

Optimal Design for Social Learning

Yeon-Koo Che*

Johannes Hörner[†]

first draft: November 25, 2012; current draft: April 28, 2015

Abstract

This paper studies the design of a recommender system for organizing social learning on a product. To improve incentives for early experimentation, the optimal design trades off fully transparent social learning by over-recommending a product (or “spamming”) to a fraction of agents in the early phase of the product cycle. Under the optimal scheme, the designer spams very little about a product right after its release but gradually increases the frequency of spamming and stops it altogether when the product is deemed sufficiently unworthy of recommendation. The optimal recommender system involves randomly triggered spamming when recommendations are public—as is often the case for product ratings—and an information “blackout” followed by a burst of spamming when agents can choose when to check in for a recommendation. Fully transparent recommendations may become optimal if a (socially-benevolent) designer does not observe the agents’ costs of experimentation.

Keywords: experimentation, social learning, mechanism design.

JEL Codes: D82, D83, M52.

1 Introduction

Most of our choices rely on recommendations by others. Whether selecting movies, picking stocks, choosing hotels or shopping online, the experiences of others can teach us to make better decisions. Internet platforms are increasingly acting as recommenders, enabling users to learn from other users on a scale never before imagined. *Amazon* (books) and *Netflix* (movies) are two well-known recommenders, but there are platforms that mediate social

*Department of Economics, Columbia University, 420 West 118th Street, 1029IAB, New York, NY 10027, USA. Email: yeonkooche@gmail.com.

[†]Department of Economics, Yale University, 30 Hillhouse Ave., New Haven, CT 06520, USA, Email: johannes.horner@yale.edu.

learning among users for almost any “experience” good, including *Yelp* (restaurants), *TripAdvisors* (hotels), *RateMD* (doctors), and *RateMyProfessors* (professors), to name just a few. Search engines such as *Google*, *Bing* and *Yahoo* organize social discovery of relevant websites based on the “search experiences” of users. Social medias such as *Facebook* and *LinkedIn* do the same for the other quintessential “experience” good, friends, linking users to new potential friends based on the experiences of other friends.

While the economics of social learning are now well appreciated, the prior literature has focused only on its positive aspects.¹ A normative perspective—namely, how best to encourage users to experiment with new products and to disseminate their experiences to others—has not yet been developed. Incentivizing early users to experiment with products is challenging precisely because they cannot internalize the benefits that later users will reap. In fact, the availability of social learning may dampen the incentives for early experimentation because users might rather wait and free ride on the information provided by others in the future.² In other words, a recommender system may crowd out information production and undermine its very foundations. The absence of sufficient initial experimentation—known as the “cold start” problem—might lead to a collapse of social learning and to the death of products (even pre-launch) that are worthy of discovery by mainstream consumers. The cold start problem is particularly relevant for new and unestablished products, which are no longer the exception but rather the rule due to the proliferation of self production.³

This paper studies a recommender system that optimally achieves the dual purposes of social learning: production and dissemination of information. To maintain realism, we focus on non-monetary tools for achieving these purposes. Indeed, monetary transfers are seldom used to compensate for experimentation perhaps because it is difficult to verify whether a reviewer has performed genuine experimentation conscientiously and submitted an unbiased review.⁴

Our key insight is that the judicious use of the recommender system itself can incentivize consumer experimentation. To illustrate the idea, suppose that a recommender system—an online movie platform—generally recommends movies to users based on the reviews of

¹For instance, the main question studied in the literature involves the circumstances under which observational learning leads to full revelation of the underlying state.

²See Chamley and Gale (1994), and Gul and Lundholm (1995) for models illustrating this idea.

³For instance, once considered vanity publishing, self-publication has expanded dramatically in recent years with the availability of easy typesetting and e-books. Bowker Market Research estimates that in 2011 more than 300,000 self-published titles were issued (New York Times, “The Best Book Review Money Can Buy,” August 25, 2012). While still in its infancy, 3D printing and other similar technologies anticipate a future that will feature an even greater increase of self-manufactured products. The popularity of self production suggests a marketplace populated by such a large number of products/titles that they will not be easily recognized, which will heighten the importance of a recommender system even further.

⁴Attempts have been made in this regard that are limited in scope. For instance, the Amazon Vine Program rewards selected reviewers with free products, and LaFourchette.com grants discounts for (verified) diners that write reviews and make reservations via their site.

past users (truthful recommendations) but occasionally recommends new movies that require experimentation (fake recommendations or “spam”).⁵ As long as the platform keeps its users uninformed as to whether a recommendation is truthful or spam—and as long as it commits not to “spam” too much—users will happily follow the recommendation and perform the necessary experimentation in the process. Indeed, we show that an optimally designed recommender system trades off fully transparent social learning by over-recommending a product in the early phase of its life cycle.

Of course, the extent to which information acquisition can be motivated in this manner depends on the agent’s cost of acquiring information and the frequency with which the platform provides truthful recommendations (as opposed to spam). Another important feature involves the dynamics of how the platform combines truthful recommendations with spam over time after the product’s initial release (*e.g.*, a movie release date). For instance, it is unlikely that many users will have experimented with an unknown product immediately after its release, so recommending such a product during its early stage is likely to be met with skepticism. Therefore, to be credible, the platform must commit to truthful recommendations with sufficiently high probability in the early stage of the product’s life cycle, which means that the designer will spam very little in the early stage and that learning will be slow. Over time, however, recommendation becomes credible, and hence the pace of learning will accelerate. This suggests that there will be nontrivial dynamics in the optimal recommendation strategy and in social learning.

The present paper explores how a recommender system optimally balances the tradeoff between experimentation and learning, what type of learning dynamics such a mechanism would entail and what implications these learning dynamics will have on welfare, particularly when compared with: (i) *no social learning* (where there is no platform supervision of learning) and (ii) *full transparency* (where the platform commits to always recommend truthfully). In the baseline analysis, we consider a platform that commits to maximize the welfare of its users. Social welfare maximization is a salient normative benchmark that is important for any market design or public policy inquiry. From a more positive perspective, social welfare maximization might also result from platforms that compete à la Bertrand by charging membership fees and providing recommendations on a collection of products of varying vintages.⁶

⁵Throughout, the term “spam” means an unwarranted recommendation, more precisely a recommendation of a product that has yet to be found worthy of recommendation.

⁶One can imagine that each such platform provides recommendations for large collections of products of varying vintages. Providing recommendations on a number of products of different vintages means that the social welfare gain from optimal experimentation will be spread evenly across users arriving at different times. For instance, a user who “sacrifices” himself or herself for future users on certain products will be over-compensated by the benefits from past learning on other products. Hence, firms competing à la Bertrand on membership fees will be forced to offer the maximum intertemporal social welfare on each product.

While the optimal recommender system we identify rests on a normative justification, its main feature is consistent with certain observed practices. For instance, it is known that search engines such as *Google* periodically shuffle their rankings of search items to provide exposure to relatively new or unknown sites, which is highly consistent with the idea of distorting recommendations for under-exposed and lesser-known items. More generally, our finding is consistent with the casual observation that ratings of many products appear to be inflated.⁷ Ratings are often inflated by sellers (as opposed to the platforms), who have every interest to promote their products even against the interests of consumers.⁸ However, platforms have instruments at their disposal to control the degree of ratings inflation, such as filters that detect false reviews, requiring verification of purchase for posting reviews, and allowing users to vote for “helpful” reviews.⁹ Our theory suggests that some tolerance for review inflation can result from the optimal supervision of social learning by a benevolent designer.

Our starting point is the standard “workhorse” model of experimentation that is borrowed from Keller, Rady and Cripps (2005). The designer offers a product to agents with unknown (binary) value. By consuming this product, a possibly costly choice, short-run agents might discover whether its value is high. Here, we are not interested in agents’ incentives to truthfully report their experience to the designer: because they consume this product only once, they are willing to do so. We thus postulate that some exogenous fraction of the agents who consume will share the information with the designer (platform); hence, the greater the number of agents who consume, the more the designer learns about the product. The learning by the designer takes the form of conclusively good news about the product, with the news arriving at a rate that is proportional to the aggregate consumption/experimentation by agents. This structure might arise from an algorithm employed by a platform that aggregates the cumulative reviews from consumers into a simple “up” or “down” signal, for instance. Importantly, agents do not directly communicate with one another. The designer mediates the transmission of information, which raises a difficult problem: How should she control social learning so as to yield the right amount of experimentation?

In the baseline model in which agents cannot delay their check-in/consumption, the optimal policy of the designer involves revealing good news truthfully but recommending the product even without the news, or “spamming,” as long as the designer’s posterior belief remains above a certain threshold. Unlike *no social learning* and *full transparency*, spamming induces experimentation, but optimal policy keeps the frequency of spam below a certain

⁷Jindal and Liu (2008) find that 60% of the reviews on Amazon have a rating of 5.0, and approximately 45% products and 59% of members have an average rating of 5.

⁸Luca and Zervas (2014) suggest that as much as 16% of Yelp reviews are suspected to be fraudulent.

⁹Mayzlin et al. (2012) find that *Expedia*’s requirement that a reviewer must verify her stay to leave a review on a hotel resulted in fewer false reviews at *Expedia* compared with *TripAdvisor*, which has no such requirement.

maximal “capacity” necessary to maintain credibility. This creates learning dynamics that are single-peaked: the designer spams very little on a product immediately after its release but gradually increases the frequency of spamming over time until a critical belief threshold is reached, at which point she stops spamming altogether.

We then extend the model in several directions. First, we study public recommendations that are common in product ratings. Unlike private recommendations, public recommendations are revealed to all agents, and this limits the scope of the designer’s ability to manipulate agents’ beliefs. The optimal policy nevertheless induces experimentation, albeit at a slower average pace than under private recommendation. Notably, the designer randomly selects a time to trigger spam (*e.g.*, a “good” rating), and the spam then lasts until the designer’s belief reaches a threshold (after which the spam is halted).

Second, we allow agents to choose when to “check-in” for a recommendation. Endogenous entry raises two novel issues. First, it is desirable to encourage early agents to wait for background information to accumulate sufficiently before making decisions. Second, those induced to experiment must be discouraged from delaying their check-in to free ride on others’ experimentation. The optimal recommender policy now involves an information “blackout” for a duration of time after release of the product, effectively inducing agents to wait until some information accumulates, followed by a massive amount of spam that induces them to experiment in a short burst. Subsequently, spamming tapers off gradually, inducing agents to experiment at rates that decrease smoothly over time until it stops altogether—a feature crucial for controlling free riding.

Third, we consider heterogeneity in the agents’ costs of experimentation. If costs are observable, then the designer can tailor her recommendations to agents’ costs levels, inducing them to experiment at different rates. If the costs are not observable, however, the designer’s ability to tailor recommendations to the agents’ costs is limited, which might lead the designer to employ a fully transparent recommendation policy, under reasonable circumstances.

In addition to the exponential bandit literature, our paper relates to several other strands of literature. First, our model can be viewed as introducing optimal design into the standard model of social learning (hence the title). In standard models (for instance, Bikhchandani, Hirshleifer and Welch, 1992; Banerjee, 1993; Smith and Sørensen, 2000), a sequence of agents takes actions myopically, ignoring their effects on the learning and welfare of agents in the future. Smith, Sørensen and Tian (2014) study altruistic agents who distort their actions to improve observational learning for posterity. Frick and Ishii (2014) examine how social learning affects the adoption of innovations of uncertain quality and explain the shape of commonly observed adoption curves. Our focus here is instead dynamic control of information to agents to incentivize their (unverifiable) experimentation. Such dynamic control of information is present in Gershkov and Szentes (2009), but that paper considers a very different environment, as there are direct payoff externalities (voting). Much more closely

related to the present paper is a recent study by Kremer, Mansour and Perry (2014). They study the optimal mechanism for inducing agents to explore over multiple products. In their model, the quality of the product is ascertained once a single consumer buys it, so learning is instantaneous. By contrast, learning is gradual in our model, and controlling its dynamic trajectory is an important aspect of the mechanism design we focus on. We also study public recommendations as well as endogenous entry by agents, which have no analogues in their model.

Our paper also contributes to the literature on Bayesian persuasion that studies how a principal can credibly manipulate agents’ beliefs to influence their behavior. Aumann, Maschler and Stearns (1995) analyze this question in repeated games with incomplete information, whereas Kamenica and Gentzkow (2011), Rayo and Segal (2010), and Ostrovsky and Schwarz (2010) study the problem in a variety of organizational settings. The current paper pursues the same question in a dynamic setting. In this regard, the current paper joins a burgeoning literature that studies Bayesian persuasion in dynamic settings (see Ely, Frankel and Kamenica (2015), Renault, Solan and Vieille (2014), Ely (2015), and Halac, Kartik, and Liu (2015)). The focus on social learning distinguishes the present paper from these other papers.¹⁰

Finally, the present paper is related to the empirical literature on the user-generated reviews (Jindal and Liu, 2008; Luca and Zervas, 2014; and Mayzlin et al. (2014)).¹¹ These papers suggest ways to empirically identify manipulations in the reviews made by users of internet platforms such as *Amazon*, *Yelp* and *TripAdvisor*. Our paper contributes a normative perspective on the extent to which the manipulation should be controlled.

2 Illustrative Example

We begin with a simple example that highlights the main themes of this paper: (1) the optimal policy trades off social learning to incentivize experimentation, and (2) the optimal policy involves slower experimentation and more transparency when agents are able to free ride.

