences when designing value-maximizing recommendation systems.

Keywords: recommendations, preference heterogeneity, optimal design

JEL Classification: D02, D47, D83

*Benkert: Department of Economics, University of Bern, jean-michel.benkert@unibe.ch. Schmutzler: Department of Economics, University of Zurich, and CEPR, armin.schmutzler@econ.uzh.ch. We are grateful for very helpful discussions with Piotr Dworzak, Johannes Johnen, Igor Letina, Marc Möller, Nick Netzer, Volker Nocke, Egor Starkov, Jakub Steiner and seminar participants in Bern, Fribourg, Zurich, at SAET (2023), Swiss Society of Economics of Statistics (2024), Swiss IO Day (2024) and Verein für Socialpolitik (2024).

1 Introduction

Recommendations play an essential role for the diffusion of economic information. Job applicants frequently rely on assessments of potential employers by their friends. Conversely, job candidates often have to provide letters of recommendation from former employers. Consumers may base their choices of travel destinations, restaurants, or doctors on what they have heard from their acquaintances. In the digital economy, product ratings, typically reflecting anonymous consumer reports, are omnipresent.

Despite their relevance, recommendations are unlikely to function smoothly. For instance, too few recommendations may be given (Che and Hörner, 2018; Kremer, Mansour, and Perry, 2014), recommenders may be positively selected (Acemoglu, Makhdoumi, Malekian, and Ozdaglar, 2022), and interested parties might interfere with the process by choosing which recommendations to publish (Bolton, Greiner, and Ockenfels, 2013; Tadelis, 2016). In the literature, it is common to make the simplifying assumption that product qualities can be objectively ranked.¹ Instead, our paper assumes that preferences are heterogeneous, in which case product qualities *cannot* be objectively ranked and recommendations may be hard to interpret, even without strategic biases. We ask whether and how recommendations can be useful in spite of preference heterogeneity. Even well-intended recommendations by non-strategic players may reflect the senders' preferences rather than the objective truth. A receiver who does not question the informational content of a recommendation may therefore make biased decisions.

To motivate our approach, consider a consumer's choice between medical doctors. This is a complex problem, as the perceived quality of medical care will typically depend on many aspects. How carefully does a doctor listen to the patient? Is the technical equipment up to date? Is the staff friendly? How crowded is the waiting room? Moreover, different people may weigh different aspects differently. Despite such preference heterogeneity, rather than randomly choosing a doctor, people often base their decision on recommendations. Importantly, senders have limited stakes in the process and therefore lack incentives to think through the effects of their recommendation carefully. By contrast, adequately using the information can be crucial for the receiver. Finally, even if the sender is willing to think through the recommendation carefully, it will typically not be possible to convey all relevant information. For all these reasons, it is not obvious why or when recommendation can be useful.

This example illustrates three key features of our theory of recommendations:

Heterogeneity of Products and Preferences: Products come in different versions and consumers value them differently.

Asymmetry of Stakes: Making sure that the recommendation is not misunderstood is more important for the receiver than for the sender.

Coarseness of Recommendations: The sender is unable to provide enough detail for all relevant aspects of the service to become fully transparent to the receiver.

These features will be even more salient for ratings provided on internet platforms, as such ratings usually do not come from close acquaintances. Hence, the receivers should consider that senders have preferences which differ from their own and are likely to give an incomplete account of their consumption experience without exerting much effort. Backing out the implications of the recommendation for the receiver's own decision is

¹See also Che and Hörner (2018) and Hui, Klein, and Stahl (2023).

therefore a complex task. Relevant considerations for the interpretation are these: To what extent does the recommendation contain objective truth? And when it does not, how likely is it that the preferences of the sender and receiver are aligned? The answers to these questions will determine whether it is in the receiver’s interest to accept the recommendation and how valuable it is.

Our model captures the above three features of recommendations in the simplest possible way. An agent must decide between several options, referred to as *products* for ease of exposition. Products are two-dimensional objects, differing only in whether they have high or low quality in each dimension. Thus, four product *versions* exist. There is a continuum of agent types. All types agree that a product is worth choosing if it has high quality in both dimensions and is thus (objectively) good. Conversely, they agree that a product is not worth choosing if it has low quality in both dimensions and is thus (objectively) bad. Accordingly, we call the products *uncontroversial* in both of these cases. By contrast, we refer to products as *controversial* when they have high quality in one dimension and low quality in the other. An agent prefers a controversial product that has high quality in the dimension that is more relevant for them, and is thus subjectively good, to a product that has high quality only in the less-relevant dimension (thus subjectively bad). Moreover, even agents who agree about which controversial product they prefer may differ in the intensity of these preferences (as captured by the continuous type distribution).

We assume that products are experience goods so that each agent only has stochastic knowledge of the quality distribution before consumption: They know the probability of each product version. The agent obtains a recommendation from a previous consumer of one of the products. The recommendation reveals whether the sender had a consumption experience that was sufficiently good (a *buy recommendation*), with payoffs above an exogenous threshold level R , or not (a *don’t-buy recommendation*). Independent of R , the sender (she) always gives a buy recommendation for a good product and a don’t-buy recommendation for a bad product. In addition, she sends a buy recommendation for a controversial product if her preferences for the product are strong enough. In line with her low stakes in the issue, the sender thus only provides a coarse description of her own consumption experience. By contrast, the receiver (he) of the recommendation, with his high stakes in the issue, will carefully evaluate the informational content of the recommendation using Bayesian inference, taking into account the distribution of product qualities, his own preferences, and those of the population.

In this framework, we address a variety of questions: How do recommendations create value for receivers, how can we interpret this value, and when is it maximal? If one could influence the payoff threshold beyond which a sender gives a buy recommendation, how high should it be? Are recommendations valuable even when uncontroversial products are unlikely? How does the value depend on how polarized the population is?

To answer these questions, we need to derive intermediate results of independent interest. Most importantly, we investigate under which conditions receivers accept all recommendations. We find that for a Bayesian expected-payoff maximizer, two key quantities measuring the objective and subjective content of the recommendation are decisive for receiver behavior. These quantities are simple functions of the model primitives. Intuitively, the objective informational content of a recommendation is high if it makes the product sufficiently more likely to be (objectively) good and less likely to be bad. The subjective content captures the effect of the recommendation on the relative probability

of the two controversial product versions. The evaluation of the subjective content thus depends on the receiver’s type. When the recommendation bears little or no subjective content, all types accept all recommendations, thus buying after a buy recommendation and choosing the outside option (one of the other alternatives) after a don’t-buy recommendation. As the subjective content increases, revealing that one of the two controversial product versions becomes more likely, only receivers without strong preferences for the other version accept all recommendations. When the recommendation threshold R approaches the maximal payoff, a buy recommendation essentially becomes objective, and similarly for a don’t-buy recommendation when R approaches the minimal payoff. As receivers are aware of this, they accept all recommendations in each of these polar cases since the subjective component vanishes.

Based on this analysis, we then study the value of recommendation systems. We think of this value as the expected-payoff increase of a randomly chosen receiver resulting from a recommendation given by a randomly chosen sender. When the recommendation is objective, or the subjective content is small, the analysis is comparatively simple: As all receivers accept all recommendations, they choose the outside option after a don’t-buy recommendation, just as if they had not received any recommendation at all. Therefore, only buy recommendations create value—they do so by inducing the receiver to purchase the product if it generates a higher expected payoff for him than the outside option. When the recommendation contains more subjective content, only a subset of the receivers accept it, rendering the analysis more complex. The value of the recommendation system then has two sources. First, receivers who accept all recommendations benefit because they expect the recommended product to be better than the outside option. Second, receivers who never accept recommendations benefit from buying the product in spite of a don’t-buy recommendation: These receivers have preferences that are not well aligned with the general population so that the don’t-buy recommendation provides good news about the chances of obtaining a subjectively good product rather than one they like less.

Next, we turn to the design of recommendation systems. Motivated by the practice of digital platforms, we consider the recommendation threshold R as a design variable. When the type distribution is symmetric, so that there is no preference bias on average, the value of the recommendation system is a monotone function of the threshold. If the value is increasing, the product optimally receives a buy recommendation only if it is good; if the value is decreasing, the product optimally receives a don’t-buy recommendation only if it is bad. The former case arises if good products are more likely ex ante than bad products are, and conversely for the latter case. Either way, the content of the recommendation essentially becomes objective. In contrast, with an asymmetrically distributed population, value maximization does not necessarily require an extreme threshold. Rather than revealing only whether a product is good or bad, it may be optimal for the sender to give a buy recommendation even if the product is controversial but sufficiently aligned with her preferences. Intuitively, giving up the exclusive focus on objective content makes sense when both controversial product versions are common and there is broad agreement on which of them is preferable. By contrast, extreme recommendation thresholds that only reveal objective content are still desirable if the product is much more likely to be objectively good than objectively bad or vice versa.

We then discuss further determinants of the value of the recommendation system, assuming that the population is symmetric. First, we find that an increase in the probability of controversial products lowers the value only if the probabilities of the two controversial

versions are similar. Second, whether a mean-preserving spread in the population distribution (corresponding to a more polarized population) increases the value of a given recommendation system depends on how strict the recommendation threshold is and on how common good products are relative to bad ones.

Finally, we show the robustness of our results. With mild adjustments, the main insights continue to hold when sender and receiver distributions are distinct—a situation that could arise because, for instance, recommendations stem from early adopters who do not represent the population at large. This also has non-trivial design implications: Our results on the relation between mean-preserving spreads can be adapted to show how a designer would want to choose the sender distribution, which may be at least partly under his control. Second, we allow for more than two types of recommendations. Again, extreme recommendation thresholds increase the probability that a recommendation has objective content and thus increase the value of the system when the population is symmetrically distributed. Third, we show that having access to many recommendations does not necessarily change receiver behavior compared to the case of a single recommendation with an optimal threshold. Then, having access to multiple recommendations does not increase the value.

Section 2 relates the paper to the literature. Section 3 introduces the model and discusses its key assumptions. In Section 4, we describe the conditions under which different receiver types accept all recommendations. Then, in Section 5 we first characterize the value of a recommendation system before investigating the determinants of the value and the design of value-maximizing systems. Section 6 discusses our results and considers extensions. Section 7 concludes. All proofs are relegated to the appendix.

2 Relation to the Literature

Our paper investigates frictions in recommendation systems arising from preference heterogeneity, emphasizing the behavior of receivers. By contrast, the literature has mainly focused on frictions on the sender side.

Public Goods Problems Research on recommendation systems often deals with the exploration-exploitation dilemma and its implications for providing recommendations, interpreted as the provision of a public good. Several authors argue that consumers usually neglect the positive externalities generated by their reviews, leading to insufficient exploration and reviewing, in particular for lesser-known niche products (e.g., Kremer et al. (2014) and Che and Hörner (2018)). Expanding on these papers, Vellodi (2022) considers the impact on market structure, showing how unbiased reviews might disadvantage entrants due to the “cold start” problem, thereby reinforcing the monopoly position of incumbents. These studies view the design of the recommendation system as the solution of a public goods problem by appropriate information management.²

Sender Selection The selection of recommenders is another source of frictions in recommendation systems. In their analysis of the learning dynamics resulting from recom-

²For example, Kremer et al. (2014) and Che and Hörner (2018) suggest reducing transparency or endorsing yet-to-be-proven products, while Vellodi (2022) supports the idea of withholding reviews from incumbents to encourage new participants.

mendations, Acemoglu et al. (2022) argue that consumers with strong initial preferences for a product are more likely to purchase it.³ Recommendations thus inevitably carry an inherent bias reflecting the predispositions of the recommender pool.⁴ Receivers should recognize this bias rather than accepting recommendations at face value. Our paper intersects with theirs by highlighting the significance of preference heterogeneity for interpreting recommendations. However, contrary to Acemoglu et al. (2022), we focus on the determinants of the value of recommendation systems and the design of those systems rather than the possibility and speed of learning.⁵

Non-truthful Recommendations and Strategic Manipulations Strategic actions of recommenders, firms, and platforms may further reduce the value of recommendations for consumers. Bolton et al. (2013) have shown that eBay users often refrain from negative feedback to avoid retaliation.⁶ Like us, Chakraborty and Harbaugh (2010) assume that the sender possesses multidimensional information. They show how this helps her influence the behavior of the receiver in spite of a common bias. Instead, we emphasize the effects of preference heterogeneity without strategic behavior, assuming that senders mechanically provide truthful recommendations. In contrast with our model with consumer recommendations, in Peitz and Sobolev (forthcoming) and Johnen and Ng (2024) firms engage in sophisticated strategies to inflate their ratings.⁷ More closely related to our paper, Bourreau and Gaudin (2022) consider incentives for firms to manipulate recommendations in a setting with heterogeneous preferences.

3 Model

We begin by introducing the assumptions of our benchmark model in Section 3.1. Then, in Section 3.2, we discuss and motivate these assumptions.

3.1 Assumptions

Products and Payoffs We consider a consumer’s choice of a product from a set of alternatives $N \in \mathbb{N}$ with $N > 1$. We assume that each product is fully characterized by a two-dimensional quality vector, where $Q_d = 1$ and $Q_d = 0$ capture high and low quality, respectively, in dimension $d \in \{1, 2\}$. Accordingly, four different product versions exist.

³In contrast to our paper, their question is essentially whether and how fast consumers find out if the product is good or not, an assessment with which all consumers agree. Beyond that, we model preference heterogeneity differently.

⁴In Section 6.1, we allow for distinct sender and receiver populations and thus for the possibility of a bias among the recommenders relative to the overall population.

⁵In Section 6.3, we introduce multiple recommendations, as do Acemoglu et al. (2022), but, in contrast to them, we ask under which conditions multiple recommendations enhance system value.

⁶Tadelis (2016) discusses the design implications of these issues—using the example of eBay deciding to refrain from allowing seller evaluations of buyers—and also touches on strategic motives in recommendations more generally.

⁷Peitz and Sobolev (forthcoming) investigate firm-driven recommendations and the choice of sales channel: targeted (with recommendation) or non-targeted. They identify motives behind “inflated recommendations,” where the firm also recommends welfare-reducing bad product matches. Johnen and Ng (2024) analyze a related phenomenon, where firms initially lower prices to attract positive ratings, later leveraging these ratings to justify price increases—effectively reducing the reliability of ratings as a quality indicator.

We assume that the products are experience goods. Thus, the consumer cannot observe product quality before consumption, so the alternatives are identical for him ex ante. However, he knows the prior distribution of the quality vectors, referred to as quality distribution for brevity. Consumers have heterogeneous preferences fully characterized by their type $i \in [-1/2, 1/2]$. A consumer of type i receives a payoff of

$$v(Q_1, Q_2, i) = (1/2 + i)Q_1 + (1/2 - i)Q_2$$

from consuming (Q_1, Q_2) .⁸ Thus, possible payoffs range from 0 to 1, with all consumers obtaining the highest payoff, 1, from consuming version $(1, 1)$ and the lowest payoff, 0, from consuming $(0, 0)$. Therefore, consumers agree that product versions $(1, 1)$ and $(0, 0)$ are (objectively) good and bad, respectively. By contrast, the payoff from consuming $(1, 0)$ or $(0, 1)$ depends on a consumer's type i so that the assessment of $(1, 0)$ and $(0, 1)$ is subjective. We thus refer to products with these quality vectors as *controversial*. Consumers of type $i > 0$ prefer $(1, 0)$ to $(0, 1)$, and vice versa for $i < 0$. Those of type $i = 0$ are indifferent between these two versions.

The distribution of consumer types is denoted by F . For now, we leave the distribution unspecified but for simplicity assume continuity and full support. Finally, the quality distribution is given as $\mathbf{q} = (q_H, q_1, q_2, q_L)$ where

$$\begin{aligned} \Pr[(Q_1, Q_2) = (1, 1)] &= q_H, & \Pr[(Q_1, Q_2) = (1, 0)] &= q_1, \\ \Pr[(Q_1, Q_2) = (0, 0)] &= q_L, & \Pr[(Q_1, Q_2) = (0, 1)] &= q_2. \end{aligned}$$

We assume that $\mathbf{q} > 0$ unless mentioned otherwise.

