Designing Social Learning*

Aleksei Smirnov[†], Egor Starkov[‡]

May 9, 2024

Abstract

This paper studies strategic communication in the context of social learning. Product reviews are used by consumers to learn product quality, but in order to write a review, a consumer must be convinced to purchase the item first. When reviewers care about welfare of future consumers, this leads to a conflict: a reviewer today wants the future consumers to purchase the item even when this comes at a loss to them, so that more information is revealed for the consumers that come after. We show that due to this conflict, communication via reviews is inevitably noisy, regardless of whether reviewers can commit to a communication strategy or have to resort to cheap talk. The optimal communication mechanism involves truthful communication of extreme experiences and pools the moderate experiences together.

Keywords: Social learning, dynamic games, strategic information transmission, experimentation.

JEL Codes: C73, D83, L15.

*We are grateful to Mikhail Drugov, Renato Gomes, Johannes Münster, Paolo Piacquadio, Marek Pycia, Armin Schmutzler, Jakub Steiner, Francesco Squintani, Bauke Visser, Nikhil Vellodi, Jidong Zhou, seminar participants at the University of Zürich and Jönköping University, and participants of ESWEM 2020 and MaCCI 2023 conferences for valuable feedback and helpful comments.

[†]Faculty of Economic Sciences, Higher School of Economics, Pokrovsky Boulevard 11, 109028, Moscow, Russia; e-mail: assmirnov@hse.ru.

[‡]Department of Economics, University of Copenhagen, Øster Farimagsgade 5, 1353 København K, Denmark; e-mail: egor@starkov.email.

1 Introduction

Whenever information is dispersed in the society, the question of social learning arises: can the society aggregate this information and achieve an efficient outcome for its members? In recent times, online customer reviews have become a powerful tool of social learning: according to multiple surveys of internet users, at least a half of respondents use ratings and online reviews "always" or "often" to inform their purchasing decisions, and most respondents find reviews to be at least "mostly reliable" (Competition & Markets Authority [2015], Mintel [2015], eMarketeer [2018]). Curiously, only about 10% of respondents to one of the aforementioned surveys say that they find product reviews "very reliable" (eMarketer [2018]). This skepticism can arise due to a variety of reasons, which mostly include various ways in which sellers can meddle with reviews, such as censorship and fake reviews.[1] However, in this paper we show that reviews can – and should – be noisy even in the absence of any intervention from sellers, with the noise arising from *how* customers write reviews.

To understand the source of this noise, one must first ask *why* customers write reviews. Surveys consistently produce a few modal answers to this question, with one of the two most popular ones being "to help other consumers" (Trustpilot [2020]).[2] Caring about other consumers making the right choice is often a sufficient incentive for people to spend their time and effort writing a review. These altruistic concerns, however, seem to only appear ex post – after the consumer has purchased and consumed the product – rather than ex ante. It is reasonable to expect that when choosing whether to buy a product and which product to buy, the consumers focus primarily on their own expected utility from consumption, rather than on their desire to provide helpful information to others.

This time-inconsistency in altruism, as we show, must lead to noise in product reviews. When product quality is uncertain, purchases have an informational externality, since they generate informative reviews, which allow future consumers to make better decisions. However, when deciding on the purchase, a self-interested consumer does not internalize this informational externality, and so her private expected value from buying a product is below the social value. This discrepancy is recognized by an altruistic reviewer, who may then want to mislead a future consumer into buying a product when it is socially, but not individually, optimal to do so.[3]

We formalize the argument above in a model of product reviews, in which a sequence of consumers decide whether to buy a product of some uncertain quality and, if they do, what kind of review to write about their experience. A consumer in our model only purchases the product if her expectation of consumption utility is high enough. The realized consumption utility is informative about the product quality, i.e., about the expected consumption utilities that other consumers would derive from this product.[4] The consumer can leave a review describing her consumption experience, and in doing so she wishes to maximize the welfare of consumers that arrive at the market after her. Note, in particular, that unlike earlier literature on experimentation in social

[1] See, e.g., Luca and Zervas [2016] for an exploration of the effects of fake reviews and Smirnov and Starkov [2022] for a model of censorship in product reviews.

[2] The other common response is "to help or punish the seller after a good or a bad experience, respectively".

[3] The point that product evaluations produce a positive externality and are hence socially underprovided was made, among others, by Avery, Resnick, and Zeckhauser [1999], who proposed cash payments to alleviate this inefficiency.

[4] Unlike most papers on experimentation, we do not restrict the model to an exponential bandits setting with binary utility outcomes, allowing instead for a rich heterogeneity in utility realizations.

learning, we impose the choice of how to communicate the consumption experiences directly on the consumers, as opposed to having a centralized recommender platform that makes these decisions [cf. Kremer, Mansour, and Perry, 2014, Che and Hörner, 2018]. This and other distinctions are discussed in more detail in Section 2.

The myopic behavior at the purchasing stage and the altruistic desire to induce some experimentation with the product at the reviewing stage conflict with each other. We show that this conflict creates noise in communication through reviews. Instead of reporting their experiences truthfully, the consumers obfuscate their reviews to foster experimentation. This is true regardless of whether consumers can commit to some communication strategy ex ante (which can be interpreted as a shared social norm in reviewing) or select their review ex post.

In particular, we show that in the commitment scenario, the optimal communication strategy features perfect communication of extreme experiences (very high or very low utilities), but remain vague about the experiences that would have put the next consumer close to indifference. The intuition is that the consumer would sometimes like to exaggerate the product quality to induce a socially optimal purchase that the next consumer would not make if she were fully informed. However, for such a recommendation to be credible, it has to be made when the product experience was, in fact, good – meaning that such a review has to be vague, only imperfectly revealing the consumer's experience with the product.

If a consumer cannot commit to a communication strategy, then despite the conflict of interest between her and a future consumer occurring only in a special set of circumstances – when buying the product is socially optimal, but not individually optimal for the future consumer, – the effects of this conflict propagate and distort communication in other cases as well. Communication is then always noisy in equilibrium, taking the interval structure often seen in the cheap talk models, when similar experiences are pooled in the same message.

Two conclusions may be drawn from our results. Firstly, coarse categories in product reviews (such as one- to five-star ratings) are almost sufficient for information transmission in the presence of the aforementioned experimentation conflict. Allowing free-form reviews in addition to – or instead of – such ratings will not necessarily increase the amount of information available to future consumers (unless the original categories were too coarse). Secondly, our paper provides a possible explanation for inflation in product reviews, namely that reviews can sometimes be inflated in order to deceive future consumers into purchasing the product they would not have bought otherwise. This complements other possible explanations, including positive ratings being sponsored or faked by the sellers. Contrary to those explanations, in our case inflation arises endogenously as a result of interaction between consumers, with no intervention from the firm. Further, noise in communication is welfare-enhancing in our model, helping not only the seller, but also the consumers.

The remainder of the paper proceeds as follows. We review the relevant literature in Section 2. We then formulate the model in Section 3. In this benchmark model, consumers are assumed to be able to commit to some communication structure. Sections 4 and 5 analyze this model, with Section 4 exploring in detail an illustrative three-period example, and Section 5 generalizing the insights to an infinite-horizon problem. Section 6 then alters the model to allow for ex post communication (cheap talk) instead of ex ante commitment. It proceeds to analyze the three-period example, and presents a weaker result for an infinite-horizon model. Section 7 concludes. The proofs of all results are relegated to the Appendix.

