Preliminary – Please do not cite or circulate!

A Theory of Recommendations

Jean-Michel Benkert and Armin Schmutzler

March 17, 2024

Abstract

We study the value of recommendations in disseminating economic information, focusing on the impact of preference heterogeneity as a key impediment. We consider Bayesian expected payoff maximizers who evaluate non-strategic recommendations that are given when the payoff of consumption exceeds or falls below some threshold. We derive conditions under which different types accept these recommendations and assess the overall value of the recommendation system. Our analysis highlights the importance of disentangling objective information from subjective preferences. We consider the design of value-maximizing and sales-maximizing recommendation systems as well as the role of a polarized population. Finally, we extend our model in several directions, including multiple recommendation levels and endogenous recommendation thresholds.

Keywords: recommendations, preference heterogeneity, optimal design
JEL Classification: D02, D47, D83

1 Introduction

Recommendations play an important role in the diffusion of economic information. Potential job applicants often rely on information from friends about the quality of potential employers. Consumers may base their own 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, there are considerable impediments to the functioning of recommendations. For instance, too few recommendations may be given (Che and Hörner, 2018; Kremer, Mansour, and Perry, 2014), senders may be positively selected (Acemoglu, Makhdoumi, Malekian, and Ozdaglar, 2022), and interested parties might interfere with the process by selecting which recommendations to publish (Bolton, Greiner, and Ockenfels, 2013; Tadelis, 2016). But even when there is no reason to believe that the senders of a recommendation have incentives to strategically bias the recommendation process, there are limitations. In this paper, we focus on the role of preference heterogeneity as an important friction for the functioning of recommendations. To some extent, even well-intended recommendations by non-strategic players are likely to reflect the sender's preferences rather than objective truth. A receiver who does not question the informational content of a recommendation may therefore make biased decisions.

To motivate our approach, consider the example of choosing between medical doctors. Rather than randomly choosing a doctor, a patient may want to rely on the recommendation of a friend who is familiar with one of the available alternatives. Importantly, the perceived quality of medical care will typically depend on many aspects. How carefully does a doctor listen to the patient? How helpful is their response? Is the technical equipment up-to-date? Is his staff friendly? How crowded is the waiting room? Different patients will value different aspects differently. Typically, recommendations can capture this multi-dimensionality only to a limited extent. While the recommendation may show whether a consumer had a good or bad experience, it will typically be much less clear how the different aspects contributed to this assessment. The potential receiver of the recommendation then faces the complex problem of backing out the implications of the recommendation for their own decision. Relevant considerations for the receiver are: To which 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? Answers to these questions will determine whether it is in the receiver's best interest to follow the recommendation and how valuable the recommendation is for the receiver.

With this in mind, we consider a setting in which economic agents must decide between several available options, which we refer to as products for ease of exposition. Each product is two-dimensional and each dimension can have high or low quality. All agents agree that a product is worth choosing if it has high quality in both dimensions and is thus *objectively good*. Similarly, they agree that a product is not worth choosing if it has low quality in both dimensions and is thus *objectively bad*. However, disagreement can occur when the two components differ in the two dimensions: Each agent will prefer a product that has high quality in the dimension that is more relevant for her and is thus *subjectively good* rather than a product that has high quality in the dimension that is less relevant for her (is *subjectively bad*). Moreover, agents are not only heterogeneous regarding which products they consider subjectively good but also regarding the intensity of these preferences. This heterogeneity is captured in a continuous type distribution.

To make the study of recommendations worthwhile, we assume that the products in question are experience goods so that each agent only has stochastic knowledge of the relevant quality distribution before consumption. In particular, she does not know whether a particular product is objectively good, objectively bad, subjectively good, or subjectively bad. However, she knows the probability of each type of product. Moreover, we assume that for one of the available products, the agent has access to a recommendation stemming from a previous consumer of that product. The recommendation provides coarse information on the product, namely 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 will always give a buy recommendation for an objectively good product and a don't-buy recommendation for an objectively bad product. In the remaining cases, buy recommendations will only come from senders whose preferences for the consumed goods are strong enough. The receiver of the recommendation must now evaluate the informational content of the recommendation, taking into account the decision-making environment and her preferences. Specifically, the decision whether to follow the recommendation will depend on the receiver's preferences as well as the distribution of product qualities and preferences in the overall population.

In this framework, we study recommendations. We are interested in a variety of questions. How do recommendations create value for the receivers and how can we interpret this value? When is this value maximal? If one could influence the quality threshold beyond which a sender gives a buy recommendation, how demanding should the requirement be? Are recommendations valuable even when the probability of products being objectively good or bad is low? How does the value for the receivers change when the population is more or less polarized?

To answer these questions, we need to take intermediate steps that are of interest in their own right. Namely, we need to understand which receivers will accept a recommendation in a particular decision environment. We find that the conditions for a Bayesian, expected payoff maximizer to follow a recommendation are defined by two key quantities measuring objective value and bias in the recommendation, which are simple functions of the model primitives. The first of these quantities captures the objective value of the recommendation. Intuitively, this is high when the recommendation makes it sufficiently more likely that the product is objectively good and sufficiently less likely that it is objectively bad. The second important quantity captures the direction and extent of the bias of the recommendation towards one or the other of the two remaining types of products, which can only be ranked subjectively. Importantly, this bias in the recommendation originates in consumers' subjective preferences. When the bias is absent or small enough, then all players accept the recommendation. That is, they buy after a buy recommendation and they choose the outside option after a don't-buy recommendation. When the bias towards one good is large relative to the objective value of the recommendation, then only those receivers who do not have too strong preferences for the other product will accept the recommendation. Notably, when the recommendation threshold R approaches 1 (the maximum), a buy recommendation becomes entirely objective, while a don't-buy recommendation becomes entirely objective when R approaches 0 (the minimum). As receivers are aware of this, they accept the recommendation in each case since, in each of these polar cases, the subjective component completely vanishes.

Against this background, we can move to the central part of our analysis, the study of the value of recommendation systems. We think of this value as the expected increase in payoffs created by receiving a recommendation from a randomly chosen sender to a randomly chosen receiver. When the recommendation is unbiased or the bias is small, the analysis is comparatively simple: In that case, we know that all receivers accept the recommendation. As only buy recommendations affect the choice of the receiver, the value of the recommendation is the probability of the buy recommendation multiplied by the expected payoff increase from following the recommendation.[1]

When a bias exists and is large enough, the analysis becomes more complex, because only a subset of receivers accepts the recommendation. The value of the recommendation system then has two components. The first part comes from receivers who accept buy recommendations. More precisely, it corresponds to the probability of a buy recommendation, multiplied by the probability that a receiver accepts a buy recommendation and the expected payoff increase from following the buy recommendation (conditional on being a receiver who accepts). The second part comes from receivers who do not accept a don't-buy recommendation. Intuitively, this happens if a receiver knows that her preferences are not well aligned with the general population. Then she takes the don't-buy recommendation as good news about the chances of obtaining the subjectively preferred product rather than the one that she likes less. Of course, this evidence still has to be weighed against the bad news that the product is now more likely to be objectively bad and less likely to be objectively good – but these objective effects may well be dominated by the higher prospects of getting the subjectively preferred product. The expected contribution to the value of recommendations that comes from this effect consists of the probability of a don't buy recommendation, multiplied by the probability that a don't buy recommendation is not accepted times the expected increase in expected payoffs coming from not accepting the recommendation.

Though the components of the value of a recommendation system have clear interpretations, it is non-trivial to understand how the total value depends on primitives. Nonetheless, we can obtain some strong results when considering the design of recommendation systems when the type distribution is symmetric, so that there is no preference bias on average in the population. In these cases, the value of the recommendation system is maximized in one of the two cases where the content of the recommendation is fully objective. In one of these extreme cases, the product will receive a buy recommendation only if it is objectively good; in the other one, it will only be accepted if it is objectively bad.

Next, we study the role of the "degree of subjectivity" in products, that is, the ex ante probability of the two products which cannot be objectively ranked. It turns out that for sufficiently well-designed recommendation systems, the value of the recommendation system is decreasing in the degree of subjectivity. Essentially, the above discussion of optimal thresholds has shown that the value of the system increases when a (relatively) objective recommendation is more likely. Thus, when thresholds are relatively extreme, less subjectivity makes both types of recommendations more objective and simultaneously increases the probability of objective recommendations, both of which increase the value. In contrast, when the system is relatively poorly designed so that the recommendation threshold is intermediate, some subjectivity can be valuable, as it can then increase the probability of a relatively objective recommendation.

[1] A receiver who obtains a don't-buy recommendation chooses the outside option and thus gets the same expected payoff as without a recommendation.

Moreover, we ask how a mean-preserving spread in the population distribution changes the value of a given recommendation system. Assuming a symmetric population, whether a recommendation system creates more value for such a more polarized population depends on the recommendation threshold and on how likely an objectively good product is relative to a bad one. Effectively, a more polarized population changes the subjectivity of the recommendations, and, depending on the relative probability of objective products, this can be a good thing or not.

We conclude by considering some extensions allowing us to showcase the robustness of our results and to make more connections to the existing literature. We first extend our model to allow for more than two different types of recommendations, now called ratings. The key implication of this change is with more than two ratings there is an intermediate rating which is purely subjective: only the lowest and highest ratings carry objective information. When the population is symmetrically distributed, and subjective products are equally likely, it follows that a value-maximizing designer once again tries to maximize the probability of an objective rating being given by opting for extreme recommendation threshold levels. Depending on the relative probability of objective products, the lowest or highest rating is turned objective. Second, we consider the possibility of receiving multiple (binary) recommendations. In this context we can show that mixed reviews, that is, when there are both buy and don't-buy recommendations, rule out both objective products. Thus, as there is more variance in the recommendations, the product is necessarily subjective. Moreover, as the number of recommendations increases, the posterior converges to the true product quality, thus fully revealing the product. Third, we endogenize the recommendation threshold so that a buy recommendation is given if and only if the payoff exceeds the prior expected payoff. Considering the case of equally likely subjective products and of a uniformly, and thus symmetrically, distributed population, we show that the recommendation system with an exogenously chosen threshold generically creates more value than this alternative system with an endogenous threshold. Essentially, the fact that subjective products are equally likely means that the endogenous recommendation threshold coincides across types. Therefore, this corresponds essentially to a fixed, interior recommendation threshold, which is, in general, not optimal. Fourth, we embed our model in a setting where there is a platform that aims to maximize sales. That is, the recommendation threshold is chosen such that the probability of buying the recommended product is maximized. We show that the optimal threshold does not in general maximize the receiver's value and that it induces a rotation in the the demand function as in Johnson and Myatt (2006). Finally, in work in progress, we consider the case of an asymmetrically distributed population. We derive the value-maximizing recommendation threshold and find that it is no longer necessarily a polar threshold but that interior solutions are now possible. Intuitively, given the population's asymmetry, the subjective content in the recommendation does not cancel out. Further, as it is, on average, valuable, it is often optimal not to have completely objective recommendations but to include some subjectivity. Nevertheless, the optimal threshold is still polar whenever the objectively good product is sufficiently more likely than the objectively bad product, or vice versa.

The remainder of the paper is organized as follows. Section 2 discusses the related literature. Section 3 introduces the model and discusses its key assumptions. In Section 4 we describe under what conditions what types accept recommendations. Building on that, we characterize the value of a recommendation system in Section 5, before turning to the design if value-maximizing and persuasive systems in Section 6. We then turn to the extensions in Section 7 and Section 8 concludes. All proofs are relegated to the appendix.

2 Related Literature

Our paper investigates the frictions in recommendation systems arising from preference heterogeneity, and it asks how such systems can be designed to function reasonably well nonetheless. As we will detail in the following, there are several papers that analyze other sources of frictions and ask how they can be mitigated.

Public goods problems: Several papers on recommendations deal with public goods problems that are related to the trade-off between exploration and exploitation, and they deal with design issues to mitigate the problem. The common element in these papers is that consumers choose between products based on the information that they have (typically from previous recommendations of other consumers or from information collected by the platform), and that they do not take into account the information that their choices provide to society. Quite generally, the choice that is optimal from a pure (individual) exploitation perspective may not reveal as much information as alternatives might do, so that insufficient exploration results. An early version of this argument was made by Kremer et al. (2014). Che and Hörner (2018) specify the point to the interesting setting where some of the available choices are niche products that are not well known and would need more exploration to find out whether they are useful or not. While these papers assume that the set of products in the market is given, Vellodi (2022) goes a step further by analyzing the implications of these mechanisms for market structure. Following up on Che and Hörner (2018), he argues that, with unbiased review systems, even if the entrant is of superior quality, she might suffer from the cold start problem whereby the lack of recommendations makes it hard or impossible for entry to be sustainable. He argues that such lack of contestability entrenches the incumbent as a monopolist.[2] Compared to this work, our paper takes the product choice of the senders and the resulting externalities as given. Instead, we focus entirely on the inference problems that receivers face when trying to extract information from recommendations.

The papers also contain closely related policy prescriptions that serve to mitigate the public goods problem. In broad terms, the authors find that information transmission should be limited or biased in a way that fosters information acquisition by consumers (exploration). More concretely, in Kremer et al. (2014) argue that the designer should be less than fully transparent, not supplying all the information she has to subsequent consumers. In Che and Hörner (2018), the designer should sometimes recommend a product that has not yet proven to be of good quality (to encourage exploitation). Vellodi (2022) emphasizes potential welfare gains from suppressing reviews for incumbents, as this can provide incentives for new firms to enter and remain active.[3] In summary, all three

[2] He summarizes his point as: “well-established firms and products often have many hundreds of reviews to their name, affording consumers unprecedented precision when making purchases. But while this stockpile of information might serve incumbent firms to great effect, it might work against a new entrant who unavoidably starts from scratch.”