Recommendations and Updating We assume that one of the products comes with a recommendation. The receiver has to choose between this product and an alternative without a recommendation. We distinguish between buy recommendations ($r = B$) and don't-buy recommendations ($r = D$). Moreover, $r = 0$ corresponds to the case without recommendation. In the benchmark model, we assume that the senders (she) and receivers (he) of recommendations are both randomly drawn from F .


To eliminate any scope for strategic behavior of the sender, we assume that she mechanically gives a buy recommendation if she obtained a payoff of at least $R \in (0, 1)$ and a don't-buy recommendation otherwise.⁹ As good products (version $(1, 1)$) yield a payoff of 1, they always result in a buy recommendation, whereas bad products (version $(0, 0)$) always yield a don't-buy recommendation. The sender's type determines the recommendation for a controversial product: After having consumed $(1, 0)$, a sender of type i gives a buy recommendation if and only if $i \geq R - 1/2$, whereas only senders with $i \leq 1/2 - R$ give a buy recommendation after consuming $(0, 1)$. All told, a recommendation can reflect objective and subjective considerations.

We assume that receivers are Bayesian, updating beliefs about product quality after having obtained a recommendation while taking into account the uncertainty about its informational content. We define $\pi^B(R)$ and $\pi^D(R)$ as the probability that a randomly

⁸For convenience, we identify a product version with the corresponding quality vector.

⁹We exclude the threshold $R = 0$ as this would render the recommendation uninformative. For symmetry, we have also excluded $R = 1$, but all results would essentially go through if we allowed $R = 1$.

FIGURE 1. ILLUSTRATION OF SENDER BEHAVIOR



Notes: The figure illustrates what types send a buy recommendation (green) and a don't-buy recommendation (red) for the two controversial product versions when the threshold is $R = 3/4$. For instance, in the top line corresponding to product version $(1, 0)$, the payoff obtained by a sender of type i is $1/2 + i$, so types such that $1/2 + i \geq 3/4 \Leftrightarrow i \geq 1/4$ give a buy recommendation.

chosen sender gives a buy or don't-buy recommendation, respectively. Clearly,

$$\begin{aligned}\pi^B(R) &:= q_H + q_1\phi_1(R) + q_2\phi_2(R), \\ \pi^D(R) &:= q_1(1 - \phi_1(R)) + q_2(1 - \phi_2(R)) + q_L,\end{aligned}$$

where $\pi^D(R) = 1 - \pi^B(R)$, and $\phi_1(R) := 1 - F(R - 1/2)$ and $\phi_2(R) := F(1/2 - R)$ denote the probability of a buy recommendation conditional on the product version being $(1, 0)$ or $(0, 1)$, respectively. In words, the probability of receiving a buy recommendation from a randomly chosen sender is given by the sum of the probability that a product is objectively good and the probability that it is subjectively good enough from a randomly chosen sender's perspective. Thus, accepting a buy recommendation the receiver's posterior $\mathbf{p}^B(R) := (p_H^B(R), p_1^B(R), p_2^B(R), p_L^B(R))$ becomes

$$\begin{aligned}p_H^B(R) &:= \frac{q_H}{q_H + q_1\phi_1(R) + q_2\phi_2(R)}, & p_2^B(R) &:= \frac{q_2\phi_2(R)}{q_H + q_1\phi_1(R) + q_2\phi_2(R)}, \\ p_1^B(R) &:= \frac{q_1\phi_1(R)}{q_H + q_1\phi_1(R) + q_2\phi_2(R)}, & p_L^B(R) &:= 0.\end{aligned}$$

A don't-buy recommendation gives rise to $\mathbf{p}^D(R) := (p_H^D(R), p_1^D(R), p_2^D(R), p_L^D(R))$, where

$$\begin{aligned}p_H^D(R) &:= 0, & p_2^D(R) &:= \frac{q_2(1 - \phi_2(R))}{q_1(1 - \phi_1(R)) + q_2(1 - \phi_2(R)) + q_L}, \\ p_1^D(R) &:= \frac{q_1(1 - \phi_1(R))}{q_1(1 - \phi_1(R)) + q_2(1 - \phi_2(R)) + q_L}, & p_L^D(R) &:= \frac{q_L}{q_1(1 - \phi_1(R)) + q_2(1 - \phi_2(R)) + q_L}.\end{aligned}$$

In the above expressions $\pi^B(R)$, $p^B(R)$, etc., we will typically drop the dependence on R to simplify notation. Finally, we use the following summary terminology:

Definition 1 $\mathcal{E} := (\mathbf{q}, F)$ defines a *decision environment*. A *recommendation system* \mathcal{R} consists of a decision environment \mathcal{E} and a recommendation threshold $R \in (0, 1)$.

3.2 Discussion of the Framework

Our theory of recommendations incorporates three key features: heterogeneity of products and preferences, coarseness of recommendations, and asymmetry of stakes between senders and receivers. We now discuss how our model addresses these and further points.

Heterogeneity of Products and Preferences Without controversial products or preference heterogeneity, the analysis would become trivial. The case without controversial products would correspond to $q_1 = q_2 = 0$. Then, for arbitrary $R \in (0, 1)$, buy recommendations would reveal that the product is good, while don't-buy recommendations would reveal that it is bad. Thus, there would be no friction in the recommendation system and no role for design considerations. With controversial products but absent preference heterogeneity, all consumers would have the same type i^* . Assuming w.l.o.g. that $i^* \geq 0$, senders would send a buy recommendation (i) only after consuming a good product if $R > \frac{1}{2} + i^*$, (ii) after consuming a good product or a controversial product of version $(1, 0)$ if $\frac{1}{2} + i^* \geq R > \frac{1}{2} - i^*$, and (iii) after consuming the good product or any of the controversial products if $\frac{1}{2} - i^* \geq R$. The first-best could be induced by the appropriate choice of the recommendation threshold, taking into account how the payoffs from controversial products compare to the expected payoff from the outside option. In sum, a meaningful analysis requires controversial products as well as preference heterogeneity.

Coarse Information Transmission With a single threshold above and below which senders give buy and don't-buy recommendations, respectively, it is impossible to always reveal the product version. Thus, as desired, recommendations provide coarse information. If, in contrast, the sender could give four different ratings depending on the consumption experience, full revelation would be possible in our parsimonious model with only four product versions. In reality, recommendations can convey more content than whether a buyer liked a product or not, for instance, by including verbal descriptions. However, given the typical complexity of products, even relatively detailed recommendations will usually only transmit coarse information.

Asymmetry of Stakes To capture the sender's limited stakes in the process, we assume that she behaves truthfully, mechanically reporting whether her consumption experience was good enough or not. More sophisticated and potentially strategic recommendations seem at odds with the lack of incentives to invest in this reporting activity, which results in the underprovision of recommendations (see the literature review).

Recommendation Threshold Our preferred way of thinking about the recommendation threshold is as a design parameter. To motivate this, we rely on Decker (2022) who argues that Airbnb advises its customers to give a five-star rating to every host that provides reasonably good services, thus reserving lower ratings exclusively to very bad hosts.¹⁰ In this spirit, we think of the threshold as reflecting the guidance of a platform such as Airbnb or as a social norm that is being actively shaped.

Other Recommendation Systems Assuming that senders mechanically provide a binary recommendation is adequate in view of their limited willingness to invest more time into giving a recommendation. Nonetheless, one could imagine more sophisticated recommendation systems. For instance, the sender might recommend a product when it is either good or when it is controversial but preferred by most of the population. This,

¹⁰In our setting, this approach corresponds to the use of a coarse recommendation system, with the vast majority of hosts receiving a buy recommendation and a small subset receiving a don't-buy recommendation. The recommendation threshold would then be close to $R = 0$ so that products with a don't-buy recommendation are typically objectively bad.

however, presupposes a sender with detailed information about the receiver’s decision environment. Instead, the sender could give a buy recommendation when the average payoff she expects receivers to obtain from consuming the product in question is higher than for the outside option. In that case, the expected value of the outside option would determine the recommendation threshold, which would then no longer be a design parameter.

Product Ranking When $q_H > 0$ and $q_L > 0$, all types agree on which product version is the best and which is the worst because their utility is fully determined by the qualities in each dimension and is increasing in each of them. In other environments, no products are objectively the best or the worst. For instance, there would be no such products in a modified version of our setting where $q_H = 0$ and $q_L = 0$, as the valuations for $(1, 0)$ and $(0, 1)$ are perfectly negatively correlated. We nonetheless keep the specific feature of the model because we believe it covers a wide range of plausible scenarios.

4 Optimal Receiver Behavior

In this section, we ask under which circumstances a receiver will accept all recommendations, depending on his type i and the recommendation system \mathcal{R} , as introduced in Definition 1. In Section 4.1, we discuss through which channels recommendations affect beliefs, distinguishing between the objective and subjective content of the recommendations. Section 4.2 then characterizes the optimal behavior of the receiver by identifying thresholds which divide the type set into receivers that accept all recommendations and those that never do.

4.1 The Effects of Recommendations on Beliefs

Recommendations contain coarse information about the product under consideration, leading the receiver to update his beliefs about its quality. We first focus on buy recommendations, the total effect of which can be decomposed as follows:¹¹

- (I) The probability of a bad product falls to zero.
- (II) A good product becomes more likely relative to the controversial products.
- (III) The odds between the controversial products change.

The first two effects in this belief-updating decomposition are objectively positive: They correspond to increases of the expected payoff for every receiver type, albeit of different sizes. By contrast, the third effect is fully type-dependent and thus subjective, resulting in an increase in the payoff for some receivers and a decrease for others.

The effects of don’t-buy recommendations are analogous to those of buy recommendations: First, the probability of a good product falls to zero; second, bad products become more likely relative to controversial products, and third, the odds between the controversial products change. Thus, again, the recommendation has two objective effects (but these are both bad now) and a subjective effect (which can be good or bad, depending on the receiver’s type).

¹¹In Appendix A.1, we provide a more formal analysis, decomposing the difference between posteriors and priors into three corresponding components.

4.2 Characterizing Optimal Receiver Behavior

We are now in a position to understand whether it is optimal for the receiver to accept recommendations. To this end, note that, after receiving a recommendation, the receiver can either buy the recommended product or not buy it. Denote by $U^B(i)$ and $U^D(i)$ player i 's expected payoff of buying a product with a buy and don't-buy recommendation, respectively, and by $U^0(i)$ the expected payoff of not buying, i.e., of buying the alternative product with no recommendation. We have

$$\begin{aligned} U^r(i) &:= p_H^r + (1/2 + i)p_1^r + (1/2 - i)p_2^r, \quad r \in \{B, D\}; \\ U^0(i) &:= q_H + (1/2 + i)q_1 + (1/2 - i)q_2. \end{aligned}$$

Now, the difference between $U^r(i)$ and $U^0(i)$ determines the behavior of an expected-payoff maximizing receiver as follows.

Definition 2 *We say that a receiver i*

- (i) *accepts a buy recommendation if $U^B(i) \geq U^0(i)$; otherwise we say that he does not accept a buy recommendation;*
- (ii) *accepts a don't-buy recommendation if $U^D(i) \geq U^0(i)$; otherwise we say that he does not accept a don't-buy recommendation*
- (iii) *accepts all recommendations if he accepts both buy and don't-buy recommendations;*
- (iv) *never accepts recommendations if he does not accept either buy and don't-buy recommendations.*

Hence, a receiver who accepts all recommendations “does as he is told”, whereas a receiver who never accepts recommendations is a “contrarian” who does the opposite of what he is told.¹² Note that the above definition does not cover all conceivable types of behavior, as a receiver i could accept buy recommendations and not accept don't-buy recommendations, which is tantamount to always buying the recommended product and thus ignoring the content of the recommendations. Analogously, a receiver i could ignore the content of the recommendations and never buy the recommended product. We will return to these cases shortly.

The condition from Definition 2(i) for receiver i to accept a buy recommendation can equivalently be written as

$$p_H^B - q_H + \frac{p_1^B - q_1}{2} + \frac{p_2^B - q_2}{2} \geq i[(p_2^B - q_2) - (p_1^B - q_1)]. \quad (1)$$

For more compact formulations, we introduce the following terminology.

Definition 3

- (i) *The **objective effect** of a recommendation r is defined as*

$$\Delta_O^r := p_H^r - q_H + \frac{p_1^r - q_1}{2} + \frac{p_2^r - q_2}{2}. \quad (2)$$

¹²The use of the term aligns with Drehmann, Oechssler, and Roeder (2005), who document contrarian behavior, i.e., investors going against the behavior of the predecessors, in a financial-markets experiment.

(ii) The **subjective effect** of a recommendation is defined as

$$\Delta_S^r := (p_2^r - q_2) - (p_1^r - q_1). \quad (3)$$

We can now restate condition (1) about accepting a buy recommendation as

$$\Delta_O^B \geq \Delta_S^B i. \quad (4)$$

We use this formulation to justify our terminology. Intuitively, the subjective effect summarizes whether and how the recommendation shifts probability between the controversial products, with $\Delta_S^r > 0$ corresponding to a shift from $(1, 0)$ to $(0, 1)$, and conversely for $\Delta_S^r < 0$.¹³ The payoff effect of this shift depends on the agent's type i and is therefore subjective. Moreover, for the unbiased receiver $i = 0$, equation (4) simplifies to $\Delta_O^B \geq 0$. As this receiver does not care about the shift in probabilities between the two controversial versions of the products, $\Delta_O^B \geq 0$ fully reflects the two objective effects (I) and (II) in the belief-updating decomposition in Section 4.1. For $i \neq 0$, the issue is more subtle. As long as $q_1 = q_2$, the expected payoff conditional on the information that a product is controversial is $1/2$ for all i and thus type-independent. Therefore, all types gain the same amount from (I) learning that a product is good or controversial instead of bad; the same applies to (II) learning that a product is good instead of controversial. For $q_1 = q_2$, Δ_O^B therefore exactly captures these two type-independent (and thus objective) effects of a buy recommendation.¹⁴ For $i \neq 0$ and $q_1 \neq q_2$, the issue is complicated further. Though the effects captured by (I) and (II) in Section 4.1 are positive for all types, their size depends on the receiver's type. Intuitively, if one of the controversial product versions (say $(1, 0)$) is more likely than the other ex ante, learning that a product is controversial rather than bad is more valuable for types who prefer $(1, 0)$ over $(0, 1)$. Similarly, learning that the product is good rather than controversial is less valuable for those types. Even in this case, the interpretation of $\Delta_O^B \geq 0$ as an objective effect is justified because of the general agreement on its sign. We can now state our first result.

Proposition 1 *For a given recommendation system \mathcal{R} , a receiver of type i*

(i) *either accepts all recommendations or never accepts recommendations;*

(ii) *accepts all recommendations if and only if $\Delta_O^B \geq i\Delta_S^B$.*

The first part of the result states that Definition 2 indeed captures all relevant forms of behavior. That is, no type of receiver will ignore the informational content of recommendations by always or never buying the recommended product. Intuitively, a receiver accepts all recommendations if he thinks that they have sufficient objective content and/or that his preferences are sufficiently aligned with the sender's. Conversely, a receiver never accepts recommendations if he thinks that they do not have sufficient objective content and/or that his preferences are not sufficiently aligned with the sender's. Whether the recommendation is positive or negative does not matter for this assessment. Rather, as laid out in the decomposition in Section 4.1, the objective effect Δ_O^B of a buy recommendation is positive, as the latter makes a good product more likely relative to the

¹³If $\Delta_S^r = 0$, there is no such shift in probabilities.

¹⁴The value of learning that the recommended product is more likely to be good is $p_H^r - q_H$, and the value of the change in the probability that it is controversial is $\frac{p_1^r + p_2^r - (q_1 + q_2)}{2}$.

controversial products and rules out a bad one. By contrast, the sign of the subjective effect Δ_S^B can be positive or negative. The second part of the result shows which types accept all recommendations, which happens unless the subjective effect is sufficiently large and the receiver's preferences are not sufficiently well aligned with those of the rest of the population. We emphasize that receivers who never accept recommendations still make use of them, by doing the opposite of what they are told.

Proposition 1 tells us that we can decompose the population into at most two different groups: those who accept all and those who never accept recommendations. Denoting the type who is indifferent between always accepting and never accepting recommendations as

$$\tilde{i} := \frac{\Delta_O^B}{\Delta_S^B} \text{ if } \Delta_S^B \neq 0,$$

we can summarize optimal receiver behavior as follows.