2 Literature Review

Our paper belongs to the literature on social experimentation and the design of social learning; notable references include Kremer et al. [2014], Mansour, Slivkins, and Syrgkanis [2015], Che and Hörner [2018], and Cohen and Mansour [2019]. This literature investigates the planner's problem of devising a communication protocol that would incentivize short-lived and/or selfish agents to experiment with a novel alternative more than they would on their own, for the sake of social benefit.[5] We explore what effectively is a decentralized version of the models mentioned above: instead of a single benevolent planner issuing recommendations, we allow each agent to communicate in the way they deem optimal. While the agents are altruistic and so have the same objective as the planner, they lack the *memory* of a central planner. In our model, once one agent withholds some of the information about their experience, this information is lost forever – unlike with the planner, who can remember privately all information that is not revealed publicly. This assumption is novel relative to the aforementioned papers and introduces a fundamentally new trade-off, since any attempt by the sender to mislead one receiver now carries not only the benefit of more social experimentation, but also the cost of subsequent receivers having worse information at their disposal. So while simple binary recommendations (whether to continue or stop experimenting/purchasing) are sufficient for the planner, in our decentralized setting a richer communication strategy is optimal.[6]

Closely related is the literature on herding and cascades in observational learning. In these models, agents receive private signals before making a choice, as opposed to after choosing to experience with the uncertain alternative, and only observe past agents' choices, which leads to inefficiencies (see Bikhchandani et al. [2021] for a more detailed overview of the topic). Wolitzky [2018] considers a setting in which agents observe only past agents' outcomes, rather than actions, while Cao, Han, and Hirshleifer [2011] and Le, Subramanian, and Berry [2016] allow for various combinations of the two. These papers conclude that imperfect communication may increase the probability of cascades, but argue that perfect communication could alleviate the inefficiency. We show that despite this, perfect communication is not welfare-maximizing. In particular, Ali and Kartik [2012] and Smith, Sørensen, and Tian [2021] consider sequential observational learning with others-regarding preferences, but assume that agents only observe past agents' actions, and not their private information or their outcomes.

Beyond experimentation, social learning with strategic information provision was explored by Swank and Visser [2015]; in their model the conflict arises from senders' career concerns, which are absent in our case. Liang and Mu [2020] consider a sequential information acquisition model, where the information that an agent acquires is also observed by future agents. They show that

[5] A separate literature explores optimal experimentation by groups of *long-lived* agents, which explores the issues of free-riding that are fundamentally similar to the core friction in our model: see Bolton and Harris [1999] and Keller, Rady, and Cripps [2005] for settings with full observability; Heidhues, Rady, and Strack [2015] explore strategic communication under private payoffs. These papers focus on the repeated game effects – i.e., devising the optimal reward and punishment strategies – which are absent from games with short-lived agents, such as ours.

[6] This trade-off is quite similar to that arising in problems of public communication to a heterogeneous audience; for some examples see Alonso and Camara [2016], Inostroza and Pavan [2022].

information acquisition driven by myopic incentives can lead to long-run inefficiencies. We show that tension between a forward-looking sender and a myopic receiver leads to noise in information provision.[7]

The version of our model without commitment belongs to the literature on dynamic cheap talk. Ambrus, Azevedo, and Kamada [2013] explore hierarchical cheap talk, in which all communication must go through the chain of biased intermediaries. In our setting, all agents are perfectly benevolent. Renault, Solan, and Vieille [2013] consider a game in which a single sender repeatedly reports the state that follows a Markov process. The repeated nature of interaction shifts focus towards the repeated game effects, which are absent in our model. The most related is the work of Chiba [2018], who looks at a committee of agents with fully aligned interests and shows that perfect communication between the committee members may be infeasible. Friction in that model comes from: (1) the agents' uncertainty about others' informedness and (2) the limited capacity of the communication channel. Communication constraints, however, play a crucial role in the result – perfect communication would be an equilibrium if agents were able to report not only their recommendation but also reveal the information this recommendation is based on. Our model, in contrast, shows that technological constraints are not necessary for communication to be imperfect in settings with highly aligned interests.

Finally, our model presents consumers as altruistic when they are writing reviews, which goes against the classical economic model, in which agents are driven exclusively by self-interest. This departure is mainly motivated by real-world consumers' own accounts on why they write reviews (Trustpilot [2020]). However, the economic literature has also argued for a long time that the "homo economicus" model does not accurately capture people's real-world behavior, which often exhibits regard for others. Various classical explanations of this departure include Andreoni [1990] (impure altruism), Fehr and Schmidt [1999] (inequality aversion) and Becker [1974] (pure altruism). The literature on altruism is surveyed by Fehr and Schmidt [2003], Konow [2003], and Meier [2006]. More recently, an attempt to provide an axiomatic foundation for such preferences has been made by Galperti and Strulovici [2017]. Of particular relevance to our argument are the experiments by March and Ziegelmeyer [2020] and Peng, Rao, Sun, and Xiao [2017], who find evidence of altruistic motives when testing standard models of observational learning. Related is work by Le Quement and Patel [2018], who explore a cheap talk model in the presence of preferences for reciprocity.

3 The Model

3.1 Primitives

Time is discrete and infinite: $t \in \{1, 2, \dots, \}$. All agents share a common discount factor $\beta < 1$.

[7]The decentralized nature of decision-making and communication in our model relates us to the literature on social learning on networks, which are inherently decentralized. Lobel and Sadler [2015] and Arieli and Mueller-Frank [2019] study sequential social learning when agents are arranged in a network or into an m-dimensional integer lattice, respectively. Campbell [2013] explores pricing and advertising in networks of friends who learn via word-of-mouth communication. Galeotti, Ghiglino, and Squintani [2013], Schopohl [2017] and Foerster [2019] analyze various games of strategic information transmission in networks. While all of these papers explore settings close in spirit to ours, the questions they focus on are different.

Seller. There is a single long-lived seller, who offers for sale a single product that he has in infinite supply at zero cost. Product quality θ, which represents the average consumption utility of the product, can be either *low* or *high*: $\theta \in \{L, H\}$, with $0 \leqslant L < H$. The price of the product is fixed at $c > 0$; to avoid triviality we assume that $L < c < H$.

Consumers. Each period a single short-lived consumer arrives at the market. All consumers are assumed to be Bayesian risk-neutral agents with lexicographic preferences, with the first-order preference being the consumer's own consumption utility, and the second-order preference being over the future consumers' expected consumption utility, discounted with factor β.

Upon entering the market, the consumer can either purchase the good at cost c, or leave the market forever. The latter option yields the reservation utility normalized to 0. In case of purchase, the consumer receives a random consumption utility v, distributed according to a quality-contingent c.d.f. F_θ with mean θ and a respective p.d.f. f_θ. We assume that both F_L and F_H have full support on the same open interval $\mathcal{V} = (\underline{v}, \overline{v}) \subseteq \mathbb{R}$.[8] Both measures are absolutely continuous on $\mathcal{V}$, and their respective densities are continuously differentiable and bounded from above. In addition, we assume that the monotone likelihood ratio property (MLRP) holds:

Assumption (MLRP)**.** *Ratio* $L(v) := \ln\left[\frac{f_H(v)}{f_L(v)}\right]$ *is a strictly increasing and continuous function of* v *on* $\mathcal{V}$*. Moreover,* $\lim_{v\to\underline{v}} L(v) = -\infty$*, and* $\lim_{v\to\overline{v}} L(v) = +\infty$.[9]

The consumer does not observe the product quality θ, so her purchasing decision is based on her belief $p := \mathbb{P}(\theta = H)$ regading the product quality, given the information available to her. In particular, the consumer purchases the product if and only if her expected consumption utility weakly exceeds the cost of purchase (we assume that the consumer purchases the product when indifferent):

$$\theta(p) := pH + (1-p)L \geqslant c \iff p \geqslant \bar{p},$$

where $\bar{p} := \frac{c-L}{H-L}$. This purchasing strategy will be taken as given in what follows.

Reviews. If the good was purchased, the consumer then sends a message (writes a review) $m \in \mathcal{M}$ to subsequent consumers, describing her experience with the product. A consumer's (pure) communication strategy is a mapping $\mathcal{V} \to \mathcal{M}$. The optimal communication strategy maximizes the expected discounted sum of consumption utilities of all future consumers. In the main model described here we assume that consumers choose their communication strategy simultaneously with their purchasing strategy. A cheap talk version of the game, in which the consumer selects message $m \in \mathcal{M}$ after observing v and not before, is explored separately in Section 6

Timing. Within a given period, the order of events is as follows:

1. Time-t consumer arrives at the market and observes all past reviews $(m_1, m_2, \ldots, m_{t-1})$ and forms belief p_t about the quality of the product.