Suppose a product, say a movie, is released at time $t = 0$, and a unit mass of agents arrive at each time $t = 1, 2$. The quality of the movie is either “good” ($\omega = 1$), in which case the movie yields a surplus of 1 to an agent, or “bad” ($\omega = 0$), in which case it yields a surplus of 0. The quality of the movie is unknown at the time of its release, with prior $p^0 := \Pr\{\omega = 1\} \in [0, 1]$. Watching the movie costs each agent $c \in (p^0, 1)$; thus, without

¹⁰Papanastasiou, Bimpikis and Savva (2014) show the insight of the current paper extends to the two-product context, but without fully characterizing the optimal mechanism.

¹¹Dai, Jin, Lee and Luca (2014) offer a structural approach to aggregate consumer ratings and apply it to restaurant reviews from *Yelp*.

further information, the agents would never watch the movie.

At time $t = 0$, the designer receives a signal $\sigma \in \{g, n\}$ (from its marketing research, for example) about the quality of the movie with probabilities:

$$\Pr\{\sigma = g|\omega\} = \begin{cases} \rho_0 & \text{if } \omega = 1; \\ 0 & \text{if } \omega = 0, \end{cases}$$

and $\Pr\{\sigma = n|\omega\} = 1 - \Pr\{\sigma = g|\omega\}$. In other words, the designer receives good news only when the movie is good; but she also may receive no news (even) when the movie is good.¹²

Suppose that the designer has received no news at the beginning of $t = 1$ but that a fraction α of agents saw the movie at $t = 1$. Then, the designer again receives conclusively good news with probability:

$$\Pr\{\sigma = g|\omega\} = \begin{cases} \alpha & \text{if } \omega = 1; \\ 0 & \text{if } \omega = 0. \end{cases}$$

The feature that the signal becomes more informative with a higher fraction α of agents experimenting at $t = 1$ captures the learning benefit that they confer to the $t = 2$ agents.

The designer chooses her recommendation policy to maximize social welfare and does so with full commitment power. Specifically, she recommends the movie to a fraction of agents in each period based on her information at that point in time.¹³ The designer discounts the welfare in period $t = 2$ by a factor $\delta \in (0, 1)$.

We consider two possible scenarios in terms of whether the agents arriving in $t = 1$ can wait until $t = 2$.

Exogenous Check-In. Suppose the agents arriving in $t = 1$ cannot delay their check-in/consumption to $t = 2$. The designer's optimal policy is then as follows. First, the

¹²Thus, it follows that the designer's posterior at time $t = 1$ on $\omega = 1$ is 1 with a probability of $\rho_0 p^0$ (in the event that she receives good news) and

$$p_1 = \frac{(1 - \rho_0)p^0}{(1 - \rho_0)p^0 + 1 - p^0},$$

with a probability of $1 - \rho_0 p^0$ (in the event that she receives no news).

¹³The designer would not gain from a stochastic recommendation policy. To see this, compare two choices: i) the designer recommends the movie to a fraction α of the agents, and ii) the designer recommends it to all agents with probability α . For agents in $t = 1$, the two options are the same in terms of welfare and thus in terms of incentives. For agents in $t = 2$, the former results in the learning of good news with probability $p^0(\rho_0 + (1 - \rho_0)\alpha)$, whereas the latter does the same because when there is no good news at the beginning of $t = 1$, all agents are recommended the movie with probability α . This equivalence means that public recommendation entails no loss in this example; however, the equivalence does not hold in our general model.

designer is truthful at time $t = 2$, as lying can only reduce welfare and can never improve the incentive for experimentation at $t = 1$.

Consider now time $t = 1$. If good news has arrived by then, the designer would again recommend the movie to all agents at $t = 1$. Suppose that no news has been received by then and that the designer nevertheless recommends—or “spams”—to a fraction α of the agents. The agents receiving the recommendation cannot determine whether the recommendation is genuine or spam; instead, they would form a posterior:

$$P_1(\alpha) := \frac{\rho_0 p^0 + \alpha p^0 (1 - \rho_0)}{\rho_0 p^0 + (1 - \rho_0 p^0) \alpha}.$$

If the designer spams to all agents (*i.e.*, $\alpha = 1$), then they will find the recommendation completely uninformative, and hence $P_1(\alpha) = p^0$. Since $p^0 < c$, they would never watch the movie. By contrast, if the designer spams rarely (*i.e.*, $\alpha \simeq 0$), then $P_1(\alpha) \simeq 1$, *i.e.*, they will be confident (and nearly certain) that the recommendation is genuine. Naturally, the agents receiving recommendations will definitely watch the movie in this case. Because the recommendation is more credible the less the designer spams, $P_i(\alpha)$ is decreasing in α . In particular, there is a maximal fraction $\hat{\alpha} =: \frac{(1-c)\rho_0 p^0}{c(1-\rho_0 p^0) - p^0(1-\rho_0)}$ of agents who can be induced to experiment.

Social welfare,

$$W(\alpha) := p^0(\rho_0 + (1 - \rho_0)\alpha)(1 - c)(1 + \delta) - \alpha(1 - p^0)c,$$

consists of the benefit from the good movie being recommended (the first term) and the loss borne by the $t = 1$ agents from a bad movie being recommended (the second term). In particular, the benefit contains the benefit that experimentation by the $t = 1$ agents confers to the $t = 2$ agents, which is captured by the term $p^0(1 - \rho_0)\alpha(1 - c)\delta$.

The optimal policy is to “spam” up to $\hat{\alpha}$, if W is increasing in α , *i.e.*, if the social value of experimentation at date 1 justifies the cost:

$$p^0 \geq \hat{p}^0 := \frac{c}{(1 - \rho_0)(1 + \delta)(1 - c) + c}. \quad (1)$$

Note that the RHS is strictly less than c when $\rho_0 < \frac{\delta}{1+\delta}$. In that case, if $p^0 \in (\hat{p}^0, c)$, the designer will “spam” the agents at $t = 1$ to consume against their myopic interest.

Endogenous Check-In and Free-Riding. Suppose next that the agents arriving at $t = 1$ can wait until $t = 2$ to check in for a recommendation. (We assume that check-in for a recommendation is costly and that an agent thus checks in only once.) Agents may delay their check-in to free ride on experimentation by other agents. We will see how the

free-riding concern affects the optimal recommendation policy.

As before, it is intuitively clear that the designer always recommends the movie in the event of good news at $t = 1$, as this improves welfare without sacrificing agents' incentives to experiment at time $t = 1$. Suppose that the designer recommends the movie to a fraction α of the agents at time $t = 1$ in the event that no news is received. Further, suppose at time $t = 2$ the designer provides a recommendation to a fraction γ_2 of agents in the event of good news and to a fraction α_2 in the event of no news. If an agent arriving at $t = 1$ checks in at $t = 1$, then his payoff is

$$U_1(\alpha) := p^0(\rho_0 + (1 - \rho_0)\alpha)(1 - c) - (1 - p^0)\alpha c.$$

If the agent waits and checks in at $t = 2$, he will receive

$$U_2(\alpha, \gamma_2, \alpha_2) := \delta p^0(\rho_0 + (1 - \rho_0)\alpha)\gamma_2(1 - c) - \delta(1 - p^0)\alpha_2 c.$$

Incentive compatibility then requires $U_1(\alpha) \geq U_2(\alpha, \gamma_2, \alpha_2)$.¹⁴

The designer's problem is then to solve

$$\max_{\alpha, \gamma_2, \alpha_2} U_1(\alpha) + U_2(\alpha, \gamma_2, \alpha_2)$$

subject to

$$U_1(\alpha) \geq U_2(\alpha, \gamma_2, \alpha_2).$$

At $t = 2$, the optimal policy is *full transparency* ($\gamma_2 = 1$ and $\alpha_2 = 0$).¹⁵ At $t = 1$, assuming (1), the optimal policy is to spam up to a level $\alpha = \tilde{\alpha}(\delta)$ that satisfies $U_1(\tilde{\alpha}(\delta)) = U_2(\tilde{\alpha}(\delta), 1, 0)$. For any $\delta < 1$, $\tilde{\alpha}(\delta) > 0$, so the designer does spam. However, for any $\delta > 0$, $\tilde{\alpha}(\delta) < \hat{\alpha}$; *i.e.*, the designer spams less than under exogenous check in. This follows from the fact that the incentive constraint is now stronger: the designer should not only incentivize agents to experiment at $t = 1$ but also keep them from free-riding on others by delaying their check-in.¹⁶ This requires the designer to spam less and induce less experimentation. In fact, as $\delta \rightarrow 1$, $\tilde{\alpha}(\delta) \rightarrow 0$. In other words, as $\delta \rightarrow 1$, the optimal policy converges to *full*

¹⁴If $\alpha = 0$, there will be no additional learning at $t = 1$, and the $t = 1$ agents will have no incentive to wait until $t = 2$. Hence, the inequality holds. If $\alpha > 0$, then the designer must be inducing some agents to check in at $t = 1$ and experiment; otherwise, $\alpha = 0$ combined with no waiting dominates that policy. To keep $t = 1$ agents from waiting, it must be that $U_1(\alpha) \geq U_2(\alpha, \gamma_2, \alpha_2)$.

¹⁵It is easy to see that at the optimum $\gamma_2 = 1$ and $\alpha_2 = 0$. Suppose not. Then, it must be the case that $\alpha > 0$; otherwise, the incentive constraint would not be binding and either raising γ_2 or lowering α_2 would increase the objective without violating the constraint. Given that $\alpha > 0$, suppose that one reduces α and simultaneously either raises γ_2 or lowers α_2 to hold $U_2(\alpha, \gamma_2, \alpha_2)$ constant. This will increase $U_1(\alpha)$ because $U_1(\cdot)$ is decreasing due to the assumption that $p^0 < c$.

¹⁶This latter result can be seen by the fact that $U_2(\tilde{\alpha}(\delta), 1, 0) > 0$, which means that the incentive constraint is stronger than simply ensuring that the agents in $t = 1$ break even.

transparency.

3 Model

Our model generalizes the example in terms of its timing and information structure. A product is released at time $t = 0$, and, for each time $t \geq 0$, a constant flow of unit mass of agents arrive and decide whether to consume the product. In the baseline model, the agents are short-lived—they make a one-time decision and then exit the market permanently. (We later extend the model to allow the agents to delay their entry after their arrival.) An agent incurs cost $c \in (0, 1)$ for consumption. This cost can be the opportunity cost of time spent or a price charged. The product is either “good,” in which case an agent derives the (expected) surplus of 1, or “bad,” in which case the agent derives the (expected) surplus of 0. The quality of a product is *a priori* uncertain but may be revealed over time.¹⁷ At time $t = 0$, the probability of the product being good, or simply “the prior,” is p^0 . We shall consider all values of the prior, although the most interesting case will be $p^0 \in (0, c)$, which makes non-consumption myopically optimal.

Agents do not observe previous agents’ decisions or their experiences. There is a **designer** who can mediate social learning by collecting information from previous agents and disclosing that information to current agents. We can think of the designer as an Internet platform, such as *Netflix*, *Google* or *Microsoft*, that has access to users’ activities and reviews and makes product recommendations based on them. As is natural with these examples, the designer may obtain information from its own marketing research and other sources including agents’ experiences themselves. For instance, there may be some flow of “fans” who experiment with the product at zero cost. We thus assume that some information arrives at a constant base rate $\rho > 0$ plus the rate at which agents experience the product. Specifically, if a flow of size μ consumes the product over some time interval $[t, t + dt)$, then the designer learns during this time interval that the movie is “good” with probability $\lambda(\rho + \mu)dt$, where $\lambda > 0$ is the arrival rate of good news and $\rho > 0$ is the rate at which the designer obtains the information regardless of the agents’ behavior. Section 7.1 extends our model to allow for news to be (conclusively) bad, showing that our qualitative results continue to hold so long as the arrival rate of the good news exceeds that of the bad news. The designer begins with the same prior p^0 , and the agents do not have access to “free” learning.

¹⁷The agents’ preferences may involve an idiosyncratic component that is realized *ex post* after consuming the product; the quality then captures only their common preference component. The presence of an idiosyncratic preference component does not affect the analysis because each agent must decide based on the expected surplus he will derive from his consumption of the product. Idiosyncratic preferences would make the designer’s interpretation of the agents’ reviews nontrivial because a good product may receive a bad review and a bad product may receive a good review, which motivates our assumption that the arrival of the news signal is gradual.

The designer provides feedback on the movie to agents at each time based on the information she has learned by that time. Since agents' decisions are binary, without loss, the designer simply decides whether to recommend the product. The designer commits to the following policy: At time t , she recommends the movie to a fraction $\gamma_t \in [0, 1]$ of agents if she learns that the movie is good, and she recommends or **spams** to fraction $\alpha_t \in [0, 1]$ if no news has arrived by t . We assume that the designer maximizes the intertemporal net surplus of the agents, discounted at rate $r > 0$, over the (measurable) functions (γ_t, α_t) .

The designer's information at time $t \geq 0$ is succinctly summarized by the **designer's belief**, which is 1 in the event that good news has arrived or some $p_t \in [0, 1]$ in the event that no news has arrived by that time. The "no news" posterior, or simply **posterior** p_t , must evolve according to Bayes' rule. Specifically, suppose for time interval $[t, t + dt)$ that there is a flow of learning by the designer at rate μ_t , which includes both "free" learning ρ and the flow α_t of agents experimenting during the period. Formally, set

$$\mu_t = \rho + \alpha_t.$$

Suppose that no news has arrived by $t + dt$, then the designer's updated posterior at time $t + dt$ must be

$$p_t + dp_t = \frac{p_t(1 - \lambda\mu_t dt)}{p_t(1 - \lambda\mu_t dt) + 1 - p_t}.$$

Rearranging and simplifying, the posterior must follow the law of motion:¹⁸

$$\dot{p}_t = -\lambda\mu_t p_t(1 - p_t), \tag{2}$$

with the initial value at $t = 0$ given by the prior p^0 . Notably, the posterior falls as time passes, as "no news" leads the designer to form a pessimistic inference regarding the movie's quality.