Corollary 1

- (i) If $|\Delta_S^B| \leq 2\Delta_O^B$, then all $i \in [-1/2, 1/2]$ accept all recommendations.
- (ii) If $\Delta_S^B < -2\Delta_O^B$, then all $i \geq \tilde{i}$ accept all recommendations.
- (iii) If $\Delta_S^B > 2\Delta_O^B$, then all $i \leq \tilde{i}$ accept all recommendations.

The intuition for Corollary 1, the proof of which is omitted, is straightforward. Without loss, consider a buy recommendation. For $\Delta_S^B = 0$, the recommendation does not shift probability between the controversial product versions. Accordingly, accepting the recommendation is optimal for all players because of its positive objective effect $\Delta_O^B \geq 0$. For $\Delta_S^B \neq 0$, the conclusion is the same as long as $|\Delta_S^B|$ remains small enough—as required by (i). In that case, even the receivers who are most adversely affected by the subjective content of the recommendation (those at $i = 1/2$ or $i = -1/2$) are convinced by its relatively objective character. As $|\Delta_S^B|$ increases, the role of the probability shift becomes apparent. As types $i > 0$ prefer (1,0) over (0,1), a shift of the probability from the latter to the former induced by a recommendation ($\Delta_S^B < 0$) increases such receivers' expected payoff, reinforcing the objective effect. By contrast, such a shift reduces the expected payoff of types $i < 0$. The condition in (ii) always holds for $i > 0$, as the objective and subjective effects reinforce each other. It also holds for $i < 0$ as long as $|i| < |\tilde{i}|$ so that the adverse subjective effect does not dominate the objective effect. Similarly, when the recommendation shifts the probability from (1,0) to (0,1) (when $\Delta_S^B > 0$), all types $i < 0$ accept. By contrast, types $i > 0$ experience a reduction in their value, which can outweigh the positive objective effect. The condition in (iii) makes sure that the latter case does not arise.¹⁵

Corollary 1 may potentially be useful to draw inferences about the relative importance of the objective and subjective effects of recommendations: If an observer sees a large share of receivers accepting recommendations, this can be interpreted as evidence for a relatively high objective effect of the recommendation.

Finally, we analyze receiver behavior for extreme recommendation thresholds.

¹⁵The effect on the expected payoffs in cases (ii) and (iii) is similar to the effect of a demand rotation as discussed in Johnson and Myatt (2006): For the types who accept the recommendation, the expected payoffs increase; for those who do not, they fall.

Proposition 2 *Suppose $R \rightarrow 1$ or $R \rightarrow 0$. Then, all types accept all recommendations irrespective of the product distribution.*

In the limits, both types of recommendations essentially contain only objective information. For instance, when $R \rightarrow 1$, buy recommendations require that the product is objectively good, while don't-buy recommendations require that it is not. Either way, the recommendation does not reveal anything about the relative probabilities of the two controversial product versions. Instead, it provides objective information that is equally valuable for all types, irrespective of their preferences for controversial products. Thus, all types accept all recommendations.

5 The Value of the Recommendation System

We now deal with the value of recommendation systems. In Section 5.1, we provide a characterization result. Section 5.2 asks how to use the recommendation threshold to influence the value. Finally, we investigate the effects of the preference and quality distributions on the value in Section 5.3.

5.1 Characterizing the Value of the Recommendation System

Before defining the value of a recommendation system, we describe the expected payoffs of a fixed type i receiver, depending on whether he has access to a recommendation system and whether he accepts its recommendations. Without a recommendation system, all products are identical ex ante, with an expected payoff based on the prior distribution of product qualities. Hence, the expected payoff of a receiver i is

$$V^0(i) := U^0(i). \tag{5}$$

Now suppose one of the products comes with a buy or don't-buy recommendation, so the receiver has to update his belief about the quality distribution. A receiver who accepts all recommendations chooses this product (rather than one of the alternatives) only if it comes with a buy recommendation. With a recommendation system, the expected payoff of a receiver i who accepts all recommendations is thus

$$V^A(i) := \pi^B U^B(i) + (1 - \pi^B) U^0(i). \tag{6}$$

The first term is the probability of receiving a buy recommendation multiplied by the expected payoff from accepting it (and thus buying the recommended product). The second term is the probability of a don't-buy recommendation multiplied by the expected payoff from buying the product without any recommendation, that is, based on priors.¹⁶

A receiver who never accepts recommendations will buy the recommended product if it comes with a don't-buy recommendation. Otherwise, he will buy one of the other products. Hence, the expected payoff of such a receiver is

$$V^N(i) := \pi^B U^0(i) + (1 - \pi^B) U^D(i). \tag{7}$$

¹⁶The notation $V^A(i)$ emphasizes the dependence of this expected payoff on the type of the receiver. Clearly, V_A also depends on R and on the decision environment; similarly for the expressions below.

A buy recommendation does not create value for receivers who do not accept it, as such types choose the alternative for which no recommendation is available. Instead, the don't-buy recommendation, which is given with probability $(1 - \pi^B)$, does influence the behavior of the types who do not accept it: Instead of buying an alternative based on prior probabilities, updating leads them to choose the product for which they received the explicit recommendation not to buy.

We now define the value of a recommendation system.¹⁷

Definition 4 $V(\mathcal{R})$, the *value of a recommendation system* \mathcal{R} , is the expected increase in payoffs resulting from the existence of recommendations, where the expectation is taken over all pairs of independently drawn senders and receivers.

Using Definition 4, we can characterize $V(\mathcal{R})$ as follows:

$$\begin{aligned} \text{(i)} \quad & \int_{-1/2}^{1/2} (V^A(i) - V_0(i)) dF(i) && \text{if } |\Delta_S^B| \leq 2\Delta_O^B; \\ \text{(ii)} \quad & \int_{-1/2}^{\tilde{i}} V^N(i) dF(i) + \int_{\tilde{i}}^{1/2} V^A(i) dF(i) - \int_{-1/2}^{1/2} V^0(i) dF(i) && \text{if } \Delta_S^B < -2\Delta_O^B; \\ \text{(iii)} \quad & \int_{-1/2}^{\tilde{i}} V^A(i) dF(i) + \int_{\tilde{i}}^{1/2} V^N(i) dF(i) - \int_{-1/2}^{1/2} V^0(i) dF(i) && \text{if } \Delta_S^B > 2\Delta_O^B. \end{aligned}$$

The simple structure of $V(\mathcal{R})$ for $|\Delta_S^B| \leq 2\Delta_O^B$ in (i) results because, by Corollary 1, all types accept a recommendation. The difference between the expressions for $\Delta_S^B < -2\Delta_O^B$ in (ii) and $\Delta_S^B > 2\Delta_O^B$ in (iii) arises because only types $i > \tilde{i}$ accept the recommendation in the former case, whereas only types $i < \tilde{i}$ do so in the latter.

We now express the value $V(\mathcal{R})$ in terms of its objective and subjective content.

Proposition 3 *The value $V(\mathcal{R})$ of the recommendation system \mathcal{R} is:*

$$\begin{aligned} \text{(i)} \quad & \pi^B [\Delta_O^B - \Delta_S^B \mathbb{E}[i]] && \text{if } |\Delta_S^B| \leq 2\Delta_O^B; \\ \text{(ii)} \quad & (1 - \pi^B) F(\tilde{i}) [\Delta_O^D - \Delta_S^D \mathbb{E}[i \mid i \leq \tilde{i}]] + \pi^B (1 - F(\tilde{i})) [\Delta_O^B - \Delta_S^B \mathbb{E}[i \mid i \geq \tilde{i}]] && \text{if } \Delta_S^B < -2\Delta_O^B; \\ \text{(iii)} \quad & \pi^B F(\tilde{i}) [\Delta_O^B - \Delta_S^B \mathbb{E}[i \mid i \leq \tilde{i}]] + (1 - \pi^B) (1 - F(\tilde{i})) [\Delta_O^D - \Delta_S^D \mathbb{E}[i \mid i \geq \tilde{i}]] && \text{if } \Delta_S^B > 2\Delta_O^B. \end{aligned}$$

Intuitively, when $|\Delta_S^B| \leq 2\Delta_O^B$, Corollary 1 implies that all receivers accept all recommendations as the objective effect dominates. Hence, by equation (6), the recommendation system creates value only in case of a buy recommendation, as a don't-buy recommendation leads to the same expected payoff as without a recommendation system. The positive payoff effect of the recommendation reflects the type-independent objective part Δ_O^B and the subjective part Δ_S^B , which is weighted by the average type. Finally, all of that only materializes with the probability of a buy recommendation, π_B .

As the absolute value of Δ_S^B increases, the analysis becomes more complex because some types never accept recommendations. We consider the case $\Delta_S^B > 2\Delta_O^B$ and reproduce the expression for $V(\mathcal{R})$ as follows:¹⁸

$$\underbrace{\pi^B F(\tilde{i}) [\Delta_O^B - \Delta_S^B \mathbb{E}[i \mid i \leq \tilde{i}]]}_{\text{types accept all recommendations}} + \underbrace{(1 - \pi^B) (1 - F(\tilde{i})) [\Delta_O^D - \Delta_S^D \mathbb{E}[i \mid i \geq \tilde{i}]]}_{\text{types never accept recommendations}}. \quad (8)$$

In this case, the buy recommendation shifts probability from $(1, 0)$ to $(0, 1)$, which is unfavorable for receivers $i > 0$, reducing their gains from accepting and possibly making

¹⁷This definition is in line with the definition of a valuable signal in Kamenica and Gentzkow (2011).

¹⁸The argument for $\Delta_S^B < -2\Delta_O^B$ is analogous.

them negative. Thus, for $\Delta_S^B > 2\Delta_O^B$, only types below \tilde{i} (a fraction $F(\tilde{i})$ of the receivers) accept all recommendations, which comes with probability π^B . In addition to the objective value of the recommendation (Δ_O^B), these receivers obtain a subjective value. In (8), the subjective value is reflected in $-\Delta_S^B \mathbb{E}[i \mid i \leq \tilde{i}]$, the average effect of the shift in the relative probabilities of (1, 0) and (0, 1) on the value of accepting the recommendation for these buyers. The sign of this term depends on the properties of F captured in $\mathbb{E}[i \mid i \leq \tilde{i}]$. It is positive if the negative contribution of types $i < 0$ dominates the positive contribution of types $i > 0$. The second component in (8), corresponding to the types who never accept recommendations, is arguably more surprising. Given their lack of alignment with the rest of the population, these types treat the recommendation not to buy as good news about the product and thus do not accept the recommendation.¹⁹ Nonetheless, its informational content creates value for them.

Going forward, we shall often invoke the following symmetry assumption on the preference distribution.

Assumption 1 $F(-i) = 1 - F(i)$ for all $i \in [-1/2, 1/2]$.

We now define the quantity

$$\beta := F(1/2 - R).$$

Due to Assumption 1, it has a simple interpretation: Given a threshold R , the parameter β captures the fraction of the population that is willing to recommend product versions (1, 0) or, equivalently (given the symmetry of the distribution), (0, 1). For a fixed type distribution F , the variables β and R are inversely related. We then obtain the following result:

Proposition 4 *Suppose Assumption 1 holds. Then, all types accept all recommendations and*

$$V(\beta) = \pi^B \Delta_O^B.$$

To understand why all types accept all recommendations, the belief-updating decomposition in Section 4.1 is helpful. Its third component captures the potential shift of probabilities between controversial products, which induces a purely type-dependent payoff effect. Given the symmetry assumption, both controversial product versions lead to a buy recommendation with probability β . Hence, a recommendation does not shift probabilities between the controversial products. As the other two effects of a buy recommendation on payoffs are positive for all types, the buy recommendation increases expected payoffs for any type, yielding the result.²⁰ Proposition 3 thus implies $V(R) = \pi^B [\Delta_O^B - \Delta_S^B \mathbb{E}[i]]$. Further, because of the symmetry of the population, $\mathbb{E}[i] = 0$, so $V(R) = \pi^B \Delta_O^B$ as required.

¹⁹In detail, a don't-buy recommendation (which happens with probability $1 - \pi^B$) carries an adverse objective content (summarized by Δ_O^D in the term in squared brackets). However, for receivers with $i > 0$ it also carries the good news that the product is more likely to be of the preferred version (1, 0) rather than (0, 1). Those with $i > \tilde{i} (> 0)$ then buy the product despite the adverse objective content, valuing the beneficial subjective effect with $-\Delta_S^D \mathbb{E}[i \mid i \geq \tilde{i}] > 0$ on average. Such receivers exist only if Δ_S^D is sufficiently large (and hence Δ_S^D is sufficiently negative relative to Δ_O^D).

²⁰An analogous argument applies for don't-buy recommendations.

5.2 Optimal Design of Recommendation Systems

We now use our characterization of the value of a recommendation system to address design issues, focusing on the optimal choice of the threshold. We start with the case of a symmetric population, then we treat asymmetric populations.

5.2.1 Symmetric Type Distribution

In the symmetric setting, we obtain a simple expression for $V(R)$, yielding transparent comparative statics results.²¹ To this end, a reparametrization is helpful:

Definition 5 *The prevalence of controversial products is given as $Q := \frac{q_1+q_2}{2}$. The odds of a good product are $\sigma := q_H/q_L$.*

Using these parameters, we immediately obtain:

Lemma 1 *Suppose Assumption 1 holds. Then,*

$$\begin{aligned}\pi^B &= \frac{(1-2Q)\sigma}{1+\sigma} + 2Q\beta; \\ \Delta_O^B &= \frac{\beta Q(\sigma+1) - 2Q\sigma + \sigma}{2\beta Q(\sigma+1) - 2Q\sigma + \sigma} - \frac{(1-2Q)\sigma}{1+\sigma} - Q; \\ V(\beta) &= \pi^B \Delta_O^B = (1-2Q) \frac{\sigma + Q(\beta - \sigma + \sigma^2(1-\beta))}{(\sigma+1)^2}.\end{aligned}$$

The value of the recommendation system depends on β as well as on the prevalence of controversial products Q and on the odds ratio σ between good and bad products. Recall that, for each of the two controversial product versions, β measures the fraction of the population willing to recommend it. Therefore, the parameter captures the influence of the distribution and the recommendation threshold on the value.²² We can interpret changes of β as either changes of the threshold R for a given distribution F or as changes of the distribution F for a given threshold R . The following result abstracts from this distinction, allowing for either interpretation. It shows that countervailing effects of β on the probability of a buy recommendation and the objective value Δ_O^B of accepting it result in a non-trivial effect on $V(\beta) = \pi^B \Delta_O^B$.

Lemma 2 *Suppose Assumption 1 holds. Then,*

- (i) π^B is strictly increasing and Δ_O^B strictly decreasing in β ;
- (ii) $\frac{\partial^2 V}{\partial \beta \partial \sigma} < 0$ and $V'(\beta) = \begin{cases} > 0 & \text{if } \sigma < 1; \\ = 0 & \text{if } \sigma = 1; \\ < 0 & \text{if } \sigma > 1. \end{cases}$

The first statement in (i) holds because as β increases a greater share of the population is willing to recommend a controversial product. Moreover, conditional on receiving a buy recommendation, the recommended product is thus more likely to be controversial.

²¹As we focus on R as the design parameter, we often write $V(R)$ instead of $V(\mathcal{R})$.

²²Observe that by varying R we can obtain any $\beta \in (0, 1)$.