[8] The support may be infinite: $\underline{v} = -\infty$ and $\overline{v} = +\infty$ are both admissible values. The common support assumption implies that no realized utility allows to rule out one of the possible quality levels θ.

[9] MLRP implies that F_H first order stochastically dominates F_L. In Smith and Sørensen [2000] terms it implies the beliefs are unbounded.

2. The consumer decides whether to purchase the product at cost c or not and selects her communication strategy.

3. After a purchase she receives random consumption utility $v_t \sim F_\theta$ and updates her belief about the product quality.

4. After a purchase, the consumer leaves review m_t about her experience, observable to all subsequent consumers. A consumer who has not purchased the product leaves no review: $m_t = \varnothing$.

Note that we assume the consumer commits to a communication strategy before observing her utility realization, as opposed to choosing ex post the review that maximizes the future consumers' welfare. This can either be taken at face value – assuming that a consumer can choose a reviewing strategy in advance and follow it through. The alternative interpretation is to see the problem as that of devising an optimal *social norm* that the consumers would use when writing reviews. Section 6 provides weaker results in the cheap talk version of the model, which drops the strategy observability assumption and replaces ex ante optimality with an requirement of interim optimality (so strategy $\mu(m|p_t, b_t)$ must be optimal given (p_t, b_t), rather than just p_t).

3.2 Histories, State Variables, Strategies

Review history $R_t := (m_1, m_2, \ldots, m_{t-1})$ is a tuple consisting of all messages sent by consumers before period t. It constitutes the public history at the beginning of period t. We denote the *public belief* about the quality of the product as

$$p_t = p(R_t) := \mathbb{P}(\theta = H \mid R_t).$$

The prior belief $p_0 = \mathbb{P}(\theta = H \mid \varnothing)$ is exogenously fixed and commonly agreed upon. The *private posterior belief* of time-t consumer in case she purchased and consumed the product is given by

$$b_t = b(p_t, v_t) := \mathbb{P}(\theta = H \mid p_t, v_t) = \frac{p_t f_H(v_t)}{p_t f_H(v_t) + (1 - p_t)\, f_L(v_t)}, \tag{1}$$

where (1) follows from the Bayes' rule.

The public belief p_t contains all payoff-relevant information available to time-t consumer at the time she decides whether to purchase the product. The pair of beliefs p_t and b_t summarizes all payoff-relevant information available to time-t consumer when she decides which message to send to subsequent consumers. Therefore, in what follows, we look at Markov equilibria (defined in Section 3.3), where p_t is the time-t *public state* and a sufficient statistic of the review history R_t, and the tuple (p_t, b_t) is the *private state* of time-t consumer and a sufficient statistic of her private history (R_t, v_t).

The consumers' purchasing decisions are myopic, and so the uniquely optimal stategy (up to indifference) is "buy if and only if $p_t \geqslant \bar{p}$". Therefore, from this point onwards we focus on the consumers' *communication* strategies. The time-t consumer's behavioral strategy is μ, where $\mu(m|p_t, b_t)$ is the probability with which the time-t consumer sends message $m \in \mathcal{M}$ conditional on private state (p_t, b_t). Let $\mathcal{M}(p_t) := \{m \in \mathcal{M} \mid \exists b_t : \mu(m \mid p_t, b_t) > 0\}$ denote the set of messages that are sent on equilibrium path at p_t. Whenever $p_t \geqslant \bar{p}$, the public belief p_{t+1} induced by an on-path message $m \in \mathcal{M}(p_t)$ is given by the Bayes' rule:

$$p_{t+1} = q(m|p_t) := \frac{p_t \cdot \int_0^1 \mu(m \mid p_t, b_t) dF_H(v_t)}{p_t \cdot \int_0^1 \mu(m \mid p_t, b_t) dF_H(v_t) + (1-p_t) \cdot \int_0^1 \mu(m \mid p_t, b_t) dF_L(v_t)} \tag{2}$$

If, on the other hand, $p_t < \bar{p}$, then message $m = \varnothing$ is sent for sure, so $p_{t+1} = p_t$: time-t consumer does not purchase the product, does not write a review, hence at $t+1$ the next consumer has exactly the same information at the time she makes her purchasing decision and does not purchase the product either. We shall refer to $p_{t+1} = q$ as the *posterior* or *induced public belief*.

Given some strategy profile, let $\mathcal{P}(p_t) := \{q(m|p_t) \mid m \in \mathcal{M}(p_t)\}$ denote the set of all public posteriors which are induced by time-t consumer. We partition this set into $\mathcal{E}(p_t) := \{q \in \mathcal{P}(p_t) \mid q \geqslant \bar{p}\}$, which includes all public posteriors q that convince the next consumer to buy the product, and $\mathcal{S}(p_t) := \{q \in \mathcal{P}(p_t) \mid q < \bar{p}\}$, which contains all q that deter her from the purchase. Note that if $p_t \geqslant \bar{p}$ then $\mathcal{E}(p_t) \neq \emptyset$, as p_t is a martingale (from the consumers' point of view). Conversely, as argued above, if $p_t < \bar{p}$ then $\mathcal{P}(p_t) = \mathcal{S}(p_t) = \{p_t\}$. The latter implies that all $q \in \mathcal{S}(p_t)$ are equivalent in the sense of shutting the market down from $t+1$ onward.

3.3 Maximization Problem and Equilibrium Definition

Given a strategy profile for all future consumers, when a consumer sends message m at private state (p_t, b_t), her continuation value (the discounted sum of future consumers' utilities) from doing so is equal to

$$V(m \mid p_t, b_t) := \mathbb{E}\left[\sum_{s=t+1}^{+\infty} \beta^{s-t-1} \cdot \mathbb{I}\left(p_s \geqslant \bar{p}\right) \cdot (v_s - c) \,\middle|\, (p_t, m), b_t\right]. \tag{3}$$

Implicit in (3) is the fact that m together with p_t determines p_{t+1} in equilibrium. It also embeds the dependence between future v_s and future p_s, stemming from the future consumers' strategies.

We are looking for the Markov Perfect Equilibria of the game, defined as follows.

Definition. *A Markov Perfect Equilibrium consists of a consumer communication strategy $\mu(m|p_t, b_t)$ and belief updating rules $b(p_t, v_t)$ and $q(p_t, m)$ such that the following conditions hold:*

- ***Belief Consistency:*** *condition* (1) *holds for all (p_t, b_t), and* (2) *holds for all p_t and all $m \in \mathcal{M}(p_t)$;*
- ***Optimality:*** *for any given p_t, and communication strategies $\{\mu(m|p_\tau, b_\tau)\}_{\tau > t}$ communication strategy $\mu(m \mid p_t, b_t)$ maximizes $\mathbb{E}_{b_t}\left[V(m \mid p_t, b_t) \mid p_t\right]$.*

The definition above implies a stationary equilibrium: the communication strategy of a period-t consumer does not depend on t, except through the consumer's belief p_t. The belief consistency condition ensures that all consumers use Bayes' rule to update their belief whenever possible. Optimality requires that every consumer chooses a communication strategy (i.e., a mapping from her private state (p_t, b_t) to messages) so as to maximize her ex ante value, as given by the expectation of (3) over v_t or, equivalently, b_t. In particular, maximizing the ex ante value means that the consumer commits to a communication strategy before observing the realized consumption utility v_t (but she can condition on the public state p_t).[10]

Without loss of generality, we restrict attention to equilibria with *direct communication*, where a consumer's review simply prescribes the belief that the next consumer must have after reading this and all other reviews: i.e., $m_t = q(p_t, m_t)$. As a consequence, it is assumed that $[0,1] \subseteq \mathcal{M}$. This assumption is made for illustrative simplicity, allowing us to ignore the distinction between message m_t and the belief p_{t+1} that it induces. Further, we assume that the set of stopping messages $\mathcal{S}(p_t)$ always consists of a single representative element whenever it is nonempty; as argued above, this is also without loss.