[3] Similarly, Decker (2022) also argues that platforms may want to encourage consumers to provide coarse information in recommendations. He starts from the observation that most ratings on Airbnb are five stars, and that this rating behavior is actively encouraged by the platform. He argues that this

papers identify public goods problems in the choices of consumers, and they show how the designer can mitigate them by suitable aggregation and transmission of information. Our paper is very different in that it does not deal with public goods issues. It is related in that it asks what determines the informational content of recommendations.

Sender Selection: An alternative source of frictions in recommendation systems comes from the selection of recommenders. As Acemoglu et al. (2022) argue, only consumers with sufficiently strong ex ante preferences for a product will buy it in the first place. Though their consumption experience may modify the ex-ante preferences in either direction, the recommendation partly reflects this upward bias in the pool of recommenders. The receivers must therefore take this bias into account rather than taking the recommendation at face value. The paper is related to ours in that it emphasizes the role of preference heterogeneity for the proper interpretation of recommendations. Contrary to us, however, it focuses exclusively on the implications for sender selection,

Non-truthful Recommendations: Positively biased recommendations may not only reflect the selection of senders, they may also result from their strategic behavior. For instance, Bolton et al. (2013) and Tadelis (2016) document that consumers on eBay used to abstain from negative comments on sellers for fear of getting negative comments themselves. Tadelis (2016) discusses the design implications of these issues – eBay decided to refrain from allowing seller evaluations on buyers.

Receiver Behavior: A related empirical literature asks whether receivers make optimal use of recommendations. Weizsäcker (2010) provides a meta-analysis of experimental studies of herding games where subjects not only receive private information on the relative value of different alternatives, but can also draw inferences from previous choices of other subjects. He finds that, conditional on being in a situation where it is empirically optimal for the participants to contradict their own private information, they choose this action, among two alternatives, in only 44 percent of the cases. In contrast, the participants are much more likely to choose optimally in cases where their own signal happens to support the empirically optimal action: there, they follow their signal nine out of ten times. Similarly, Ronayne and Sgroi (2019) document an unwillingness to follow expert advice even where that would be beneficial for receivers to do, similarly Yaniv and Kleinberger (2000). Our paper is related to this literature in that it provides a theoretical benchmark that gives testable predictions on receiver behavior.

3 Model

We introduce the assumption of our model in Section 3.1, followed by a brief discussion in Section 3.2.

3.1 Assumptions

In what follows we introduce our model in two steps, first introducing products and payoffs before turning to recommendations and the updating of beliefs associated therewith.

deliberate reduction in the informativeness of the system may help to induce more consumers to use the platform and that this is consistent with profit-maximizing behavior. He also takes the model to the data in a structural model. In spite of the similarity in prescriptions (suppression of information), the reasoning is different, as it does not rely on public goods arguments (and is driven by profit rather than welfare maximization).

Products and Payoffs A product is two-dimensional, where $Q_d = 1$ and $Q_d = 0$ capture high and low quality, respectively, in dimension $d \in \{1, 2\}$. 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 product (Q_1, Q_2), and possible payoffs range from 0 to 1.[4] Thus, all consumers obtain the highest possible payoff 1 from consuming product $(1, 1)$ and the lowest possible payoff 0 from consuming $(0, 0)$. They, therefore, agree that products with $(1, 1)$ and $(0, 0)$ are *objectively* good and bad, respectively. By contrast, the value from consuming $(1, 0)$ or $(0, 1)$ depends on a consumer's type i, so that the assessment of products $(1, 0)$ and $(0, 1)$ is *subjective*. More specifically, consumers $i > 0$ prefer product $(1, 0)$ over product $(0, 1)$ and vice versa for consumers $i < 0$. Consumers' types are distributed according to some distribution F. For now, we leave the distribution unspecified but assume continuity and full support for simplicity.

We assume that the products are experience goods, so that a consumer cannot observe the quality vectors of a product before consumption. However, she knows the prior distribution of the quality vectors, which is identical and independent across products. We set

$$\Pr[(Q_1, Q_2) = (1, 1)] = q_H, \qquad \Pr[(Q_1, Q_2) = (1, 0)] = q_1,$$
$$\Pr[(Q_1, Q_2) = (0, 0)] = q_L, \qquad \Pr[(Q_1, Q_2) = (0, 1)] = q_2,$$

and write $\mathbf{q} = (q_H, q_1, q_2, q_L)$.

Recommendations and Updating Throughout the paper, we will consider a situation in which a consumer has the choice between a product that comes with a recommendation and a product that does not. We distinguish between buy and don't-buy recommendations. Formally, the recommendations are defined as $r \in \{B, D, 0\}$, where 0 captures that no recommendation was received. We refer to our consumer of interest as the receiver (of a recommendation) and to the consumer who made the recommendation as the sender.

To eliminate any scope for strategic behavior on the part of the sender, we assume that the sender is randomly chosen from the population according to the distribution F and mechanically gives a buy or don't-buy recommendation depending on the payoff they have received from consuming some product. In particular, the sender sends a buy recommendation if and only if she obtained a payoff of at least $R \in [0, 1]$. Hence, a recommendation can be the result of an objective or a subjective assessment of the product. For instance, the objectively good product (1,1) yielding a payoff of 1 will always result in a buy recommendation, while the objectively bad product (0,0) will always yield a don't-buy recommendation. Whether the subjective products (1,0) and (0,1) yield a buy or a don't-buy recommendation depends on the sender's type. For instance, following the consumption of a (1,0) product, a sender of type i gives a buy recommendation if and only if $i \geq R - 1/2$.

Receiving a buy or a don't-buy recommendation leads the receiver to update her belief about the distribution of the quality vector of that product, taking into account the uncertainty about the source of the recommendation and its informational content. To

[4] For convenience, we identify a product with the corresponding quality vector.

describe the updated probabilities, we denote the probability of a buy recommendation conditional on the product being $(1,0)$ or $(0,1)$, respectively, as

$$\phi_1(R) := 1 - F(R - 1/2), \qquad \phi_2(R) := F(1/2 - R),$$

Note that $\phi_1(R)$ is the probability of a sender having a type of at least $R - 1/2$, whereas $\phi_2(R)$ is the probability of a sender having a type of at most $1/2 - R$. Thus, having received a buy recommendation, the posterior reads

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

Similarly, a don't-buy recommendation results in a posterior of

$$p_H^D(R) = 0, \qquad p_2^D(R) = \frac{q_2(1-\phi_2(R))}{q_1(1-\phi_1(R)) + q_2(1-\phi_2(R)) + q_L}$$
$$p_1^D(R) = \frac{q_1(1-\phi_1(R))}{q_1(1-\phi_1(R)) + q_2(1-\phi_2(R)) + q_L}, \quad p_L^D(R) = \frac{q_L}{q_1(1-\phi_1(R)) + q_2(1-\phi_2(R))) + q_L}.$$

We will use the notation $\pi_B \equiv \pi_B(R)$ for the probability that a randomly chosen sender gives a buy recommendation and $\pi_D \equiv \pi_D(R)$ for the probability that a randomly chosen sender gives a don't-buy recommendation. Clearly,

$$\pi_B(R) = q_H + q_1\phi_1(R) + q_2\phi_2(R), \tag{1}$$
$$\pi_D(R) = q_1(1-\phi_1(R)) + q_2(1-\phi_2(R)) + q_L, \tag{2}$$

and $\pi_D = 1 - \pi_B$.

Finally, we use the following summary terminology.

Definition 1 *The* ***decision environment*** *is given by* $\mathcal{E} := (\boldsymbol{q}, F)$. *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

In the following, we discuss some aspects of the framework. We first highlight the role of certain assumptions, and then we deal with interpretation.

First, when $q_H > 0$ and $q_L > 0$, the decision environment is such that all types agree which product is the best and which is the worst (at least weakly). This is an immediate implication of the fact that, for all consumers, utility is fully determined by the qualities in each component and increasing in each of these qualities. Of course, there are conceivable environments where there are no products which are objectively the best or obviously the worst. For instance, there are clearly no such products in a modified version of our setting where $q_H = 0$ and $q_L = 0$, as the valuations for the remaining products $(1,0)$ and $(0,1)$ are perfectly negatively correlated. We nonetheless keep the specific feature of the model, because we want to emphasize the complications arising from the coexistence of the objective and subjective aspects of recommendations.

Second, there are some other aspects of the framework that we shall deal with separately as extensions in Section 7. In the benchmark model of the current section, there are only two types of recommendations ("buy" or "don't buy"), we will handle the case of multiple levels of recommendations in Section 7.1. In Section 7.2, we return to the case of two types of recommendations, but, in contrast to the benchmark, we allow for the possibility that the receiver has access to more than one recommendation. Finally, the recommendation threshold in the benchmark model is exogenous. In Section 7.3, we handle the alternative that the threshold is endogenous, determined by the expected payoffs that the receiver would obtain without a recommendation.

Third, the interpretation of the recommendation threshold in the benchmark model needs some discussion. Our preferred way of thinking about it is as a design parameter. To motivate this, we follow Decker (2022) who analyzes the five-star rating system of Airbnb in a structural model. He argues that the platform advises its customers to give a five-star rating to every host that provides reasonably good services and reserve lower ratings exclusively to hosts that are very bad. One way to think of this using our setting would be to argue that the platform is essentially using a coarse recommendation system, with the vast majority of hosts receiving a buy recommendation and a small subset receiving a don't-buy recommendation. In our language, the recommendation threshold would then be close to $R = 0$, so that the goods that are not recommended are typically objectively bad. In that spirit, the threshold should not be taken too literally but rather thought of as the guidance a platform such as Airbnb gives or as a social norm that has evolved or is being actively shaped.

Finally, with a single threshold above and below which buy and don't-buy recommendations, respectively, are given, it is impossible to always reveal the product in question. If, in contrast, the sender could simply give four different ratings depending on the consumption experience, full revelation would be possible by associating a rating with each product type. However, in practice the multi-dimensionality of products, which we model here as simply as possible using four different product types, would typically render such a full-revelation system infeasible, resulting in coarse information transmission as in our model.

4 Optimal Receiver Behavior

In this section, we ask under which circumstances a receiver will accept recommendations, depending on his type i and the decision environment, comprising the probability distribution over products as well as the type distribution. Finally, it includes the recommendation threshold, which, as argued above, can usefully be interpreted as a design parameter. In Section 4.1, we first provide further terminology and some useful auxiliary results. Section 4.2 then characterizes the optimal behavior of the receiver by identifying thresholds dividing the type set into receivers that accept and do not accept the recommendation.

4.1 Notation and Auxiliary Results

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,

$$\mathbf{p}(B) - \mathbf{q} = (\mathbf{p}(B) - \mathbf{q}'') + (\mathbf{q}'' - \mathbf{q}') + (\mathbf{q}' - \mathbf{q}), \tag{3}$$

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 it 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 subjectively good and bad 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 changes the odds between the subjective products.* 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 $\mathbf{p}(B)$.

The first two steps of this decomposition are positive for any type of receiver and thus capture objectively positive effects of a buy recommendation: The receiver knows that the product is not objectively bad, and the probability of an objectively good product is higher than without a recommendation. By contrast, an analogous decomposition for the case of a don't-buy recommendation reveals that after a don't-buy recommendation, the receiver knows that the product is not objectively good, and the probability of an objectively bad product is higher than without a recommendation. Finally, the third step of the decomposition captures heterogeneous effects on receivers. Moreover, the effect of a recommendation on the odds of the two subjective products will depend on the parameters of the model. Therefore, we introduce the following self-explanatory distinction.

Definition 2

(i) A recommendation r ***shifts probability from*** *$(0, 1)$* ***to*** *$(1, 0)$ if $p_2^r - q_2 < p_1^r - q_1$.*

(ii) A recommendation r ***shifts probability from*** *$(1, 0)$* ***to*** *$(0, 1)$ if $p_2^r - q_2 > p_1^r - q_1$.*

(iii) Otherwise, a recommendation r ***does not affect the probability*** *between $(0, 1)$ and $(1, 0)$.*

Note that the terminology applies both to buy and don't-buy recommendations. Intuitively, a recommendation shifts probability towards one of the two products that cannot be compared objectively if it makes it relatively more likely than the other one of these goods. In the case of a buy recommendation, this will happen if the ex-post probability of the product after such a recommendation (relative to the other one of these products) is high compared to the ex-ante probability.

We are now in a position to understand the effects of a recommendation on expected payoffs and, thus, whether it is optimal for the receiver to *accept a recommendation.* To this end, denote by U_i^B and U_i^D the expected payoff of player i of buying a product with a buy and don't-buy recommendation, respectively, and U_i^0 the expected payoff of buying the alternative product with no recommendation.[5] We have

$$U_i^r = p_H^r + p_1^r(1/2 + i) + p_2^r(1/2 - i), \quad r \in \{B, D\}, \tag{4}$$

and

$$U_i^0 = q_H + q_1(1/2 + i) + q_2(1/2 - i). \tag{5}$$

For an expected payoff maximizer, it is optimal to accept a buy recommendation if $U_i^B - U_i^0 \geq 0$, that is, if the expected payoff from buying the recommended product is higher than the expected payoff from buying the other product (that has no recommendation). Similarly, it is optimal to accept a don't-buy recommendation if $U_i^D - U_i^0 \leq 0$, that is, if the expected payoff from buying the product without any recommendation is higher than the expected payoff from buying the product with the don't-buy recommendation. Whether it is optimal to accept a recommendation depends on the receiver's type i and the parameters of the model, that is, the probability vector $\mathbf{q}$ and the type distribution F. As our first result shows, it turns out that the condition for accepting a recommendation is the same for buy and don't-buy recommendations.