Therefore, the objective payoff component Δ_O^B decreases as β increases, implying the second statement in (i). The effect of β on $V(\beta)$ in (ii) reflects the effect on its two components. As σ increases (good products become relatively more likely than bad ones), a buy recommendation becomes more likely. Hence, a high Δ_O^B is particularly valuable when σ is high so that increasing σ and lowering β are complementary ($\frac{\partial^2 V}{\partial \beta \partial \sigma} < 0$).²³ The value-maximizing level of β must thus be weakly decreasing in σ . The symmetry assumption turns out to render the problem linear in β , so the second result in (ii) follows. Using $\beta = F(1/2 - R)$, Lemma 2 directly implies the following result.²⁴

Proposition 5 *Suppose Assumption 1 holds. Then, $V(R)$ is*

- (i) *decreasing in R if $\sigma < 1$;*
- (ii) *increasing in R if $\sigma > 1$;*
- (iii) *constant in R if $\sigma = 1$.*

By Lemma 2, an increase in R reduces the probability of a buy recommendation, while it increases Δ_O^B , the expected average payoff conditional on a buy recommendation. These effects exactly offset each other if and only if $\sigma = 1$. Otherwise, one effect dominates the other. Based purely on uncontroversial products, a buy recommendation would be more likely than a don't-buy recommendation when the good product is relatively more likely than the bad one ($\sigma > 1$). It is thus worthwhile to increase the expected average payoff conditional on a buy recommendation by increasing R at the expense of the probability of receiving such a recommendation. Put differently, the designer chooses between revealing whether the product is good or whether it is bad. He moves the threshold in the direction which yields full revelation with higher probability, which hinges only on the odds ratio of the uncontroversial product versions.

5.2.2 Asymmetric Type Distribution

For a symmetric population, the subjective component cancels out in the value of the recommendation system, allowing for a focus on the role of the objective content. Moreover, symmetry simplifies the analysis by ensuring that all types accept the recommendation. However, the assumption of a symmetric population may prove restrictive in some contexts, and it does not allow us to study the role of the subjective component for the optimal design of the recommendation system. To make headway in that direction, we introduce the following assumption.

Assumption 2 (i) $F(i) = (i + 1/2)^a$ for $a > 0$; (ii) $q_1 = q_2 = Q$.

We thus capture asymmetry by means of a simple but flexible parameterization. By assuming that the controversial products are equally likely a priori, we ensure that any effect of a recommendation exclusively reflects the population asymmetry. Corollary 1 implies the following result:

²³More technically, repeated application of the product rule shows that $\frac{\partial^2 V}{\partial \beta \partial \sigma} = \frac{\partial \pi^B}{\partial \beta} \frac{\partial \Delta_O^B}{\partial \sigma} + \frac{\partial \pi^B}{\partial \sigma} \frac{\partial \Delta_O^B}{\partial \beta} + \frac{\partial^2 \pi^B}{\partial \beta \partial \sigma} \Delta_O^B + \frac{\partial^2 \Delta_O^B}{\partial \beta \partial \sigma} \pi^B$. The (incomplete) intuition just given corresponds to the first term being positive.

²⁴Result (iii) follows because $V(\beta)$ is independent of β in this case.

Corollary 2 *Suppose Assumption 2 holds. Then, we have:*

$$(i) \Delta_S^B = 0 \text{ if } a = 1; (ii) \Delta_S^B < 0 \text{ if } a > 1; (iii) \Delta_S^B > 0 \text{ if } a < 1.$$

Thus, whether and how a recommendation shifts the probabilities of the controversial product versions depends entirely on the type distribution. When it is symmetric ($a = 1$), there is no shift. When it is convex ($a > 1$), so that more people prefer version $(1, 0)$ to $(0, 1)$, a buy recommendation shifts probability from $(0, 1)$ to $(1, 0)$, as it is more likely to come from a person who prefers $(1, 0)$ to $(0, 1)$. Similarly, when the type distribution is concave ($a < 1$), the probability is shifted to $(0, 1)$.


We can now analyze the extent to which the result in Proposition 5 that, generically, extreme thresholds are desirable is an artifact of the symmetry assumption. This is challenging because, contrary to the symmetric case, the set of types that accept all recommendations depends on the threshold R . Reflecting this complication, there is no direct counterpart to Proposition 5. Nevertheless, the following two results provide some insights into the optimal recommendation system in the case of an asymmetrically distributed population.

Proposition 6 *Suppose Assumption 2 holds. An optimal threshold $R^* \in (0, 1)$ exists if $Q > 1 - \max\{\frac{a}{a+1}, \frac{1}{a+1}\}$ or if (i) or (ii) holds:*

$$(i) a < 1 \text{ and } \frac{1 - Q(a+1)}{a - Q(a+1)} \geq \sigma \geq \frac{a - Q(a+1)}{1 - Q(a+1)};$$

$$(ii) a > 1 \text{ and } \frac{1 - Q(a+1)}{a - Q(a+1)} \leq \sigma \leq \frac{a - Q(a+1)}{1 - Q(a+1)}.$$

FIGURE 2. ILLUSTRATION OF PROPOSITIONS 6 AND 7



Notes: An interior solution arises between the respective top and bottom curve segments (Proposition 6). For $a \leq 1/9 \approx 0.11$ and $a \geq 9$, an interior solution arises for all values of σ . Below the respective top and bottom curve segments, extreme thresholds are optimal if Q is sufficiently small (Proposition 7); for example, $Q = 0.1$, as depicted.

As illustrated in Figure 2, Proposition 6 provides sufficient conditions for interior solutions. It shows that an interior optimum exists when (i) controversial products are likely (Q is

large) and the population is very asymmetric (a is far away from 1) or otherwise when (ii) the uncontroversial product versions are similarly likely (σ is close to 1).²⁵ The result starkly contrasts the symmetric case, where an interior solution is only possible in the knife-edge case $\sigma = 1$. In that case, recommendations do not reveal any information about the relative likelihood of the controversial products beyond the prior, so the subjective payoff effect cancels out and the optimal threshold depends only on the odds ratio σ of the uncontroversial products. In contrast, the average subjective payoff effect does not cancel out in the asymmetric case. Thus, the choice of the optimal threshold no longer only depends on σ but also on the asymmetry parameter a and the prevalence of controversial products Q . We explore under which circumstances the conditions given in Proposition 6 are necessary as well as sufficient:

Proposition 7 *Suppose Assumption 2 holds and fix σ . Then, as $Q \rightarrow 0$, $V(R)$ is*

$$(i) \text{ decreasing in } R \text{ if } \sigma \leq \min \left\{ \frac{1-Q(a+1)}{a-Q(a+1)}, \frac{a-Q(a+1)}{1-Q(a+1)} \right\};$$

$$(ii) \text{ increasing in } R \text{ if } \sigma \geq \max \left\{ \frac{1-Q(a+1)}{a-Q(a+1)}, \frac{a-Q(a+1)}{1-Q(a+1)} \right\};$$

(iii) maximized for some $R^* \in (0, 1)$ otherwise.

The proof essentially proceeds in two steps. First, all types accept all recommendations when Q is sufficiently small, since then the subjective content in the recommendation vanishes.²⁶ This simplifies $V(R)$ by eliminating the case distinctions. Leveraging this, we can prove the quasiconcavity of the function, ensuring that the sufficient conditions in Proposition 6 are necessary, too.

The requirement that $Q \rightarrow 0$ is much stronger than necessary for this result. Even for higher values of Q , Figure 2 quite accurately separates the parameter range under which an interior solution exists from the area where it does not.²⁷ In a nutshell, the result shows that an interior solution exists if (i) controversial products are sufficiently prevalent, (ii) the population asymmetry is substantial, and (iii) the odds ratio of the uncontroversial products is sufficiently close to 1. To understand this intuitively and to contrast it with the symmetric case, suppose that $\sigma > 1$. Hence, the value of the recommendation system is increasing in R in the symmetric case, as the gain in the expected payoff conditional on receiving a buy recommendation outweighs the loss from the lower probability of a buy recommendation. This follows because the subjective effect in $V(R) = \pi^B [\Delta_O^B - \Delta_S^B \mathbb{E}[i]]$ cancels out with symmetry. With an asymmetric population, the term $-\Delta_S^B \mathbb{E}[i]$ does not vanish. To illustrate this, consider a very asymmetric preference distribution such as $F(x) = (x + 1/2)^{10}$, for which the probability of a type $i \leq 0$ is less than 0.1%. Then, $\mathbb{E}[i] > 0$ and, by Corollary 2, $\Delta_S^B < 0$ so that $-\Delta_S^B \mathbb{E}[i] > 0$. Hence, the subjective effect is positive on average. Intuitively, though this population is distributed very asymmetrically

²⁵The figure might suggest that the parameter region where $R^* \rightarrow 1$ is in some sense larger than the area where $R^* \rightarrow 0$. This is an artefact of the otherwise quite convenient parameterization of the preference distribution but does not have any significant economic meaning.

²⁶ $Q \rightarrow 0$ implies that $|\Delta_S^B| \leq 2\Delta_O^B$, ensuring that all types accept all recommendations by Corollary 1. The key step is that $\Delta_S^B \rightarrow 0$ as $Q \rightarrow 0$, while $\Delta_O^B \rightarrow 1 - q_H > 0$ when $Q \rightarrow 0$. The proof shows that, even for larger values of Q , $|\Delta_S^B| \leq 2\Delta_O^B$ often holds, explaining why the conclusion of Proposition 7 holds not only in the limit.

²⁷The figure is drawn for $Q = 0.1$, in which case the if-and-only-if statement is correct for $\sigma \in [1/60, 60]$ and $a \in [1/10000, 10000]$.

around zero, it is essentially homogeneous in the sense that almost all types agree on the ranking of the controversial products: If such a product results in a buy recommendation, it almost surely has quality vector $(1,0)$, as almost all types prefer this to $(0,1)$. As such, these products are not that controversial anymore, and even the subjective effect is “fairly objective.” Now, as R increases, the probability of a buy recommendation still falls, and the “purely objective” payoff part Δ_O^B increases. However, for R high enough, this “fairly objective” payoff part $-\Delta_S^B \mathbb{E}[i]$ also decreases, as an extreme threshold R impedes the ability to distinguish between a good product and a widely preferred controversial product. While consumers are homogeneous regarding which controversial product version they prefer, they are still heterogeneous with respect to the intensity with which they do so. Hence, disentangling the good and the preferred controversial product is worthwhile. Therefore, an interior optimal threshold results.

5.3 Preference Heterogeneity and Controversial Products

As discussed after Proposition 7, preference heterogeneity only matters when controversial products exist. Moreover, the less homogeneous the population, the more likely it is that extreme thresholds are desirable so that receivers learn solely about uncontroversial products. Hence, one might conjecture that greater preference heterogeneity and a higher prevalence of controversial products reduce the value of recommendation systems. In the following, we address both issues in turn, demonstrating that these conjectures do not hold without qualification. For tractability, we focus on symmetric populations.

5.3.1 Preference Heterogeneity

In this section, we analyze how changes in the (exogenously given) preference distribution affect the value of the recommendation system.²⁸ Under Assumption 1, we can analyze increases in the polarization of the distribution of preferences by keeping the threshold R fixed and varying β .²⁹ This corresponds to a mean-preserving spread of F (see Figure 3), coinciding with an increase in β for $R > 1/2$ and a decrease for $R < 1/2$. We obtain the following result.

Proposition 8 *Suppose Assumption 1 holds.*

1. *Suppose $\sigma < 1$. Then, a mean-preserving spread of the type distribution increases the value of the recommendation system for $R > 1/2$ and decreases it for $R < 1/2$.*
2. *Suppose $\sigma > 1$. Then, a mean-preserving spread of the type distribution decreases the value of the recommendation system for $R > 1/2$ and increases it for $R < 1/2$.*


To gain some intuition, suppose $\sigma < 1$ and $R > 1/2$. Figure 1 shows that with $R > 1/2$, only extreme types give buy recommendations for controversial products.³⁰ Increasing

²⁸When we consider distinct sender and receiver distributions in Section 6.1, we will build further on this and take a design perspective instead of the comparative statics approach here.

²⁹For $R = 1/2$, we must have $\beta = 1/2$ because the symmetry of F requires $F(0) = 1/2$, so this exercise is only meaningful for thresholds $R \neq 1/2$. Further, for $R > 1/2$ we must have $\beta \leq 1/2$ because $\beta = F(1/2 - R) \leq F(0) = 1/2$. Accordingly, for $R < 1/2$ we must have $\beta \geq 1/2$.

³⁰Conversely, for $R < 1/2$ only relatively extreme types give don't-buy recommendations for controversial products. This can be seen by flipping the colors in Figure 1 so that it corresponds to $R = 1/4$.

FIGURE 3. POPULATION POLARIZATION



Notes: We have fixed $R = 3/4$. The dotted and dashed lines both correspond to symmetric type distributions F , where $\beta = 1/2$ for the dotted line and $\beta = 0$ for the dashed line. A move from $\beta = 0$ to $\beta = 1/2$ corresponds to a mean-preserving spread of the distribution of types.

β corresponds to a mean-preserving spread of F given R . Consequently, there are more such extreme senders and, hence, buy recommendations are increasingly likely to result from controversial products. This increases the probability of buy recommendations while lowering their objective value. The resolution of this trade-off, and thus whether the value of the recommendation system goes up or down, is determined by the odds ratio σ (see Proposition 5). For instance, if $\sigma < 1$, good products are rare compared to bad ones, so the value of the recommendation system is increasing in β . Therefore, the positive effect of revealing bad products by means of don't-buy recommendations outweighs the negative effect of a reduced value conditional on receiving a buy recommendation. Thus, given a high recommendation threshold ($R > 1/2$), a more polarized population benefits more from the recommendation system, as this means only bad products generate don't-buy recommendations. Conversely, when the recommendation threshold is relatively low ($R < 1/2$), a less polarized population centered around the middle gives don't-buy recommendations only to bad products.

5.3.2 Prevalence of Controversial Products


Without controversial products, preference heterogeneity would not matter: All receivers would accept all recommendations, as they would reveal whether the product is good or bad. Hence, the value of the recommendation system would be independent of the preference distribution. More broadly, we now ask how an increasing prevalence of controversial products (an increase in Q) affects the value of a recommendation system. A natural conjecture is that recommendations become less valuable when controversial products are more likely. We find that this conjecture needs to be qualified.

Corollary 3 *Suppose Assumption 1 holds.*

- (i) *If $3\sigma - \beta - \sigma^2 + \sigma^2\beta > 0$, then the value of the recommendation system is decreasing in Q . In particular, this is the case when $\sigma = 1$.*
- (ii) *If $3\sigma - \beta - \sigma^2 + \sigma^2\beta < 0$, then $Q^* := \frac{3\sigma - \beta - \sigma^2 + \sigma^2\beta}{4\sigma - 4\beta - 4\sigma^2 + 4\sigma^2\beta} \in (0, \frac{1}{2})$ maximizes $V(R)$.*

Figure 4 illustrates the result. In the white area in Panel A, $V(\beta)$ is decreasing in β ; in the shaded area, an interior maximizer exists. When good and bad products are similarly

FIGURE 4. ILLUSTRATION OF COROLLARY 3



Notes: In Panel A, the shaded area indicates the parameter space for which an interior solution for Q arises by Corollary 3. In Panel B, the shaded area indicates where π^B is increasing in Q . In Panel C, the shaded area indicates where Δ_O^B is increasing in Q , illustrated for $Q = 0.1$.

likely (σ is close to 1), a higher prevalence of controversial products indeed reduces $V(\beta)$. For very high or low values of σ , however, this need not be true. This surprising result reflects the multiplicative structure of $V(\beta)$. Under Assumption 1, $V(\beta)$ is the product of the buying probability π^B and the objective value of the recommendation per buyer, Δ_O^B . The effect of Q on π^B is negative only if σ is sufficiently large (see Panel B). Intuitively, as Q increases, good products are less likely, reducing the probability of buy recommendations. However, there are now more buy recommendations stemming from controversial products, as these have become more likely. β and σ determine which of these effects dominates. When β is high, the greater likelihood that controversial products receive buy recommendations increases the total probability of buy recommendations, strengthening the positive effect. When σ is high, the lower share of good products dominates, strengthening the negative effect. Similarly, one can show that the effect of Q on Δ_O^B becomes negative only for small σ (see Panel C). Again, this reflects countervailing forces. The negative effect results from the increased probability that a product with a buy recommendation is controversial. The positive effect reflects a change in the value of the outside option. The values of σ and β determine which effect dominates.³¹

6 Extensions

Our benchmark model relies on several assumptions. First, the preference distribution of senders and receivers is identical. Second, there are only two levels of recommendations (“buy” and “don’t-buy”). Third, each receiver only has access to a single recommendation. We now relax each of these assumptions in turn.

³¹When β is small, the negative effect is weak, as only few types give buy recommendations for controversial products. When σ is small, the outside option becomes more valuable, as bad products become less likely, weakening the positive effect.