4 Three-Period Example

This section demonstrates the main insights in a simple three-period setting: $t \in \{1, 2, 3\}$ (the current period t is included as a part of both public and private states for the purposes of equilibrium definition – i.e., strategies are allowed to be time-dependent). We shall denote the three consumers as C_1, C2, and C3, respectively. Suppose that there is no discounting, i.e., C1 treats welfare of both C2 and C3 equally. We solve the example by backward induction. Fix some prior belief $p_0 \geqslant \bar{p}$. For this example, assume that consumption utilities v_t are normally distributed with mean θ and variance σ^2, and that they are mutually independent.

In period $t = 3$, C3 purchases the product if and only if $\theta(p_3) = Hp_3 + L(1 - p_3) \geqslant \bar{p}$, and her messaging strategy is irrelevant, since no consumers arrive at the market after her. We hence continue straight to $t = 2$.

4.1 Second Period

In the second period, if $p_2 < \bar{p}$ then, as mentioned in the model setup, the game effectively ends: C2 does not buy the product, so writes no review, so $p_3 = p_2 < \bar{p}$, and C3 does not buy the product either. Payoffs of C2 and C3 are zero in this case. Conversely, if $p_2 \geqslant \bar{p}$ then C2's continuation value equals C3's expected consumption utility: $V(m|p_2, b_2) = \theta(q(m|p_2)) - c$. Therefore, at the review stage C2 would prefer to act in the best interest of C3. Truthful communication, where C2 reports $m_2 = p_2$, is thus optimal.

However, perfect communication is not necessary to achieve the maximal payoff for C3. Note that the only piece of information relevant to C3 is whether to buy the product or not. She cannot make use of more precise information to make better recommendations to future consumers because there are no future consumers. Therefore, a simple binary communication strategy that sends $m \geqslant \bar{p}$ whenever $b_2 \geqslant \bar{p}$ and $m \leqslant \bar{p}$ whenever $b_2 \leqslant \bar{p}$ is also optimal for C2.

[10]To work around the problem of the equilibrium multiplicity, we assume that the whole strategy μ_t chosen by consumer t is observable to all future consumers in addition to her review m_t. Without this assumption multiple equilibria could arise, in which a specific communication strategy could be supported by some specific on- and off-the-equilibrium-path beliefs. It is straightforward that the equilibrium we find Pareto-dominates all such self-reinforcing equilibria.

4.2 First Period: Payoffs

As shown above, perfect communication is optimal for C2. We now show that the same is not true for C1. We begin in this section by analyzing C1's continuation value $V(m_1 \mid p_1, b_1)$ as a function of the message she sends or, equivalently, of the public belief $p_2 = m_1 = q(m_1|p_1)$ she induces.

Recall that we assumed in the example setup that $p_1(= p_0) \geqslant \bar{p}$, otherwise C1 does not buy the product and all values are zero. C1 then buys the good and receives utility v_1. If she sends $m_1 \leqslant \bar{p}$, then neither C2, nor C3 will buy the product, and C1's continuation value is zero. If she sends $m_1 \geqslant \bar{p}$, then C2 purchases the product and obtains utility v_2. Following that, C3 purchases the product if and only if $p_3 \geqslant \bar{p} \Leftrightarrow b_2 = b(p_2, v_2) \geqslant \bar{p}$ (since, as argued above, it is optimal for C2 to reveal her posterior b_2 truthfully). The threshold utility level $\bar{v}_2$ which would render C2 indifferent between issuing either recommendation – i.e., the one such that $b_2 = b(p_2, \bar{v}_2) = \bar{p}$, – can be found from

$$\frac{b_2}{1-b_2} = \frac{p_2}{1-p_2} \cdot \frac{f_H(\bar{v}_2)}{f_L(\bar{v}_2)} = \frac{\bar{p}}{1-\bar{p}},$$

where the first equality follows from (1). Since v_2 was assumed to be normally distributed, we have $\frac{f_H(\bar{v}_2)}{f_L(\bar{v}_2)} = \exp\left\{\frac{H-L}{\sigma^2}\left(\bar{v}_2 - \frac{H+L}{2}\right)\right\}$, and thus

$$\bar{v}_2(p_2) = \frac{\sigma^2}{H-L}\left[\ln\left(\frac{\bar{p}}{1-\bar{p}}\right) - \ln\left(\frac{p_2}{1-p_2}\right)\right] + \frac{H+L}{2}. \tag{4}$$

Therefore, if C2 buys the product, C1's continuation value from inducing belief $p_2 \geqslant \bar{p}$ is given by

$$\hat{V}(p_2 \mid p_1, b_1) := \mathbb{E}\left[(v_2 - c) + (v_3 - c) \cdot \mathbb{I}\{v_2 \geqslant \bar{v}_2(p_2)\} \mid b_1\right]. \tag{5}$$

Given C2's sequential rationality, C1's continuation value is

$$V(m_1 \mid p_1, b_1) = \begin{cases} \hat{V}(m_1 \mid p_1, b_1) & \text{if } m_1 \geqslant \bar{p}, \\ 0 & \text{if } m_1 < \bar{p}. \end{cases}$$

From the point of view of C1, the good is of high quality with probability b_1. In that case C3 buys the good with probability $1 - F_H(\bar{v}_2)$ and receives $H - c$ in expectation. Similarly, with probability $1 - b_1$ the good is of low quality, and then C3 gets $L - c$ conditional on purchase which occurs with probability $1 - F_L(\bar{v}_2)$. In the end, expression (5) can be rewritten as

$$\hat{V}(p_2 \mid p_1, b_1) = \theta(b_1) - c + b_1 \cdot (1 - F_H(\bar{v}_2(p_2)))\,(H - c) + (1 - b_1) \cdot (1 - F_L(\bar{v}_2(p_2)))\,(L - c)\,, \tag{6}$$

where $\bar{v}_2$ is given by (4).

Analyzing (6), we can identify several important properties of $\hat{V}(p_2 \mid p_1, b_1)$. Firstly, it is strictly positive at $b_1 = \bar{p}$ for all p_2, which follows from the fact that $F_H(\bar{v}_2(p_2)) < F_L(\bar{v}_2(p_2))$. Function $\hat{V}(p_2 \mid p_1, b_1)$ is continuous in b_1, hence it is also strictly positive in some neighborhood of $b_1 = \bar{p}$. This implies that C1 strictly prefers to induce $p_2 \geqslant \bar{p}$ for at least some posteriors $b_1 < \bar{p}$, as compared to inducing $p_2 < \bar{p}$: she wants C2 to purchase the product despite herself believing that this is not myopically optimal. This is due to the social value of experimentation (i.e., of information generated

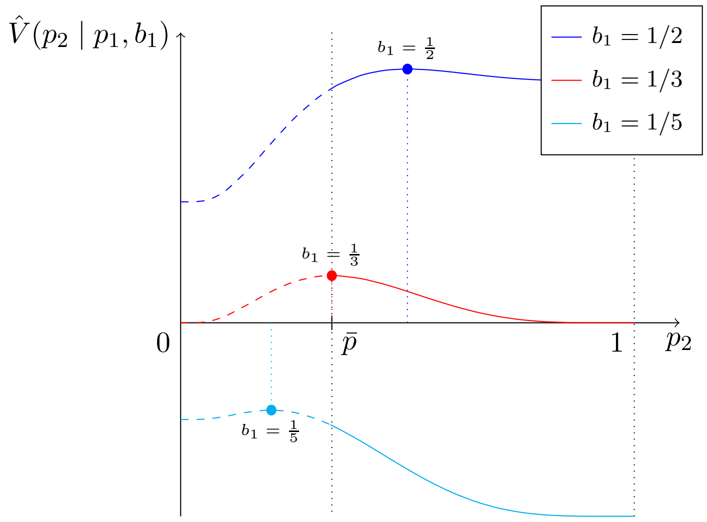

Figure 1: $\hat{V}(p_2 \mid p_1, b_1)$ as a function of p_2.
Note: the parameter values are $H = 3, L = 0, c = 1$ (so $\bar{p} = 1/3$), $\sigma = 4$.

by C2's purchase), which is internalized by C1 in her communication strategy, but not by C2 in her purchasing strategy. There is thus a conflict between the two consumers, which leads to noise in communication between the two.