Proposition 1 *For given probability vector $\mathbf{q}$ and the type distribution F, a receiver of type i accepts a buy recommendation if and only if she accepts a don't-buy recommendation.*

Unless noted differently, all proofs are relegated to the appendix. Proposition 1 establishes that the conditions under which a receiver i accepts a buy or a don't-buy recommendation are the same. Intuitively, following the above discussion, a receiver will accept a recommendation if he thinks that it has sufficient objective content and/or that his tastes are sufficiently aligned with those of the sender. Whether the recommendation is positive or negative does not matter for this assessment.

4.2 Characterizing Optimal Receiver Behavior

Next, we ask which types i will accept a recommendation or not. Proposition 1 allows us to consider only buy recommendations. We have

$$\begin{aligned} & U_i^B - U_i^0 \geq 0 \\ \Leftrightarrow & p_H^B(R) - q_H + (p_1^B(R) - q_1)(1/2 + i) + (p_2^B(R) - q_2)(1/2 - i) \geq 0, \\ \Leftrightarrow & (p_H^B(R) - q_H) + \frac{p_1^B(R) - q_1}{2} + \frac{p_2^B(R) - q_2}{2} \geq i[(p_2^B(R) - q_2) - (p_1^B(R) - q_1)]. \end{aligned} \tag{6}$$

To formulate this condition in compact form, we use the following terminology.

[5] There are different ways to interpret this setting. First, the consumer could face the choice between two a priori identical (in terms of the distribution of the possible quality vectors) experience goods and receive additional information about one of them in the form of a recommendation. Second, the consumer faces the choice between an experience good and an outside option, where the expected payoff of the experience good before receiving a recommendation coincides with the value of the outside option.

Definition 3

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

$$\Delta_O^r := p_H^r(R) - q_H + \frac{p_1^r(R) - q_1}{2} + \frac{p_2^r(R) - q_2}{2} \tag{7}$$

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

$$\Delta_S^r := (p_2^r(R) - q_2) - (p_1^r(R) - q_1) \tag{8}$$

Δ_O^r is independent of i, capturing the effect of following the recommendation on the average player ($i = 0$). Because of this feature, we can think of it as the objective part of the effect of following the recommendation. By contrast, Δ_S^r captures how the recommendation affects the probabilities of products 1 and 2, about which the players disagree – it is non-zero only if the recommendation affects these probabilities in different ways. A positive (negative) value of Δ_S^r corresponds to a positive (negative) effect of r on the likelihood of product 2 relative to product 1. Whether each type of effect means that a receiver of type i will feel more inclined to buy the product will obviously depend on the sign of i, and is therefore subjective. Inequality (6) can thus be rewritten as

$$\Delta_O^B \geq i\Delta_S^B. \tag{9}$$

Thus, the acceptance of the recommendation hinges on the interplay of the objective and subjective effects of the recommendation on the expected payoffs, captured by Δ_O^B and Δ_S^B, respectively. To improve the intuition further, the next result is helpful.

Lemma 1 *For any prior distribution* $\mathbf{q}$ *and any recommendation threshold R we have* $\Delta_O^B \geq 0$ *and*

(i) $\Delta_S^B > 0$ *if a buy recommendation shifts the probability from* $(1, 0)$ *to* $(0, 1)$*;*

(ii) $\Delta_S^B < 0$ *if a buy recommendation shifts the probability from* $(0, 1)$ *to* $(1, 0)$*;*

(iii) $\Delta_S^B = 0$ *otherwise.*

Intuitively, as laid out in the decomposition of the recommendation's effect in equation (3), the objective part Δ_O^B of the recommendation must be positive, as it makes the objectively good product more likely relative to the two subjectively good and bad products and rules out the objectively bad product. By contrast, the sign of subjective content $\Delta_S^B > 0$ can be positive or negative, corresponding to Step 3 in the above decomposition, with the result of Lemma 1 following directly from the definition.

It will turn out to be useful to also consider the case when a receiver rejects a don't-buy recommendation, that is, buys the product despite it receiving a negative recommendation. Accordingly, player i rejects a don't-buy recommendation if and only if

$$\begin{aligned} & U_i^D - U_i^0 \geq 0 \\ \Leftrightarrow & p_H^D(R) - q_H + (p_1^D(R) - q_1)(1/2 + i) + (p_2^D(R) - q_2)(1/2 - i) \geq 0, \\ \Leftrightarrow & \Delta_O^D \geq i\Delta_S^D. \end{aligned} \tag{10}$$

With this in hand, we obtain the following counterpart to Lemma 1.

Lemma 2 *For any prior distribution* $\mathbf{q}$ *and any recommendation threshold* R *we have* $\Delta_O^D \leq 0$*. Further,* $\pi_B \Delta_O^B = -\pi_D \Delta_O^D$ *and* $\operatorname{sgn}(\Delta_S^B) = -\operatorname{sgn}(\Delta_S^D)$ *.*

Intuitively, decomposing the effect of a don't-buy recommendation as we did above for a buy recommendation, Δ_O^D captures the objectively bad effect on the expected product quality, as the don't-buy recommendation rules out the objectively good product and makes the objectively bad one more likely than the subjective products. This decrease in the recommended product's objective value corresponds to the increase resulting from a buy recommendation, yielding $\pi_B \Delta_O^B = -\pi_D \Delta_O^D$. Hence, a receiver will only reject a don't-buy recommendation if the subjective effect Δ_S^D is sufficiently positive. Finally, this subjective effect of a don't-buy recommendation goes in the opposite direction of the subjective effect of a buy recommendation. We can now determine how the behavior of the receivers in terms of accepting or rejecting a recommendation depends on their type, where, using equation (6), we define the indifferent type $\tilde{i}$ as

$$\tilde{i} := \frac{\Delta_O^B}{\Delta_S^B} \text{ if } \Delta_S^B \neq 0. \tag{11}$$

Proposition 2

(i) If $|\Delta_S^B| \leq 2\Delta_O^B$*, then all* $i \in [-1/2, 1/2]$ *accept the recommendation.*

(ii) If $\Delta_S^B < -2\Delta_O^B$*, then all* $i \geq \tilde{i}$ *accept the recommendation.*

(iii) If $\Delta_S^B > 2\Delta_O^B$*, then all* $i \leq \tilde{i}$ *accept the recommendation.*

Note that Lemma 2 implies that in case (i), both buy and don't-buy recommendations leave the probability between the subjective products unchanged. Further, we could also formulate cases (ii) and (iii) in terms of Δ_S^D, in which case the signs of Δ_S^D would just be flipped according to Lemma 2.

The intuition for Proposition 2 is straightforward. For $\Delta_S^B = 0$, the recommendation leaves the probability of the products 1 and 2 unaffected. Accordingly, accepting the recommendation is optimal for all players because of its positive objective value $\Delta_O^B \geq 0$. For $\Delta_S^B \neq 0$, this conclusion still holds as long as $|\Delta_S^B|$ remains small enough – as required by (i): In that case, even the players who are most negatively affected by the subjective content of the recommendation (those at $i = 1/2$ or $i = -1/2$) are convinced by its objective character.

As $|\Delta_S^B|$ increases, the role of the shifting of the probability becomes apparent. Type $i \geq 0$ prefers product (1,0) over product (0,1). Hence, when the recommendation shifts the probability from the latter to the former ($\Delta_S^B < 0$), this yields a positive effect on that receiver's expected payoff, reinforcing the objective effect. By contrast, the expected payoff of a type $i \leq 0$ is adversely affected by the shift of the probability. The condition in (ii) is always fulfilled for $i > 0$, as the objective and subjective effects reinforce each other. It also holds when $i < 0$, as long as $|i|$ is small enough 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 $\Delta_S^B > 0$), 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.

The following simple observation follows immediately from Proposition 2.

Corollary 1

(i) If a recommendation increases the probability of (1,0), then all $i \geq 0$ accept it.

(ii) If a recommendation increases the probability of (0,1), then all $i \leq 0$ accept it.

The proof is omitted, as the result follows directly from showing that $\tilde{i} \leq 0$ in case (i) and $\tilde{i} \geq 0$ in case (ii). Intuitively, a consumer who is indifferent between products $(1,0)$ and $(0,1)$, only cares about the objective content of the information that the recommendation conveys, and therefore always accepts a recommendation. A fortiori, the receivers for whom the information increases the probability of the subjective product that they prefer will also accept a recommendation.

Finally, we provide a result on the effects of changes in the recommendation threshold R on receiver behavior.

Proposition 3 *Suppose $R \to 1$ or $R \to 0$. Then, all types accept the recommendation irrespective of the product distribution.*

Essentially, in the limits, one of the recommendations contains only objective information. When $R \to 1$, buy-recommendations are only given when the product is objectively good, while don't-buy recommendations contain both subjective and objective informational content.[6] When $R \to 0$, every product receives a buy recommendation except the objectively bad one so that the don't-buy recommendation contains only objective informational content.[7] We know from Proposition 1 that the equivalence condition is the same for buy and don't-buy recommendations. Thus, it does not matter which recommendation is purely objective for all types to follow as long as one of them is.

Finally, the fact that all types follow the recommendation for these extreme thresholds does not imply that the value of the recommendation system is maximal in these cases, as we will see in the Section 6.

5 The Value of the Recommendation System

The previous section has characterized the optimal receiver behavior under a given recommendation system. In the following, we characterize the value of recommendation systems for a given economic environment. Before describing this value, we first describe the expected payoffs of a fixed type i, depending on whether she has access to a recommendation system and whether she accepts it or not.

Without a recommendation system, the expected payoff of a receiver i is given by

$$V_0(i) = q_H + (1/2 + i)q_1 + (1/2 - i)q_2, \tag{12}$$

as this is the a priori expected payoff in the absence of a recommendation. Next, the probability that the sender gives a buy recommendation is π_B. In the presence of a

[6] Note that for $R = 1$ the probability of a subjective product leading to a buy recommendation is a probability-zero event.

[7] Once more, note the probability-zero events leading to this observation.

recommendation system, the expected payoff of a receiver i who accepts a recommendation from a random sender i is thus given by

$$\begin{aligned} V_A(i) =& \pi_B(R)\big(p_H^B(R) + (1/2+i)p_1^B(R) + (1/2-i)p_2^B(R)\big) \\ &+ (1-\pi_B(R))\big(q_H + (1/2+i)q_1 + (1/2-i)q_2\big). \end{aligned} \tag{13}$$

Here the first term captures the probability of receiving a buy recommendation multiplied by the expected payoff from accepting it (and thus buying the recommended product), and 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.[8] Analogously, we obtain the expected payoff of a receiver i who does not accept the recommendation as

$$\begin{aligned} V_N(i) =& \pi_B(R)\big(q_H + (1/2+i)q_1 + (1/2-i)q_2\big) \\ &+ (1-\pi_B(R))\big(p_H^D(R) + (1/2+i)p_1^D(R) + (1/2-i)p_2^D(R)\big). \end{aligned} \tag{14}$$

Naturally, the probabilities of the respective recommendations are the same as above. However, the values for receivers now result from not accepting the recommendation, that is, from buying the product without a recommendation after having received a buy recommendation for another one and conversely, from buying a product for which one has received a don't-buy recommendation.

We now define the value of a recommendation system.

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

Using this definition, we can characterize the value of a recommendation system as follows.

$$V(\mathcal{R}) = \begin{cases} \int_{-1/2}^{1/2}(V_A(i) - V_0(i))dF(i) & \text{if } |\Delta_S^B| \leq 2\Delta_0^B \\ \int_{-1/2}^{\tilde{i}} V_N(i)dF(i) + \int_{\tilde{i}}^{1/2} V_A(i)dF(i) - \int_{-1/2}^{1/2} V_0(i)dF(i) & \text{if } \Delta_S^B < -2\Delta_0^B \\ \int_{-1/2}^{\tilde{i}} V_A(i)dF(i) + \int_{\tilde{i}}^{1/2} V_N(i)dF(i) - \int_{-1/2}^{1/2} V_0(i)dF(i) & \text{if } \Delta_S^B > 2\Delta_0^B \end{cases} \tag{15}$$

To understand this result, recall from Proposition 2 that the change in probability from products (1,0) to (0,1) essentially determines which types accept a recommendation: Depending on the sign and size of Δ_S^B, either (i) all types accept a recommendation, (ii) all types above the threshold $\tilde{i}$ accept a recommendation, or (iii) all types below the threshold $\tilde{i}$ accept a recommendation. The difference between the expressions for $\Delta_S^B < -2\Delta_0^B$ and $\Delta_S^B > 2\Delta_0^B$, respectively, comes from the fact that only types $i > \tilde{i}$ accept the recommendation in the former case, whereas only types $i < \tilde{i}$ accept it in the latter case. In our central result, we express the value of recommendation system in terms of its objective and subjective content. Let $M := \max\{-1/2, \tilde{i}\}$ and $m := \min\{1/2, \tilde{i}\}$. Then, after simple rearrangements of (15) we obtain:

[8] We choose the notation $V_A(i)$ to emphasize the dependence of this expected payoff on the type of the receiver. Clearly, the expression also depends on R and on the decision environment; similarly, for the expressions below.