6.1 Distinct Sender and Receiver Distributions

We modify the benchmark model by assuming that the sender and the receiver are drawn from different distributions F and G , respectively. For instance, this enables us to take into account that the “early adopters” of a new product may have different preferences than the general population (similar to Acemoglu et al. (2022)). We show that the results of the benchmark case easily generalize to this setting. More importantly, the general setting allows us to treat the choice of the sender distribution as a potentially interesting design variable.

The sender distribution F fully determines the probabilities of buy and don’t-buy recommendations, the posterior probabilities, and thus the objective and subjective effects and the expected payoffs, V_A and V_N . Hence, all results on receiver behavior in Section 4 apply verbatim if F is taken to be the sender distribution, as the receiver’s behavior is independent of the receiver distribution. The value of the recommendation system, however, depends on the receiver distribution G as well, as it involves integration over this distribution. The next result generalizes Proposition 3.³²

Proposition 9 *Consider the model with distinct preference distributions of senders and receivers. Suppose the sender distribution F satisfies Assumption 1. Then, all types accept all recommendations. Further, $V(R) = \pi_F^B [\Delta_{O,F}^B - \Delta_{S,F}^B \mathbb{E}_G[i]]$ is*

- (i) decreasing in R if $\sigma < \frac{Q - (q_1 - q_2) \mathbb{E}_G[i]}{Q - (q_2 - q_1) \mathbb{E}_G[i]}$,
- (ii) increasing in R if $\sigma > \frac{Q - (q_1 - q_2) \mathbb{E}_G[i]}{Q - (q_2 - q_1) \mathbb{E}_G[i]}$,
- (iii) constant in R if $\sigma = \frac{Q - (q_1 - q_2) \mathbb{E}_G[i]}{Q - (q_2 - q_1) \mathbb{E}_G[i]}$.

Recall that in the benchmark model symmetry of F implies that all receivers accept all recommendations and the subjective effect of the recommendation cancels out. This is still true if F is distinct from G as long as G is also symmetric and hence $\mathbb{E}_G[i] = 0$. For more general distributions G , all receivers still accept all recommendations for the same reason as above. However, as the receiver distribution is not symmetric, the subjective effect does not cancel out anymore. Even so, the result qualitatively mirrors Proposition 5, as changing the threshold still leads to the same trade-off: Reducing the threshold increases the probability of a buy recommendation but decreases its value conditional on that recommendation. However, in the latter effect, the value is no longer purely objective but consists of the objective part Δ_O^B and a subjective part given by $-\Delta_S^B \mathbb{E}_G[i]$, reflecting the receiver distribution G . Further, as the subjective effect now influences the value of the system, so does the prevalence of controversial products by affecting the tipping point at which the designer opts for a high or a low threshold. In any case, when $\mathbb{E}_G[i] = 0$, we obtain the result from the baseline model as a special case. Similarly, the other results in the benchmark model remain (qualitatively) valid. With asymmetrically distributed populations, Propositions 6 and 7 hold for arbitrary receiver distributions as long as we impose Assumption 2 on the sender distribution.

Allowing for differences in sender and receiver distributions makes it possible to address another important design issue. In some cases, it is natural to assume that a designer, for instance a platform, can influence the sender distribution or at least the receiver’s

³²The notation for the quantities in Proposition 9 reflects their dependence on F and G , respectively.

perception of it. Assuming that the platform has better knowledge about sender types, it may choose to share some of this information with the receivers or even restrict the option of sending recommendations to certain types of senders (and inform the receivers accordingly). For instance, many platforms such as Amazon inform browsing customers that “customers who bought this item also bought” another item.

There are different ways to model this. First, suppose that F is symmetric and that the designer can render the sender distribution more or less precise using mean-preserving contractions or mean-preserving spreads, respectively. Intuitively, one would expect more precise knowledge about the sender to increase the value of the recommendation system. However, the logic of Proposition 9 shows that this is not necessarily the case. As in Section 5.3.1, the effects of mean-preserving contractions and spreads can be captured as changes in $\beta = F(1/2 - R)$, just like changes in R . With mild adaptations, the proof of Proposition 9 can thus be used to show that whether more or less precise information about the sender increases the value of recommendations depends on the parameters.

Second, suppose that the sender and receiver distributions are fixed but that the designer can provide the receivers with a signal about the sender that allows for inference about the sender’s type. To make the point as starkly as possible, suppose the signal fully reveals the sender’s type. If the sender is revealed to be type $i = 0$, all types accept all recommendations, as both types of recommendations contain only objective information. In contrast to Proposition 2—where this behavior is shown to ensue with extreme thresholds—this would arise for all thresholds $R \in (0, 1)$. Intuitively, as the type $i = 0$ is indifferent between controversial products, her recommendation never reveals subjective information. In contrast, if the signal reveals a sender of type $i \neq 0$, the receiver will be able to learn about controversial products, too.³³ Thus, having a sender who is partial to some controversial product is not necessarily a bad thing.

Finally, suppose that the sender and receiver distributions are not independent but somehow correlated. Recommendations would then be “tailored” to the receiver \bar{i} , as they would be drawn from a sender distribution $F(i|\bar{i})$, resulting in type-specific posteriors $p_j^r(i)$ for $j \in \{H, 1, 2, L\}$. In such a setting, an analogous result to Proposition 1 still holds, establishing that either all or no recommendations are accepted. Further analysis would give rise to similar complications as in Section 5.2.2 with asymmetric distributions, as the correlation between sender and receiver distribution would affect the set of types who accept all recommendations.

6.2 Multiple Recommendation Thresholds

We now analyze the effects of reducing the coarseness of the recommendation system. We maintain all assumptions of the benchmark case, except that there are two threshold levels $R_1, R_2 \in (0, 1)$ with $R_2 \geq R_1$ and three types of recommendations $r \in \{D, N, B\}$ which we interpret as “don’t-buy,” “neither buy nor don’t-buy,” and “buy.”³⁴ The sender gives a buy recommendation if the payoff is at least R_2 , a don’t-buy recommendation if it is below R_1 , and neither if it is in $[R_1, R_2)$. We maintain Assumption 1.

The proof of Proposition 10 in Appendix A.3 provides the receiver’s posteriors. For

³³For instance, if $i = 1$ and $R \in (0, 1)$, a buy recommendation reveals that the product is either of version $(1, 1)$ or $(1, 0)$

³⁴Such a two-threshold structure is found to be optimal in models of strategic communication in the context of social learning by Smirnov and Starkov (2024) and Bénabou and Vellodi (2024).

buy and don't-buy recommendations, these posteriors are analogous to the benchmark, except that the different thresholds in each case have to be taken into account. As in the benchmark model, the highest and lowest recommendations rule out the bad and good product versions, respectively. The intermediate recommendation $r = N$, however, rules out *both* of these versions, so the receiver knows that the recommended product is controversial.

We say that a receiver accepts all recommendations if he buys if $r = B$ and does not buy if $r = D$, allowing for either behavior if $r = N$. Assumption 1 again implies that the receiver always buys a product with $r = B$ and never buys if $r = D$. As noted above, $r = N$ rules out both uncontroversial products. Given the symmetric population, $r = N$ leaves the relative likelihood of the controversial products unchanged. Hence, whether a type i wants to buy a product with $r = N$ depends on (i) how likely a good product is and (ii) how much he (dis-)likes the more likely controversial product. As in the benchmark model, we can derive the indifferent type so that, depending on parameters, either all, only sufficiently aligned, or no types buy the product with an intermediate recommendation.

$V(R)$ now consists of two components. First, as in the benchmark, all types accept $r = B$, resulting in the component $\pi^B \Delta_{\mathcal{O}}^B$.³⁵ Moreover, $r = N$ can create value for the types that buy following that neutral recommendation. As noted above, parameter values determine which types do this. Despite this ambiguity, we can formulate the following result.

Proposition 10 *Consider the model with multiple recommendation thresholds. Suppose Assumption 1 holds. $V(R)$ is (i) decreasing in R_1 if $\sigma < 1$ and (ii) increasing in R_2 if $\sigma > 1$.*

The result provides an incomplete characterization of the optimal recommendation system: It establishes that one of the two thresholds should be extreme so that one recommendation essentially becomes objective. Thus, as in Proposition 5, the designer maximizes the probability of an objective recommendation, considering which uncontroversial product version is more likely a priori. The optimal value of the threshold that is not pinned down by Proposition 10 depends on the model parameters unless $q_1 = q_2$. In this degenerate case, any value for the other threshold is optimal, but it is also possible that both the high and the low recommendations are essentially turned objective.³⁶

6.3 Multiple Recommendations

We now show that the value of a recommendation system does not always increase if receivers obtain more than one recommendation: The additional recommendations do not add value if the ex ante probabilities of the good and bad products are similar while those of the two controversial product versions are not.

We assume that receivers have access to a number b of buy recommendations and a number d of don't-buy recommendations. As in the benchmark model, recommendations come from randomly drawn senders. We add the assumption that they are drawn independently.³⁷ Just as before, a single buy recommendation is enough to rule out a bad product, and a single don't-buy recommendation is enough to rule out a good product.

³⁵As before, accepting don't-buy recommendations does not create value.

³⁶This arises for instance for $F(i) = 1/2 + i$ with parameters $q_1 = 0.7$, $q_2 = 0.1$, and $q_H = 0.03$.

³⁷In Appendix A.4, we provide the posteriors, depending on b and d .

Consequently, when the receiver obtains mixed recommendations, that is, both buy and don't-buy recommendations, he rules out both uncontroversial product versions.³⁸

To make our main point that multiple recommendations need not increase the value of the recommendation system as starkly as possible, we focus on the case of “infinite learning” where the receiver has access to an unlimited number of recommendations.

Proposition 11 *Consider the model with multiple recommendations. Suppose Assumption 1 holds. Let $\lambda := q_1/q_2$. Compared with the optimal recommendation system in the benchmark model, infinite learning does not increase the value of the recommendation system if and only if either (i), (ii), or (iii) holds:*

$$(i) \lambda = 1; (ii) \lambda > 1 \text{ and } \sigma \notin \left[\frac{1}{\lambda}, \lambda \right]; (iii) \lambda < 1 \text{ and } \sigma \notin \left[\lambda, \frac{1}{\lambda} \right].$$

To grasp the result intuitively, recall the receiver's behavior in the single-recommendation system. With a threshold of $R \rightarrow 0$, the receiver buys the recommended product unless it is bad. For $R \rightarrow 1$, he does so only if it is good. In the benchmark model, $V(R)$ is decreasing in R if $\sigma < 1$ and increasing if $\sigma > 1$. Lemma 4 in Appendix A.4 describes optimal receiver behavior with infinite learning. It shows that behavior coincides with the optimally designed single-recommendation system if the product is uncontroversial: Good products are bought but bad ones are not. Thus, the infinite-learning system potentially differs from the benchmark only for controversial products, which generate mixed recommendations. When σ is very low, bad products are relatively likely, so the outside option looks relatively bleak. A consumer facing a product with mixed recommendations will thus buy this (controversial) product. In contrast, when σ is very high, any such consumer will buy the outside option, as it is likely to be good. Thus, when the probabilities of the uncontroversial product versions are not too similar, the behavior with a single recommendation and with multiple recommendations coincides, and multiple recommendations do not increase the value of the recommendation system. However, multiple recommendations do affect behavior if the probabilities of the good and bad products are similar. For instance, when $\lambda > 1$, low types who value the less likely controversial product version $(0, 1)$ more strongly than product version $(1, 0)$ will opt for the outside option rather than the controversial product. Then, the system's value with multiple recommendations is higher than that of the single-recommendation system, as it provides more granular information. Inspecting the conditions on σ in (ii) and (iii), we find that the more extreme (further away from 1) the odds ratio of the controversial products is, the more extreme the odds ratio of the uncontroversial product versions needs to be for there to be no value in additional recommendations.

7 Conclusion

We have set out to study the role of preference heterogeneity in recommendations, asking questions such as: Under what circumstances do people accept a recommendation? What determines the value of a recommendation system and when is it maximized? Our

³⁸In Sun (2011), a higher variance in the recommendations corresponds to a niche product that some consumers like and some don't. Similarly, in our paper, mixed recommendations (and thus a higher variance) result from controversial products that, indeed, some consumers like more than others.

analysis highlights the importance of disentangling objective information from subjective preferences. We have seen that, depending on the preference distribution, the subjective content can reinforce the positive news of a buy recommendation or reverse it so that a buy recommendation is treated as a bad signal. Conversely, the subjective content can turn a don't-buy recommendation into good news for some people. The design of a recommendation system, as summarized in the threshold R , should thus take into account how it influences the relative objectivity of recommendations.

We find that extreme recommendation systems may be valuable, as they convey objective information even when preferences are heterogeneous. However, when preferences are not distributed symmetrically and controversial products are sufficiently prevalent, with enough consensus on their ranking, the optimal recommendation threshold is set at an intermediate level. This allows receivers to not only draw inferences about the uncontroversial products but also about the controversial ones. For asymmetrically distributed populations, subjective content in the recommendation is thus valuable when controversial products are common and not too controversial.

Throughout the paper, we have analyzed the value of the recommendation system from a consumer perspective. Arguably, online recommendation systems employed on an e-commerce platform are more likely to be geared towards the interests of the platform. Depending on the details of the fee system that it employs, the platform could, for instance, be interested in maximizing the probability of purchases. It is possible to look at such questions in the context of our model, too. A naïve conjecture would be that it is then optimal to set as low a recommendation threshold as possible so that all products except the bad one yield buy recommendations. However, recall that it is not necessarily the case that all receivers accept recommendations. The analysis of this problem is then non-trivial.³⁹

Finally, we have maintained a number of classical assumptions such as risk neutrality, linear probability weights, and Bayesian updating. Preliminary results suggest that a promising avenue for future research is to relax these assumptions so as to study their impact on the receiver's tendency to accept a recommendation or not and, thus, on the value of the recommendation system.

³⁹In the case of a symmetric population, in which all recommendations are accepted, the problem is simpler, and a sales-maximizing recommendation system indeed sets as low a threshold as possible. Moreover, our results show that this may even be optimal for consumers. This latter part stands in contrast to, for instance, the welfare-reducing inflated ratings in Peitz and Sobolev (forthcoming).

A Formal Details and Proofs

A.1 Proofs and Technical Details for the Benchmark Model

The Effects of Recommendations on Beliefs

To understand how beliefs are updated after a recommendation, it is useful to decompose the difference between beliefs with and without a recommendation into several components. We illustrate this for buy recommendations. Let $\mathbf{p}(r) = (p_H^r, p_1^r, p_2^r, p_L^r)$ denote the vector of posterior probabilities following a recommendation $r \in \{B, D\}$. The total effect of the buy recommendation B on the probability vectors is the sum of three parts,

$$\mathbf{p}(B) - \mathbf{q} = (\mathbf{p}(B) - \mathbf{q}'') + (\mathbf{q}'' - \mathbf{q}') + (\mathbf{q}' - \mathbf{q}), \quad (9)$$

where the three components of the right-hand side will be explained in Steps 1 to 3 below.

Step 1 The first effect of a buy recommendation is that *the probability of an objectively bad product falls to zero*. We isolate this effect by assuming that the odds ratios between the remaining events remain unchanged so that the effect corresponds to a change from the prior probability \mathbf{q} to $\mathbf{q}' := (\frac{q_H}{1-q_L}, \frac{q_1}{1-q_L}, \frac{q_2}{1-q_L}, 0)$. This effect of the buy recommendation states that all three quality profiles other than $(0, 0)$ become more likely at the expense of quality profile $(0, 0)$.

Step 2 Next, the buy recommendation means that *the objectively good product becomes more likely relative to the controversial products*. We isolate this effect by assuming that (i) the probability of $(1, 1)$ increases to p_H and (ii) the odds ratios between the remaining events remain unchanged so that the effect corresponds to a change from $\mathbf{q}' := (\frac{q_H}{1-q_L}, \frac{q_1}{1-q_L}, \frac{q_2}{1-q_L}, 0)$ to $\mathbf{q}'' := (p_H, k \frac{q_1}{1-q_L}, k \frac{q_2}{1-q_L}, 0)$, where $k = \frac{(1-p_H)(1-q_L)}{q_1+q_2}$ to guarantee that \mathbf{q}'' is a probability vector. Clearly, $p_H > \frac{q_H}{1-q_L}$.