In our model, agents do not directly observe the designer's information or her belief. However, they can form a rational belief about the designer's belief. They know that the designer's belief is either 1 or p_t , depending on whether good news has been received by time t . Let g_t denote the probability that the designer has received good news by time t . This belief g_t is pinned down by the martingale property, *i.e.*, that the designer's posterior must

¹⁸Subtracting p_t from both sides and rearranging, we obtain

$$dp_t = -\frac{\lambda\mu_t p_t(1 - p_t)dt}{p_t(1 - \lambda\mu_t dt) + 1 - p_t} = -\lambda\mu_t p_t(1 - p_t)dt + o(dt),$$

where $o(dt)$ is a term such that $o(dt)/dt \rightarrow 0$ as $dt \rightarrow 0$.

on average equal the prior:

$$g_t \cdot 1 + (1 - g_t)p_t = p^0. \quad (3)$$

Notably, g_t rises as p_t falls, *i.e.*, the agents find it increasingly probable that the news has arrived as time progresses.

The designer chooses a (measurable) policy (α, γ) to maximize social welfare, namely,

$$\mathcal{W}(\alpha, \beta, \chi) := \int_{t \geq 0} e^{-rt} g_t \gamma_t (1 - c) dt + \int_{t \geq 0} e^{-rt} (1 - g_t) \alpha_t (p_t - c) dt,$$

where (p_t, g_t) must follow the required laws of motion: (2) and (3), and $\mu_t = \rho + \alpha_t$ is the total experimentation rate and r is the discount rate of the designer.¹⁹ The welfare consists of the discounted value of consumption— $1 - c$ in the event of good news and $p_t - c$ in the event of no news—for those the designer induces to consume.

In addition, for the policy (α, γ) to be implementable, the agents must have an incentive to follow the recommendation. Since the exact circumstances of the recommendation (whether the agents are recommended because of good news or despite no news) is kept hidden from the agents, their incentives for following the recommendation depend on their posterior regarding the information held by the designer. Specifically, an agent will have the incentive to consume the product, if and only if the posterior that the movie is good is no less than the cost:

$$\frac{g_t \gamma_t + (1 - g_t) \alpha_t p_t}{g_t \gamma_t + (1 - g_t) \alpha_t} \geq c. \quad (4)$$

(There is also an incentive constraint for the agents not to consume the product when it is not recommended by the designer. Because this constraint will not be binding throughout—as the designer typically desires more experimentation than do the agents—we shall ignore it.)

Our goal is to characterize the optimal policy of the designer and the pattern of social learning it induces. To facilitate this characterization, it is useful to consider three benchmarks.

- **No Social Learning:** In this regime, the agents receive no information from the designer, and hence they decide solely based on the prior p^0 . When $p^0 < c$, no agent ever consumes.
- **Full Transparency:** In this regime, the designer discloses her information—or her

¹⁹We allow the designer to randomize over (γ, α) although we show in the Appendix that such a policy is never optimal.

beliefs—truthfully to the agents. In our framework, full disclosure is implemented by the policy of $\gamma_t \equiv 1$ and $\alpha_t = \mathbf{1}_{\{p_t \geq c\}}$.

- **First-Best:** In this regime, the designer optimizes her policy without having to satisfy the incentive compatibility constraint (4).

To distinguish the current problem from the *first-best*, we call the optimal policy maximizing \mathcal{W} subject to (2), (3) and (4), the **second-best** policy.

Before proceeding, we observe that it never pays the designer to lie if the news actually arrives.

LEMMA 1. *It is optimal for the designer to disclose the good news immediately. That is, an optimal policy has $\gamma_t \equiv 1$.*

Proof. If one increases γ_t , it can only increase the value of objective \mathcal{W} and relax (4) (and does not affect other constraints). \square

Lemma 1 reduces the scope of optimal intervention by the designer to choosing α , the recommendation policy following “no news.” In the sequel, we shall thus fix $\gamma_t \equiv 1$ and focus on α as the sole policy instrument.

4 Optimal Recommendation Policy

We begin by further characterizing the process by which the designer’s posterior, and the agents’ beliefs about the designer’s posterior, evolve under arbitrary policy α . To understand how the designer’s posterior evolves, it is convenient to work with the likelihood ratio $\ell_t = \frac{p_t}{1-p_t}$ of the posterior p_t . Given the one-to-one correspondence between the two variables, we shall refer to ℓ simply as a “posterior.” It then follows that (2) can be restated as:

$$\dot{\ell}_t = -\ell_t \lambda \mu_t, \quad \ell_0 = \ell^0 := \frac{p^0}{1-p^0}. \quad (5)$$

The agents’ belief g_t concerning the arrival of news is determined via (3):

$$g_t = \frac{\ell^0 - \ell_t}{1 - p^0}. \quad (6)$$

Immediately following the release of the product, news is very unlikely to have arrived, and hence $g_t \simeq 0$. As time elapses, the agents find it increasingly probable that news has arrived.

To see how this affects the designer’s ability to “spam,” substitute g_t into (4) to obtain:

$$\alpha_t \leq \bar{\alpha}(\ell_t) := \min \left\{ 1, \frac{\ell^0 - \ell_t}{k - \ell_t} \right\}, \quad (7)$$

if the normalized cost $k := c/(1 - c)$ exceeds ℓ_t and $\bar{\alpha}(\ell_t) := 1$ otherwise. It is convenient to interpret $\bar{\alpha}(\ell_t)$ as the designer’s *capacity* for spamming. This capacity depends on the prior p^0 . If $p^0 \geq c$, then the designer can “freeze” the agents’ belief at the prior p^0 by always spamming, and the prior is high enough to convince the agents to consume the product. Hence, $\bar{\alpha}(\ell_t) = 1$ for all ℓ_t . By contrast, if $p^0 < c$, then the spamming capacity is initially zero and increases gradually over time. Immediately after the product’s release, the designer has nearly no ability to spam because good news could never have arrived instantaneously, and the agents’ prior is unfavorable. Over time, however, $\bar{\alpha}(\ell_t)$ increases. In other words, even when no news is received, and ℓ_t falls as a result, the arrival of good news becomes increasingly probable, which allows the designer to develop her credibility over time and expands her capacity to spam.

Effectively, spamming “pools” recommendations across two very different circumstances: one in which the good news has arrived and one in which no news has arrived. Although the agents in the latter circumstance will never knowingly follow the recommendation, pooling the two circumstances for recommendation enables the designer to siphon the slack incentives from the former circumstance to the latter and to incentivize the agents to experiment, so long as the recommendation in the latter circumstance is kept sufficiently infrequent/improbable. Since the agents do not internalize the social benefit of experimentation, spamming becomes a useful tool for the designer’s *second-best* policy.

Substituting (6) into the objective function and using $\mu_t = \rho + \alpha_t$ —and normalizing by ℓ —the *second-best* program is restated as follows:

$$[SB] \quad \sup_{\alpha} \int_{t \geq 0} e^{-rt} (\ell^0 - \ell_t - \alpha_t (k - \ell_t)) dt$$

subject to

$$\dot{\ell}_t = -\lambda(\rho + \alpha_t)\ell_t, \forall t, \text{ and } \ell_0 = \ell^0, \quad (8)$$

$$0 \leq \alpha_t \leq \bar{\alpha}(\ell_t), \forall t. \quad (9)$$

Obviously, the *first-best* program, labeled $[FB]$, is the same as $[SB]$, except that the upper bound for $\bar{\alpha}(\ell_t)$ is replaced by 1. We next characterize the optimal recommendation policy.

PROPOSITION 1. (i) *The first-best policy prescribes experimentation*

$$\alpha^{**}(p_t) = \begin{cases} 1 & \text{if } p_t \geq p^*; \\ 0 & \text{if } p_t < p^*, \end{cases}$$

where

$$p^* := c \left(1 - \frac{rv}{\rho + r(v + \frac{1}{\lambda})} \right),$$

and $v := \frac{1-c}{r}$ denotes the continuation payoff upon the arrival of good news.

(ii) *The second-best policy prescribes experimentation at*

$$\alpha^*(p_t) = \begin{cases} \bar{\alpha}(\frac{p_t}{1-p_t}) & \text{if } p_t \geq p^*; \\ 0 & \text{if } p_t < p^*. \end{cases}$$

(iii) *If $p^0 \geq c$, then the second-best policy implements the first-best, and if $p^0 < c$, then the second-best induces slower experimentation/learning than the first-best. Whenever $p^0 > p^*$, the second-best induces strictly more experimentation/learning than either no social learning or full transparency.*

The *first-best* and *second-best* policies have a cutoff structure: They induce maximal feasible experimentation, which equals 1 under *first-best* and equals the spamming capacity $\bar{\alpha}$ under the *second-best*, so long as the designer's posterior remains above the threshold level p^* and there is no experimentation otherwise. The optimal policies induce interesting learning trajectories, which are depicted in Figure 1 for the case of $p^0 < c$.

The optimality of a cutoff policy and the optimal level of the cutoff posterior can be explained by the main tradeoff facing the designer, namely the marginal benefit and cost of inducing additional experimentation at any given belief p :

$$\underbrace{\lambda p v \left(\frac{1}{(\lambda \rho / r) + 1} \right)}_{\text{value of experimentation}} - \underbrace{c - p}_{\text{cost of experimentation}}. \quad (10)$$

The cost is the flow cost borne by the experimenting agents (the second term). The benefit is the social learning that the additional experimentation may generate (the first term): With probability p , the product is good, and experimentation will reveal this information at rate λ , which will enable the future generation of agents to collect the benefit of $v = (1 - c)/r$. The term $\frac{1}{(\lambda \rho / r) + 1}$ discounts this benefit by the rate at which the good news will be learned by “background learning” even with no experimentation. The optimal threshold posterior, p^* , which equates the benefits and the costs, is the same under either *first-best* and *second-best*




Figure 1: Path of α for $(c, \rho, p^0, r, \lambda) = (2/3, 1/4, 1/2, 1/10, 4/5)$.

because the opportunity cost of experimentation, *i.e.*, relying solely on background learning, is the same under both regimes.

If $p^0 \geq c$, the designer can implement the *first-best* policy by simply “spamming” *all* agents if and only if full experimentation is warranted under *first-best*. The agents comply with the recommendation because their belief is frozen at $p^0 \geq c$ under the policy. Admittedly, informational externalities are not particularly severe in this case because early agents will have an incentive to consume on their own. Note, however, that *full transparency* does not implement the *first-best* in this case, since agents will stop experimenting once p_t reaches c . In other words, spamming—or “obfuscation”—is crucial to achieving the *first-best*, even in this case.

In the more interesting case with $p^0 < c$, the *second-best* policy cannot implement the *first-best*. In this case, the spamming constraint for the designer is binding. As can be seen in Figure 1, spamming capacity is initially zero and increases gradually. Consequently, experimentation initially takes off very slowly and builds up gradually over time until the posterior reaches the threshold p^* , at which point the designer abandons experimentation. Throughout, the experimentation rate remains strictly below 1. In other words, learning is always slower under the *second-best* than under the *first-best*. Since the threshold belief is the same under both regimes, the agents are induced to experiment longer under the *second-best* than under the *first-best* regime, as Figure 1 shows.




Figure 2: Optimal (*second-best*) spamming as a function of ρ (here, $(k, \ell_0, r, \lambda) = (2/5, 1/3, 1/2, 1)$). The dots on the x -axis indicate stopping times under *first-best*.

In either case, as long $p^0 > p^*$, the *second-best* policy implements strictly higher experimentation/learning than either *no social learning* or *full transparency*, strictly dominating both of these benchmarks.

Comparative statics reveal further implications. The values of (p^0, ρ) parameterize the severity of the cold start problem facing the designer. The lower these values, the more severe the cold start problems are. One can see how these parameters affect optimal experimentation policies and induced social learning.

COROLLARY 1. (i) *As p^0 rises, the optimal threshold remains unchanged and total experimentation/social learning increases under both the first-best and second-best policies. The learning speed remains the same in the first-best policy but rises in the second-best policy.*

(ii) *As ρ rises, the optimal threshold p^* rises and total experimentation/social learning declines under both the first-best and the second-best policies. The learning speed remains the same in the first-best policy but rises in the second-best policy.*

Unlike the *first-best* policy, the severity of the cold start problem affects the flow rate of experimentation under the *second-best* policy. Specifically, the more severe the cold start problem is, in the sense of (p^0, ρ) being smaller, the more difficult it becomes for the designer to credibly spam the agents, thereby reducing the flow rate of experimentation that the designer can induce. This has certain policy implications. Internet recommenders such as

Netflix and *Amazon* have the ability to increase ρ by investing resources in product research. Corollary 1-(ii) shows the sense in which such an investment “substitutes” for agents’ experimentation: an increased free learning raises the opportunity cost of experimentation, calling for its termination at a higher threshold, under both *first-best* and *second-best* policies. In the *second-best* world, however, there is also a sense in which this investment “complements” experimentation: free learning makes spamming credible, and this allows the designer to induce a higher level of experimentation at each t . Figure 2 shows that this indirect effect can accelerate the social learning significantly: as ρ rises, the time it takes to reach the threshold belief is reduced much more dramatically under the *second-best* policy than under the *first-best*.