Secondly, the expression in (6) is strictly increasing in p_2 on $[0, b_1]$ and is strictly decreasing on $[b_1, 1]$, i.e., it is single-peaked with a peak at $p_2 = b_1$. This means that when $b_1 \geqslant \bar{p}$, C1 would prefer to tell the truth to C2 and induce the correct belief $p_2 = b_1$. To see this, observe that

$$\frac{\partial \hat{V}(p_2 \mid p_1, b_1)}{\partial p_2} = (1 - b_1)\, \bar{p} f_L(\bar{v}_2(p_2)) \cdot \frac{\sigma^2}{(1-\bar{p})p_2(1-p_2)} \left(\frac{b_1}{1-b_1} \cdot \frac{1-\bar{p}}{\bar{p}} \cdot \frac{f_H(\bar{v}_2(p_2))}{f_L(\bar{v}_2(p_2))} - 1 \right).$$

Since $\frac{f_H(\bar{v}_2(p_2))}{f_L(\bar{v}_2(p_2))}$ is strictly decreasing in p_2, and the fraction multiplying the bracket is positive, we get that $\hat{V}(p_2 \mid p_1, b_1)$ is single-peaked in p_2. We can find the peak from the condition $\frac{\partial V(p_2|p_1,b_1)}{\partial p_2} = 0$, which can be rewritten as

$$\bar{v}_2 = \frac{\sigma^2}{H-L} \left[\ln\left(\frac{\bar{p}}{1-\bar{p}} \right) - \ln\left(\frac{b_1}{1-b_1} \right) \right] + \frac{H+L}{2},$$

thereby implying together with (4) that the peak is at $p_2 = b_1$.

Function $\hat{V}(p_2|p_1, b_1)$ as given by (6) is plotted in Figure 1 (as a function of p_2 given different values of b_1). Since $\hat{V}(p_2|p_1, b_1)$ only coincides with $V(p_2|p_1, b_1)$ when $p_2 \geqslant \bar{p}$ (and $V(p_2|p_1, b_1) = 0$ otherwise), we use dashed lines for values $p_2 < \bar{p}$.

The two properties of $\hat{V}(p_2 \mid p_1, b_1)$ outlined above – that it is positive for $b_1 = \bar{p} - \varepsilon$ for at least some $\varepsilon > 0$, and that it peaks at $p_2 = b_1$ – will be used heavily in the analysis that follows.

4.3 First Period: Communication

Under direct communication ($m_1 := q(m_1|p_1)$), instead of writing an arbitrary review, which is then somehow incorporated into C2's prior belief, C1 directly communicates in her review what C2's belief should be after observing the review. Then by committing to truthful communication, C1 with private belief $b_1 = b(p_1, v_1)$ can achieve value

$$V(m_1 \mid p_1, b_1) = \begin{cases} \hat{V}(b_1 \mid p_1, b_1) & \text{if } b_1 = b(p_1, v_1) \geqslant \bar{p}, \\ 0 & \text{if } b_1 < \bar{p}. \end{cases}$$

Our goal in this section is to demonstrate that C1 can do better than this (in expectation over b_1 for a given p_1).

On the one hand, obfuscating information by pooling different posteriors b_1 above $\bar{p}$ is costly for C1. To see this, first note that $\hat{V}(m_1 \mid p_1, b_1)$ as given by (6) is linear in b_1 given p_1 and m_1. Hence $\max_{m_1} \hat{V}(m_1 \mid p_1, b_1)$ for a given p_1 is an upper envelope of a set of functions that are linear in b_1, and is thus itself convex in b_1. It is then immediate that adding a new inducable posterior q' to $\mathcal{E}(p_1)$ would increase $\max_{m_1} \hat{V}(m_1 \mid p_1, b_1)$, since it expands the set of linear functions over which the maximum is taken (and it is easy to verify that neither of these functions dominates any other).

On the other hand, it is also strictly optimal to stop experimentation for all private beliefs b_1 lower than some $\bar{\bar{p}} < \bar{p}$ given by the condition $\hat{V}(\bar{\bar{p}} \mid p_1, \bar{\bar{p}}) = 0$.[11] At that point, C1's opinion of the product is so low that the product is not worth experimenting with any further, even after accounting for the informational externality of C2's purchase. Revealing truthfully that $b_1 < \bar{\bar{p}}$ is hence optimal in that case. The two points above support the idea that truthful communication is optimal for $b_1 \geqslant \bar{p}$ and $b_1 < \bar{\bar{p}}$. However, it is for $b_1 \in (\bar{\bar{p}}, \bar{p})$ that garbling information may be beneficial.

Indeed, suppose that instead of being perfectly informative, C1 sends the same message $\bar{p}$ after all b_1 in some ε-neighborhood of $\bar{p}$. As argued above, reporting b_1 inaccurately for $b_1 > \bar{p}$ yields a welfare loss. However, the magnitude of this loss is only of the order ε^2, since $\hat{V}(\bar{p} \mid p_1, b_1)$ is tangent to $\hat{V}(b_1 \mid p_1, b_1)$ at $b_1 = \bar{p}$. Intuitively, this little lies does not affect C2's expected payoff, but may lead C2 to discourage C3 from buying a product that actually has positive expected payoff. Both the expected payoff that C3 would be missing out on and the probability of such a mistake are roughly proportional to ε, hence the total loss is of order ε^2. On the other hand, misreporting $b_1 = \bar{p} - \varepsilon$ as $\bar{p}$ yields benefit approximately equal to $\varepsilon \cdot \hat{V}(\bar{p} \mid p_1, \bar{p})$, due to inducing a purchase from C2 in cases where C2 would have passed if she knew b_1. As was previously established, $\hat{V}(\bar{p} \mid p_1, \bar{p}) > 0$, hence the benefit is of order ε, and it outweighs the cost for at least some small ε. Therefore, garbling information when b_1 is close to $\bar{p}$ is indeed optimal for C1.

Finally, belief p_2 induced by the pooling message must be exactly equal to $\bar{p}$. It can not be lower, since then C2 would not buy the product, which is C1's goal. If the induced posterior is higher than $\bar{p}$, then this only means that C1 send this message after too many private posteriors $b_1 > \bar{p}$ and thus incurs an unnecessary loss, which can not be optimal. In the end, we argue that the optimal communication strategy for C1 is to send message $\bar{p}$ when b_1 is in some neighborhood of $\bar{p}$, and to report b_1 truthfully otherwise. Figure 2 plots the value attained by C1 under this communication strategy.

[11] One can see this formally from combining two inequalities: first, $\hat{V}(b_1 \mid p_1, \bar{\bar{p}}) < \hat{V}(\bar{\bar{p}} \mid p_1, \bar{\bar{p}}) = 0$ for any $b_1 \leqslant \bar{\bar{p}}$ due to single-peakedness of $\hat{V}$ in the first argument. Second, $\hat{V}(b_1 \mid p_1, b_1) < \hat{V}(b_1 \mid p_1, \bar{\bar{p}})$, since $\hat{V}$ is strictly increasing in its last argument. Telling the truth at b_1 then yields negative continuation value to C1, and we know from single-peakedness of $\hat{V}$ that telling the truth is optimal conditional on inducing a purchase by C2.