Proposition 4 *$V(R)$, the value of the recommendation system with threshold R, is:*

$$V(\mathcal{R}) = \begin{cases} \pi_B \Delta_O^B & \text{if } |\Delta_S^B| \le 2\Delta_0^B \\ \pi_B F(m)\left[\Delta_O^B - \Delta_S^B \mathbb{E}[i \mid i \le m]\right] + (1-\pi_B)(1-F(m))\left[\Delta_O^D - \Delta_S^D \mathbb{E}[i \mid i \ge m]\right] & \text{if } \Delta_S^B < -2\Delta_0^B \\ (1-\pi_B)F(M)\left[\Delta_O^D - \Delta_S^D \mathbb{E}[i \mid i \le M]\right] + \pi_B(1-F(M))\left[\Delta_O^B - \Delta_S^B \mathbb{E}[i \mid i \ge M]\right] & \text{if } \Delta_S^B > 2\Delta_0^B \end{cases}$$

The above expressions can be interpreted as follows. When $\Delta_S^B = 0$, the recommendation does not affect the relative probabilities of $(0,1)$ and $(1,0)$. Accordingly, it follows from equation (6) that all receivers accept the recommendation and benefit from the objective increase Δ_O^B in expected payoffs after a buy recommendation (that arises with probability π_B). When $\Delta_S^B \neq 0$, the analysis is more complex. We will interpret the term for $V(R)$ for $\Delta_S^B > 0$, the argument for $\Delta_S^B < 0$ is analogous.

If $\Delta_S^B > 0$, 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 it negative. Indeed, if Δ_S^B is sufficiently large that $\tilde{i} < 1/2$ and thus $m = \min\{\tilde{i}, 1/2\} < 1/2$, Proposition 2 implies that only a fraction $F(m)$ of the receivers accepts a buy recommendation. The term $-\Delta_S^B \mathbb{E}[i \mid i \le m]$ captures 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 effect depends on the distributional properties of F captured in $\mathbb{E}[i \mid i \ge M]$: It will be positive if the positive contribution of types $i < 0$ dominates the negative contributions of types $i \in (0,1)$.

The second component of the value of the recommendation system for $\Delta_S^B > 0$ is arguably more surprising. In brief, it comes from receivers who do not accept a don't-buy recommendation because, given their lack of alignment with the rest of the population, they treat the recommendation not to buy as good news about the product. In more detail, a don't-buy recommendation (which happens with probability $1 - \pi_B$) carries an adverse objective content (as summarized in the term in squared brackets by Δ_O^D). However, for buyers with $i > 0$ the don't-buy recommendation also carries some good news, namely that the product is more likely to be of the preferred type $(1,0)$ rather than $(0,1)$. Those with $i > m (> 0)$ then buy the product in spite of the adverse objective content, valuing the beneficial subjective effect with $-\Delta_S^D \mathbb{E}[i \mid i \ge m] > 0$ on average.[9] Such receivers exist only if Δ_S^B is sufficiently large (and hence Δ_S^D is sufficiently negative relative to Δ_O^D).

In the following, we shall often invoke the following symmetry assumption.

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

With this assumption in place, the following quantity has a simple interpretation.

$$\beta := F(1/2 - R). \tag{16}$$

It captures the fraction of the population that is willing to recommend product $(1,0)$ or, equivalently (given the symmetry of the distribution), product $(0,1)$. For fixed type distribution F, β and R are inversely related. In the symmetric setting, it is possible to provide a simple expression for $V(R)$ that allows for transparent comparative statics results. To this end, a reparametrization is helpful:

[9] Note that $-\Delta_S^D \mathbb{E}[i \mid i \ge m]$ is positive: By Lemma 2, $\Delta_S^B > 0$ implies $\Delta_S^D < 0$.

Definition 5 *The **degree of subjectivity** of the recommendations is given as* $Q := \frac{q_1+q_2}{2}$. *The **odds of an objectively good product** are* $\sigma := q_H/q_L$.

Using these parameters, we can rewrite the value of the recommendation system as follows.

Proposition 5 *Suppose Assumption 1 holds. Then, irrespective of the change in probabilities between (1,0) and (0,1), all types follow the recommendation. Moreover,*

$$V(\beta) = \pi_B \Delta_O^B = (1-2Q)\frac{\sigma + Q\left(\beta - \sigma + \sigma^2(1-\beta)\right)}{(\sigma+1)^2}$$

To prove this, observe that, under Assumption 1, $\min\{\Delta_O^B + \frac{\Delta_S^B}{2}, \Delta_O^B - \frac{\Delta_S^B}{2}\} \geq 0$. Then, according to Proposition 2, for any shift in probabilities between (1,0) and (0,1), all types follow. Thus, using Proposition 4, we can write the value of the recommendation system as $\pi_B\left[\Delta_O^B - \Delta_S^B \mathbb{E}[i]\right]$. Because of the symmetry of the population, we have $\mathbb{E}[i] = 0$, so the value of the recommendation system is given by $\pi_B \Delta_O^B$, which we can rearrange as written in the proposition.

In this symmetric setting, $\beta \in [0,1]$ is a sufficient statistic capturing the influence of the distribution and the recommendation threshold on the value. In addition, the value of the recommendation system depends on the degree of subjectivity Q and the odds of the objectively good product σ.[10] Studying the value of the recommendation system as a function of β is particularly interesting, as it allows to capture different aspects of the model. We can interpret changes of β as either changing the threshold R for a given distribution F of as changes of in the distribution F for a given threshold R.

6 Optimal Design of Recommendation Systems

In the previous section, we have obtained a general and intuitive characterization of the value of a recommendation system with threshold R. We now turn to design issues, focusing on the role of the recommendation threshold as a design parameter. In doing so, we begin by deriving the optimal recommendation threshold for the symmetric model, before looking into the implications of relaxing this assumption.

We start by proving the following auxiliary result, showing that the effect of changes in β on the probability of a buy recommendation and the average value of accepting it may be countervailing, resulting in a non-trivial relation between β and $V(\beta) = \pi_B \Delta_O^B$.

Lemma 3 *Suppose that Assumption 1 holds. Then*

(i) π_B *is strictly increasing in* β*;*

(ii) Δ_O^B *is strictly decreasing in* β*;*

(iii) $\frac{\partial^2 V}{\partial \beta \partial \sigma} < 0$.

(iv) the value of the recommendation system is maximized for $\beta^* = 1$ *if* $\sigma < 1$ *and for* $\beta^* = 0$ *if* $\sigma > 1$.

[10] Observe that by varying R we can obtain any $\beta \in [0,1]$.

The results reflect the assumption of a symmetrically distributed population. Under symmetry, all receivers follow the recommendation, so that the buying probability is identical to the likelihood of a buying recommendation. Therefore, result (i) is immediately intuitive given the interpretation of β as the fraction of the population that recommends products $(1, 0)$ or the fraction of the population that recommends $(0, 1)$. Turning to (ii), as β increases, a greater fraction of the population is willing to recommend $(1, 0)$ and $(0, 1)$. Thus, the likelihood that a buy recommendation refers to an objectively good product falls. Put differently, the informational content of the recommendation becomes less objective and the objective part of the recommendation Δ_O^B falls.

Next, the average subjective effect of the recommendation cancels out as $\mathbb{E}[i] = 0$. Therefore, $V(R) = \pi_B \Delta_O^B$. To understand (iii) and (iv), we thus need to understand the effect of β on this product. From part (ii), we know that the objective component Δ_O^B is strictly decreasing in β. Moreover, as σ increases (high states become more likely relative to low states), the probability of a buy recommendation becomes larger. Therefore, having a high share of objective goods is particularly valuable when β is low. This explains result (iii). While this result immediately implies that value-maximizing level of β must be weakly increasing in σ, result (iv) is much stronger, stating that only boundary solutions can arise. Intuitively, the symmetry assumption renders the problem linear in β, yielding a corner solution, the nature of which is in line with the implication of (iii) that the value-maximizing level of β must be increasing in σ.

With Lemma 3 in hand, it is easy to derive the threshold that maximizes the value of the recommendation system.

Proposition 6 *Suppose Assumption 1 holds. Then, the value of the recommendation system is maximized for*

(i) $R^* = 0$ *if* $\sigma < 1$*;*

(ii) $R^* = 1$ *if* $\sigma > 1$*;*

(iii) any $R^* \in [0, 1]$ *if* $\sigma = 1$*.*

The result is an immediate implication of Lemma 3(iv), as $R = 0$ corresponds to $\beta = 1$ and $R = 1$ corresponds to $\beta = 0$. The above intuition for the value-maximizing choice of β thus translates directly to the choice of R: for $R = 1$ the informational content of a buy recommendation is purely objective, while for $R = 0$ it is purely objective for a don't-buy recommendation. Hence, by choosing either a threshold $R = 1$ or $R = 0$ the designer is choosing in what case the recommendation system is purely objective and when it is rather subjective. Following the above logic, therefore, objectively good products are more likely than objectively bad ones (i.e., when $\sigma > 1$), an objective buy recommendation maximizes the value of the recommendation system and vice versa when bad products are more likely ($\sigma < 1$).

The assumption of a symmetric population allows for a relatively clean interpretation of the role of objective information in determining the (optimal) value of the recommendation system since the subjective component cancels out due to the symmetry in the population. However, the assumption of a symmetric population may prove restrictive depending on the context and does not allow us to study the role of the subjective component in the optimal design of the recommendation system. To make headway in that direction, we make the following assumption.

Assumption 2 *We have* $q_1 = q_2 = Q$ *and let* $F(x) = (x + 1/2)^a$ *for* $a > 0$*.*

By assuming that the subjective products are a priori equally likely, we ensure that any shift in the probabilities between the subjective products following a recommendation is purely driven by the asymmetry in the population. Together with the parameterization of the distribution of types in the population, this allows for a relatively tractable analysis. Importantly, we can model situations in which there are more or less people who prefer product 1 over product 2 and vice versa.

Before we turn to the optimal design of the threshold, we obtain the following corollary of Proposition 2.

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

1. $\Delta_S^B = 0$ *if* $a = 1$*.*
2. $\Delta_S^B < 0$ *if* $a > 1$*.*
3. $\Delta_S^B > 0$ *if* $a < 1$*.*

Thus, as mentioned above, any shift in the probabilities between the subjective products is driven by the distribution of the types. More precisely, when types are symmetrically distributed ($a = 1$) we have no shift. When the distribution of types is convex ($a > 1$), so that more people prefer product 1 over product 2, the probability is shifted to product 1. The reason is that a buy recommendation is more likely to come from a person who prefers product 1, so that product 1 is relatively more likely than product 2 after a buy recommendation. Analogously, when the distribution of types is concave ($a < 1$), the probability is shifted to product 2.

We can now gauge the extent to which the corner solution in Proposition 6 is an artefact of the symmetry assumption.

Proposition 7 *Suppose Assumption 2 holds. The optimal threshold* R^* *is interior, i.e.,* $R^* \in (0, 1)$*, if*

$$\frac{a}{a+1} \gtrless Q + q_H \gtrless \frac{1}{a+1} \text{ for } a \lesseqgtr 1.$$

We can infer two things from this result, which is illustrated in Figure 1. First, when a is sufficiently close to 1, the parameter space for which an interior solution exists, vanishes. Hence, the corner solution in Proposition 6 is not simply the result of a knife-edge condition. Second, even when the distribution of the types is very asymmetric, i.e., when a is either very large or very close to 0, an corner solution may exist for some parameters.

It is difficult to derive the optimal threshold under Assumption 2, since changes in the in the threshold may change the set of types that accept a recommendation. However, it is possible to derive the solution for specific numerical examples, that are insightful and representative of the solution more generally. Suppose $a = 2$, $\sigma \in [0.2, 5]$ and $Q \in (0, 0.3]$. For these parameters, we can show that all types accept recommendations so that the value of the recommendation system can be written as

$$\begin{aligned} V(R) &= \pi_B \left[\Delta_O^B - \Delta_S^B \mathbb{E}[i]\right] \\ &= \frac{3\sigma - 6Q^2(1 - R - \sigma + R\sigma^2) - Q(9\sigma - 3 + R^2(1+\sigma)^2 + 2R(1 - \sigma - 2\sigma^2))}{3(1+\sigma)^2}. \end{aligned}$$

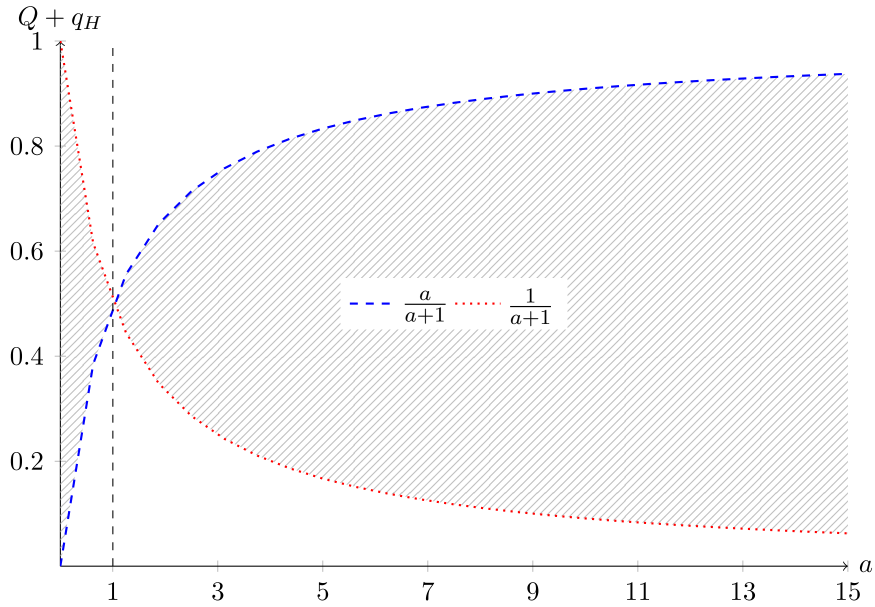

Figure 1 Illustration of the result in Proposition 7. The shaded area captures the parameter space in which an interior solution arises. Clearly, near $a = 1$, this space colappses and no interior solution exists for $a = 1$, while the the space expans fully for $a \to 0, \infty$.