5 Public Recommendations

Thus far, we have considered private recommendations that may differ across agents and are not shared amongst them. Such personalized private recommendations are an important part of the Internet recommender system; *Netflix* and *Amazon* make personalized recommendations based on its users’ past purchase/usage histories. Search engines are known to rank search items differently across users based on their past search behavior. At the same time, platforms also make public recommendations that are commonly observable to all users. Product ratings provided by *Amazon*, *Netflix*, *Yelp*, *Michelin*, and *Parker* are public and unpersonalized, to our knowledge.

Personalized recommendations benefit consumers when they have heterogenous tastes, but do personalized recommendations matter even when the users have the same preferences, as assumed herein? In particular, do they play any role in incentivizing experimentation? We answer this question by studying the (optimal) public recommendation policy. Plainly, any public recommendation can be made privately without any loss. Hence, a public recommendation cannot be strictly better than private recommendations. It is not obvious, however, that private recommendations are strictly better than public recommendations. Our main result is that although the designer can still induce agents to experiment under public recommendation, the experimentation is random and slower (on average) and thus welfare is strictly lower than under the optimal (private) recommendation.

To begin, consider the baseline model with $p^0 < c$. Recall *full transparency* induces no experimentation in this case. We first show that a public recommendation policy can induce some experimentation and improve upon *full transparency*. Consider a policy with the following structure:

- At $t = 0$, the designer begins by making no recommendation, which means that the only source of learning is free learning at rate ρ .

- If good news has arrived by $t > 0$, then the designer recommends the product and does so forever.
- If good news has not arrived by time $t > 0$ and the designer’s posterior is above p^* (our *first-best/second-best* cutoff), then the designer “spams” (or recommends the product without good news) according to a distribution function $F(t)$ (to be defined). Once she recommends the product, she continues to recommend until her posterior reaches p^* .
- If good news has not arrived by t and the posterior is below p^* , then the designer recommends that the agents do not consume.

Proposition 2 below demonstrates that the optimal public recommendation policy has this structure. As with private recommendations, the designer reveals the good news when it arrives, but the way she “spams” is different. Given the public nature, the designer must effectively recommend to “all” agents when she does so at all. After the recommendation is made successfully, meaning that all agents choose to consume, the designer can “freeze” the agents’ posterior by simply continuing to recommend the product, which allows the designer to implement the *first-best* policy from that point forward. Of course, the first time that a public recommendation is made cannot be arbitrary. For instance, if the designer begins the recommendation at time 0 with certainty (which is what she would like to do), then the agents’ posterior will be p^0 , and they will thus not follow the recommendation. The designer instead randomizes the time at which she begins spamming on the product.

We now study how the random spamming distribution $F(t)$ can be constructed to satisfy incentive compatibility. For ease of exposition, we assume $F(t)$ to be atomless. (This will be justified formally in the proof of Proposition 2 in the Appendix.) Let $h(t) := f(t)/(1 - F(t))$ denote the hazard rate; *i.e.*, $h_t dt$ is the probability that the designer spams for the first time during $[t, t + dt)$. Then, by Bayes’ rule, the posterior of an agent who receives the public recommendation to consume for the first time at t is given by:

$$P_t = \frac{p^0 e^{-\lambda \rho t} (\lambda \rho + h(t))}{p^0 e^{-\lambda \rho t} (\lambda \rho + h(t)) + (1 - p^0) h(t)}.$$

The incentive constraint requires that $P_t \geq c$, which in turn yields:

$$h(t) \leq \frac{\lambda \rho \ell^0}{k - e^{\lambda \rho t} - \ell^0}.$$

It is intuitive (and formally shown) that the incentive constraint is binding at the optimal policy, which gives rise to a differential equation for F , alongside the boundary condition

$F(0) = 0$. Its unique solution is

$$F(t) = \frac{\ell^0(1 - e^{-\lambda\rho t})}{k - \ell^0 e^{-\lambda\rho t}}, \quad (11)$$

for all $t < t^* := -\frac{1}{\lambda\rho} \ln(\frac{k}{\ell^0} \frac{\lambda\rho+r}{\lambda(1+\rho)+r})$, the time at which the designer's posterior reaches the threshold belief p^* .

Clearly, $F(t) > 0$, for any $t \in (0, t^*)$, so the designer spams and the agents experiment with positive probability. Meanwhile, $F(t^*) = \frac{\ell^0 - \ell^*}{k - \ell^*} < 1$, and hence there is a positive probability that the designer never spam—a property crucial for the designer to maintain her credibility. Learning under this policy can be either faster or slower depending on the realization of the random triggering of spam, but it is on average slower under public recommendation than under the optimal private recommendation. To see this more clearly, fix any $t \leq t^*$. Under optimal public recommendation, spam is triggered at s according to $F(s)$ and lasts until t , unless the posterior reaches p^* . Let $T(s)$ be the time at which the latter event occurs if spam was triggered at s . Then, the expected level of experimentation performed by time t under public experimentation is:

$$\begin{aligned} \int_0^t (\min\{t, T(s)\} - s) dF(s) &\leq \int_0^t (t - s) dF(s) \\ &= \int_0^t F(s) ds = \int_0^t \frac{\ell^0 - \ell^0 e^{-\lambda\rho s}}{k - \ell^0 e^{-\lambda\rho s}} ds < \int_0^t \frac{\ell^0 - \ell_s}{k - \ell_s} ds = \int_0^t \bar{\alpha}(\ell_s) ds, \end{aligned}$$

where ℓ_s is the likelihood ratio at time s under the optimal private recommendation. The first equality follows from integration by parts, and the inequality holds because $\ell_s = \ell^0 e^{-\lambda \int_0^s (\bar{\alpha}(\ell_{s'}) + \rho) ds'} < \ell^0 e^{-\lambda\rho s}$.

We now state our findings.

PROPOSITION 2. *Under the optimal public recommendation policy, the designer recommends the product at time t if good news is received by that time. If good news is not received and a recommendation is not made by t , she triggers spam at a random time given by $F(t)$ in (11) and the spam lasts until t^* in the event that no good news arrives by that time. The experimentation speeds up until t^* on each realization of the policy (and thus on average). The induced experimentation under optimal public recommendation is on average slower—and the welfare attained is strictly lower—than under optimal private recommendation.*

Proof. The proof in the Appendix formulates the mechanism design problem fully generally, allowing the designer to communicate arbitrary messages upon receiving good news and no news. Although the argument is involved, the key observation is that once the agents are induced to consume the product, the designer can “freeze” the posteriors of the agents arriving subsequently and induce them to continue, which means that the *first-best* policy

can be implemented following any history that credibly convinces the agents to consume the product. This fact allows us to formulate the optimal mechanism design problem as a stopping problem that is indexed by two stopping times, one specifying how soon after good news is received that the designer should recommend the product to the agents and one specifying how soon she should recommend it to the agents despite having received no news. The unusual feature of the problem is that the incentive constraint must be obeyed, which forces the optimal stopping time (*i.e.*, triggering spam) to be random. We show that the above public recommendation policy solves this general mechanism design problem. \square

6 Endogenous Entry

Our model has thus far considered short-lived agents whose only decision is *whether* to consume a product given a recommendation. In practice, consumers also choose *when* to consume. In particular, they may delay their consumption decision until after more accurate recommendations become available. A recommender facing such patient consumers must guide consumers not only toward the right experimentation decision—and must make the recommendation persuasive in the sense discussed above—but also toward the right time for doing so. In particular, overcoming a possible free-rider problem becomes a new challenge.

To study this issue, we extend our baseline model in Section 4 as follows. As in our baseline model, agents arrive at a unit flow rate in each period, but upon arrival, say at t , an agent is free to choose any time $\tau \geq t$ to “check in” for a recommendation. In principle, agents may check in multiple times to obtain an improved recommendation, but such an incentive can be muted by ensuring that an agent receives the same recommendation each time he checks in, which is possible under personalized/private recommendation.²⁰ Agents would then never delay their consumption decision after checking in. As in Section 3, we focus on private recommendation.

The designer now controls two instruments: the mass X_t of agents the designer induces to check in by time t and the fraction α_t of agents to whom the designer spams among those who check in at t . Specifically, the designer faces the following problem:

$$[SB'] \quad \sup_{\{\alpha_t, X_t\}} \int_{t \geq 0} e^{-rt} (\ell^0 - \ell_t - \alpha_t (k - \ell_t)) dX_t$$

²⁰For instance, an Internet recommender may program the same recommendation to appear at its Internet portal for any user who logs in with the same membership ID or the same IP address. Although motivated users may circumvent the system at some inconvenience, the stakes or benefits from this behavior are often not large enough for it to be worthwhile. However, multiple check-ins might be potentially beneficial from a welfare perspective because agents may reduce their wait time after good news has arrived; for instance, the designer may have the agents “continuously” check in for news. In practice, checking in for a recommendation can be costly. We assume that this cost is high enough for multiple check-ins to be undesirable.

subject to (7),

$$\dot{\ell}_t = -\lambda(\rho + \alpha_t dX_t)\ell_t, \forall t, \text{ and } \ell_0 = \ell^0, \quad (12)$$

$$X_t \leq t, \forall t \quad (13)$$

$$t \in \arg \max_{t' \geq t} (\alpha_t (\ell_t - k) - \ell_t), \forall t \in \text{supp}(X). \quad (14)$$

The problem $[SB']$ has a similar objective function to $[SB]$, and involves the incentive constraint (7) and the law of motion (12), much as in $[SB]$. The difference is that the designer now controls the agents' check-in time according to a measure X_t . She does so subject to two constraints: (13) ensures that the mass of agents who check in by time t cannot exceed the mass who have arrived by that time, and (3) requires that agents should have no incentive to delay their check-in beyond that desired by the designer.²¹

The following proposition characterizes the optimal recommendation policy and *full transparency* in the presence of endogenous entry.

PROPOSITION 3. *Suppose that agents can delay checking in for a recommendation, and suppose that $\ell^0 < k$.*

- (i) *Under full transparency, agents who arrive before $t^{FT} = -\frac{1}{\rho\lambda} \ln \frac{r}{r+\lambda\rho}$ all check in at t^{FT} , and those who arrive after t^{FT} check in as they arrive. Upon checking in, agents consume if and only if news has arrived by that time. In other words, no agent experiments under full transparency.*
- (ii) *The second-best policy coincides with full transparency if ρ is sufficiently large for any p . The second-best policy differs from full transparency if $\rho \leq \frac{1}{\lambda} \left(\frac{r\ell^0 + \sqrt{r\ell^0} \sqrt{4k\lambda + \ell^0}}{2k} - r \right)$.*
- (iii) *In case the second-best differs from full transparency, there exist $t^* > 0$ and $T > t^*$ such that the optimal mechanism induces the following behavior from the agents:²²*
 - a. *(Blackout followed by a burst of spam) The designer offers no recommendation (even with good news) until t^* , which induces the agents who arrive before t^* to*

²¹Here, $\text{supp}(X)$ is support of X . More precisely,

$$\text{supp}(X) := \{t \in [0, \infty) \mid X_{t+\epsilon} - X_{\max\{t-\epsilon, 0\}} > 0, \forall \epsilon > 0\}.$$

We ignore the incentive for the agents never to check in too soon relative to what the designer intends, since the designer can simply withhold all recommendations to prevent agents from checking in for a period of time.

²²As before, as soon as good news arrives, it is immediately shared with the agents, including during the process of atom split, which is described in a.

check in and experiment at t^* but sequentially with probability $\alpha(\ell)$ that increases along the locus $\alpha(\ell) := \frac{\ell^+ + \alpha_2(k - \ell^+) - \ell}{k - \ell}$, for some (α^+, ℓ^+) with $\ell^+ < \ell^{FT} := \frac{r\ell^0}{r + \lambda\rho}$.

- b. (*Smoothly tapered spam*) The agents who arrive after t^* check in as they arrive and experiment with probability α_t which declines over time and reaches 0 at $t = T$. The agents who arrive after T never experiment.

Proof. The proof is available in the Appendix and the Supplementary Material. \square

The intuition behind the results can be explained as follows. There are two conflicting welfare considerations involved in an agent’s timing of check-in.

First, there is a social benefit from delaying agents’ check-in time. By delaying the time at which agents obtain a recommendation, the designer allows an agent to make a decision based on more updated information. This benefit exists even without any experimentation because of free learning but declines over time due to discounting. Therefore, absent experimentation, there is an interior time $t^{FT} > 0$ such that the designer would seek to delay agents’ check-in up to that time. This is precisely what happens under *full transparency*, which is socially optimal, assuming that no agent is induced to experiment.

The second consideration arises when the designer wishes to induce positive experimentation. Since experimentation is costly for agents who experiment and beneficial for later agents who do not experiment, the former have incentives to delay their checking in to free ride on other experimenters (even when delaying is not socially beneficial). This incentive problem can be seen clearly by the *second-best* policy in Section 4. In Figure 1, all agents who arrive before $t \simeq 2.4$ are told to experiment at a rate that leaves them with zero expected payoff, whereas those arriving slightly after 2.4 enjoy a strictly positive payoff; thus, all the former agents would wait until after that time. To overcome this free-rider problem, the designer must “bribe” the experimenting agents by reducing their experimentation and reducing the externalities enjoyed by later agents by prolonging experimentation. In particular, experimentation cannot fall discontinuously as in our baseline case, and all experimenting agents must enjoy strictly positive payoffs, which means that incentive constraint (7) is no longer binding.