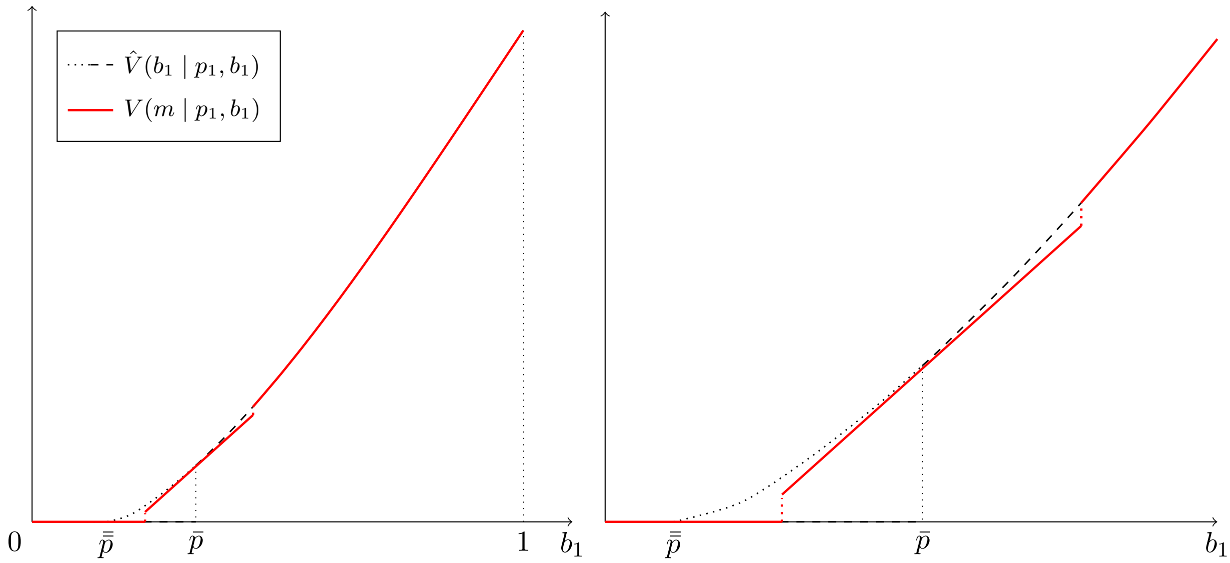

Figure 2: $\hat{V}(b_1 \mid p_1, b_1)$ and $V(m \mid p_1, b_1)$.

Note: $H = 3, L = 0, c = 1, \sigma = 2$. To illustrate convexity, it is also assumed for this graph that C1 cares about third consumer's utility 9 times as much as about C2's (one can think that 9 consumers arrive in the third period). The left panel plots functions for all $b_1 \in [0, 1]$; the right panel focuses on the neighboorhood of $\bar{p}$.

5 Infinite-Horizon Model

This section describes the equilibrium of the infinite-horizon version of the game, assuming it exists. We show that perfect communication is not the optimal communication strategy. Instead, it is always optimal for the consumer to send a vague message when her private posterior belief b is close to $\bar{p}$.

The statement of the main result below and the argument behind it mirror the conclusions from the three-period example (Section 4.3): the reviewer's desire to inflate the review of a marginally-bad item for sake of social experimentation results in garbled communication being optimal. The reviewer ends up issuing the same review for a product that she believes is barely good enough and for a product that is subpar but not bad enough for her to outright reject. After reading such a review, the next consumer is exactly indifferent between buying the product and not, so buys it for sake of writing a review and benefitting the future generations. Conversely, if a reviewer is sufficiently confident in her judgement of the product quality, then it is optimal for her to report her impressions truthfully.

Theorem 1. *In equilibrium, at every public state $p_t \geqslant \bar{p}$, the t-consumer's communication strategy is characterized by cutoffs $l(p_t) < \bar{p} < r(p_t)$ such that:*

1. *For all $b_t \in [0, l(p_t)]$, the consumer sends message $m_t \in \mathcal{S}(p_t)$, i.e., the experimentation stops.*
2. *For any $b_t \in [r(p_t), 1]$, the consumer truthfully reveals her private belief b_t: $m_t = b_t$.*
3. *For all $b_t \in (l(p_t), r(p_t))$, the consumer sends message $m_t = \bar{p}$.*

This conclusion is driven by the implicitly lexicographic nature of the consumers' preferences. When buying the product, a consumer maximizes own expected utility, but when writing a review, she cares about all future generations. No consumer is willing to sacrifice her consumption utility for the sake of other consumers. This creates a conflict, since time-t reviewer would like the consumer at $t+1$ to sometimes buy the product when it is myopically suboptimal for her, so that more information about product quality is generated.

The reviewer thus faces an incentive to sometimes send message $m \geqslant \bar{p}$ when $b_t < \bar{p}$, i.e., the reviewer herself thinks the product yields negative expected utility. The benefit of such an upwards distortion for an interval of beliefs $b_t \in (\bar{p}-\varepsilon, \bar{p})$ is approximately equal to $\varepsilon V(\bar{p} \mid p_t, \bar{p}) \sim \mathcal{O}(\varepsilon)$. However, for such a message to be credible, in the sense of it inducing a posterior belief $q(m|p_t) \geqslant \bar{p}$, the upwards distortion must be balanced by a proportionate downwards distortion. That is, the same message m must be sent after some $b_t > \bar{p}$. Such a downwards distortion is costly, since it not only conceals some information about the state from all future consumers, but it also decreases the amount of experimentation from $t+2$ onwards. However, the losses from distorting the review downwards after $b_t \in (\bar{p}, \bar{p}+\varepsilon)$ are of order $\mathcal{O}(\varepsilon^2)$, because $V(p_{t+1} \mid p_t, b_t)$ is continuous and smooth in $p_{t+1} \geqslant \bar{p}$. In the end, for small enough ε, the gains outweigh the costs, so it is optimal to send a vague "it's okay" review when the reviewer is sufficiently uncertain of whether is product is worth buying, as opposed to revealing the private posterior b_1 exactly.

The formal proof of Theorem 1 in the appendix proceeds in two main steps. First, we show that $V(b_t \mid p_t, b_t)$ is continuously differentiable and (weakly) convex in b_t above $\bar{p}$. This implies that pooling is only beneficial around the cutoff. The second step shows that gains from pooling over an arbitrarily small interval of posteriors will be of first order, while the losses will be of second order, meaning that some noise is always optimal in equilibrium.

Interestingly, Theorem 1 can be related to empirical evidence suggesting that consumers are most likely to leave reviews after either very positive, or very negative experiences (Trustpilot [2020]). Our result suggests that this reviewing strategy is socially optimal, since leaving no review can be treated as a pooling (catch-all) message $m = \bar{p}$. Of course, in the real world, many consumers do not leave a review regardless of their experience with the product, but our model can be extended to account for that without affecting the main result.

It is also worth pointing out that while the optimal communication strategy induces more social experimentation than perfect communication, it features less experimentation than the first best (in which consumers can transmit their information perfectly while also having perfect control over the future consumers' actions). This is because inducing experimentation is costly: to provide incentives for future consumers to experiment with the product after $b_t < \bar{p}$, the sender must distort the information transmitted after $b_t \geqslant \bar{p}$, meaning that some information is permanently destroyed, leading to consumers after $t+1$ making worse choices that they could have. In particular, if there existed an intermediary, such as a recommender platform, which could privately collect the consumer experiences and make private purchase recommendations to other consumers, such an intermediary could induce more experimentation and get closer to the first-best. This is because such an intermediary could distort the information available to any consumer t without losing the original knowledge, meaning that this knowledge could still be used to make more informed recommendations to other consumers. The information design problem of such an intermediary is explored by Che and Hörner [2018].

However, we can show that if the planner does not have access to private memory, but rather has to select a communication strategy for all consumers that trades off the experimentation incentives against the information loss, then such a planner's strategy would coincide with the equilibrium communication strategy chosen by the consumers.

In particular, let us define the solution that such a planner would choose as follows:

Definition. *Given public belief* p_t *a* first-best solution *consists of communication strategies* $\{\mu(m|p_\tau, b_\tau)\}_{\tau\geqslant t}$ *and belief updating rules* $\{b(p_\tau, v_\tau)\}_{\tau\geqslant t}$ *and* $\{q(p_\tau, m)\}_{\tau\geqslant t}$ *such that the following conditions hold:*

- ***Belief Consistency:*** *condition* (1) *holds for all* (p_t, b_t)*, and* (2) *holds for all* p_t *and all* $m \in \mathcal{M}(p_t)$*;*

- ***Planner's Optimality:*** *communication strategies* $\{\mu(m|p_\tau, b_\tau)\}_{\tau\geqslant t}$ *maximize* $\mathbb{E}_{b_t}[V(m \mid p_t, b_t) \mid p_t]$.