Further, we can show that this function is quasiconcave in R so that applying the Karusch-Kuhn-Tucker Theorem (e.g., Sundaram, 1996, Theorem 8.13) gives the solution

$$R^*(\sigma, Q) = \begin{cases} 0 & \text{if } \sigma < \frac{3Q-1}{3Q-2} \\ \frac{2\sigma-1-3Q(\sigma-1)}{\sigma+1} & \text{if } \frac{3Q-1}{3Q-2} \leq \sigma \leq \frac{3Q-2}{3Q-1} \\ 1 & \text{if } \sigma > \frac{3Q-2}{3Q-1}. \end{cases} \tag{17}$$

Comparing this to the solution in the symmetric case in Propositio 6 is instructive. There, the optimal threshold was $R^* = 1$ for $\sigma < 1$. In the asymmetric example above we still get the corner $R^* = 1$ when σ is sfficiently large. Analogously, we still get a an optimal threshold of $R^* = 0$ when σ is sufficiently small. However, there is a wedge around $\sigma = 1$ for which the interior solution obtains. Just as in Proposition 6, setting $R = 1$ makes a buy recommendation fully objective. Consequently, the subjective value of the recommendation Δ_S^B is zero. Further, the probability of the buy recommendation is minimized at $\pi_B = q_H$. However, since the objectively good product is so likely, it' is still optimal to proceed in this way. As σ falls, however, the objective value of the buy recommendation falls, while the subjective value increases for any $R \neq 1, 0$. Hence, at some point the value of the objective content no longer dominates the value of the subjective content. Thus, it becomes optimal to lower the threshold, thus also including (on average positive) subjective information in the buy recommendation by lowering the threshold. Moreover, this increases the probability of a buy recommendation, which is also increasing the value of the recommendation system. As σ falls further, the objective value of the buy recommendation falls further, while the subjective value increases further eventually leading to the case where the subjective value dominates the objective one. Finally, when the objectively good product has a sufficiently high probability, i.e., when σ is sufficiently high, it is optimal to set $R = 0$, thus fully objectifying the now relatively likely don't-buy recommendation.

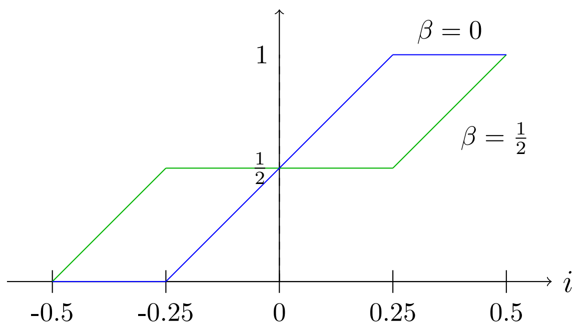

Figure 2 We have fixed $R = 3/4$. The green and blue line both correspond to symmetric type distributions. For F given by the green line, we have $\beta = 1/2$, whereas $\beta = 0$ for the blue line. All probability mass is on $[-1/4, 1/4]$ for $\beta = 0$ while all mass is on $[-1/2, -1/4] \cup [1/4, 1/2]$ for $\beta = 1/2$. Hence, as we move from $\beta = 0$ to $\beta = 1/2$ we are employing a mean-preserving spread of the distribution of types.

6.1 The Effect of the Degree of Subjectivity and of Polarized Populations

Having gained a clear understanding of the optimal design of the recommnendation threshold, we turn our attention to two comparative-static exercises. We first study the effect the degree of subjectivity Q has on the value of the recommendation system and then consider changes in the polarization of the population. For this to be tractable, we restrict attention throughout this section to the symmetric-population case.

6.1.1 The Effect of a Mean-preserving Spread

In the discussion so far, we have kept the distribution of types F fixed, and by varying β, we have varied the threshold R with the goal of reaching some objective. We will now consider the opposite case, where we fix the threshold R and vary the distribution F.[11] Specifically, we ask how a mean-preserving spread of the distribution of types affects the value of the recommendation, thus conducting comparative statics on the polarization of the population. For the symmetric population, such a mean-preserving spread can also be captured by varying β. When $R > 1/2$, an increase in β corresponds to a mean-preserving spread of the distribution of types, as illustrated in Figure 2. When $R < 1/2$, a decrease in β corresponds to a mean-preserving spread of the distribution of types. Together, we obtain the following result.

Proposition 8 *Suppose Assumption 1 holds.*

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

[11]Note that for $R = 1/2$ 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 $\beta \geq 1/2$.

Consider the case $\sigma < 1$ with a threshold $R > 1/2$ to gain some intuition for the result. From Proposition 6 we know that in that case the value of the recommendation system is maximized for $\beta = 1$, which yields an objective don't-buy recommendation, which is observed with probability $\pi_D = q_L$. Now, with a threshold of $R > 1/2$, a buy recommendation is given for objectively good products and for subjective products when the sender is sufficiently extreme, i.e., for very low or very high i. As noted above, increasing β corresponds to a mean-preserving spread of the distribution of types given the threshold R, which means that the buy recommendations are more likely to be the result of subjective products, thus increasing the probability of buy recommendations while lowering their objective value. However, for $\sigma < 1$ this is the value-maximizing thing to do, as objectively good products are relatively unlikely in that case. Thus, given a high recommendation threshold of $R > 1/2$, a more polarized population benefits more from the recommendation system. Conversely, when the recommendation threshold is relatively low with $R < 1/2$, a less polarized population centered around the middle benefits more from the recommendation system than a more polarized one.

6.1.2 The Effect of the Degree of Subjectivity

Finally, we consider the role of the degree of subjectivity in the product distribution. Given the discussion above, one might expect that a recommendation system is more valuable when subjective products are relatively unlikely, that is, for low degrees of subjectivity Q. The following result shows that this notion is only partly correct.

Corollary 3

(i) If $3\sigma-\beta-\sigma^2+\sigma^2\beta > 0$, then the value of the recommendation system Q 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 the maximal value of the recommendation system is achieved for the interior value $Q^* := \frac{3\sigma-\beta-\sigma^2+\sigma^2\beta}{4\sigma-4\beta-4\sigma^2+4\sigma^2\beta} \in (0, \frac{1}{2})$.

To gain some intuition for this result, it is useful to rewrite the condition to obtain an interior solution to

$$\frac{\sigma(3-\sigma)}{1-\sigma^2} < \beta \quad \text{if } \sigma < 1,$$
$$\frac{\sigma(3-\sigma)}{1-\sigma^2} > \beta \quad \text{if } \sigma > 1.$$

Recall from Proposition 6 that for $\sigma > 1$ it is optimal to set $\beta = 0$, corresponding to $R = 1$. That is, it is optimal to make the buy recommendation fully objective. The above condition tells us that there is an interior solution if and only if the β is sufficiently small and thus sufficiently close to that optimal solution. Suppose β is sufficiently small so that the value of the recommendation system is maximized at an interior Q and consider the corner cases of $Q \to 0$ and $Q \to 1/2$. In the former case, subjective products are virtually non-existent, and the only thing a buy recommendation does is confirm that the product is objectively good, which was already a priori more likely. As such, the receiver does not learn much from the recommendation system, reducing its value. In the latter case of $Q \to 1/2$ with a sufficiently small β and thus sufficiently high R, very few (albeit relatively objective in content) buy recommendations are sent, so that the system is once more not very valuable. At an interior Q, these effects balance out so that buy recommendations are sent sufficiently often and convey relatively objective and, thus, valuable information.

To complete the picture, suppose β is not sufficiently small, so that the value of the system is decreasing in Q. In that case, the relatively high β and thus low R mean that the buy recommendation is fairly subjective in its informational content. Given that $\sigma > 1$, this is not desirable for the receiver. As Q decreases, the subjective products become less likely, thus making buy recommendations more objective and more likely and thereby increasing the value of the recommendation system.

7 Extensions

Our benchmark model relied on several assumptions. First, there were only two levels of recommendations (”buy“ and ”don't buy“). Second, the recommendation threshold was exogenously given. Third, each receiver only has access to one recommendation. We now relax each of these assumptions in turn.

7.1 Multiple levels of recommendations

We now consider an alternative recommendation system that nests the benchmark model. There are two threshold levels $R_2 \geq R_1$ both in $[0, 1]$ and three types of recommendations $r \in \{D, N, B\}$ which we interpret as “don't buy”, “neither buy or don't buy” and “buy”. A buy recommendation obtains if the payoff was above R_2, a don't buy recommendation obtains if the payoff was below R_1 and neither obtains for payoffs in (R_1, R_2).

In Appendix A.1, we provide the resulting posteriors. For buy and don't-buy recommendations, they 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 rating rule out the objectively bad and the objectively good product, respectively. The intermediate rating $r = N$, however, rules out *both* objective products and is thus fully subjective in its informational content.

In the modified setting, it is natural to say that a receiver accepts a recommendation if he buys after a buy recommendation, does not buy after a don't buy recommendation, whereas the behavior in the intermediate case is arbitrary.

To maintain tractability in the face of the high number of different cases that can arise, we assume a symmetrically distributed population and equally likely subjective products. Thus, we assume that Assumption 1 holds and that $q_1 = q_2 = Q$. With this in hand, the posteriors above simplify substantially (see Section A.1), and it is easy to verify that the receiver will always buy a product with the highest rating $r = B$, and never buy one with the lowest rating $r = D$. A product with the middle rating $r = N$ is bought if and only if $1/2 \geq q_H + Q \Leftrightarrow q_L \geq q_H$. Intuitively, the intermediate rating $r = N$ rules out both objective products. Whether this is overall good or bad news depends on the relative likelihood of the objectively good and bad products. Indeed, ruling out the objective products is good news if and only if the bad product is more likely than the good one. Finally, observe that because of the two symmetry assumptions, all of this was independent of the receiver's type, so that all types will behave in the same way. Thus, the value of this recommendation system can be written as

$$V(\beta,\gamma) = (q_H + 2Q\beta)\left(\frac{q_H + q\beta}{q_H + 2q\beta} - (q_H + Q)\right) + [1 - (q_H + 2Q\beta) - (1 - q_H - 2Q\gamma)]\max\{1/2 - (q_H + Q), 0\}.$$

The term in the first line corresponds to $\pi_B \Delta_O^B$ in our benchmark model and captures the value that is generated when the highest rating $r = B$ is received. The second term captures the additional value that may come from the intermediate rating $r = N$, which only arises if $q_L \geq q_H$. As such, the problem closely mirrors that in the benchmark model. Before we state the corresponding result, observe that R_1 and R_2 are implicitly determined by β and γ and since $R_2 \geq R_1$ we must have $\gamma \geq \beta$.

Proposition 9 *The optimal thresholds R_1^* and R_1^*, are such that*

$$\gamma^* = 1 \text{ if } \sigma < 1$$
$$\beta^* = 0 \text{ if } \sigma > 1$$

Thus, the optimal recommendation system, in many ways, mimics that in the benchmark model, where Proposition 6 is the point of comparison. Essentially, the designer again try to maximize the probability of an objective rating. Depending on what objective product is a priori more likely, it is either the high or the low ranking which is turned fully objective by choosing the right thresholds. Since the intermediate rating is always subjective, it can effectively be bundled together with the not-objective rating, so that one of the two thresholds can always be chosen freely without affecting the value of the recommendation system. This is, of course, an artifact of the symmetric population and equally likely subjective products, where the intermediate rating conveys no information beyond ruling out the objective products.

7.2 Multiple recommendations

We now consider the situation where the receiver obtains not only a single but multiple recommendations. Just as before, they are sent from randomly drawn individuals from the population.

In Appendix , we provide the posteriors, depending on the number of buy recommendations (b) and don't-buy recommendations (d). Just as before, a single buy recommendation is enough to rule out the objectively bad product and a single don't-buy recommendation is enough to rule out the objectively good product. Consequently, when the receiver receives mixed recommendations, that is, both buy and don't-buy recommendations, she rules out *both* objective products. We are first interested in determining what the receiver can learn about the product quality as the number of recommendations increases.

Proposition 10 *Consider the case of a symmetric population with any thresholds $R \in (0,1)$. If the receiver obtains only buy or don't-buy recommendations, respectively, then her posterior belief converges to the objectively good or bad product, respectively, as the number of recommendations increases. 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 $b, d > 0$.*

Thus, in the case of a symmetric population, it is never possible to learn which of the two subjective products is being recommended. This is a direct result of the symmetry assumption of the population: irrespective of the threshold, the probability of receiving a (don't-) buy recommendation is the same for both subjective products so that the receiver cannot learn anything about the relative likelihood of the two products. Moreover, the ability to fully learn about the objective products in the case of homogenous recommendation is tightly connected to the recommendation system in place, as evidenced by the following corollary.

Corollary 4 *Consider the case of a symmetric population with threshold $R \in \{0, 1\}$. Then, the receiver will never get mixed recommendations and obtaining more than one recommendation does not lead to additional learning. Moreover, the product is only revealed if it is objectively good or bad, respectively, when the threshold is $R = 1$ or $R = 0$, respectively.*

With the extreme thresholds, the subjective products are indistinguishable from the objectively good or bad product, respectively, when the threshold is $R = 0$ or $R = 1$. Hence, in those cases, everything the receiver can learn is learned from the first recommendation.

The above results suggest, not very surprisingly, that receiving more recommendations allows the receiver to learn more about the recommended product quality. However, it is not clear how this translates into the value of the recommendation system. To make the point as starkly as possible, we consider the case of "infinite learning" and compare it to the case of receiving a single recommendation as in the main text. To do so, we first need to characterize the value of the recommendation system in the limit with infinitely many recommendations. It follows from Proposition 10 that the receiver can learn most from multiple recommendations when the threshold is not extreme. In that case, the objective products are revealed with probability 1. In the case of subjective products, nothing can be inferred beyond the fact that the product is subjective. Thus, it is not clear whether a receiver i would want to buy the product with mixed recommendations or the outside option. As we will see, the question of which receivers buy the product is quite relevant for the value of this "infinite learning" recommendation system. First, let us characterize the receiver's behavior. Let

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

so we can formulate the following result.