Both of these considerations reduce the value of experimentation, particularly when ρ is large (significant free learning). Indeed, for a sufficiently large ρ , the *second-best* policy is reduced to *full transparency*. When ρ is sufficiently low, however, *full transparency* is again not optimal (part (ii) of Proposition 3). In this case, the *second-best* policy involves “spamming”—much as in our baseline model—but with qualitatively new features. Figure 3 depicts how experimentation and belief evolve over time under the *second-best* policy. As the belief ℓ_t declines over time, the relevant outcome (α_t, ℓ_t) moves from right to left in the figure.




Figure 3: The probability of a recommendation as a function of ℓ (here, $(r, \lambda, \rho, k, \ell^0) = (1/2, 2, 1/5, 15, 9)$). In dotted line, the spamming probability α^* in the baseline model with exogenous entry, as a function of ℓ .

First, the agents who arrive before $t^* > 0$ are induced to delay their checking in for recommendation until t^* . This can be implemented by an information “blackout”—or the recommender refusing to recommend a product—for a duration of time after its release. This can be interpreted in practical terms as the recommender refusing to take a stance on the product. Given the initially unfavorable prior, all agents will then choose to wait during the blackout.

At time t^* , it is optimal for the agents to all check in at that moment; however, it is not optimal to have them experiment simultaneously. If the designer were to “split” the mass into two smaller masses and have one of them move immediately after the other, the former will benefit from the experimentation of the latter, and can be asked to experiment more without violating their incentives. Repeating this argument, one finds it optimal to divide the mass of agents into flows of them who check in sequentially—but without taking any real time—and to experiment with increasingly higher probability, in a manner that leaves them indifferent across this menu of $(\tilde{\ell}, \tilde{\alpha})$ offered sequentially at that instant.²³ Doing so

²³The need to allow for sequential moves that do not take any real time arises from the feature of the continuous time: the soonest next time after any time t is not well defined. The Supplementary Material resolves this issue by enriching the policy space, specifically by adding a “virtual” clock that is used to sequence the check-in of agents who arrive at the same “real” time. See Simon and Stinchcombe (1989)

sequentially is optimal because it increases the overall experimentation that can be performed with the agents; Indifference among these agents gives rise to a locus $(\tilde{\ell}, \tilde{\alpha}) = (\tilde{\ell}, \alpha(\tilde{\ell}))$ that begins at (α^-, ℓ^-) and ends at (α^+, ℓ^+) , as depicted in Figure 3. The mass of agents are all dissipated at a point (ℓ^+, α^+) .

After t^* , the experimentation tapers off along a smooth locus, as depicted in Figure 3. Specifically, these agents check in as they arrive, and they experiment with probability that declines continuously at a rate that eliminates incentives for delaying entry. At T , the experimentation stops, and *full transparency* prevails from this point forward.

In sum, the optimal policy under endogenous entry exhibits qualitatively the same features as that from the baseline model (such as spamming and single-peaked experimentation), but it also introduces new features such as an information blackout and a gradual phase-out of experimentation. Endogenous entry also affects the amount of experimentation performed by the agents. Under exogenous entry, the agents who are induced to experiment do so at the level that leaves them with no rents. By contrast, under endogenous entry, the designer must leave them with positive rents to prevent them from free-riding on others. This means that, for each $\ell > \ell^*$, agents experiment less under endogenous entry than under exogenous entry; see the optimal experimentation α^* induced under the latter regime for comparison in Figure 3.²⁴ Meanwhile, agents are induced to experiment even when $\ell < \ell^*$ with endogenous entry (whereas no more experimentation occurs when $\ell < \ell^*$ with exogenous entry). This feature is in turn attributed to the free-riding problem. Inducing agents to experiment when $\ell < \ell^*$ entails net social loss (as was shown before); yet this distortion is necessary to prevent early agents who are induced to experiment when $\ell > \ell^*$ from delaying their check-in. In this sense, the distortion at a low belief is what makes experimentation at a higher belief implementable.²⁵

and Ausubel (2004) for adopting similar enrichment of strategies to guarantee existence of an equilibrium. Intuitively, we can think of the optimal policy in the enriched space as the supremum over policies in which the designer can only provide one recommendation at each time t , or the limit of the optimal outcomes of the sequence of discrete time models as its time interval shrinks to zero. In such a model, the designer will divide the mass into an increasing number of smaller masses as the time interval shrinks to zero, but with total real time approaching zero in the limit. Note, however, that our optimal policy is a well-defined maximum if the optimal control problem is written with ℓ as the “time” index.

²⁴This does not imply that learning is necessarily slower here in the temporal sense than under exogenous entry. The agents who arrive before t^* delay their check-in, and the free learning up to that point can motivate them to experiment more at $t \geq t^*$. This effect may offset the otherwise reduced incentive for experimentation.

²⁵This feature is reminiscent of optimal second-degree price discrimination wherein a seller distorts a low-value buyer’s purchase to extract rents from a high-value buyer.

7 Extensions

The model can be further extended to incorporate additional features. The detailed analysis is provided in the Supplementary Material; here, we illustrate the main ideas and results.

7.1 General Signal Structure

Thus far, our model assumed a simple signal structure which features only good news. This is a reasonable assumption for many products whose priors are initially unfavorable but can be improved dramatically through social learning. For some other products, however, social learning may involve discovery of poor quality. Our signal structure can be extended to allow for such a situation via “bad” news.²⁶ Specifically, news can be either good or bad, where good news means $\omega = 1$ and bad news reveals $\omega = 0$, and the arrival rates of the good news and bad news are respectively $\lambda_g > 0$ and $\lambda_b > 0$ conditional on the state.²⁷

The relative arrival rate $\Delta := \frac{\lambda_g - \lambda_b}{\lambda_g}$ of good news proves crucial for both the evolution of the designer’s belief and the optimal recommendation policy. Thus, we say that the signal structure is “good” news if $\Delta > 0$ and “bad” news if $\Delta < 0$. The posterior of the designer in the absence of any news evolves according to:

$$\dot{\ell}_t = -\ell_t \Delta \lambda_g \mu_t, \quad \ell_0 = \ell^0. \quad (15)$$

Intuitively, the belief becomes pessimistic in the good news environment and optimistic in the bad news environment as time progresses with no news.

As before, the designer reveals both good news and bad news truthfully. The more important question is whether the designer would recommend the product despite having no news. Here, the main thrust of our earlier analysis continues to hold. That is, the designer would like to spam the agents up to the maximum capacity allowed by incentive compatibility as long as her posterior remains above some threshold. In particular, the optimal policy in the good news case is virtually identical to that described in Proposition 1, in fact with λ in the optimal threshold p^* being replaced by λ_g .²⁸

²⁶See Keller and Rady (2015) for the standard bad news model of strategic experimentation.

²⁷More precisely, if a flow of size μ consumes the product over some time interval $[t, t + dt)$, then during this time interval the designer learns that the movie is “good” with probability $\lambda_g(\rho + \mu)dt$ and “bad” with probability $\lambda_b(\rho + \mu)dt$.

²⁸At first glance, it may be surprising that the arrival rate λ_b of the bad news does not affect the optimal threshold p^* in the case of good news, and vice versa. The reason is that the tradeoff captured in (10) does not depend on the arrival rate of bad news. At the same time, the arrival rate of bad news affects both the duration and the rate of incentive-compatible experimentation. As (15) shows, as λ_b rises (toward λ_g), it slows down the decline of the posterior. Hence, it takes a longer time for the posterior to reach the threshold level, which means that the agents are induced to experiment for longer (until news arrives or the threshold

The optimal recommendation policy in the bad news environment has the same cutoff structure, although the threshold posterior is somewhat different. The formal result, the proof of which is available in the Supplementary Material, is as follows:

PROPOSITION 4. *The first-best policy (absent any news) prescribes no experimentation until the posterior p rises to p_b^{**} and then full experimentation at the rate of $\alpha^{**}(p) = 1$ thereafter, for $p > p_b^{**}$, where*

$$p_b^{**} := c \left(1 - \frac{rv}{\rho + r(v + \frac{1}{\lambda_b})} \right).$$

*The second-best policy implements the first-best if $p_0 \geq c$ or if $p_0 \leq \hat{p}_0$ for some $\hat{p}_0 < p_b^{**}$. If $p_0 \in (\hat{p}_0, c)$, then the second-best policy prescribes no experimentation until the posterior p rises to p_b^* and then experimentation at the maximum incentive-compatible level thereafter for any $p > p_b^*$,²⁹ where $p_b^* > p_b^{**}$. In other words, the second-best policy triggers experimentation at a later date and at a lower rate than the first-best policy.*

Although the structure of the optimal recommendation policy is similar between the good news and bad news cases, the intertemporal trajectory of experimentation is quite different. Figure 4 depicts an example with $\Delta < 0$ and a sufficiently low prior belief. Initially, the designer finds the prior to be too low to trigger recommendation, and she never spams as a result. However, as time progresses without receiving any news (good or bad), her belief improves gradually, and as her posterior reaches the optimal threshold, she begins spamming at the maximal capacity allowed by incentive compatibility. One difference here is that the optimal *second-best* threshold differs from that of the *first-best*. The designer has a higher threshold, so she waits longer to trigger experimentation under the *second-best* policy than she would under the *first-best* policy. This is due to the difference in the tradeoffs at the margin between the two regimes. Although the benefit from not triggering experimentation is the same between the two regimes, the benefit from triggering experimentation is lower in the *second-best* regime due to the constrained experimentation that follows in the regime.

7.2 Heterogenous Costs

Thus far, agents have been assumed to be homogenous, which has greatly simplified the analysis and made the question of the observability of the agents' costs moot. To discuss

p^* is reached), holding constant the per-period experimentation rate. Furthermore, the spamming capacity can be seen to increase with λ_b , as the commitment never to recommend in the event of bad news means that a recommendation is more likely to have been the result of good news. Hence, experimentation increases in two different senses when λ_b rises.

²⁹The maximally incentive-compatible level is $\bar{\alpha}(\ell_t) := \min \left\{ 1, \frac{\left(\frac{\ell_t}{\ell_0}\right)^{-\frac{1}{\Delta}} - 1}{k - \ell_t} \ell_t \right\}$.




Figure 4: Path of α for $\Delta < 0$ and $(c, \rho, p^0, r, \lambda_g, \lambda_b) = (1/2, 1, 2/7, 1/10, 1, 2)$.

how our findings generalize, consider first a simple two-type example. Suppose that there are two cost types, c_L and c_H , with $c_H > c_L > p^0$, arising with probabilities q_L and $q_H = 1 - q_L$, respectively. For concreteness, we assume that

$$c_H/(1 - c_H) = 12/10, c_L/(1 - c_L) = 11/10, p^0 = 1/2, \text{ and } q_H = q_L = 1/2.$$

Furthermore, $\rho = r = \lambda = 1$. We contrast two scenarios. In the first, the types are observable. In the second, they are not.

7.2.1 Observable types

Here, we assume that the designer perfectly observes or infers the type of the agent. In practice, such an inference may be possible from a user's past consumption history. For instance, the frequencies of downloading or streaming movies may indicate a user's (opportunity) cost of experimentation, and the online platform may use that information to tailor its recommendation to the user. With observable costs, the designer can spam an agent with a maximum probability that depends on his type, namely,

$$\bar{\alpha}_i(\ell_t) := \min \left\{ 1, \frac{\ell^0 - \ell_t}{k_i - \ell_t} \right\},$$

for $i = L, H$, where $k_i := c_i/(1 - c_i)$.

The optimal policies, described fully in the Supplementary Material, are characterized by two cutoffs, $p_2 < p_1 < c$, in the designer's belief such that the designer induces both types

of agents to experiment if $p > p_1$, only low-cost type agents to experiment if $p \in (p_2, p_1]$, and no agents to experiment if $p < p_2$. The agents induced to experiment do so up to the maximal probability, which is one under the *first-best* regime and $\bar{\alpha}_i(\ell)$ for type $i = L, H$ under the *second-best* regime.

As in Section 4, the last threshold is equal to

$$p_2 = p_2^{**} = c_L \left(1 - \frac{rv}{\rho + r(v + 1/\lambda)} \right),$$

under both *first-best* and *second-best* policies. In the current example, $p_2^{**} \simeq .425$. The first threshold differs between under the two regimes. In our example, the threshold is $p_1^{**} = 0.451$ under the *first-best* policy and $p_1^* = .448$ under the *second best* policy. In other words, the high-cost agents are induced to experiment longer (in terms of the belief) under the *second-best* policy than under the *first-best* policy. This is attributed to the differing consequences of triggering the phase in which only low-cost agents experiment. In the *first-best* case, the designer can rely on all low-cost agents to experiment from that point on; in the *second-best*, she can only credibly ask a fraction of them to experiment. Her alternative channel for experimentation appears bleaker, and as a result, she is forced to ask high-cost agents to experiment for lower beliefs than in the *first-best* case. As before, since experimentation occurs at a lower rate under the *second-best* policy, it takes longer to reach each belief threshold: it takes $t_1^* \simeq 1.15$ and $t_2^* \simeq 1.22$ to reach the first and second threshold under *second-best*, whereas it takes $t_1^{**} \simeq .1$ and $t_2^{**} \simeq .17$ under the *first-best* policy.

This feature is not specific to the particular numerical example we picked. In the Supplementary Material, we prove that, for the case of good news, and two types, the high-cost agents are spammed at beliefs below the point at which they should stop under the *first-best* policy.