In other words, the planner's solution replaces the [consumer] optimality with the requirement for the communication strategy to maximize the ex ante welfare. The difference between optimality and planner's optimality is that players select μ for a given p_t to maximize the continuation payoff from that p_t, whereas the planner selects all strategies in order to maximize the continuation value from p_0. Note that the planner can only choose the communication strategies, but has no power to control the actual purchasing decisions of the consumers – we maintain that the period-t consumer buys the product if and only if $p_t \geqslant \bar{p}$. The equivalence of the two settings can then be stated as follows.

Proposition 1. *Assume* $\frac{f_L(v)}{L'(L(v))}$ *is bounded for* $v \to -\infty$*. Then for any* $p_1 \in [\bar{p}, 1)$ *there exists a first-best solution with strategies* $\{\mu^*(m|p_t, b_t)\}_{t\geqslant 1}$*. Moreover,*

1. *for any public belief* p_t*, strategy profile* $\{\mu^*(m|p_\tau, b_\tau)\}_{\tau\geqslant t}$ *also constitutes a first-best solution at* p_t *when coupled with the respective belief updating rules;*

2. *this first-best solution constitutes a Markov Perfect Equilibrium.*

6 Cheap Talk

In this section we relax the assumption that reviewers can verifiably commit to a communication strategy and assume instead that consumers choose a message *after* observing their utility realization v_t. We also drop the assumption that a consumer's communication *strategy* is observed by future consumers. This section begins by introducing payoff functions and the equilibrium concept used in the analysis, then proceeds to analyzing the three-period model, and concludes by obtaining some (weaker) general results in the infinite-horizon model.

6.1 Definitions and Preliminaries

We stick to the model introduced in Section 3 is used, with the only modification being that in any period t, the current consumer selects her communication strategy at stage 4 of the timing (i.e., after observing v_t, as opposed to stage 2 in the baseline model). We will continue to use $\mu(m \mid p_t, b_t)$ to refer to time-t consumer's communication strategy. The continuation value that time-t consumer obtains by sending message m in private state (p_t, b_t) is still given by $V(m \mid p_t, b_t)$, as defined in (3).

Definition. *A Markov Perfect Equilibrium of the cheap talk game consists of a strategy profile $\mu(m|p_t, b_t)$ and belief updating rules $b(p_t, v_t)$ and $q(m|p_t)$ such that the following conditions hold:*

- ***Belief Consistency:*** *condition* (1) *holds for all* (p_t, b_t)*, and* (2) *holds for all* p_t *and* $m \in \mathcal{M}(p_t)$*;*
- ***Ex Post Optimality:*** *if* $\mu(m^* \mid p_t, b_t) > 0$ *then* $m^* \in \arg\max_m V(m \mid p_t, b_t)$*.*

In particular, the original (ex ante) Optimality condition has been replaced with Ex Post Optimality: message sent by the time-t consumer must now maximize her continuation value given her private state. To pin down the off-path beliefs, we assume that $q(m|p_t) = \varepsilon$ for all p_t and all $m \notin \mathcal{M}(p_t)$ for some $\varepsilon < \bar{p}$.[12]

Finally, we call a strategy profile (and any corresponding equilibrium) *babbling at* p_t if $|\mathcal{P}(p_t)| = 1$. This means that either time-t consumer sends the same message regardless of v_t, or all the different messages she sends induce the same public posterior belief $q(m|p_t)$. The martingale property of beliefs then implies that $\mathcal{P}(p_t) = \{p_t\}$.

6.2 Three-Period Model

In this section we analyze the three-period cheap talk model. As in Section 4, we include current period t in both private and public states, allowing strategies to be time-dependent. The first thing to note is that the analysis of periods $t = 3$ and $t = 2$ is completely analogous to what is presented in Section 4 for the model with commitment. Specifically, in period 3, communication is irrelevant, since no one else arrives at the market to read C3's review. Then in period 2, C2 has no reason to misreport her experience, and hence truthful reporting is an equilibrium.[13] Furthermore, the analysis of C1's continuation payoffs in Section 4.2 carries over to this setting as well. Therefore, we focus on analyzing the equilibrium first-period communication strategy and show that the equilibrium communication strategy must have an interval structure: i.e., there exists a partition $0 = \Delta_0 < \Delta_1 < \Delta_2 < \ldots = 1$ and messages $m_1, m_2, \ldots$ such that if $b_1 \in (\Delta_{j-1}, \Delta_j)$ then $\mu(m_j \mid p_1, v_1) = 1$. For simplicity and without loss of generality we restrict attention to equilibria with direct communication: for any m sent in equilibrium at p_t, $m = q(m|p_t)$. I.e., the reviewer tells the next consumer what belief she must have upon reading this (and all other) reviews, and this message is credible.

Suppose that $\mathcal{S}(p_1)$ is nonempty for the empty review history R_1, i.e., there exists a review $m_0 \in \mathcal{S}(p_1)$ that will prevent C2 from buying the product. Then this review will be used by C1 if her private posterior b_1 is low enough. To see this, observe that for $b_1 \approx 0$, expression (6) (which is still a valid representation of C1's continuation value conditional on C2 buying the item) reduces to $\hat{V}(p_2 \mid p_1, b_1) = (L - c) \cdot (1 + 1 - F_L(\bar{v}_2(p_2)))$, which is negative because $L < c$, whereas sending m_0 yields $V(m_0 \mid p_1, b_1) = 0$.

Consider now the smallest posterior belief among those available in equilibrium that lead C2 to purchase the product, $m_1 = \min \mathcal{E}(p_1)$. After which experiences v_1 or, equivalently, for which

[12] This belief is admissible in equilibrium, since v_t has full support, and thus at every p_t there exists v deemed possible on equilibrium path such that $b(p_t, v) = \varepsilon$.

[13] Unlike in Section 4, there now exist continuation equilibria at $t = 2$ that are *payoff-distinct* from the perfect communication equilibrium. However, for the purpose of exposition, in this section we select continuation equilibria with perfect communication in period 2 in order to simplify the analysis of period-1 communication.

private posteriors b_1 will C1 send this message? Let Δ_1 denote the private posterior b_1 s.t. C1 is indifferent between sending reviews m_0 and m_1:

$$0 = \qquad V(m_0 \mid p_1, \Delta_1) = V(m_1 \mid p_1, \Delta_1) \qquad = \hat{V}(m_1 \mid p_1, \Delta_1).$$

Since $\hat{V}(p_2 \mid p_1, b_1)$ is increasing in b_1 for given p_1, p_2, it follows that between m_0 and m_1, C1 will prefer to send m_0 when $b_1 < \Delta_1$ and send m_1 when $b_1 > \Delta_1$. Further, recall from Section 4.2 that $\hat{V}(p_2 \mid p_1, b_1) > 0$ when $b_1 = \bar{p}$, i.e., C1 prefers m_1 (a recommendation to buy) when her belief is $b_1 \sim \bar{p}$, so long as m_1 is not too large. On the other hand, she also prefers m_1 to all other reviews $m_j > m_1$ in this case, because $\hat{V}(p_2 \mid p_1, b_1)$ is decreasing in p_2 when $p_2 > b_1$ (due to single-peakedness shown in 4.2). In the end, there exists a non-trivial range of private posteriors b_1, for which it is strictly optimal for C1 to send m_1. The intuition is the same as in Section 4: if there exists a way to make the "most cautious recommendation to buy", then C1 would like to adopt that phrasing for at least some range of posteriors $b_1 \sim \bar{p}$. This is because she wants C2 to purchase the product, thus generating information, even when it is not myopically optimal for C2 – but does not want to distort the information that C2 passes onwards too much.