Step 3 Finally, the buy recommendation may change the odds between the controversial product versions. We isolate this effect as a change in the probability vector from $\mathbf{q}'' := (p_H, k \frac{q_1}{1-q_L}, k \frac{q_2}{1-q_L}, 0)$ to $\mathbf{p}(B)$.

Proof of Proposition 1

Part (i): Accepting the don't-buy recommendation is optimal if and only if $U(0) - U(D) \geq 0$ or, equivalently,

$$q_H + q_1(1/2 + i) + q_2(1/2 - i) \geq p_1^D(1/2 + i) + p_2^D(1/2 - i). \quad (10)$$

Similarly, accepting the buy recommendation is optimal if and only if $U(B) - U(0) \geq 0$ or, equivalently,

$$p_H^B + p_1^B(1/2 + i) + p_2^B(1/2 - i) \geq q_H + q_1(1/2 + i) + q_2(1/2 - i). \quad (11)$$

We shall show that conditions (10) and (11) are equivalent, thus proving the claim that receivers either accept all recommendations or never accept recommendations. To do so, we prove the following two steps.

Step 1: (10) is equivalent with

$$1 - p_1^B - p_2^D \geq (1/2 + i)[p_2^B - p_2^D + p_1^D - p_1^B]. \quad (12)$$

Step 2: (11) is equivalent with (12).

The proofs of both steps will rely on the identities

$$q_s = \pi^B p_s^B + \pi^D p_s^D = (1 - \pi^D) p_s^B + \pi^D p_s^D \text{ for } s \in \{H, 1, 2\} \quad (13)$$

and

$$p_H^B = 1 - p_1^B - p_2^B, \quad (14)$$

Proof of Step 1: Replacing q_H , q_1 and q_2 in (10) using (13), we obtain

$$\begin{aligned} & (1 - \pi^D) p_H^B + ((1 - \pi^D) p_1^B + \pi^D p_1^D) \left(\frac{1}{2} + i \right) \\ & + ((1 - \pi^D) p_2^B + \pi^D p_2^D) \left(\frac{1}{2} - i \right) \geq p_1^D \left(\frac{1}{2} + i \right) + p_2^D \left(\frac{1}{2} - i \right). \end{aligned}$$

This simplifies to

$$p_H^B + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(\frac{1}{2} - i \right) \geq \left(p_1^D \left(\frac{1}{2} + i \right) + p_2^D \left(\frac{1}{2} - i \right) \right).$$

Using (14) and then rearranging further gives

$$\begin{aligned} & 1 - p_1^B - p_2^B + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(\frac{1}{2} - i \right) \geq p_1^D \left(\frac{1}{2} + i \right) + p_2^D \left(\frac{1}{2} - i \right) \Leftrightarrow \\ & 1 - p_1^B - p_2^D - p_2^B + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(\frac{1}{2} - i \right) \geq p_1^D \left(\frac{1}{2} + i \right) - p_2^D \left(\frac{1}{2} + i \right) \Leftrightarrow \\ & 1 - p_1^B - p_2^D + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(-\frac{1}{2} - i \right) \geq (p_1^D - p_2^D) \left(\frac{1}{2} + i \right) \Leftrightarrow \\ & 1 - p_1^B - p_2^D \geq (1/2 + i)[p_2^B - p_2^D + p_1^D - p_1^B]. \end{aligned}$$

Proof of Step 2: Replacing q_H and q_1 and q_2 in (11) using (13) gives

$$\begin{aligned} & p_H^B + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(\frac{1}{2} - i \right) \geq \\ & (1 - \pi^D) p_H^B + ((1 - \pi^D) p_1^B + \pi^D p_1^D) \left(\frac{1}{2} + i \right) + ((1 - \pi^D) p_2^B + \pi^D p_2^D) \left(\frac{1}{2} - i \right). \end{aligned}$$

Rearranging this expression gives

$$p_H^B + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(\frac{1}{2} - k \right) - p_1^D \left(\frac{1}{2} + i \right) - p_2^D \left(\frac{1}{2} - i \right) \geq 0$$

Adding $-p_1^B - p_2^D + p_2^D + p_1^B$ on the left-hand side gives:

$$p_H^B - p_1^B - p_2^D + p_2^D + p_1^B + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(\frac{1}{2} - i \right) - p_1^D \left(\frac{1}{2} + i \right) - p_2^D \left(\frac{1}{2} - i \right) \geq 0$$

and thus

$$p_H^B - p_1^B - p_2^D \geq -p_2^D - p_1^B - p_1^B \left(\frac{1}{2} + i \right) - p_2^B \left(\frac{1}{2} - i \right) + p_1^D \left(\frac{1}{2} + i \right) + p_2^D \left(\frac{1}{2} - i \right).$$

Using (14) gives

$$\begin{aligned} & 1 - p_1^B - p_2^B - p_1^B - p_2^D \\ & \geq -p_2^D - p_1^B - p_1^B \left(\frac{1}{2} + i\right) - p_2^B \left(\frac{1}{2} - i\right) + p_1^D \left(\frac{1}{2} + i\right) + p_2^D \left(\frac{1}{2} - i\right) \end{aligned}$$

or equivalently

$$1 - p_1^B - p_2^B \geq -p_1^B \left(\frac{1}{2} + i\right) - p_2^B \left(\frac{1}{2} - i\right) + p_1^D \left(\frac{1}{2} + i\right) + p_2^D \left(\frac{1}{2} - i\right)$$

Adding $-p_2^D + p_2^B$ on both sides gives:

$$\begin{aligned} & 1 - p_1^B - p_2^D - p_2^B + p_2^B \\ & \geq -p_1^B \left(\frac{1}{2} + i\right) - p_2^B \left(\frac{1}{2} - i\right) + p_1^D \left(\frac{1}{2} + i\right) + p_2^D \left(\frac{1}{2} - i\right) - p_2^D + p_2^B, \end{aligned}$$

which is equivalent with (12).

Part (ii): This follows directly from plugging in Δ_O^r and Δ_S^r into equation (1) for the buy recommendation the analogous equation for a don't-buy recommendation.

Proof of Proposition 2

Recall from equation (1) the optimal acceptance condition

$$(p_H^B - q_H) + \frac{p_1^B - q_1}{2} + \frac{p_2^B - q_2}{2} \geq i[(p_2^B - q_2) - (p_1^B - q_1)].$$

First, observe that

$$\lim_{R \rightarrow 1} p_1^B = \lim_{R \rightarrow 1} p_2^B = 0$$

and thus $\lim_{R \rightarrow 1} p_H^B = 1$. Hence, in the case $R \rightarrow 1$ the optimality condition simplifies to

$$1 - q_H - \frac{q_1 + q_2}{2} \geq i(q_1 - q_2). \quad (15)$$

Equation (15) always holds for $i \geq 0$. Further, for $q_1 > q_2$ the condition is most stringent for $i = 1/2$. In this case, it is equivalent to $1 \geq q_1 + q_H$, which always holds. The case $i < 0$ is analogous. In the case $R \rightarrow 0$ we get

$$\lim_{R \rightarrow 0} p_s(B) = \frac{q_s}{1 - q_L}$$

for $s \in \{H, 1, 2\}$ and thus the optimality condition simplifies to

$$\begin{aligned} & \frac{q_H q_L}{1 - q_L} + \frac{q_1 q_L}{2(1 - q_L)} + \frac{q_2 q_L}{2(1 - q_L)} \geq i \left[\frac{q_2 q_L}{1 - q_L} - \frac{q_1 q_L}{1 - q_L} \right] \\ & \Leftrightarrow 2q_H + q_1 + q_2 \geq 2i[q_2 - q_1]. \end{aligned}$$

For $i \geq 0$, the last condition always holds if $q_1 \geq q_2$. Further, for $q_2 > q_1$ the condition is most stringent for $i = 1/2$. In this case, it is equivalent to $q_H + q_1 \geq 0$, which is always satisfied. The case $i < 0$ is analogous.

Proof of Proposition 3

(i) If $|\Delta_S^B| \leq 2\Delta_O^B$, inserting (5) and (6) yields

$$\begin{aligned}
V(R) &= \int_{-1/2}^{1/2} (V^A(i) - V_0(i))dF(i) \\
&= \pi^B \int_{-1/2}^{1/2} (p_H^B + (1/2 + i)p_1^B + (1/2 - i)p_2^B) - (q_H + (1/2 + i)q_1 + (1/2 - i)q_2)dF(i) \\
&= \pi^B \int_{-1/2}^{1/2} \underbrace{(p_H^B - q_H) + \frac{p_1^B - q_1}{2} + \frac{p_2^B - q_2}{2}}_{=\Delta_O^B} dF(i) \\
&\quad - \pi^B \int_{-1/2}^{1/2} i \underbrace{((p_2^B - q_2) - (p_1^B - q_1))}_{=\Delta_S^B} dF(i) \\
&= \pi^B [\Delta_O^B - \Delta_S^B \mathbb{E}[i]]
\end{aligned}$$

(ii) If $\Delta_S^B < -2\Delta_O^B$, inserting (5), (6) and (7) yields

$$\begin{aligned}
V(R) &= \int_{-1/2}^{\tilde{i}} V^N(i)dF(i) + \int_{\tilde{i}}^{1/2} V^A(i)dF(i) - \int_{-1/2}^{1/2} V^0(i)dF(i) \\
&= \int_{-1/2}^{\tilde{i}} \pi^B \left(q_H + q_1 \left(\frac{1}{2} + i \right) + q_2 \left(\frac{1}{2} - i \right) \right) + (1 - \pi^B) \left(p_H^D + p_1^D \left(\frac{1}{2} + i \right) + p_2^D \left(\frac{1}{2} - i \right) \right) dF(i) \\
&\quad + \int_{\tilde{i}}^{1/2} \pi^B \left(p_H^B + p_1^B \left(\frac{1}{2} + i \right) + p_2^B \left(\frac{1}{2} - i \right) \right) + (1 - \pi^B) \left(q_H + q_1 \left(\frac{1}{2} + i \right) + q_2 \left(\frac{1}{2} - i \right) \right) dF(i) \\
&\quad - \int_{-1/2}^{\tilde{i}} q_H + q_1 \left(\frac{1}{2} + i \right) + q_2 \left(\frac{1}{2} - i \right) dF(i) - \int_{\tilde{i}}^{1/2} q_H + q_1 \left(\frac{1}{2} + i \right) + q_2 \left(\frac{1}{2} - i \right) dF(i) \\
&= (1 - \pi^B) \int_{-1/2}^{\tilde{i}} \underbrace{(p_H^D - q_H) + \frac{p_1^D - q_1}{2} + \frac{p_2^D - q_2}{2}}_{=\Delta_O^D} - i \underbrace{((p_2^D - q_2) - (p_1^D - q_1))}_{=\Delta_S^D} dF(i) \\
&\quad + \pi^B \int_{\tilde{i}}^{1/2} \underbrace{(p_H^B - q_H) + \frac{p_1^B - q_1}{2} + \frac{p_2^B - q_2}{2}}_{=\Delta_O^B} - i \underbrace{((p_2^B - q_2) - (p_1^B - q_1))}_{=\Delta_S^B} dF(i) \\
&= (1 - \pi^B) F(\tilde{i}) [\Delta_O^D - \Delta_S^D \mathbb{E}[i | i \leq \tilde{i}]] + \pi^B (1 - F(\tilde{i})) [\Delta_O^B - \Delta_S^B \mathbb{E}[i | i \geq \tilde{i}]].
\end{aligned}$$

(iii) If $\Delta_S^B > 2\Delta_O^B$, then $V(R) = \int_{-1/2}^{\tilde{i}} V^A(i)dF(i) + \int_{\tilde{i}}^{1/2} V^N(i)dF(i) - \int_{-1/2}^{1/2} V^0(i)dF(i)$. Inserting (5), (6) and (7) and proceeding analogously as in case (ii) gives the expression.

Proof of Proposition 4

According to Corollary 1, all types accept the recommendation if

$$\min \left\{ \Delta_O^B + \frac{\Delta_S^B}{2}, \Delta_O^B - \frac{\Delta_S^B}{2} \right\} \geq 0. \tag{16}$$

To see that this condition is fulfilled, first observe that

$$\begin{aligned}
\Delta_O^B + \frac{\Delta_S^B}{2} &= p_H^B - q_H + p_2^B - q_2 \\
&= \frac{q_H}{q_H + \beta(q_1 + q_2)} - q_H + \frac{\beta q_2}{q_H + \beta(q_1 + q_2)} - q_2 \\
&= \frac{q_H + \beta q_2 - (q_H + q_2)(q_H + \beta(q_1 + q_2))}{q_H + \beta(q_1 + q_2)} \\
&= \frac{\beta(q_2 - (q_H + q_2)(q_1 + q_2)) + q_H(1 - q_H - q_2)}{q_H + \beta(q_1 + q_2)}
\end{aligned}$$

This term is positive if $q_2 - (q_H + q_2)(q_1 + q_2) \geq 0$. If $q_2 - (q_H + q_2)(q_1 + q_2) < 0$, then

$$\begin{aligned}
&\frac{\beta(q_2 - (q_H + q_2)(q_1 + q_2)) + q_H(1 - q_H - q_2)}{q_H + \beta(q_1 + q_2)} \\
&\geq \frac{(q_2 - (q_H + q_2)(q_1 + q_2)) + q_H(1 - q_H - q_2)}{q_H + \beta(q_1 + q_2)} = \frac{(q_H + q_2)(1 - q_1 - q_2 - q_H)}{q_H + \beta(q_1 + q_2)} \geq 0.
\end{aligned}$$

Thus, $\Delta_O^B + \frac{\Delta_S^B}{2} \geq 0$ even in this case. Proceeding analogously, one obtains $\Delta_O^B - \frac{\Delta_S^B}{2} = p_H^B - q_H + p_1^B - q_1 \geq 0$. Corollary 1 thus implies that all types accept the recommendation. Assumption 1 implies $\mathbb{E}[i] = 0$, so that the result follows from Proposition 3.

Proof of Lemma 1

The result on π^B follows from inserting $q_H = (1 - 2Q)\sigma/(1 + \sigma)$, $\phi_1(R) = \phi_2(R) = \beta$ and $q_1 + q_2 = 2Q$ into $\pi^B = q_H + q_1\phi_1(R) + q_2\phi_2(R)$. Next, equation (2) gives

$$\Delta_O^B = \frac{\beta Q(\sigma + 1) - 2Q\sigma + \sigma}{2\beta Q(\sigma + 1) - 2Q\sigma + \sigma} - \frac{(1 - 2Q)\sigma}{1 + \sigma} - Q \tag{17}$$

The result on $V(\beta)$ then follows immediately from Proposition 4.

Proof of Lemma 2

(i) Note that $\pi^B = (1 - 2Q)\sigma/(1 + \sigma) + 2Q\beta$, so that π_B is strictly increasing. Moreover,

$$\frac{\partial \Delta_O^B}{\partial \beta} = -\frac{(Q(1 - 2Q)\sigma(1 + \sigma))}{(\sigma - 2Q\sigma + 2\beta Q(1 + \sigma))^2} < 0.$$

(ii) Lemma 1 implies

$$V'(\beta) = \frac{Q(1 - 2Q)(1 - \sigma)}{1 + \sigma},$$

and hence

$$\frac{\partial^2 V}{\partial \beta \partial \sigma} = -\frac{2Q(1 - 2Q)}{(1 + \sigma)^2} < 0,$$

$V'(\beta)$ is negative for $\sigma > 1$, positive for $\sigma < 1$ and zero for $\sigma = 1$.