PROPOSITION 5. *Both the first-best and second-best policies are characterized by a pair of thresholds $0 \leq \ell_L \leq \ell_H \leq \ell^0$, such that (i) all agents are asked to experiment with maximum probability for $\ell \geq \ell_H$, (ii) only low-cost agents experiment (with maximum probability) for $\ell \in [\ell_L, \ell_H)$, and (iii) no agent experiments for $\ell < \ell_L$. Furthermore, the belief thresholds are $\ell_L^{**} = \ell_L^*$ and $\ell_H^{**} \geq \ell_H^*$ under first-best and second-best policies, respectively, with a strict inequality whenever $\ell^0 > \ell_H^{**}$.*

Alternatively, one could allow for a continuum of types. Although we are unable to solve for the optimal policy with full generality, it is possible to solve for the limit policy as $r \rightarrow 0$ in the case of the uniform distribution (and pure good news). Figure 5 depicts the optimal policies for some parameter values. As time passes, the belief decreases, as does the threshold cost type above which it is optimal for the designer not to spam (highest curve in Figure 5). Types below this threshold are spammed to different extents: those whose cost




Figure 5: Optimal policy with a continuum of cost (here, $p^0 = 1/2, k \sim U[0, 2]$).

is below the prior ($c \leq p^0 = 0.5$) are spammed for sure, while those whose cost is above the prior are spammed to the maximum extent that makes them still willing to follow the recommendation. Hence, the total rate of experimentation is lower than the highest curve and equal to the lower curve: at the initial instant, agents with $c \leq p^0$ are “fully” spammed whereas agents with higher costs cannot be spammed at all. As time passes and the belief decreases, higher cost agents can be spammed (although fewer of them are targeted), and experimentation increases. Eventually, however, the designer focuses on spamming only lower and lower types only, leading to a decline in rate of experimentation. When this threshold coincides with the cost type that can be spammed with a probability of one (the type with $c = .5$), the *first-best* policy can be followed. Until then, however, the *second-best* threshold is above the *first-best* threshold, as indicated by the dotted line: a given cost type (such that $c > .5$) is spammed at beliefs below those at which he would no longer be asked to experiment in the *first-best*.³⁰

7.2.2 Unobservable Costs

If the designer cannot infer agents’ costs, then her ability to tailor recommendations is severely limited. Specifically, consider our two-type example but assume that the type is unobservable to the agent. Because the designer cannot use transfers to elicit truth-telling, and because all types of agents prefer a more honest recommendation, the designer cannot

³⁰Keep in mind, however, that under the *first-best* policy, types below the dotted threshold are experimenting with probability 1, while under the *second-best* policy, this is not the case for cost types $c \geq 0.5$.

effectively screen agents’ types. In fact, the designer can only choose one probability of spamming α and has effectively three options regarding its value at any given point in time.³¹

1. No spamming: The designer chooses $\alpha = 0$.
2. Spamming only to the low-cost agents: The designer spams to a fraction $\alpha = \bar{\alpha}_L$ of agents. A low-cost agent follows that recommendation, while a high-cost agent does not. Thus, this means, in particular, that at any such time a high-cost agent will not consume even when the designer sends a genuine recommendation after good news.
3. Spamming to both types of agents: The designer spams to fraction $\alpha = \bar{\alpha}_H$ of agents. Both types of agents then follow the recommendation, but since $\bar{\alpha}_H < \bar{\alpha}_L$, experimentation will proceed at a relatively slow rate—although it will be spread across all agents, as opposed to only the low-cost agents.

Intuitively, the designer triggers the options in the following order: First, she spams to both types, then only to the low type, and finally to none. It turns out that in this example the designer skips the first phase and begins with $\alpha = \bar{\alpha}_L$ incentivizing only the low-cost type to experiment. She does that until her belief reaches $p_2 := .468$, which occurs at time $t_2 = .109$, assuming that no news has been received by that time. Compared to the observable case, the designer stops experimenting much sooner (.109 vs. 1.22!), and at a higher belief (.468 vs. .425) than when types are observable. Furthermore, she stops spamming high-cost agents altogether—or rather, they do not listen to the designer’s recommendation and will not consume.

Clearly, experimentation is much costlier when types are unobservable. Given spamming at $\bar{\alpha}_H$, not all low-cost agents are induced to experiment and in fact some low-cost agents receive a “do not consume” recommendation while some high-cost agents receive a “consume” recommendation. This mismatch makes this option costly (making it suboptimal to use it at least in our example). Meanwhile, spamming at rate $\bar{\alpha}_L$ is costly because the designer must then forego consumption by high-cost agents altogether—even if he has found out that the product is of high quality. This makes it more expensive to prolong experimentation with low-cost agents and results in a shorter experimentation spell than when types are observable.

In sum, unobservability makes spamming costly and pushes the designer toward greater transparency. In fact, as we show in the Supplementary Material, this push toward transparency is exacerbated with more types. Namely,

³¹While this claim sounds intuitive, it requires proof, which follows from the general result stated below.

PROPOSITION 6. *When agents’ costs are uniformly distributed and unobservable to the designer, full transparency is optimal.*³²

8 Conclusion

Early experimentation is crucial for users to discover and adopt potentially valuable products on a large scale. The purpose of the present paper has been to understand how a recommendation policy can be designed to promote such early experimentation. There are several takeaways from the current study.

First, “spamming” on a product—namely, recommending a product that has yet to be found worthy of recommendation—can turn users’ belief favorably toward the product and can thus incentivize experimentation by early users. Accordingly, spamming can be part of the socially optimal design of a recommendation system.

Second, spamming is effective only when it is properly underpinned by genuine learning. Spamming can leverage genuine learning to amplify the incentive but cannot work without genuine learning. This insight has two useful implications. First, the recommender would be well-advised to scale up spamming only gradually, beginning at a low level upon the release of a product, when an agent could not credibly believe that much had been learned about the product. Excessive spamming or highly inflated reviews at an early stage can backfire and harm the recommender’s reputation. Second, a recommender could raise the level of genuine learning by investing in independent reviews of a product. Such an investment in product reviews can directly contribute to the learning on a product (which can then be shared with future consumers). But it can also substantially increase the credibility with which the designer can persuade agents to experiment. This indirect benefit could very well be as important as the direct benefit.

The optimal recommendation policy involves randomly triggered spamming when recommendation is public, as is often the case with product ratings. And in case users can choose when to check in for a recommendation, the optimal policy features an early blackout (or a suppressed recommendation) followed by a flurry of spamming and its gradual phase-out—a dynamics that encourages early agents to wait for sufficient information to accumulate but not to free ride on other agents’ experimentation. Several extensions show that the general thrust of our findings is robust to heterogeneity among users with respect to their costs of experimentation and to the nature of the information learned from experimentation.

³²It is important to note that this result is obtained without any restriction on the (finite) number of messages the designer may employ and that the designer is allowed to randomize over *paths* of recommendation. That is, she is allowed to “flip a coin” at time 0—unknownst to the agents—and decide on a particular function $(\alpha_t)_t$ as a function of the outcome. Such “chattering” controls are often useful but not in this instance.

Throughout, we have assumed that the recommender pursues a benevolent policy and does so with full commitment. These assumptions served to identify a social optimal policy and can be justified if a recommender competes on a fixed fee (*e.g.*, membership fee) and offers recommendations on a mix of products with varying vintages. Nevertheless, it is worth exploring what happens if these assumptions are violated. If the recommender is benevolent but unable to commit, she will spam excessively to the point of losing credibility, and users will then simply ignore the recommendation and not consume. Consequently, no experimentation occurs, and the outcome coincides with *full transparency*. If the recommender can commit but maximizes consumption (in the event that the recommender earns revenue proportional to consumption),³³ then the recommender will follow a policy similar to our optimal policy except that she will never stop recommending even after a critical level of belief is passed. The fuller implications of relaxing these assumptions and other aspects of a recommendation policy remains a subject for future research.

References

- [1] Aumann, R.J., M. Maschler and R.E. Stearns, 1995. *Repeated Games With Incomplete Information*, Cambridge: MIT Press.
- [2] Ausubel, L., 2004. “An Efficient Ascending-Bid Auction for Multiple Objects,” *American Economic Review*, **94**, 1452–1475.
- [3] Banerjee, A., 1993. “A Simple Model of Herd Behavior,” *Quarterly Journal of Economics*, **107**, 797–817.
- [4] Bikhchandani S., D. Hirshleifer and I. Welch, 1992. “A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades,” *Journal of Political Economy*, **100**, 992–1026.
- [5] Cesari, L., 1983. *Optimization-Theory and Applications. Problems with ordinary differential equations, Applications of Mathematics* 17, Berlin-Heidelberg-New York: Springer-Verlag.
- [6] Chamley, C. and D. Gale, 1994. “Information Revelation and Strategic Delay in a Model of Investment,” *Econometrica*, **62**, 1065–1085.
- [7] Dai, W., G. Jin, J. Lee and M. Luca, 2014. “Optimal Aggregation of Consumer Ratings: An Application to Yelp.com,” working paper, Harvard Business School.

³³Revenue may be tied to consumption, because a platform may charge a fee for each book or movie downloaded but also from the advertising profit associated with streaming or downloading.

- [8] Ely, J., 2015. “Beeps,” mimeo, Northwestern University.
- [9] Ely, J., A. Frankel and E. Kamenica, 2015. “Suspense and Surprise,” *Journal of Political Economy*, **123**, 215–260.
- [10] Frick, M. and Y. Ishii, 2014. “Innovation Adoption by Forward-Looking Social Learners,” working paper, Harvard.
- [11] Gershkov A. and B. Szentes, 2009. “Optimal Voting Schemes with Costly Information Acquisition,” *Journal of Economic Theory*, **144**, 36–68.
- [12] Gul, F. and R. Lundholm, 1995. “Endogenous Timing and the Clustering of Agents’ Decisions,” *Journal of Political Economy*, **103**, 1039–1066.
- [13] Halac, M., N. Kartik, and Q. Liu, 2014. “Contests for Experimentation,” mimeo, Columbia University.
- [14] Jindal, N. and B. Liu, 2008. “Opinion Spam and Analysis,” *Proceedings of the 2008 International Conference on Web Search and Data Mining, ACM*, 219–230.
- [15] Kamenica E, and M. Gentzkow, 2011. “Bayesian Persuasion,” *American Economic Review*, **101**, 2590–2615.
- [16] Keller, G. and S. Rady, 2015. “Breakdowns,” *Theoretical Economics*, **10**, 175–202.
- [17] Keller, G., S. Rady and M. Cripps, 2005. “Strategic Experimentation with Exponential Bandits,” *Econometrica*, **73**, 39–68.
- [18] Kremer, I., Y. Mansour, and M. Perry, 2014. “Implementing the “Wisdom of the Crowd,”” *Journal of Political Economy* **122**, 988–1012.
- [19] Luca, M. and G. Zervas, 2014. “Fake It Till You Make It: Reputation, Competition and Yelp Review Fraud,” mimeo, Harvard Business School.
- [20] Mayzlin, D., Y. Dover and J. Chevalier, 2014. “Promotional Reviews: An Empirical Investigation of Online Review Manipulation,” *American Economic Review*, **104**, 2421–2455.
- [21] Ostrovsky, M. and M. Schwarz, 2010. “Information Disclosure and Unraveling in Matching Markets,” *American Economic Journal: Microeconomics*, **2**, 34–63.
- [22] Papanastasiou, Y., K. Bimpikis and N. Savva, 2014. “Crowdsourcing Exploration,” mimeo, London Business School.

- [23] Rayo, L. and I. Segal, 2010. “Optimal Information Disclosure,” *Journal of Political Economy* **118**, 949–987.
- [24] Renault, J., E. Solan, and N. Vieille, 2014. “Optimal Dynamic Information Provision,” arXiv:1407.5649 [math.PR].
- [25] Seierstad, A. and K. Sydsæter, 1987. *Optimal Control Theory with Economic Applications*, Amsterdam: North-Holland.
- [26] Simon, L.K. and M. Stinchcombe, 1989. “Extensive Form Games in Continuous Time: Pure Strategies,” *Econometrica*, **57**, 1171–1214.
- [27] Smith, L. and P. Sørensen, 2000. “Pathological Outcomes of Observational Learning,” *Econometrica*, **68**, 371–398.
- [28] Smith, L., P. Sørensen, and J. Tian, 2014. “Informational Herding, Optimal Experimentation, and Contrarianism,” mimeo, University of Wisconsin.

Appendix

Proof of Proposition 1. To analyze this tradeoff precisely, we reformulate the designer’s problem to conform to the standard optimal control framework. First, we switch the roles of variables so that we treat ℓ as a “time” variable and $t(\ell) := \inf\{t | \ell_t \leq \ell\}$ as the state variable, interpreted as the time it takes for a posterior ℓ to be reached. Up to constant (additive and multiplicative) terms, the designer’s problem is written as: For problem $i = SB, FB$,

$$\sup_{\alpha(\ell)} \int_0^{\ell^0} e^{-rt(\ell)} \left(1 - \frac{k}{\ell} - \frac{\rho \left(1 - \frac{k}{\ell} \right) + 1}{\rho + \alpha(\ell)} \right) d\ell$$

s.t. $t(\ell^0) = 0$,

$$t'(\ell) = -\frac{1}{\lambda(\rho + \alpha(\ell))\ell},$$

$$\alpha(\ell) \in \mathcal{A}^i(\ell),$$

where $\mathcal{A}^{SB}(\ell) := [0, \bar{\alpha}(\ell)]$, and $\mathcal{A}^{FB} := [0, 1]$.