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 objectively good;

(ii) no receiver type $i \in [-1/2, 1/2]$ buys the product if it is objectively bad;

(iii) if the product is subjective,

(a) and $q_1 > q_2$, all receiver types $i \geq \tilde{i}_\infty$ buy the product;

(b) *and* $q_1 < q_2$*, all receiver types* $i \leq \tilde{i}_\infty$ *buy the product.*

(c) *and* $q_1 = q_2$*, all receiver types* $i \in [-1/2, 1/2]$ *buy the product if* $q_L \geq q_H$ *and none buy it if* $q_L < q_H$*.*

The first two cases of the lemma are obvious. To understand the last case, observe that the receiver knows that the recommended product is either (1,0) or (0,1). If the probability of the former is sufficiently high (Case, iii.a), types who care sufficiently strongly about the first dimension will buy it. Note, however, that it might be that not even type $i = 1/2$ buys it, if q_H is sufficiently high, in which case the outside option may be more attractive. It is precisely this objective effect that leads to Case iii.c when $q_1 = q_2$. Finally, an analogous interpretation applies to Case iii.b. With this in hand, we obtain that the value of the "infinite learning" recommendation system is given by

$$V_\infty = q_H \int_{-1/2}^{1/2} (1 - V_0(i)) dF(i) + (q_1 + q_2) \int_I \left(\frac{1}{2} + i \frac{q_1 - q_2}{q_1 + q_2} - V_0(i) \right) dF(i),$$

where I is the set of types who buy the recommended good according to Lemma 4. We can now determine when infinite learning is (not) valuable.

Proposition 11 *Suppose the population distribution is symmetric. Then, having multiple recommendations does not increase the value of the recommendation system if* $q_1 = q_2$ *or if* $q_1 \neq q_2$ *and*

$$q_H \notin \left[\min \left\{ \frac{1 - q_1 - q_2}{2}, \frac{-3q_2^2 - 4q_1q_2 + 3q_2 - q_1^2 + q_1}{4(q_1 + q_2)} \right\}, \max \left\{ \frac{1 - q_1 - q_2}{2}, \frac{-3q_1^2 - 4q_1q_2 + 3q_1 - q_2^2 + q_2}{4(q_1 + q_2)} \right\} \right]$$

To grasp this result intuitively, consider first the receiver's behavior in the single-recommendation system. With a threshold of $R = 0$, the receiver buys the recommended product when it is objectively good or subjective. In contrast, with a threshold of $R = 1$, the receiver buys the recommended product only when it is objectively good. In the optimal system, the former happens when $q_H < q_l$ and the latter if $q_H \geq q_L$. Now, in the "infinite learning" system, behavior coincides when the product is objective in quality. Thus, the key point lies in behavior regarding subjective products, which corresponds to the case of mixed recommendations. When q_L is sufficiently high, the outside option looks relatively bleak and so any type will buy either of the subjective products. Analogously, when q_H is sufficiently high, any type will buy the recommended product. Thus, when the objective probabilities are sufficiently extreme, the behavior coincides and receiving multiple recommendations does not increase the value of the recommendation system. Only when the objective probabilities are sufficiently close to each other, behavior do differ. For instance, when $q_1 > q_2$, low types who value the second dimension more strongly, which is less likely to be high in that case, will opt for the outside option rather than the subjective product, so that the value of the system with multiple recommendations is higher than that of the single-recommendation system, as it provides more granular information.

7.3 Endogenous Recommendation Threshold

The recommendation threshold was so far an exogenously given value. As we have argued in Section 3.2, the threshold should not be taken too literally but rather interpreted as something being shaped by the designer of the platform or as a social norm that has evolved. Such a norm could be to give a buy recommend a product when it was better than initially expected. As such, the recommendation threshold is endogenous, which we consider in this extension. Thus, an agent i sends a buy recommendation whenever the payoff is at least the prior expected payoff $q_H + q_1(1/2 + i) + q_2(1/2 - i)$. In that case, posteriors take the following form ....

It is straightforward to verify that Propositions 1 and 2 still hold in this case. We are interested in how the value created by this recommendation system compares to those of a fixed threshold as in the main model.

For the sake of tractability, we assume equally likely subjective products $q_1 = q_2 = Q$ and consider a uniformly (and thus symmetrically) distributed population. Then, we obtain that $i_1 = \frac{2q_H+2q-1}{2}$ and $i_2 = \frac{1-2q_H-2q}{2} = -i_1$ thus implying that $F(i_2)-(1-F(i_1)) = (-i_1 + 1/2) - 1 + (i_1 + 1/2) = 0$. Using this we get

$$
\begin{aligned}
\Delta_S^B &= [(p_2^B(R) - Q) - (p_1^B(R) - Q)] \\
&= Q\left(\frac{F(i_2)}{q_H + Q(1 - F(i_1)) + F(i_2)} - \frac{(1 - F(i_1))}{q_H + Q(1 - F(i_1)) + qF(i_2)}\right) \\
&= 0
\end{aligned}
$$

so that all types follow the recommendation according to Proposition 2. Further, by Proposition 4 we obtain that the value of the recommendation system is given by

$$
V_E = (2Q - 1)\left(Q^2 + Q(2q_H - 1) - (1 - q_H)q_H\right).
$$

Naturally, since the recommendation threshold is endogenously determined, this expression does not depend on R. The question is how it compares to the system where the threshold is exogenous and optimally chosen as in the benchmark model.

Proposition 12 *Suppose $q_1 = q_2 = Q$ and $F(x) = x + 1/2$. Then the value of the recommendation system in the endogenous-threshold model is strictly lower than in the benchmark with the optimally chosen threshold whenever $\sigma \neq 1$ and equal when $\sigma = 1$.*

The result can be understood quite easily by observing that with equally likely subjective products, the outside option is type-independent and given by $q_H + Q$. Thus, the endogenously determined reference point corresponds to the model with an exogenously given reference point of $R = q_H + Q$. Since the solution to the benchmark model is to generically have a corner solution, the only case in which the value of the system with an endogenously determined threshold is optimal is when $\sigma = 1$, as then any threshold is optimal.

7.4 Revenue-maximizing platform

We now embed our recommendation system into a simple model of a revenue-maximizing platform. Suppose that one of the products the consumer can buy is being offered on a platform that designs the recommendation system. Consumers' expected payoff should thus be interpreted as net of the product's price. We assume that for every unit of the product sold via the platform, the platform gets a revenue cut of $\alpha > 0$. The platform's objective is to maximize its revenue. In the absence of a recommendation system, the product offered via the platform is a priori identical to the outside option and we assume that a fraction $\gamma \in [0,1]$ of consumers buy the product on the platform. In this setting, should the platform introduce a recommendation system and if so, what is the optimal threshold?

In general, we can deduce from Propositions 2 and 4 that, following the introduction of a recommendation system with a threshold R, the demand for the recommended product will be

$$D(R) = \begin{cases} \pi_B & \text{if } \Delta_S^B = 0 \\ \pi_B F(m) + (1-\pi_B)(1-F(m)) & \text{if } \Delta_S^B > 0 \\ (1-\pi_B)F(M) + \pi_B(1-F(M)) & \text{if } \Delta_S^B < 0, \end{cases}$$

where $M := \max\{-1/2, \tilde{i}\}$ and $m := \min\{1/2, \tilde{i}\}$. Further, we know from Proposition 3 that all types accept a recommendation for $R = 0$ and none do for $R = 1$, so that $D(R) = \pi_B$ in those cases. More precisely, we have $D(0) = q_H + q_1 + q_2$ and $D(1) = q_1 + q_2 + q_L$. For intermediate thresholds and depending on the distribution of types as well as the product probabilities, the demand will be somewhere in between. For instance, in the case of a symmetric population, we will always have $D(R) = \pi_B$, as all types accept a recommendation in that case. Then, it is optimal to set $R = 0$, as it ensures the highest probability of buy recommendations and thus yields the highest demand. For asymmetrically distributed populations, it could in principle be that the demand is higher for intermediate thresholds. Intuitively, while a higher threshold decreases the probability of a buy recommendation, it could lead to some types who reject the recommendation and thus buy the product whenever they receive a don't-buy recommendation. If this happens sufficiently often for sufficiently many types, an interior threshold could be optimal. The following result provides a sufficient condition for the extreme threshold $R = 0$ to be optimal for any population distribution.

Proposition 13 *Suppose $q_H \geq 1/2$. Then, the optimal threshold for the platform is given by $R = 0$, yielding a demand of $D(0) = q_H + q_1 + q_2$. Moreover, such a recommendation system is implemented if and only if $q_H + q_1 + q_2 \geq \gamma$ so that the induced demand exceeds the demand without a recommendation system.*

The result tells us that, in a sense, setting the lowest possible threshold yields the most persuasive recommendation system in terms of the probability of the recommended product being bought. Somewhat remarkably, this is at least sometimes also the most valuable recommendation system from the perspective of the population of receivers, as we know from Proposition 6 that the extreme threshold $R = 0$ can be optimal.[12]

Moreover, the result allows for a nice connection to Johnson and Myatt (2006), who study demand rotations which are distinct from demand shifts. Following a demand rotation, as the name suggests, the demand for a product for some types will increase while it decreases for others. In contrast, a demand shift yields an increase (or decrease) for all types. In our model, a buy recommendation may lead to a demand rotation: Following a buy recommendation with intermediate threshold r, some type may buy with more conviction than before while others may no longer buy the product. Proposition 13

[12]Note that in Proposition 6 the threshold of $R = 0$ is optimal when $\sigma < 1 \Leftrightarrow q_H < q_L$, which clearly does not meet the sufficient condition in Proposition 13. However, as noted in the text preceding the result, in the case of a symmetric population sales are always maximized with $R = 0$.

tells us, however, that it will never be optimal to implement a recommendation system that leads to a demand rotation when an objectively good product is relatively likely.

Finally, we can interpret the platform's problem as one of constrained Bayesian persuasion (Kamenica and Gentzkow, 2011). The sender (the platform) wants to persuade the receiver (population of consumers) to buy the product. However, the platform is constrained in what signal structures and thus distributions of posteriors it can induce, as they can only be generated through the recommendation system.

8 Conclusion

We have set out to study the role of preference heterogeneity in recommendations. Under what circumstances do people follow a recommendation? What determines the value of a recommendation system and when is it maximized? Our analysis highlights the importance of disentangling objective information from subjective preferences. We have seen that, depending on the distribution of preferences in the population, 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. It is thus important for the designer of a recommendation system to take into account how the choice of system, boiling down to the recommendation threshold in our model, influences the relative objectivity of recommendations.

Our results show that extreme recommendation systems can often be the most valuable, as they then manage, despite the presence of preference heterogeneity, to convey objective information with a high probability. Moreover, this value-maximizing recommendation system is sometimes the most persuasive model in terms of maximizing the probability that receivers buy the product in question. In particular, when objectively bad products are relatively likely, the objectives of a designer who wants to maximize the receivers' expected payoffs and of a designer interested in maximizing sales coincide.

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

A Proofs

Proof of Proposition 1

Accepting the don't-buy recommendation means not buying the recommended good but the product without recommendation. Observe that the condition for accepting a don't-buy recommendation, $U(0) - U(D) \geq 0$, is equivalent with

$$q_H + q_1(1/2 + i) + q_2(1/2 - i) \geq p_1^D(R)(1/2 + i) + p_2^D(R)(1/2 - i). \tag{18}$$

Replacing

$$q_s = \pi_B p_s(B) + \pi_D p_s(D) = (1 - \pi_D)p_s(B) + \pi_D p_s(D) \text{ for } s \in \{H, 1, 2\}$$

and using

$$p_H^B(R) = 1 - p_1^B(R) - p_2^B(R),$$

equation (18) simplifies to

$$1 - p_1^B(R) - p_2^D(R) \geq (1/2 + i)[p_2^B(R) - p_2^D(R) + p_1^D(R) - p_1^B(R)]. \tag{19}$$

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

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

Proceeding as above, one finds that this condition is equivalent to equation (19).