Proof of Corollary 2

Using (3), applying the q_H notation for ease of exposition and imposing Assumption 2 and $q_1 = q_2$ we obtain

$$\Delta_S^B = -\frac{Q(1 - R^a - (1 - R)^a)}{q_H + Q(1 - R^a + (1 - R)^a)}. \quad (18)$$

For $a = 1$, this term simplifies to $\Delta_S^B = 0$. Hence, it suffices to show that Δ_S^B is strictly increasing in a . To this end, note that the derivative of Δ_S^B with respect to a is given by

$$\frac{\partial \Delta_S^B}{\partial a} = \frac{Q[\log(R)R^a(2Q(1 - R)^a + q_H) + \log(1 - R)(1 - R)^a(2Q(1 - R^a) + q_H)]}{(Q(1 - R^a + (1 - R)^a) + q_H)^2}$$

Since $R \in (0, 1)$ this expression is negative. Further, given the assumptions, the indifferent type can be expressed as (for $a \neq 1$)

$$\begin{aligned} \tilde{i} &= \frac{2Q^2(1 - R^a + (1 - R)^a) - 2(1 - q_H)q_H}{2Q(1 - R^a - (1 - R)^a)} \\ &\quad - \frac{Q(1 - R^a + (1 - R)^a - 2q_H(2 - R^a + (1 - R)^a))}{2Q(1 - R^a - (1 - R)^a)} \end{aligned}$$

where

$$\frac{\partial \tilde{i}}{\partial Q} = \frac{Q^2(1 - R^a + (1 - R)^a) + (1 - q_H)q_H}{Q^2(1 - R^a - (1 - R)^a)}$$

so that \tilde{i} is increasing in Q when $a > 1$ and decreasing in Q when $a < 1$.

Proof of Proposition 6

We provide the proof for $a < 1$.⁴⁰ In this case, Proposition 3 implies

$$V(R) = \begin{cases} \pi^B [\Delta_O^B - \Delta_S^B \mathbb{E}[i]] & \text{if } \tilde{i} > 1/2 \\ \pi^B F(\tilde{i}) [\Delta_O^B - \Delta_S^B \mathbb{E}[i \mid i \leq \tilde{i}]] & \text{if } \tilde{i} \leq 1/2 \\ +(1 - \pi^B)(1 - F(\tilde{i})) [\Delta_O^D - \Delta_S^D \mathbb{E}[i \mid i \geq \tilde{i}]] & \end{cases}$$

Proposition 2 implies $\tilde{i} > 1/2$ if $R \rightarrow 0$ or $R \rightarrow 1$. Thus, it suffices to consider

$$\begin{aligned} V(R) &= (q_H + Q(1 - R^a + (1 - R)^a)) \left(\frac{2q_H + Q(1 - R^a + (1 - R)^a)}{2(q_H + Q(1 - R^a + (1 - R)^a))} - q_H - Q \right) \\ &\quad + (q_H + Q(1 - R^a + (1 - R)^a)) \frac{Q(1 - R^a - (1 - R)^a)}{q_H + Q(1 - R^a + (1 - R)^a)} \frac{a - 1}{2(a + 1)} \\ &= \left(\frac{2q_H + Q(1 - R^a + (1 - R)^a)}{2} - (q_H + Q(1 - R^a + (1 - R)^a))(q_H + Q) \right) \\ &\quad + Q(1 - R^a - (1 - R)^a) \frac{a - 1}{2(a + 1)} \end{aligned}$$

⁴⁰The proof for $a > 1$ is analogous.

for R sufficiently close to 0 and 1. Using $q_H = \frac{(1-2Q)\sigma}{1+\sigma}$ and rearranging gives

$$V(R) = \frac{(1-2Q)\sigma}{1+\sigma} + \frac{aQ(1-R^a) + Q(1-R)^a}{a+1} \\ - \left(\frac{(1-2Q)\sigma}{1+\sigma} + Q(1-R^a) + Q(1-R)^a \right) \left(\frac{(1-2Q)\sigma}{1+\sigma} + Q \right)$$

We then obtain the corresponding derivative

$$V'(R) = \frac{aQ \left((a+1)(R^{a-1} + (1-R)^{a-1}) \left(Q + \frac{(1-2Q)\sigma}{1+\sigma} \right) - aR^{a-1} - (1-R)^{a-1} \right)}{a+1}$$

We can then write

$$V'(R) > 0 \\ \Leftrightarrow (a+1) \left(R^{a-1} + (1-R)^{a-1} \right) \left(Q + \frac{(1-2Q)\sigma}{1+\sigma} \right) - aR^{a-1} - (1-R)^{a-1} > 0 \\ \Leftrightarrow (a+1) \left(Q + \frac{(1-2Q)\sigma}{1+\sigma} \right) > \frac{aR^{a-1} + (1-R)^{a-1}}{R^{a-1} + (1-R)^{a-1}} = \frac{a(1-R)^{1-a} + R^{1-a}}{(1-R)^{1-a} + R^{1-a}},$$

so that

$$\lim_{R \rightarrow 0} V'(R) > 0 \Leftrightarrow (a+1) \left(Q + \frac{(1-2Q)\sigma}{1+\sigma} \right) > a$$

and analogously

$$\lim_{R \rightarrow 1} V'(R) < 0 \Leftrightarrow (a+1) \left(Q + \frac{(1-2Q)\sigma}{1+\sigma} \right) < 1.$$

Thus, a sufficient condition for an interior solution when $a < 1$ is given by

$$\frac{1}{a+1} > \left(Q + \frac{(1-2Q)\sigma}{1+\sigma} \right) > \frac{a}{a+1}. \quad (19)$$

Rearranging the first inequality, we obtain

$$\left(\frac{1}{a+1} - Q \right) (1+\sigma) > (1-2Q)\sigma \\ \Leftrightarrow \frac{1}{a+1} - Q > (1-2Q)\sigma - \left(\frac{1}{a+1} - Q \right) \sigma \\ \Leftrightarrow \frac{1}{a+1} - Q > \sigma \left(\frac{a}{1+a} - Q \right)$$

For $a < 1$, the condition $Q \geq 1 - \max\{\frac{a}{a+1}, \frac{1}{a+1}\}$ is equivalent to $Q \geq \frac{a}{1+a}$ so that the right-hand side is negative while the left-hand side is positive. Hence, this inequality is always satisfied. Proceeding analogously for the second inequality, we obtain

$$\left(\frac{a}{a+1} - Q \right) (1+\sigma) < (1-2Q)\sigma,$$

which is also implied by $Q \geq 1 - \max\{\frac{a}{a+1}, \frac{1}{a+1}\}$. Hence, the first statement of Proposition 6 for the case $a < 1$ follows. For the second statement when $Q < 1 - \max\{\frac{a}{a+1}, \frac{1}{a+1}\}$, we can rewrite the two inequalities in (19) as

$$\frac{1 - Q(a+1)}{a - Q(a+1)} > \sigma > \frac{a - Q(a+1)}{1 - Q(a+1)}$$

giving the second condition in the statement.

Proof of Proposition 7

Using (17) and (18), we obtain

$$\lim_{Q \rightarrow 0} \Delta_S^B = \lim_{Q \rightarrow 0} -\frac{Q(1 - R^a - (1 - R)^a)}{q_H + Q(1 - R^a + (1 - R)^a)} = 0,$$

and

$$\lim_{Q \rightarrow 0} \Delta_O^B = \lim_{Q \rightarrow 0} \frac{2q_H + Q(1 - R^a + (1 - R)^a)}{2(q_H + Q(1 - R^a + (1 - R)^a))} - q_H - Q = 1 - q_H > 0.$$

Therefore, $2\Delta_O^B \geq |\Delta_S^B|$ so that all types accept a recommendation according to Corollary 1. Thus, Proposition 3 implies

$$V(R) = \pi^B[\Delta_O^B - \Delta_S^B \mathbb{E}[i]]$$

Rearranging this equation gives

$$\begin{aligned} V(R) &= \frac{(1 - 2Q)\sigma}{1 + \sigma} + \frac{aQ(1 - R^a) + Q(1 - R)^a}{a + 1} \\ &\quad - \left(\frac{(1 - 2Q)\sigma}{1 + \sigma} + Q(1 - R^a) + Q(1 - R)^a \right) \left(\frac{(1 - 2Q)\sigma}{1 + \sigma} + Q \right) \\ &= \frac{(1 - 2Q)\sigma}{1 + \sigma} \left(1 - \left(\frac{(1 - 2Q)\sigma}{1 + \sigma} + Q \right) \right) \\ &\quad + \frac{Q}{a + 1} (a(1 - R^a) + (1 - R)^a) \\ &\quad - \frac{Q}{a + 1} ((1 - R^a) + (1 - R)^a) (a + 1) \left(\frac{(1 - 2Q)\sigma}{1 + \sigma} + Q \right) \\ &= \frac{(1 - 2Q)\sigma(Q(\sigma - 1) + 1)}{(\sigma + 1)^2} \\ &\quad + \frac{Q}{a + 1} (1 - R^a) \left(a - (a + 1) \left(\frac{(1 - 2Q)\sigma}{1 + \sigma} + Q \right) \right) \\ &\quad + \frac{Q}{a + 1} (1 - R)^a \left(1 - (a + 1) \left(\frac{(1 - 2Q)\sigma}{1 + \sigma} + Q \right) \right) \\ &= c_0 + \frac{Q}{a + 1} [c_1(a)(1 - R^a) + c_2(a)(1 - R)^a] \end{aligned}$$

with $c_0 = \frac{(1-2Q)\sigma(Q(\sigma-1)+1)}{(\sigma+1)^2}$, $c_1(a) = a + (a+1)\frac{Q(\sigma-1)-\sigma}{\sigma+1}$ and $c_2(a) = 1 + (1+a)\frac{Q(\sigma-1)-\sigma}{\sigma+1}$. We thus obtain

$$V'(R) = -\frac{Qa}{a+1} [R^{a-1}c_1(a) + (1-R)^{a-1}c_2(a)],$$

$$V''(R) = -\frac{Qa(a-1)}{a+1} [R^{a-2}c_1(a) - (1-R)^{a-2}c_2(a)],$$

For $c_1(a), c_2(a) \geq 0$ and $c_1(a), c_2(a) \leq 0$ the function $V(R)$ is monotonically decreasing and increasing, respectively. Further, for $a \geq 1$, $c_1(a) > 0$ and $c_2(a) < 0$ it is concave. Conversely, for $a < 1$ and $c_1(a) < 0$ and $c_2(a) > 0$ it is concave. Further,

$$c_1(a) - c_2(a) \begin{cases} < 0 & a < 1 \\ = 0 & a = 1 \\ > 0 & a > 1, \end{cases}$$

Therefore, for $a < 1$ either (i) $c_1(a), c_2(a) \leq 0$, (ii) $c_1(a) \leq 0 \leq c_2(a)$, or (iii) $0 \leq c_1(a), c_2(a)$. Hence, in all three cases $V(R)$ is quasiconcave. Similarly, for $a > 1$ either (i) $c_1(a), c_2(a) \leq 0$, (ii) $c_2(a) \leq 0 \leq c_1(a)$, or (iii) $0 \leq c_1(a), c_2(a)$. Hence, in all three cases $V(R)$ is quasiconcave. Finally, for $a = 1$ quasiconcavity follows because $c_1(a) = c_2(a)$ and thus V is monotone. Putting all of this together, it follows that V is quasiconcave in R . Hence, the sufficient conditions in Proposition 6 for an interior solution are also necessary.

Proof of Proposition 8

As argued in the text, a mean-preserving spread corresponds to an increase in β for $R > 1/2$, while for $R < 1/2$ it corresponds to a decrease in β . From the proof of Lemma 2, we can see that the value of the recommendation system is increasing in β for $\sigma < 1$ and decreasing in β for $\sigma > 1$. Thus, the statement follows.

Proof of Corollary 3

We first take the derivative of V with respect to Q to obtain

$$\frac{\partial V}{\partial Q} = \frac{1}{(\sigma+1)^2} (\beta - 3\sigma + 4Q\sigma - 4Q\beta + \sigma^2 - 4Q\sigma^2 - \sigma^2\beta + 4Q\sigma^2\beta)$$

Setting this equal to zero and solving for the candidate solution we obtain

$$Q^* = \frac{3\sigma - \beta - \sigma^2 + \sigma^2\beta}{4\sigma - 4\beta - 4\sigma^2 + 4\sigma^2\beta}$$

Next, the SOC reads

$$4\sigma - 4\beta - 4\sigma^2 + 4\sigma^2\beta < 0,$$

where the l.h.s. of the SOC is the denominator of Q^* . Thus, a necessary condition for a positive interior solution is

$$3\sigma - \beta - \sigma^2 + \sigma^2\beta < 0.$$

This proves (i). To prove (ii), note that $3\sigma - \beta - \sigma^2 + \sigma^2\beta < 0$ implies $4\sigma - 4\beta - 4\sigma^2 + 4\sigma^2\beta < 0$. Hence, the candidate solution is positive. It remains to show that $Q^* < \frac{1}{2}$. Recall that, by the requirements that the SOC holds and $Q^* > 0$, we can focus on the case that both the numerator and the denominator in Q^* are negative. Accordingly, the requirement that $\frac{3\sigma - \beta - \sigma^2 + \sigma^2\beta}{4\sigma - 4\beta - 4\sigma^2 + 4\sigma^2\beta} < 1/2$ becomes

$$2\sigma - 2\beta - 2\sigma^2 + 2\sigma^2\beta < 3\sigma - \beta - \sigma^2 + \sigma^2\beta.$$

Equivalently,

$$(\sigma + 1)(\sigma(1 - \beta) + \beta) > 0$$

which always holds for $\beta \leq 1$.

A.2 Distinct Sender and Receiver Distributions

Proof of Proposition 9

The value of the recommendation system can be written as

$$V(\mathcal{R}) = \begin{cases} \int_{-1/2}^{1/2} (V_{A,F}(i) - V_{0,F}(i))dG(i) & \text{if } |\Delta_{S,F}^B| \leq 2\Delta_{0,F}^B \\ \int_{\tilde{i}}^{1/2} V_{N,F}(i)dG(i) + \int_{\tilde{i}}^{1/2} V_{A,F}(i)dG(i) - \int_{-1/2}^{1/2} V_{0,F}(i)dG(i) & \text{if } \Delta_{S,F}^B < -2\Delta_{0,F}^B \\ \int_{-1/2}^{\tilde{i}} V_{A,F}(i)dG(i) + \int_{\tilde{i}}^{1/2} V_{N,F}(i)dG(i) - \int_{-1/2}^{1/2} V_{0,F}(i)dG(i) & \text{if } \Delta_{S,F}^B > 2\Delta_{0,F}^B \end{cases}$$

and the result from Proposition 3 translates accordingly so that we can rewrite $V(\mathcal{R})$ as

$$\begin{aligned} & \pi^B[\Delta_{O,F}^B - \Delta_{S,F}^B \mathbb{E}_G[i]] && \text{if } |\Delta_{S,F}^B| \leq 2\Delta_{O,F}^B \\ & (1 - \pi^B)G(\tilde{i}) [\Delta_O^D - \Delta_{S,F}^D \mathbb{E}_G[i \mid i \leq \tilde{i}]] + \pi^B(1 - G(\tilde{i})) [\Delta_{O,F}^B - \Delta_{S,F}^B \mathbb{E}_G[i \mid i \geq \tilde{i}]] && \text{if } \Delta_{S,F}^B < -2\Delta_{O,F}^B \\ & \pi^B G(\tilde{i}) [\Delta_{O,F}^B - \Delta_{S,F}^B \mathbb{E}_G[i \mid i \leq \tilde{i}]] + (1 - \pi^B)(1 - G(\tilde{i})) [\Delta_O^D - \Delta_{S,F}^D \mathbb{E}_G[i \mid i \geq \tilde{i}]] && \text{if } \Delta_{S,F}^B > 2\Delta_{O,F}^B. \end{aligned}$$

Since F satisfies Assumption 1, all receivers will accept the recommendation so that we are in the case $|\Delta_{S,F}^B| \leq 2\Delta_{O,F}^B$. Thus, we have

$$\begin{aligned} V(\beta) &= \pi^B[\Delta_{O,F}^B - \Delta_{S,F}^B \mathbb{E}_G[i]] \\ &= q_H + \beta Q - (q_H + 2\beta Q)(q_H + Q) + (q_2 - q_1)(q_H - \beta(1 - 2Q))\mathbb{E}_G[i] \end{aligned}$$

and therefore

$$V'(\beta) = Q - 2Q(q_H + Q) + (q_2 - q_1)(1 - 2Q)\mathbb{E}_G[i]$$

Replacing $q_H = ((1 - 2Q)\sigma)/(\sigma + 1)$ and rearranging the inequalities $V'(\beta) > 0$ and $V'(\beta) < 0$ yields the statement in the proposition.