This transformation enables us to focus on the optimal recommendation policy directly as a function of the posterior ℓ . Given the transformation, the admissible set no longer depends on the state variable (since ℓ is no longer a state variable), thus conforming to the standard specification of the optimal control problem.

Next, we focus on $u(\ell) := \frac{1}{\rho + \alpha(\ell)}$ as the control variable. With this change of variable, the designer's problem (both *second-best* and *first-best*) is restated, up to constant (additive and multiplicative) terms: For $i = SB, FB$,

$$\sup_{u(\ell)} \int_0^{\ell^0} e^{-rt(\ell)} \left(1 - \frac{k}{\ell} - \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) u(\ell) \right) d\ell, \quad (16)$$

s.t. $t(\ell^0) = 0$,

$$t'(\ell) = -\frac{u(\ell)}{\lambda \ell},$$

$$u(\ell) \in \mathcal{U}^i(\ell),$$

where the admissible set for the control is $\mathcal{U}^{SB}(\ell) := [\frac{1}{\rho + \bar{\alpha}(\ell)}, \frac{1}{\rho}]$ for the *second-best* problem and $\mathcal{U}^{FB}(\ell) := [\frac{1}{\rho+1}, \frac{1}{\rho}]$. With this transformation, the problem becomes a standard linear optimal control problem (with state t and control α). A solution exists by the Filippov-Cesari theorem (Cesari, 1983).

We shall thus focus on the necessary condition for optimality to characterize the optimal recommendation policy. To this end, we write the Hamiltonian:

$$\mathcal{H}(t, u, \ell, \nu) = e^{-rt(\ell)} \left(1 - \frac{k}{\ell} - \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) u(\ell) \right) - \nu \frac{u(\ell)}{\lambda \ell}. \quad (17)$$

The necessary optimality conditions state that there exists an absolutely continuous function $\nu : [0, \ell^0]$ such that, for all ℓ , either

$$\phi(\ell) := \lambda e^{-rt(\ell)} \ell \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) + \nu(\ell) = 0, \quad (18)$$

or else $u(\ell) = \frac{1}{\rho + \bar{\alpha}(\ell)}$ if $\phi(\ell) > 0$ and $u(\ell) = \frac{1}{\rho}$ if $\phi(\ell) < 0$.

Furthermore,

$$\nu'(\ell) = -\frac{\partial \mathcal{H}(t, u, \ell, \nu)}{\partial t} = r e^{-rt(\ell)} \left(\left(1 - \frac{k}{\ell} \right) (1 - \rho u(\ell)) - u(\ell) \right) \quad (\ell - \text{a.e.}) \quad (19)$$

Finally, transversality at $\ell = 0$ implies that $\nu(0) = 0$ (since $t(\ell)$ is free) .

Note that

$$\begin{aligned} \phi'(\ell) &= -rt'(\ell) \lambda e^{-rt(\ell)} \ell \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) \\ &+ \lambda e^{-rt(\ell)} \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) + \frac{\rho k \lambda e^{-rt(\ell)}}{\ell} + \nu'(\ell), \end{aligned}$$

or using the formulas for t' and ν' ,

$$\phi'(\ell) = \frac{e^{-rt(\ell)}}{\ell} (r(\ell - k) + \rho\lambda k + \lambda(\rho(\ell - k) + \ell)), \quad (20)$$

so ϕ cannot be identically zero over some interval, as there is at most one value of ℓ for which $\phi'(\ell) = 0$. Every solution must be “bang-bang.” Specifically,

$$\phi'(\ell) \begin{matrix} \geq \\ \leq \end{matrix} 0 \Leftrightarrow \ell \begin{matrix} \geq \\ \leq \end{matrix} \tilde{\ell} := \left(1 - \frac{\lambda(1 + \rho)}{r + \lambda(1 + \rho)}\right) k > 0.$$

Also, $\phi(0) = -\lambda e^{-rt(\ell)} \rho k < 0$. So $\phi(\ell) < 0$ for all $0 < \ell < \ell^*$, for some threshold $\ell^* > 0$, and $\phi(\ell) > 0$ for $\ell > \ell^*$. The constraint $u(\ell) \in \mathcal{U}^i(\ell)$ must bind for all $\ell \in [0, \ell^*)$ (a.e.), and every optimal policy must switch from $u(\ell) = 1/\rho$ for $\ell < \ell^*$ to $1/(\rho + \bar{\alpha}(\ell))$ in the *second-best* problem and to $1/(\rho + 1)$ in the *first-best* problem for $\ell > \ell^*$. It remains to determine the switching point ℓ^* (and establish uniqueness in the process).

For $\ell < \ell^*$,

$$\nu'(\ell) = -\frac{r}{\rho} e^{-rt(\ell)} \ell^{\frac{1}{\Delta}-1}, \quad t'(\ell) = -\frac{1}{\rho\lambda\ell},$$

so that

$$t(\ell) = C_0 - \frac{1}{\rho\lambda} \ln \ell, \text{ or } e^{-rt(\ell)} = C_1 \ell^{\frac{r}{\rho\lambda}},$$

for some constants $C_1, C_0 = -\frac{1}{r} \ln C_1$. Note that $C_1 > 0$; or else $C_1 = 0$ and $t(\ell) = \infty$ for every $\ell \in (0, \ell^*)$, which is inconsistent with $t(\ell^*) < \infty$. Hence,

$$\nu'(\ell) = -\frac{r}{\rho} C_1 \ell^{\frac{r}{\rho\lambda}},$$

and so (using $\nu(0) = 0$),

$$\nu(\ell) = -\frac{r\lambda}{r + \rho\lambda} C_1 \ell^{\frac{r}{\rho\lambda}+1},$$

for $\ell < \ell^*$. We now substitute ν into ϕ , for $\ell < \ell^*$, to obtain

$$\phi(\ell) = \lambda C_1 \ell^{\frac{r}{\rho\lambda}} \ell \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) - \frac{r\lambda}{r + \rho\lambda} C_1 \ell^{\frac{r}{\rho\lambda}+1}.$$

We now see that the switching point is uniquely determined by $\phi(\ell) = 0$, as ϕ is continuous and C_1 cancels. Simplifying,

$$\frac{k}{\ell^*} = 1 + \frac{\lambda}{r + \rho\lambda},$$

which leads to the formula for p^* in the Proposition (via $\ell = p/(1 - p)$ and $k = c/(1 - c)$). We have identified the unique solution to the program for both *first-best* and *second-best*,

and shown in the process that the optimal threshold p^* applies to both problems.

The *second-best* implements the *first-best* if $p^0 \geq c$, since then $\bar{\alpha}(\ell) = 1$ for all $\ell \leq \ell^0$. If not, then $\bar{\alpha}(\ell) < 1$ for a positive measure of $\ell \leq \ell^0$. Hence, the *second-best* implements a lower and thus a slower experimentation than does the *first-best*.

As for sufficiency, we use Arrow sufficiency theorem (Seierstad and Sydsæter, 1987, Theorem 5, p.107). This amounts to showing that the maximized Hamiltonian $\hat{\mathcal{H}}(t, \ell, \nu(\ell)) = \max_{u \in \mathcal{U}^i(\ell)} \mathcal{H}(t, u, \ell, \nu(\ell))$ is concave in t (the state variable), for all ℓ . To this end, it suffices to show that the terms inside the big parentheses in (17) are negative for all $u \in \mathcal{U}^i$, $i = FB, SB$. This is indeed the case:

$$\begin{aligned} & 1 - \frac{k}{\ell} - \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) u(\ell) \\ & \leq 1 - \frac{k}{\ell} - \min \left\{ \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) \frac{1}{1 + \rho}, \left(\rho \left(1 - \frac{k}{\ell} \right) + 1 \right) \frac{1}{\rho} \right\} \\ & = - \min \left\{ \frac{k}{(1 + \rho)\ell}, \frac{1}{\rho} \right\} < 0, \end{aligned}$$

where the inequality follows from the linearity of the expression in $u(\ell)$ and the fact that $u(\ell) \in \mathcal{U}^i \subset [\frac{1}{\rho+1}, \frac{1}{\rho}]$, for $i = FB, SB$. The concavity of maximized Hamiltonian in t , and thus sufficiency of our candidate optimal solution, then follows. \square

Proof of Proposition 2. Write h_t^P for the public history up to time t , and h_t for the private history of the designer — which includes whether or not she received positive feedback by time t . Let $p(h_t)$ denote the designer's belief given her private history.

- Suppose that, given some arbitrary public history h_t^P , the agent is willing to buy at t . Then, they are willing to buy if nothing more is said afterwards. To put it differently, the designer can receive her incentive-unconstrained *first-best* after such a history, and since this is an upper bound on her payoff, we might assume that this is what she does (full experimentation as long as she wishes after such a history).
- It follows that the only public histories that are non-trivial are those after which the agents are not yet willing to buy. Given h_t , the designer chooses (possibly randomly) a stopping time τ , which is the time at which she first tell the agent to buy (she then gets her *first-best*). Let $F(\tau)$ denote the distribution that she uses to tell them to buy at time τ conditional on her not having had good news by time τ ; let $F_t(\tau)$ denote the distribution that she uses if she had positive news precisely at time $t \leq \tau$. We will assume for now that the designer emits a single “no buy” recommendation at any given time; we will explain why this is without loss as we proceed.

- Note that, as usual, once the designer's belief $p(h_t)$ drops below p^* , she might as well use "truth-telling:" tell agents to abstain from buying unless she has received conclusive news. This policy is credible, as the agent's belief is always weakly above the designer's belief who has not received positive news, conditional on h_t^P . And again, it gives the designer her *first-best* payoff, so given that this is an upper bound, it is the solution. It follows immediately that $F(t^*) > 0$, where t^* is the time it takes for the designer's belief to reach p^* absent positive news, given that $\mu_t = \rho$ until then. If indeed $F(t) = 1$ for some $t \leq t^*$, then the agent would not be willing to buy conditional on being told to do so at some time $t \leq \max\{t' : t' \in \text{supp}(F)\}$. (His belief would have to be no more than his prior for some time below this maximum, and this would violate $c > p^0$.) Note that $F_t(t^*) = 1$ for all $t \leq t^*$: conditional on reaching time t^* at which, without good news, the designer's belief would make telling the truth optimal, there is no benefit from delaying good news if it has occurred. Hence, at any time $t > t^*$, conditional on no buy recommendation (so far), it is common knowledge that the designer has not received good news.
- The final observation: whenever agents are told to buy, their incentive constraint must be binding (unless it is common knowledge experimentation has stopped and the designer has learnt that the state is good). If not at some time t , the designer could increase $F(t)$ (the probability with which she recommends to buy at that date conditional on her not having received good news yet): she would get her *first-best* payoff from doing so; keeping the hazard rate $F(dt)/(1 - F(t))$ fixed at later dates, future incentives would not change.

Let

$$H(\tau) := \int_0^\tau \int_0^t \lambda \rho e^{-\lambda \rho s} (1 - F(s)) ds F_s(dt).$$

This (non-decreasing) function is the probability that the agent is told to buy for the first time at some time $t \leq \tau$ given that the designer has learnt that the state is good at some earlier date $s \leq t$. Note that H is constant on $\tau > t^*$, and that its support is the same as that of F . Because $H(0) = 0$, $F(0) = 0$ as well.

Let $P(t)$ denote the agent's belief conditional on the (w.l.o.g., unique) history h_t^P such that he is told to buy at time t for the first time. We have, for any time t in the support of F ,

$$P(t) = \frac{p^0 (H(dt) + e^{-\rho \lambda t} F(dt))}{p^0 (H(dt) + e^{-\rho \lambda t} F(dt)) + (1 - p^0) F(dt)}.$$

Indifference states that

$$P(t) = c, \text{ or } L(t) = k,$$

where $L(t)$ is the likelihood ratio

$$L(t) = \ell^0 \frac{H(dt) + e^{-\rho\lambda t} F(dt)}{F(dt)}.$$

Combining, we have that, for any t in the support of F ,

$$(k/\ell^0 - e^{-\rho\lambda t}) F(dt) = H(dt).^{34} \quad (21)$$

This also holds for any $t \in [0, t^*]$, as both sides are zero if t is not in the support of F . Note that, by definition of H , using integration by parts,

$$H(\tau) = \int_0^\tau \lambda \rho e^{-\lambda \rho t} (1 - F(t)) F_t(\tau) dt.$$

Integration by parts also yields that

$$\int_0^\tau (k/\ell^0 - e^{-\rho\lambda t}) F(dt) = (k/\ell^0 - e^{-\rho\lambda \tau}) F(\tau) - \int_0^\tau \lambda \rho e^{-\lambda \rho t} F(t) dt.$$

Hence, given that $H(0) = F(0) = 0$, we may rewrite the incentive compatibility constraint as, for all $t \leq t^*$,

$$(k/\ell^0 - e^{-\rho\lambda \tau}) F(\tau) = \int_0^\tau \lambda \rho e^{-\lambda \rho t} ((1 - F(t)) F_t(\tau) + F(t)) dt,$$

and note that this implies, given that $F_t(\tau) \leq 1$ for all $t, \tau \geq t$, that

$$(k/\ell^0 - e^{-\rho\lambda \tau}) F(\tau) \leq \int_0^\tau \lambda \rho e^{-\lambda \rho t} dt = 1 - e^{-\lambda \rho \tau},$$

so that

$$F(t) \leq \frac{1 - e^{-\lambda \rho t}}{k/\ell^0 - e^{-\rho\lambda t}}, \quad (22)$$

an upper bound that is achieved for all $t \leq t^*$ if, and only if, $F_t(t) = 1$ for all $t \leq t^*$.