Proof of Lemma 1

The statement regarding the sign of Δ_S^B follows immediately. For the sign of Δ_O^B note that we can write

$$\begin{aligned}\Delta_O^B = & \frac{2q_H(1 - (q_H + q_1\phi_1(R) + q_2\phi_2(R)))}{2(q_H + q_1\phi_1(R) + q_2\phi_2(R))} \\ & + \frac{q_1(\phi_1(R) - (q_H + q_1\phi_1(R) + q_2\phi_2(R)))}{2(q_H + q_1\phi_1(R) + q_2\phi_2(R))} \\ & + \frac{q_2(\phi_2(R) - (q_H + q_1\phi_1(R) + q_2\phi_2(R)))}{2(q_H + q_1\phi_1(R) + q_2\phi_2(R))},\end{aligned}$$

the sign of which depends only on the numerator. Further, the derivative of the numerator with respect to $\phi_2(R)$ is given by $q_2(q_L - q_H)$ and the derivative of the numerator with respect to $\phi_1(R)$ is given by $q_1(q_L - q_H)$. Thus, suppose $q_L \geq q_H$ in which case the numerator is smallest for $\phi_1(R) = \phi_2(R) = 0$. It then reads

$$2q_H(1 - q_H) - q_1 q_H - q_2 q_H \geq q_H(1 - q_H - q_1 - q_2) \geq q_H q_L \geq 0.$$

Conversely, when $q_L < q_H$ the numerator is smallest when $\phi_1(R) = \phi_2(R) = 1$, in which case it reads

$$(2q_H + q_1 + q_2)(1 - q_L) \geq 0.$$

Proof of Lemma 2

The proof for $\Delta_O^D \leq 0$ is analogous to that part of the proof in Lemma 1 and thus omitted. For the second part, recall that

$$q_s = \pi_B p_s(B) + \pi_D p_s(D) = (1-\pi_D)p_s(B) + \pi_D p_s(D) \text{ for } s \in \{H,1,2\}.$$

With this, we show that $\pi_B \Delta_O^B = -\pi_D \Delta_O^D$. We have

$$\begin{aligned}
\pi_B \Delta_O^B &= \pi_B \left(p_H^B(R) - q_H + \frac{1}{2}\left(q_1(B) - q_1 + q_2(B) - q_2\right) \right) \\
&= \pi_B \Bigg(p_H^B(R) - \pi_B p_H^B(R) - \pi_D p_H^D(R) + \frac{1}{2}\left(p_1^B(R) - \pi_B p_1^B(R) - \pi_D p_1^D(R)\right) \\
&+ \frac{1}{2}\left(p_2^B(R) - \pi_B p_2^B(R) - \pi_D p_2^D(R)\right) \Bigg) \\
&= \pi_B \pi_D \left(p_H^B(R) - p_H^D(R) + \frac{1}{2}\left(p_1^B(R) - p_1^D(R) + p_2^B(R) - p_2^D(R)\right) \right) \\
&= -\pi_D \left(\pi_B(p_H^D(R) - p_H^B(R)) + \frac{1}{2}\left(\pi_B(p_1^D(R) - p_1^B(R)) + \pi_B(p_2^D(R) - p_2^B(R))\right) \right) \\
&= -\pi_D \Bigg((1-\pi_D)p_H^D(R) - \pi_B p_H^B(R) + \frac{1}{2}\left((1-\pi_D)p_1^D(R) - \pi_B p_1^B(R)\right) \\
&+ \frac{1}{2}\left((1-\pi_D)p_2^D(R) - \pi_B p_2^B(R)\right) \Bigg) \\
&= -\pi_D \left(p_H^D(R) - q_H + \frac{1}{2}\left(p_1^D(R) - q_1 + q_2(D) - p_2\right) \right) \\
&= -\pi_D \Delta_O^D.
\end{aligned}$$

Finally, to prove $\Delta_S^B = \Delta_S^D \Leftrightarrow \Delta_S^B = 0$ we once more use the identity

$$q_s = \pi_B p_s(B) + \pi_D p_s(D) = (1-\pi_D)p_s(B) + \pi_D p_s(D) \text{ for } s \in \{H,1,2\}.$$

We need to show that the two values are equal if and only if they are zero. To see this, note that

$$\begin{aligned}
&\Delta_S^D = \Delta_S^B \\
\Leftrightarrow & p_2^B(R) - q_2 - p_1^B(R) + q_1 = p_2^D(R) - q_2 - p_1^D(R) + q_1 \\
\Leftrightarrow & p_2^B(R) - p_1^B(R) = p_2^D(R) - p_1^D(R) \\
\Leftrightarrow & p_2^B(R)(1-\pi_B) - p_1^B(R)(1-\pi_B) = \pi_D p_2^D(R) - \pi_D p_1^D(R) \\
\Leftrightarrow & p_2^B(R) - p_1^B(R) = \pi_B p_2^B(R) + \pi_D p_2^D(R) - \pi_B p_1^B(R) - \pi_D p_1^D(R) \\
\Leftrightarrow & p_2^B(R) - p_1^B(R) = q_2 - q_1 \\
\Leftrightarrow & \Delta_S^B = 0
\end{aligned}$$

Next, observe that

$$\begin{aligned}
\frac{\partial \Delta_S^B}{\partial(1-\phi_1(R))} &= \frac{q_1(q_H + 2\phi_2(R)q_2)}{\pi_B^2} \geq 0 \\
\frac{\partial \Delta_S^D}{\partial(1-\phi_1(R))} &= -\frac{q_1(q_L + 2q_2(1-\phi_2(R))}{(1-\pi_B)^2} \leq 0.
\end{aligned}$$

Thus, Δ_S^B is increasing in $(1-\phi_1(R))$ and Δ_S^D is decreasing in $(1-\phi_1(R))$ and they intersect at 0, so that whenever one is positive, the other is negative.

Proof of Proposition 2

We proceed case by case and suppress the dependence on the recommendation to simplify notation.

(i) In this case the inequality in equation (6) reduces to

$$\begin{aligned}
&0 \leq (p_H - q_H) + (p_2 - q_2)\\
\Leftrightarrow\ &0 \geq q_H - \frac{q_H}{q_H + q_1\phi_1(R) + q_2\phi_2(R)} + q_2 - \frac{q_2\phi_2(R)}{q_H + q_1\phi_1(R) + q_2\phi_2(R)}\\
\Leftrightarrow\ &0 \geq q_H(q_H + q_1\phi_1(R) + q_2\phi_2(R) - 1) + q_2(q_H + q_1\phi_1(R) + q_2\phi_2(R)) - q_2\phi_2(R)\\
\Leftrightarrow\ &0 \geq q_1q_2 - q_Hq_L - q_2(q_1+q_L)F(1/2-R) - q_1(q_2+q_H)(1-\phi_1(R))
\end{aligned}$$

Further, observe that the inequality is most stringent when $\phi_2(R) = (1-\phi_1(R)) = 0$. Hence, for the statement to be true we need

$$0 \geq q_1q_2 - q_Hq_L.$$

Further, observe that the condition for leaving the odds between (1,0) and (0,1) unaffected can be written as

$$\begin{aligned}
&p_2 - q_2 = p_1 - q_1\\
\Leftrightarrow\ &\frac{q_2\phi_2(R)}{q_H + q_1\phi_1(R) + q_2\phi_2(R)} - q_2 = \frac{q_1\phi_1(R)}{q_H + q_1\phi_1(R) + q_2\phi_2(R)} - q_1
\end{aligned}$$

which, when $\phi_2(R) = (1-\phi_1(R)) = 0$, simplifies to

$$-q_2 = \frac{q_1}{q_H + q_1} - q_1 \Leftrightarrow q_2 = -\frac{q_1(1-q_1-q_H)}{q_1+q_H}$$

Plugging this into the above inequality yields

$$0 \geq q_1q_2 - q_Hq_L \Leftrightarrow 0 \geq -\frac{q_1^2(1-q_1-q_H)}{q_1+q_H} - q_Hq_L,$$

which is satisfied.

(ii) Since the recommendation shifts the odds from $(0,1)$ to $(1,0)$, we have $p_2 - q_2 < p_1 - q_1$. Note that

$$p_H - q_H + p_2 - q_2 \geq 0 \Leftrightarrow \Delta_O \geq -\Delta_S/2 \Leftrightarrow \tilde{i} \leq -1/2.$$

Thus, given the condition $p_H - q_H + p_2 - q_2 \geq 0$ in case (a), all types accept the recommendation. It follows immediately, that in case (b) all types $i \geq \tilde{i}$ accept the recommendation.

(iii) Since the recommendation shifts the odds from $(1,0)$ to $(0,1)$, we have $p_2 - q_2 > p_1 - q_1$. Note that

$$p_H - q_H + p_1 - q_1 \geq 0 \Leftrightarrow \Delta_O \geq \Delta_S/2 \Leftrightarrow \tilde{i} \geq 1/2.$$

Thus, given the condition $p_H - q_H + p_1 - q_1 \geq 0$ in case (a), all types accept the recommendation. It follows immediately, that in case (b) all types $i \leq \tilde{i}$ accept the recommendation.

Proof of Proposition 3

Recall from equation (6) the optimal acceptance condition

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

First, observe that

$$\lim_{R\to 1} p_1^B(R) = \lim_{R\to 1} p_2^B(R) = 0$$

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

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

Consider $i \geq 0$ and note that for $q_2 \geq q_1$ this is always satisfied. Further, for $q_1 > q_2$ the condition is most stringent for $i = 1/2$ in which case it is equivalent to $1 \geq q_1 + q_H$, which is always satisfied. The case for $i < 0$ is analogous.

In the case $R \to 0$ we get

$$\lim_{R\to 0} p_i(B) = \frac{q_i}{1 - q_L}$$

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

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

Consider $i \geq 0$ and note that for $q_1 \geq q_2$ this is always satisfied. Further, for $q_2 > q_1$ the condition is most stringent for $i = 1/2$ in which case it is equivalent to $q_H \geq q_1$, which is always satisfied. The case for $i < 0$ is analogous.

Proof of Proposition 4

(i) If $\Delta_S^B = 0$, then $V(R) = \int_{-1/2}^{1/2} (V_A(i) - V_0(i))dF(i)$. Inserting (12) and (13) yields

$$\begin{aligned} V(R) &= \pi_B \int_{-1/2}^{1/2} \left(p_H^B(R) + (1/2 + i)p_1^B(R) + (1/2 - i)p_2^B(R)\right) - (q_H + (1/2 + i)q_1 + (1/2 - i)q_2)dF(i) \\ &= \pi_B \int_{-1/2}^{1/2} \underbrace{(p_H^B(R) - q_H) + \frac{p_1^B(R) - q_1}{2} + \frac{p_2^B(R) - q_2}{2}}_{=\Delta_O^B} dF(i) \\ &+ \pi_B \int_{-1/2}^{1/2} i(\underbrace{[(p_2^B(R) - q_2) - (p_1^B(R) - q_1)]}_{=\Delta_S^B = 0})dF(i) \\ &= \pi_B \Delta_O^B \end{aligned}$$

(ii) If $\Delta_S^B > 0$, then $V(R) = \int_{-1/2}^{m} V_A(i)dF(i) + \int_{m}^{1/2} V_N(i)dF(i) - \int_{-1/2}^{1/2} V_0(i)dF(i)$. Inserting (12), (13) and (14) and proceeding analogously as in case (i) gives the expression.

(iii) If $\Delta_S^B < 0$, then $V(R) = \int_{-1/2}^{M} V_N(i)dF(i) + \int_{M}^{1/2} V_A(i)dF(i) - \int_{-1/2}^{1/2} V_0(i)dF(i)$. Inserting (12), (13) and (14) and proceeding analogously as in case (i) gives the expression.

Proof of Proposition 5

According to Proposition 2, all types accept the recommendation if

$$\min\{\Delta_O^B + \frac{\Delta_S^B}{2}, \Delta_O^B - \frac{\Delta_S^B}{2}\} \geq 0. \tag{20}$$

To see that this condition is fulfilled, first observe that

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

Now, if $q_2 - (q_H + q_2)(q_1 + q_2) \geq 0$, then this is clearly positive. If, in contrast, we have $q_2 - (q_H + q_2)(q_1 + q_2) < 0$, we obtain

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

so that $\Delta_O^B + \frac{\Delta_S^B}{2} \geq 0$ even in this case. Proceeding analogously, one also obtains that $\Delta_O^B - \frac{\Delta_S^B}{2} = p_H^B(R) - q_H + p_1^B(R) - q_1 \geq 0$. Together with Proposition 2 this implies that all types accept the recommendation, so that

$$\begin{aligned}
V(\beta) &= \int_{-1/2}^{1/2} (q_H + \beta((1/2 + i)q_1 + (1/2 - i)q_2))\, dF(i) \\
&\quad - \int_{-1/2}^{1/2} (q_H + \beta(q_1 + q_2))\,(q_H + (1/2 + i)q_1 + (1/2 - i)q_2)\, dF(i) \\
&= q_H + \frac{\beta(q_1 + q_2)}{2} - (q_H + \beta(q_1 + q_2))\left(q_H + \frac{(q_1 + q_2)}{2}\right) \\
&\quad + (q_1 - q_2)\,[\beta - (q_H + \beta(q_1 + q_2))] \int_{-1/2}^{1/2} i\, dF(i) \\
&= q_H + \frac{\beta(q_1 + q_2)}{2} - (q_H + \beta(q_1 + q_2))\left(q_H + \frac{q_1 + q_2}{2}\right),
\end{aligned}$$

because $F(-x) = 1 - F(x)$ implies $\int_{-1/2}^{1/2} i\, dF(i) = 0$. Finally, replacing $(q_1 + q_2)/2$ by Q and q_H by $(1 - 2Q)\sigma/(1 + \sigma)$ yields the result.

Proof of Lemma 3

(i) Note that $\pi_B = q_H + 2Q\beta$ so that the result follows immediately.

(ii) We have

$$\begin{aligned}\Delta_O^B &= (p_H^B(R) - q_H) + \frac{p_1^B(R) - q_1}{2} + \frac{p_2^B(R) - q_2}{2} \\ &= \frac{q_H + \beta Q}{q_H + 2\beta Q} - q_H - Q\end{aligned}$$

so that

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

(iii) We have

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

Proof of Proposition 6

We have

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

which is negative for $\sigma > 1$, positive for $\sigma < 1$ and zero for $\sigma = 1$. Thus, V is maximized for $\beta = 0$ and $\beta = 1$, respectively, in the first two cases and for any β in the last case. The result then follows by the continuity and full-support assumption on F together with $\beta = F(1/2 - R)$.