A.3 Multiple Recommendation Levels

Proof of Proposition 10

The posteriors are given as follows. Following a buy recommendation,

$$\begin{aligned} p_H^B(R) &= \frac{q_H}{q_H + q_1(1 - F(R_2 - 1/2)) + q_2F(1/2 - R_2)}, \\ p_1^B(R) &= \frac{q_1(1 - F(R_2 - 1/2))}{q_H + q_1(1 - F(R_2 - 1/2)) + q_2F(1/2 - R_2)}, \\ p_2^B(R) &= \frac{q_2F(1/2 - R_2)}{q_H + q_1(1 - F(R_2 - 1/2)) + q_2F(1/2 - R_2)} \\ p_L^B(R) &= 0. \end{aligned}$$

A don't-buy recommendation results in a posterior of

$$\begin{aligned} p_H^D(R) &= 0, \\ p_1^D(R) &= \frac{q_1F(R_1 - 1/2)}{q_1F(R_1 - 1/2) + q_2(1 - F(1/2 - R_1)) + q_L}, \\ p_2^D(R) &= \frac{q_2(1 - F(1/2 - R_1))}{q_1F(R_1 - 1/2) + q_2(1 - F(1/2 - R_1)) + q_L} \\ p_L^D(R) &= \frac{q_L}{q_1F(R_1 - 1/2) + q_2(1 - F(1/2 - R_1)) + q_L}. \end{aligned}$$

Finally, the recommendation neither to buy nor not to buy yields

$$\begin{aligned} p_H^N(R) &= 0 & p_2^N(R) &= \frac{q_2\Gamma_2}{q_1\Gamma_1 + q_2\Gamma_2} \\ p_1^N(R) &= \frac{q_1\Gamma_1}{q_1\Gamma_1 + q_2\Gamma_2} & p_L^N(R) &= 0. \end{aligned}$$

where $\Gamma_1 \equiv F(R_2 - 1/2) - F(R_1 - 1/2)$ and $\Gamma_2 \equiv F(1/2 - R_1) - F(1/2 - R_2)$. For the symmetric case of Assumption 1, setting $F(1/2 - R_2) = 1 - F(R_2 - 1/2) = \beta_2$ and $F(1/2 - R_1) = 1 - F(R_1 - 1/2) = \beta_1$, this reduces to

$$\begin{aligned} p_H^B &= \frac{q_H}{q_H + 2Q\beta_2}, & p_H^N &= 0, & p_H^D &= 0, \\ p_1^B &= \frac{q_1\beta_2}{q_H + 2Q\beta_2}, & p_1^N &= \frac{q_1}{q_1 + q_2}, & p_1^D &= \frac{q_1(1 - \beta_1)}{1 - q_H - 2Q\beta_1}, \\ p_2^B &= \frac{q_2\beta_2}{q_H + 2Q\beta_2}, & p_2^N &= \frac{q_2}{q_1 + q_2}, & p_2^D &= \frac{q_2(1 - \beta_1)}{1 - q_H - 2Q\beta_1}, \\ p_L^B &= 0, & p_L^N &= 0, & p_L^D &= \frac{q_L}{1 - q_H - 2Q\beta_1}. \end{aligned}$$

Next, we determine the indifferent type following an intermediate recommendation. Type i will buy the product with the intermediate recommendation whenever

$$\begin{aligned} \frac{q_1}{q_1 + q_2}(1/2 + i) + \frac{q_2}{q_1 + q_2}(1/2 - i) &\geq q_H + q_1(1/2 + i) + q_2(1/2 - i) \\ \Leftrightarrow i \left(\frac{(q_1 - q_2)(q_H + q_L)}{q_1 + q_2} \right) &\geq \frac{1}{2}(q_H - q_L) \end{aligned}$$

If $q_1 = q_2$, all types buy if $q_L \geq q_H$ and none buy otherwise. If $q_1 > q_2$, all types

$$i \geq \frac{1}{2} \frac{(q_H - q_L)(q_1 + q_2)}{(q_1 - q_2)(q_H + q_L)}$$

buy the product if $r = N$.⁴¹ Taken together, we can define

$$\tilde{i}_M := \begin{cases} 1/2 & q_1 = q_2, q_L < q_H \\ -1/2 & q_1 = q_2, q_L \geq q_H \\ \frac{1}{2} \frac{(q_H - q_L)(q_1 + q_2)}{(q_1 - q_2)(q_H + q_L)} & q_1 \neq q_2. \end{cases}$$

The value of the recommendation system therefore reads

$$\begin{aligned} V(\beta_2, \beta_1) = & (q_H + 2Q\beta_2) \left(\frac{q_H + Q\beta_2}{q_H + 2Q\beta_2} - (q_H + Q) \right) \\ & + 2Q(\beta_1 - \beta_2) \int_{\tilde{i}_M}^{1/2} 1/2(q_L - q_H) + i \frac{(1 - 2Q)(q_1 - q_2)}{2Q} dF(i), \end{aligned}$$

where the first part corresponds to the value of accepting the highest recommendation and the second part comes from buying following an intermediate recommendation. Depending on parameters, \tilde{i}_M changes so that some, all or no types buy following an intermediate recommendation. Taking the respective derivatives we obtain

$$\begin{aligned} \frac{\partial V}{\partial \beta_2} &= Q(q_L - q_H)F(\tilde{i}_M) - (1 - 2Q)(q_1 - q_2)(1 - F(\tilde{i}_M))\mathbb{E}[i \mid i \geq \tilde{i}_M] \\ \frac{\partial V}{\partial \beta_1} &= Q(q_L - q_H)(1 - F(\tilde{i}_M)) + (1 - 2Q)(q_1 - q_2)(1 - F(\tilde{i}_M))\mathbb{E}[i \mid i \geq \tilde{i}_M] \end{aligned}$$

Suppose $\sigma > 1$ so that $q_L < q_H$. Then, $\frac{\partial V}{\partial \beta_2} < 0$. Next, consider $\sigma < 1$ so that $q_L > q_H$. Then, $\frac{\partial V}{\partial \beta_1} > 0$.

Finally, suppose $q_1 = q_2$. If $q_L < q_H$, then $\tilde{i}_M = 1/2$, so that the derivatives read

$$\begin{aligned} \frac{\partial V}{\partial \beta_2} &= Q(q_L - q_H) < 0 \\ \frac{\partial V}{\partial \beta_1} &= 0. \end{aligned}$$

For $q_L > q_H$, we obtain $\tilde{i}_M = -1/2$, so that the derivatives read

$$\begin{aligned} \frac{\partial V}{\partial \beta_2} &= 0 \\ \frac{\partial V}{\partial \beta_1} &= Q(q_L - q_H) > 0. \end{aligned}$$

Thus, for $q_L < q_H$ any β_1 and thus any R_1 is optimal, whereas for $q_L > q_H$ any β_2 and thus any R_2 is optimal.

⁴¹The argument for $q_1 < q_2$ is analogous.

A.4 Multiple Recommendations

After b buy recommendations and d don't-buy recommendations, where $b + d > 0$, the posterior reads

$$\begin{aligned}
 p_H(b, d) &= \begin{cases} \frac{q_H}{q_H + q_1(\phi_1(R))^b + q_2(\phi_2(R))^b} & \text{if } d = 0 \\ 0 & \text{if } b = 0 \end{cases} \\
 p_1(b, d) &= \frac{q_1(\phi_1(R))^b(1 - \phi_1(R))^d}{q_1(\phi_1(R))^b(1 - \phi_1(R))^d + q_2(\phi_2(R))^b(1 - \phi_2(R))^d} \\
 p_2(b, d) &= \frac{q_2(\phi_2(R))^b(1 - \phi_2(R))^d}{q_1(\phi_1(R))^b(1 - \phi_1(R))^d + q_2(\phi_2(R))^b(1 - \phi_2(R))^d} \\
 p_L(b, d) &= \begin{cases} \frac{q_L}{q_1(1 - \phi_1(R))^d + q_2(1 - \phi_2(R))^d + q_L} & \text{if } b = 0 \\ 0 & \text{if } d = 0 \end{cases}
 \end{aligned}$$

Lemma 3 *Consider the model with multiple recommendations. Suppose Assumption 1 holds and consider any threshold $R \in (0, 1)$.*

- (i) *If the receiver obtains only buy (don't buy) recommendations, then, as the number of recommendations increases, her posterior belief converges to $p_H = 1$ ($p_L = 1$).*
- (ii) *If the receiver obtains mixed recommendations, then $p_1(b, d)/p_2(b, d) = q_1/q_2$ and $p_H(b, d) = p_L(b, d) = 0$ for any $b, d > 0$.*

Proof. (i) Suppose the receiver gets only buy recommendations and that $b \rightarrow \infty$. Then, $\phi_1(R) = \phi_2(R) \in (0, 1)$ if $R \in (0, 1)$, so that $\lim_{b \rightarrow \infty} (\phi_i(R))^b = 0$ for $i = 1, 2$. Hence, $\lim_{b \rightarrow \infty} p_H(b, 0) = 1$. The argument for don't-buy recommendations is analogous.

(ii) For any $b, d > 0$, $p_H(b, d) = p_L(b, d) = 0$ obviously holds. Further, because $\phi_1(R) = \phi_2(R)$ by Assumption 1,

$$p_i(b, d) = \frac{q_i}{q_1 + q_2}$$

for $i = 1, 2$, yielding statement (ii). ■

Optimal Receiver Behavior with Infinite Learning

Lemma 4 *Consider the “infinite learning” recommendation system with interior threshold $R \in (0, 1)$. Then,*

- (i) *any receiver type $i \in [-1/2, 1/2]$ buys the product if it is good;*
- (ii) *no receiver type $i \in [-1/2, 1/2]$ buys the product if it is bad;*
- (iii) *if the product is controversial,*
 - (a) *and $\lambda > 1$, all receiver types $i \geq \tilde{i}_\infty$ buy the product;*
 - (b) *and $\lambda < 1$, all receiver types $i \leq \tilde{i}_\infty$ buy the product.*
 - (c) *and $\lambda = 1$, all receiver types $i \in [-1/2, 1/2]$ buy the product if $\sigma \leq 1$ and none buy it if $\sigma > 1$.*

Proof. First, Lemma 3 implies that infinite learning fully reveals the good and bad products, whereas the posteriors of the controversial product versions are $p_i(b, d) = q_i/(q_1 + q_2)$ for $b, d > 0$. Hence, the expected payoff of buying the recommended good with mixed recommendations reads

$$\frac{q_1}{q_1 + q_2}(1/2 + i) + \frac{q_2}{q_1 + q_2}(1/2 - i) = \frac{1}{2} + i \frac{q_1 - q_2}{q_1 + q_2}.$$

Therefore, receiver i buys the product with mixed recommendations if and only if

$$\frac{1}{2} + i \frac{q_1 - q_2}{q_1 + q_2} \geq V^0(i) = q_H + q_1(1/2 + i) + q_2(1/2 - i). \quad (20)$$

The indifferent consumer satisfies

$$\frac{1}{2} + \tilde{i}_\infty \frac{q_1 - q_2}{q_1 + q_2} = q_H + q_1(1/2 + \tilde{i}_\infty) + q_2(1/2 - \tilde{i}_\infty) \text{ and thus}$$

$$\tilde{i}_\infty = \frac{(q_1 + q_2)(q_1 + q_2 + 2q_H - 1)}{2(q_1 - q_2)(1 - (q_1 + q_2))}$$

Using $q_1 = \frac{2Q\lambda}{\lambda+1}$, $q_2 = \frac{2Q}{\lambda+1}$ and $q_H = ((1 - 2Q)\sigma)/(\sigma + 1)$ gives

$$\tilde{i}_\infty = \frac{(\sigma - 1) \lambda + 1}{2(1 + \sigma) \lambda - 1}$$

Inserting this into (20) gives the result for part (iii) when $\lambda \neq 1$. For the case $\lambda = 1$, the condition to accept the recommendation simplifies to

$$\frac{1}{2} \geq q_H + Q \Leftrightarrow 1 \geq 2q_H + (1 - q_H - q_L) \Leftrightarrow q_L \geq q_H \Leftrightarrow \sigma \leq 1.$$

Finally, parts (i) and (ii) follow trivially. ■

Proof of Proposition 11

(i) If $\lambda = 1$ the value of the infinite learning recommendation system is given by

$$V_\infty = q_H(1 - q_H - Q) + 2Q \max \left\{ \frac{1}{2} - q_H - Q, 0 \right\}.$$

The first term corresponds to the value generated when the product is good, so that the sender always buys it. The second term captures the value generated from a controversial product if it is bought, where the purchasing decision depends on the parameters. Hence, V_∞ coincides with the value of the single-recommendation system for the optimal β , depending on whether q_L or q_H is bigger.

(ii) If $\lambda > 1$, only $i \geq \tilde{i}_\infty$ buy the product by Lemma 4. Hence, no type buys a controversial product in the infinite-learning recommendation system if $\tilde{i}_\infty \geq 1/2$, which is equivalent to $\sigma \geq \lambda$. This inequality implies $\sigma > 1$. Hence, in that case all receivers behave in the same way in both the optimal single-recommendation system (Proposition 5) and the infinite-learning system, buying only objectively

good products. Conversely, if $\tilde{i}_\infty \leq -1/2$ all types buy a controversial product with the infinite-learning recommendation system. This can be rewritten as $\sigma \leq 1/\lambda$, which implies $\sigma < 1$, so that all receivers buy all products except those that are objectively bad and therefore behave in the same way in the optimal single-recommendation system and the infinite-learning system.

(iii) The proof is analogous to case (ii) and thus omitted.

References

- ACEMOGLU, D., A. MAKHDOUMI, A. MALEKIAN, AND A. OZDAGLAR (2022): “Learning from reviews: The selection effect and the speed of learning,” *Econometrica*, 90, 2857–2899.
- BÉNABOU, R. AND N. VELLODI (2024): “(Pro-) Social Learning and Strategic Disclosure,” NBER Working Paper 32483, National Bureau of Economic Research.
- BOLTON, G., B. GREINER, AND A. OCKENFELS (2013): “Engineering trust: reciprocity in the production of reputation information,” *Management Science*, 59, 265–285.
- BOURREAU, M. AND G. GAUDIN (2022): “Streaming Platform and Strategic Recommendation Bias,” *Journal of Economics and Management S*, 31, 25–47.
- CHAKRABORTY, D. AND R. HARBAUGH (2010): “Persuasion by Cheap Talk,” *American Economic Review*, 100, 2361–2382.
- CHE, Y.-K. AND J. HÖRNER (2018): “Recommender systems as mechanisms for social learning,” *The Quarterly Journal of Economics*, 133, 871–925.
- DECKER, C. (2022): *What’s in an Airbnb Five-Star Rating? An Empirical Model of Bayesian Persuasion*, Working Paper.
- DREHMANN, M., J. OECHSSLER, AND A. ROIDER (2005): “Herding and Contrarian Behavior in Financial Markets: An Internet Experiment,” *American Economic Review*, 95, 1403–1426.
- HUI, X., T. KLEIN, AND K. STAHL (2023): *Learning from Online Ratings*, Working Paper.
- JOHNEN, J. AND R. NG (2024): “Harvesting Ratings,” Working Paper, CORE/LIDAM, UCLouvain and University of Mannheim.
- JOHNSON, J. P. AND D. P. MYATT (2006): “On the Simple Economics of Advertising, Marketing, and Product Design,” *American Economic Review*, 96, 756–784.
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615.
- KREMER, I., Y. MANSOUR, AND M. PERRY (2014): “Implementing the “wisdom of the crowd,”” *Journal of Political Economy*, 122, 988–1012.
- PEITZ, M. AND A. SOBOLEV (forthcoming): “Inflated Recommendations,” *Rand Journal of Economics*.
- SMIRNOV, A. AND E. STARKOV (2024): “Designing Social Learning,” Working Paper.
- SUN, M. (2011): “How Does the Variance of Product Ratings Matter?” *Management Science*, 58, 696–707.
- TADELIS, S. (2016): “Reputation and feedback systems in online platform markets,” *Annual Review of Economics*, 8, 321–340.
- VELLODI, N. (2022): “Ratings design and barriers to entry,” *Working Paper*.