Before writing the designer's objective, let us work out some of the relevant continuation

³⁴If multiple histories of “no buy” recommendations were considered, a similar equation would hold after any history h_t^P for which “buy” is recommended for the first time at date t , replacing $F(dt), H(dt)$ with $\tilde{F}(h_t^P), \tilde{H}(h_t^P)$; $\tilde{F}(h_t^P)$ is then the probability that such a history is observed without the designer having received good news by then, while $\tilde{H}(h_t^P)$ stands for the probability that it does with the designer having observed good news by then, yet producing history h_t^P . Define then $F, H : \mathbf{R}_+ \rightarrow \mathbf{R}_+$ as (given t) the expectation $F(t)$ (resp. $H(t)$) over all public histories h_t^P for which t is the first time at which “buy” is recommended. Taking expectations over such histories h_t^P gives (21). The remainder of the proof is unchanged.

payoff terms. First, t^* is given by our familiar threshold defined by the belief $\ell_{t^*} = k \frac{\lambda\rho+r}{\lambda(1+\rho)+r}$; given that, until t^* , conditional on no buy recommendation, experimentations occurs at rate ρ , it holds that $e^{-\lambda\rho t^*} = \ell_{t^*}/\ell^0$.

From time t^* onward, if the designer has not recommended to buy yet, there cannot have been good news. Experimentation only occurs at rate ρ from that point on. This history contributes to the expected total payoff the amount

$$p^0(1 - F(t^*))e^{-(r+\lambda\rho)t^*} \frac{\lambda\rho}{r + \lambda\rho} \frac{1 - c}{r}.$$

Indeed, this payoff is discounted by the factor e^{-rt^*} ; it is only positive if the state is good, and it is reached with probability $p^0(1 - F(t^*))e^{-\lambda\rho t^*}$: the probability that the state is good, the designer has not received any good news and has not given a buy recommendation despite not receiving any good news. Finally, conditional on that event, the continuation payoff is equal to

$$\int_0^\infty \lambda\rho e^{-rs-\lambda\rho s} ds \cdot \frac{1 - c}{r} = \frac{\lambda\rho}{r + \lambda\rho} \frac{1 - c}{r}.$$

Next, let us consider the continuation payoff if the designer emits a buy recommendation at time $\tau \leq t^*$, despite not having received good news. As mentioned, she will then experiment at maximum rate until her belief drops below p^* . The stopping time $\tau + t$ that she will pick must maximize her expected continuation payoff from time τ onward, given her belief p_τ , that is,

$$W(\tau) = \max_t \left\{ p_\tau \left(1 - \frac{r}{\lambda\rho + r} e^{-(\lambda(1+\rho)+r)t} \right) \frac{1 - c}{r} - (1 - p_\tau)(1 - e^{-rt}) \frac{c}{r} \right\}.$$

The second term is the cost incurred during the time $[\tau, \tau + t]$ on agents when the state is bad. The first is the sum of three terms, all conditional on the state being good: (i) $(1 - e^{-rt})(1 - c)/r$, the flow benefit on agents from experimentation during $[\tau, \tau + t]$; (ii) $(1 - e^{-\lambda(1+\rho)t})e^{-rt}(1 - c)/r$, the benefit afterwards in case good news has arrived by time $\tau + t$; (iii) $e^{-(r+\lambda(1+\rho))t} \frac{\lambda\rho}{r+\lambda\rho} (1 - c)/r$, the benefit from the free experimentation after time $\tau + t$, in case no good news has arrived by time $\tau + t$. Taking first-order conditions, this function is uniquely maximized by

$$t(\tau) = \frac{1}{\lambda(1 + \rho)} \ln \left(\frac{\ell_\tau}{k} \frac{\lambda(1 + \rho) + r}{\lambda\rho + r} \right).$$

Note that we can write $W(\tau) = p_\tau W_1(\tau) - (1 - p_\tau) W_0(\tau)$, where $W_1(\tau)$ ($W_0(\tau)$) is the benefit (resp., cost) from the optimal choice of t given that the state is good (resp., bad). Plugging

in the optimal value of t gives that

$$w_1(\tau) := rW_1(\tau)/(1-c) = 1 - \frac{r}{\lambda\rho + r} \left(\frac{\ell_\tau}{k} \frac{\lambda(1+\rho) + r}{\lambda\rho + r} \right)^{-1 - \frac{r}{\lambda(1+\rho)}},$$

and

$$w_0(\tau) := rW_0(\tau)/c = 1 - \left(\frac{\ell_\tau}{k} \frac{\lambda(1+\rho) + r}{\lambda\rho + r} \right)^{-\frac{r}{\lambda(1+\rho)}}.$$

Note that, given that absent any good news by time t , we have $\ell_t = \ell^0 e^{-\rho t}$. It follows that

$$\begin{aligned} k(1 - w_0(t)) - \ell^0 e^{-\lambda\rho t}(1 - w_1(t)) &= k \left(1 - \frac{r}{\lambda(1+\rho) + r} \right) \left(\frac{k}{\ell_t} \frac{\lambda\rho + r}{\lambda(1+\rho) + r} \right)^{\frac{r}{\lambda(1+\rho)}} \\ &= K e^{\frac{r\rho}{1+\rho}t}, \end{aligned} \quad (23)$$

with

$$K := k \frac{\lambda(1+\rho)}{\lambda(1+\rho) + r} \left(\frac{k}{\ell^0} \frac{\lambda\rho + r}{\lambda(1+\rho) + r} \right)^{\frac{r}{\lambda(1+\rho)}}.$$

For future reference, note that, by definition of ℓ_{t^*} ,

$$K e^{\frac{r\rho}{1+\rho}t^*} = k \frac{\lambda(1+\rho)}{\lambda(1+\rho) + r} \left(\frac{k}{\ell_{t^*}} \frac{\lambda\rho + r}{\lambda(1+\rho) + r} \right)^{\frac{r}{\lambda(1+\rho)}} = k \frac{\lambda(1+\rho)}{\lambda(1+\rho) + r}. \quad (24)$$

We may finally write the objective. The designer wishes to choose $\{F, (F_s)_{s=0}^{t^*}\}$ so as to maximize

$$\begin{aligned} J &= p^0 \int_0^{t^*} e^{-rt} \left(\frac{1-c}{r} H(dt) + e^{-\rho\lambda t} W_1(t) F(dt) \right) \\ &\quad - (1-p^0) \int_0^{t^*} e^{-rt} W_0(t) F(dt) + p^0 (1 - F(t^*)) e^{-(r+\lambda\rho)t^*} \frac{\lambda\rho}{r + \lambda\rho} \frac{1-c}{r}. \end{aligned}$$

The first two terms are the payoffs in case a buy recommendation is made over the interval $[0, t^*]$ and is split according to whether the state is good or bad; the third term is the benefit accruing if no buy recommendation is made by time t^* .

Multiplying by $\frac{r}{1-c} \frac{e^{rt^*}}{1-p^0}$, this is equivalent to maximizing

$$\int_0^{t^*} e^{-r(t-t^*)} \left(\ell^0 H(dt) + \ell^0 e^{-\rho\lambda t} w_1(t) F(dt) - k w_0(t) F(dt) \right) + \ell^0 (1 - F(t^*)) e^{-\lambda\rho t^*} \frac{\lambda\rho}{r + \lambda\rho}.$$

We may use (21) (as well as $\ell^0 e^{-\lambda \rho t^*} = \ell_{t^*} = k \frac{\lambda \rho + r}{\lambda(1+\rho) + r}$) to rewrite this as

$$\int_0^{t^*} e^{-r(t-t^*)} (k(1-w_0(t)) - \ell^0 e^{-\rho \lambda t} (1-w_1(t))) F(dt) + (1-F(t^*)) \frac{\lambda \rho k}{\lambda(1+\rho) + r}.$$

Using (23), and ignoring the constant term $\frac{\lambda \rho k}{\lambda(1+\rho) + r}$ (irrelevant for the maximization) this gives

$$e^{rt^*} K \int_0^{t^*} e^{-\frac{r}{1+\rho} t} F(dt) - \frac{\lambda \rho k}{\lambda(1+\rho) + r} F(t^*),$$

or, integrating by parts and using that $F(0) = 0$, as well as (24),

$$e^{rt^*} \frac{rK}{1+\rho} \int_0^{t^*} e^{-\frac{r}{1+\rho} t} F(t) dt + \left(k \frac{\lambda(1+\rho)}{\lambda(1+\rho) + r} - k \frac{\lambda \rho}{\lambda(1+\rho) + r} \right) F(t^*),$$

or finally (using (24) once more to eliminate K)

$$\frac{\lambda k}{\lambda(1+\rho) + r} \left(\int_0^{t^*} r e^{-\frac{r}{1+\rho} (t-t^*)} F(t) dt + F(t^*) \right).$$

Note that this objective function is increasing pointwise in F . Hence, it is optimal to set F as given by its upper bound given by (22), namely, for all $t \leq t^*$,

$$F(t) = \frac{\ell^0 (1 - e^{-\lambda \rho t})}{k - \ell^0 e^{-\rho \lambda t}},$$

and for all $t \leq t^*$, $F_t(t) = 1$. □

Proof of Proposition 3. Part (i): Under *full transparency*, agents who check in at t realize payoff

$$\mathcal{U}_t = e^{-rt} (\ell_0 - \ell_t).$$

Note that the timing at which agents check in is irrelevant for belief updating (because those who check in never experiment), so

$$\ell_t = \ell_0 e^{-\lambda \rho t}.$$

The function \mathcal{U}_t is quasi-concave in t , with a maximum achieved at the time

$$t^{FT} = -\frac{1}{\rho \lambda} \ln \frac{\ell^*}{\ell_0}, \quad \ell^{FT} := \frac{r \ell_0}{r + \lambda \rho}.$$

Part (ii): This is a perturbation argument around *full transparency*. Starting from this policy, consider the following modification. At some time t_2 (belief ℓ_2), the designer

is fully transparent ($\alpha_2 = 0$). An instant $\Delta > 0$ before, however, he recommends to buy with probability α_1 to some fraction κ of the queue $Q_{t_1} = t_1$, so that the agent is indifferent between checking in and waiting until time $t_2 = t_1 + \Delta$:

$$\ell^0 - \ell_1 - \alpha_1(k - \ell_1) = e^{-r\Delta}(\ell^0 - \ell_2), \quad (25)$$

where

$$\ell_1 = \ell^0 e^{-\lambda \rho t_1},$$

and

$$\ell_2 = \ell_1 e^{-\lambda(\rho\Delta + \kappa\alpha_1 t_1)}.$$

We solve (25) for κ (given ℓ_2), and insert into the payoff from this policy:

$$\mathcal{W}_\kappa = e^{-rt_2} \left((\ell^0 - \ell_2)t_2 + \frac{\ell^0}{r} - \frac{\ell_2}{r + \lambda\rho} \right).$$

Transparency is the special case $\kappa = \Delta = 0$, $t_1 = t^*$, and we compute a Taylor expansion of the gain for small enough Δ with $\ell_1 = \ell^* + \tau\Delta$ and $\alpha_1 = a_1\Delta^2$, with τ, a_1 to be chosen. We pick a_1 so that $\kappa = 1$, which gives

$$a_1 = \frac{\rho(r + \lambda\rho)(\lambda\ell^0\rho r - 2\tau(r + \lambda\rho))}{2\rho(k(\lambda\rho + r) - \ell^0 r) - 2\ell^0 r \ln\left(\frac{r}{r + \lambda\rho}\right)},$$

and choose τ to maximize the first-order term from the expansion, namely, we set

$$\tau = \frac{\lambda\ell^0\rho^2 r (k(\lambda\rho + r)^2 - \ell^0 r(\lambda\rho + r) - \ell^0 r\lambda)}{(\lambda\rho + r)^2 \left(\rho(k(\lambda\rho + r) - \ell^0 r) - \ell^0 r \ln\left(\frac{r}{\lambda\rho + r}\right) \right)}.$$

Plugging back into the expansion, we obtain

$$\mathcal{W}_\kappa - \mathcal{W}_0 = \frac{\lambda^2\ell^{03}\rho^3 r^3 \left((\lambda\rho + r) \ln\left(\frac{r}{\lambda\rho + r}\right) - \lambda\rho \right) (k(\lambda\rho + r)^2 - \ell^0 r(\lambda\rho + r) - \ell^0 r\lambda)}{(\rho + 1)(\lambda\rho + r)^5 \left(\rho(k(\lambda\rho + r) - \ell^0 r) - \ell^0 r \ln\left(\frac{r}{\lambda\rho + r}\right) \right)} \Delta + \mathcal{O}(\Delta^2),$$

and the first term is of the same sign as

$$\ell^0 r(\lambda\rho + r) + \ell^0 r\lambda - k(\lambda\rho + r)^2.$$

Note that this expression is quadratic and concave in $(\lambda\rho + r)$, and positive for $(\lambda\rho + r) = 0$. Hence it is positive if and only if it is below the higher of the two roots of the polynomial,

i.e., if and only if

$$\rho \leq \frac{1}{\lambda} \left(\frac{r\ell^0 + \sqrt{r\ell^0}\sqrt{4k\lambda + \ell^0}}{2k} - r \right).$$

Part (iii) is more involved and is provided in Supplementary Material.

□