Proof of Proposition 7

Suppose $a < 1$, in which case we have by Proposition 4

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

In the case $\tilde{i} > 1/2$ we can write this as

$$V(R) = q_H + \frac{aQ(1-R^a) + Q(1-R)^a}{a+1} - \big(q_H + Q(1-R^a) + q(1-R)^a\big)(q_H + Q)$$

Thus, we can write

$$V'(R) = \begin{cases} \frac{aQ\left((a+1)\left(R^{a-1} + (1-R)^{a-1}\right)(Q+q_H) - aR^{a-1} - (1-R)^{a-1}\right)}{a+1} & \text{if } \tilde{i} > 1/2 \\ \ldots & \text{if } \tilde{i} \leq 1/2 \end{cases}$$

where ... is not written out for simplicity as we do not need it. Observe that Proposition 3 implies $\tilde{i} > 1/2$ if $R \in \{0, 1\}$, so that for R sufficiently close to 0 and 1, respectively, $V'(R)$ is given by the first line above. We can then write

$$\begin{aligned} V'(R) > 0 &\Leftrightarrow (a+1)\left(R^{a-1} + (1-R)^{a-1}\right)(Q + q_H) - aR^{a-1} - (1-R)^{a-1} > 0 \\ &\Leftrightarrow (a+1)(Q+q_H) > \frac{aR^{a-1} + (1-R)^{a-1}}{R^{a-1} + (1-R)^{a-1}} = \frac{a(1-R)^{1-a} + R^{1-a}}{(1-R)^{1-a} + R^{1-a}}, \end{aligned}$$

so that

$$\lim_{R \to 0} V'(R) > 0 \Leftrightarrow (a+1)(Q+q_H) > a$$

and analogously

$$\lim_{R \to 1} V'(R) < 0 \Leftrightarrow (a+1)(Q+q_H) < 1.$$

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

$$\frac{1}{a+1} > Q + q_H > \frac{a}{a+1}.$$

Proceeding analogously for the case of $a > 1$ and noting that Proposition 3we obtain the corresponding result, completing the proof.

Proof of Proposition 8

We have argued in the text that for $R > 1/2$ a mean-preserving spread corresponds to an increase in β while for $R < 1/2$ it corresponds to a decrease in β. From the proof of Proposition 6 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

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

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 already proves the first part of (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,

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

or

$$\sigma > \frac{\beta}{\beta - 1},$$

which always holds for $\beta - 1 < 0$. For the case of $\beta = 1$, one can directly verify that for any σ satisfying the condition for case (ii), we also obtain a strictly interior solution.

A.1 Multiple Levels of Recommendations

The posteriors are given as follows. Following a buy recommendation we have

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

while 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, neither buy nor don't buy yields

$$p_1(N) = 0 \qquad p_3(N) = \frac{q_2\Gamma_2}{q_1\Gamma_1 + q_2\Gamma_2}$$

$$p_2(N) = \frac{q_1\Gamma_1}{q_1\Gamma_1 + q_2\Gamma_2} \qquad p_4(N) = 0.$$

where $\Gamma_1 \equiv F(R_2 - 1/2) - F(R_1 - 1/2)$ and $\Gamma_2 \equiv F(1/2 - R_1) - F(1/2 - R_2)$.

Symmetric Case

Let $F(1/2 - R_2) = 1 - F(R_2 - 1/2) = \beta$ and $F(1/2 - R_1) = 1 - F(R_1 - 1/2) = \gamma$. We have

$$\begin{aligned}
p_1^B(R) &= \frac{q_H}{q_H + 2Q\beta}, & p_1(N) &= 0, & p_1^D(R) &= 0, \\
p_2^B(R) &= \frac{Q\beta}{q_H + 2Q\beta}, & p_2(N) &= \frac{1}{2}, & p_2^D(R) &= \frac{Q(1-\gamma)}{1 - q_H - 2Q\gamma}, \\
p_3(B) &= \frac{Q\beta}{q_H + 2Q\beta}, & p_3(N) &= \frac{1}{2}, & p_3(D) &= \frac{Q(1-\gamma)}{1 - q_H - 2Q\gamma}, \\
p_4(B) &= 0, & p_4(N) &= 0, & p_4(D) &= \frac{q_L}{1 - q_H - 2Q\gamma}.
\end{aligned}$$

Proof of Proposition 9

We need to consider two cases. First, suppose $\sigma < 1$ so that $1/2 > q_H + Q$, yielding

$$\begin{aligned}
V(\beta,\gamma) =& (q_H + 2Q\beta)\left(\frac{q_H + q\beta}{q_H + 2q\beta} - (q_H + Q)\right) \\
&+ [1 - (q_H + 2Q\beta) - (1 - q_H - 2Q\gamma)](1/2 - (q_H + Q))
\end{aligned}$$

Taking the respective derivatives we obtain

$$\begin{aligned}
\frac{\partial V}{\partial \beta} &= 0 \\
\frac{\partial V}{\partial \gamma} &= 2Q(1/2 - q_H - Q) \geq 0.
\end{aligned}$$

In contrast, for $\sigma \geq 1$ so that $1/2 \leq q_H + Q$ we have

$$V(\beta,\gamma) = (q_H + 2Q\beta)\left(\frac{q_H + q\beta}{q_H + 2q\beta} - (q_H + Q)\right)$$

Taking the respective derivatives we obtain

$$\begin{aligned}
\frac{\partial V}{\partial \beta} &= Q(q_L - q_H) \leq 0 \\
\frac{\partial V}{\partial \gamma} &= 0.
\end{aligned}$$

The result follows directly from here.

Endogenous Recommendation Threshold: Proof of Proposition 12

In this model, the posteriors are given as follows:

$$\begin{aligned}
p_H^B(R) &= \frac{q_H}{q_H + q_1(1 - F(i_1)) + q_2 F(i_2)}, \\
p_1^B(R) &= \frac{q_1(1 - F(i_1))}{q_H + q_1(1 - F(i_1)) + q_2 F(i_2)}, \\
p_2^B(R) &= \frac{q_2 F(i_2)}{q_H + q_1(1 - F(i_1)) + q_2 F(i_2)} \\
p_L^B(R) &= 0,
\end{aligned}$$

where $i_1 = \frac{2q_H+q_1+q_2-1}{2(1-q_1+q_2)}$ and $i_2 = \frac{1-2q_H-q_1-q_2}{2(1-q_2+q_1)}$.[13] Analogously, we obtain the posteriors of a don't-buy recommendation as

$$\begin{aligned}
p_H^D(R) &= 0,\\
p_1^D(R) &= \frac{q_1F(i_1)}{q_1F(i_1)+q_2(1-F(i_2))+q_L},\\
p_2^D(R) &= \frac{q_2(1-F(i_2))}{q_1F(i_1)+q_2(1-F(i_2))+q_L},\\
p_L^D(R) &= \frac{q_L}{q_1F(i_1)+q_2(1-F(i_2))+q_L},.
\end{aligned}$$

Choosing the optimal threshold for the model with endogenous threshold the value of the system reads

$$V(R^*) = \begin{cases} q_H(1-q_H-q) & \text{if } q_H \geq q_L \\ (q_H+q)(1-q_H-2q) & \text{if } q_H < q_L. \end{cases}$$

We need to compare this to (see Lemma **??**)

$$V(R) = q_H + Q(1-R) - \big(q_H + 2Q(1-R)\big)(q_H+Q).$$

Consider first $q_H \geq q_L$ so that $V(R^*) = q_H(1-q_H-q)$. Then, we have

$$\begin{aligned}
V(R^*) - V^E &= q_H(1-q_H-q) - (q_H + q(1-q_H-q)) + (q_H + 2q(1-q_H-q))(q_H+q)\\
&= -q(1+2q^2-3q_H+2q_H^2+q(4q_H-3))
\end{aligned}$$

Now, take the derivative of this with respect to q_H to obtain $-q(4q_H-3+4q)$ and observe that this is negative whenever $q_H \geq 3/4-q$ which is satisfied for $q_H \geq q_L$. Thus, for $q_H = 1-2q$ we have $V(R^*) - V^E \geq -q^2(4q-3) \geq 0$. Proceeding analogously for the case of $q_H < q_L$ completes the proof.

A.2 Multiple Recommendations

The posterior when having received the set of recommendations (b,d) with $b+d>0$ then reads

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

[13] To obtain for instance i_1 we need to solve $1/2+i_1 = q_H + q_1(1/2+i_1) + q_2(1/2-i_1)$.

Proof of Proposition 10

Suppose the receiver gets only buy recommendations and that $b \to \infty$. Then, for $R \in (0,1)$ we must have $\phi_1(R) = \phi_2(R) \in (0,1)$ so that $\lim_{b\to\infty}(\phi_i(R))^b = 0$ for $i = 1,2$. Hence, $\lim_{b\to\infty} p_H(b,0) = 1$. Proceeding analogously we obtain the statement for the case of only don't-buy recommendations.

Assuming mixed recommendations, we clearly have $p_H(b,d) = p_L(b,d) = 0$ for any $b,d > 0$. Further, because $\phi_1(R) = \phi_2(R)$ we get for $i = 1,2$ that

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

yielding the statement.

Proof of Corollary 4

To see this, consider for instance the posterior for the case $R = 0$ noting that mixed reviews are not possible in this case. We have:

$$p_H(b,d) = \begin{cases} \frac{q_H}{q_H+q_1+q_2} & \text{if } d = 0 \\ 0 & \text{if } b = 0 \end{cases}$$

$$p_1(b,d) = \begin{cases} \frac{q_1}{q_H+q_1+q_2} & \text{if } d = 0 \\ 0 & \text{if } b = 0 \end{cases}$$

$$p_2(b,d) = \begin{cases} \frac{q_2}{q_H+q_1+q_2} & \text{if } d = 0 \\ 0 & \text{if } b = 0 \end{cases}$$

$$p_L(b,d) = \begin{cases} 0 & \text{if } d = 0 \\ 1 & \text{if } b = 0 \end{cases}$$

Proof of Lemma 4

Observe that the expected payoff of buying the recommended good with mixed recommendations reads

$$\frac{1}{2} + i\frac{q_1 - q_2}{q_1 + q_2}$$

so that the result obtains from appropriately rearranging the inequality

$$\frac{1}{2} + i\frac{q_1 - q_2}{q_1 + q_2} \geq V_0(i).$$

Proof of Proposition 11

Suppose $q_1 = q_2 = Q$. Then, the value of the "infinite learning" recommendation system is given by

$$V_\infty = q_H(1 - q_H - Q) + 2Q\mathbb{I}_{q_L \geq q_H}\left(\frac{1}{2} - q_H - Q\right),$$

which coincides with the value of the single-recommendation system for the optimal β, depending on whether q_L or q_H is bigger.

Suppose $q_1 > q_2$ so that by Lemma 4 all types $i \geq \tilde{i}_\infty$ buy the product. Further, assume $q_H \geq q_L$, which is equivalent to $q_H \geq (1 - q_1 - q2)/2$. Then, no type buys a subjective, recommended product if $\tilde{i}_\infty \geq 1/2$, which is equivalent to $q_H \geq (-3q_1^2 - 4q_1q_2 + 3q_1 - q_2^2 + q_2)/(4(q_1 + q_2))$. Thus, if

$$q_H \geq \max\left\{\frac{1 - q_1 - q_2}{2}, \frac{-3q_1^2 - 4q_1q_2 + 3q_1 - q_2^2 + q_2}{4(q_1 + q_2)}\right\}$$

no type buys a subjective, recommended product. Observe that in that case, the value of the "infinite learning" recommendation system, is given by

$$q_H - q_H\left(q_H + \frac{q_1 + q_2}{2}\right),$$

which coincides with the value of the single-recommendation system for $\beta = 0$.

Now, suppose $q_H < q_L$, which is equivalent to $q_H < (1 - q_1 - q_2)/2$. Then, every type buys a subjective, recommended product if $\tilde{i}_\infty \leq -1/2$, which is equivalent to $q_H \leq (-3q_1^2 - 4q_1q_2 + 3q_1 - q_2^2 + q_2)/(4(q_1 + q_2))$. Thus, if

$$q_H \leq \min\left\{\frac{1 - q_1 - q_2}{2}, \frac{-3q_2^2 - 4q_1q_2 + 3q_2 - q_1^2 + q_1}{4(q_1 + q_2)}\right\}$$

all types buy a subjective, recommended product. In that case, the value of the "infinite learning" recommendation system is given by

$$q_H + \frac{(q_1 + q_2)}{2} - (q_H + (q_1 + q_2))\left(q_H + \frac{q_1 + q_2}{2}\right),$$

which coincides with the value of the single-recommendation system for $\beta = 1$.

Proceeding analogously, we obtain the same for the case of $q_1 < q_2$, which concludes the proof.

Revenue-Maximizing Platform: Proof of Proposition 13

Note that $q_H \geq 1/2$ implies $\pi_B \geq 1/2$, as $\pi_B = q_H + q_1\phi_1(R) + q_2\phi_2(R)$. Thus, we have for the case of $\Delta_S > 0$

$$1/2 \leq \pi_B \Leftrightarrow \pi_B F(m) + (1 - \pi_B)(1 - F(m)) \leq \pi_B \leq q_H + q_1 + q_2.$$

and similarly for the case of $\Delta_S < 0$.

$$1/2 \leq \pi_B \Leftrightarrow (1 - \pi_B)F(M) + \pi_B(1 - F(M)) \leq \pi_B \leq q_H + q_1 + q_2.$$

Together, these two inequalities imply the optimality of $R = 0$. The case $\Delta_S^B = 0$ is trivial.

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.

Bolton, G., B. Greiner, and A. Ockenfels (2013): "Engineering trust: reciprocity in the production of reputation information," *Management Science*, 59, 265–285.

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*, WorkingPaper.

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.

Ronayne, D. and D. Sgroi (2019): *Ignoring good advice*, University of Warwick, Centre for Competitive Advantage in the Global ....

Sundaram, R. K. (1996): *A first course in optimization theory*, Cambridge university press.

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*.

Weizsäcker, G. (2010): "Do we follow others when we should? A simple test of rational expectations," *American Economic Review*, 100, 2340–2360.

Yaniv, I. and E. Kleinberger (2000): "Advice taking in decision making: Egocentric discounting and reputation formation," *Organizational behavior and human decision processes*, 83, 260–281.