Since C2 is rational and Bayesian, her inference from m_1 must be consistent with C1's equilibrium strategy. Since $m_1 \geqslant \bar{p}$ and it is sent by some types $b_1 < \bar{p}$, there must exist types $b_1 > \bar{p}$ that also send message m_1. Consider the supremum of such types b_1 and denote it by Δ_2. C1 with posterior $b_1 = \Delta_2$ must (by continuity of $\hat{V}$) be indifferent between leaving review m_1 and a higher review $m_2 > m_1$. However, we know that $\hat{V}(p_2 \mid p_1, b_1)$ is single peaked in p_2 with a peak at $p_2 = b_1$, hence the indifference condition $\hat{V}(m_1 \mid p_1, \Delta_2) = \hat{V}(m_2 \mid p_1, \Delta_2)$ implies that $m_2 > \Delta_2$. Then C1 strictly prefers to leave review m_2 whenever $b_1 \in [\Delta_2, m_2]$, but to make this review credible, C1 must also leave review m_2 for some $b_1 > m_2$, etc. In the end, by iterating the argument, we get that

$$... < \Delta_j < m_j < \Delta_{j+1} < m_{j+1} < ...$$

with strict inequalities at every step. To clarify, this sequence need not be infinite. But this means that equilibrium communication at $t = 1$ necessarily has interval structure: instead of communicating her private belief b_1 (or, equivalently, consumption utility v_1 she received) perfectly, C1 only indicates which interval $[\Delta_j, \Delta_{j+1}]$ her posterior b_1 belongs to. Intuitively, the fact that the aforementioned "most cautious recommendation to buy" is noisy and not perfectly revealing of b_1 implies that all other messages must be noisy as well. Notably, perfect communication is thus impossible even for high posteriors b_1, when there is no conflict between the sender and the receiver. The latter stands in contrast to the classic result of Crawford and Sobel [1982], who also show that cheap talk communication has interval structure, but in their model the conflict is present throughout the state space.

Figure 3 illustrates the payoffs in an example cheap talk equilibrium. It plots the continuation payoff of C1 in an interval equilibrium with three messages: $m_0 < \bar{p}$, $m_1 = \bar{p}$, and $m_2 > \bar{p}$.[14] This payoff coincides with the first-best (when C1 can choose any p_2 and force C2 to purchase the item) whenever $b_1 \in \{m_1, m_2\}$, but is strictly lower for all other posteriors. The noise in communication thus hurts C1 by making the purchasing decision of the *third* consumer less efficient, but the losses

[14] An equilibrium with three messages may not exist; equilibria with more or less than three messages will typically exist, as well as equilibria with $m_1 \neq \bar{p}$.

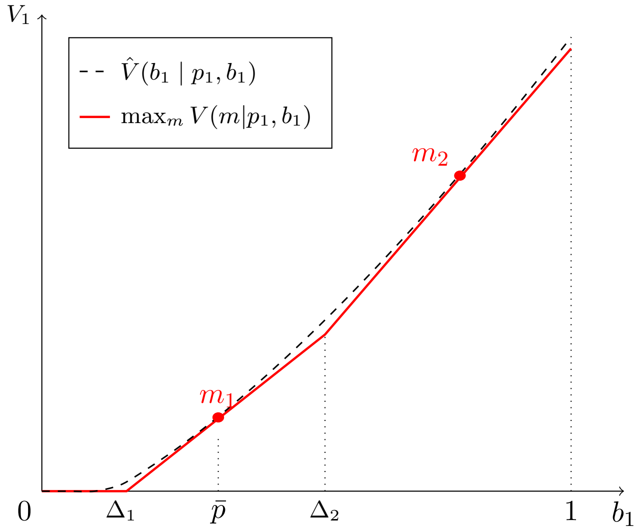

Figure 3: $\hat{V}(b_1 \mid p_1, b_1)$ and $\max_{m \in \mathcal{P}(p_1)} V(m|p_1, b_1)$ in a cheap talk equilibrium.
Note: $H = 3, L = 0, c = 1, \sigma = 2$. To illustrate convexity, it is also assumed for this graph that C1 cares about third consumer's utility 9 times as much as about C2's.

from this noise are offset to some extent by the gain from more social experimentation conducted by C2 (as compared to if C1 communicated truthfully).

6.3 Infinite-Horizon Model

We now generalize the intuition from the three-period model to infinite horizon. The equilibrium multiplicity problem is greatly exacerbated in such a setting, which prevents us from providing a tight characterization of equilibria. However, we can still make some general statements. Specifically, we show that perfect communication can not be an equilibrium outcome in a given period, unless communication in all subsequent period is uninformative.

To formulate the result, we first need to introduce the notion of a cascade. Cascades are prominent in the observational learning literature, where this label is used whenever the society gets locked into one of the available alternatives (possibly at a loss to efficiency).

Definition. *Message $m \in \mathcal{E}(p_t)$ at public state p_t starts a* cascade *if after such a message, $p_s \geqslant \bar{p}$ for all $s > t$.*

In other words, we say that some recommendation to purchase issued at p_t leads to all future consumers buying the product, regardless of any of the interim consumers' experiences and reviews. Once a cascade starts, no new reviews can change the future consumers' behavior. There are two things to note in relation to cascades. First, any message $m \in \mathcal{S}(p_t)$ at any p_t necessarily starts a cascade as well, in the sense that no future consumers buy the product again, as discussed in Section 3.2.[15] Second, with cheap talk, there always exists a continuation equilibrium in which any given $m \in \mathcal{E}(p_t)$ starts a cascade. One example is the babbling equilibrium, the one in which all future reviews are uninformative and are perceived as such, and thus the public belief remains frozen at $q(m|p_t)$.[16] However, in general, a cascade need not shut down the information transmission

[15] The definition above relates to positive cascades, while any $m \in \mathcal{S}(p_t)$ inevitably starts a negative cascade.

[16] Babbling is prominent in cheap talk literature. To see that it is an equilibrium note that neither player has

completely: reviews may be informative and affect the public belief p_t as long as they do not affect future consumers' actual purchasing decisions.

Proposition 2. *In the cheap talk model, at any public state p_t: $[\bar{p}, 1] \subset \mathcal{P}(p_t)$ only if any message $m \in \mathcal{M}(p_t)$ starts a cascade.*

Proposition 2 demonstrates that the conflict between the sender and the receiver of a review precludes perfect communication. It claims that unless *any* message m available in period t starts a cascade, time-t consumer cannot have access to all possible public posteriors $[\bar{p}, 1]$. This implies that some information is inevitably lost in period t, unless all information after period t is ignored. The idea is that if the reviewer's (time-t consumer's) posterior b_t is just below the myopic cutoff $\bar{p}$, she generally wants the next consumer to purchase the product and generate information about quality for the sake of future generations. The receiver (consumer at $t + 1$), however, would not buy the product if she learned that given all available information, the product is good only with probability $b_t < \bar{p}$. The reviewer thus wants to misrepresent her posterior as if it was barely above $\bar{p}$. As argued in the previous subsection, in the absence of commitment, this noisiness of communication unravels. In particular, even though the sender-receiver conflict only exists for some $b_t < \bar{p}$, the noise propagates to all $b_t > \bar{p}$.

The proposition above also illustrates that the noise arises exactly from the reviewer's regard for consumers beyond time $t+1$. In particular, if no informative communication is possible at $t+1$ or afterwards, then time-t consumer has no reason to induce experimentation that the consumer at $t + 1$ is trying to avoid, because the information from these experiments would not be conveyed to the subsequent generations either way.

Proposition 2 is a negative statement, claiming that perfectly informative equilibria are unattainable in the cheap talk game. Theorem 2 below is, conversely, a positive statement, providing a partial (if weak) characterization of how all equilibria *should* look like. It claims that experiences b_t in some neighborhood of the myopic cutoff $\bar{p}$ are always pooled together into a single review. This statement applies to any equilibrium of the game.

Theorem 2. *In any equilibrium of the cheap talk game, for any p_t for which there exists $m \in \mathcal{M}(p_t)$ that does not start a cascade, there exist $l(p_t)$ and $r(p_t)$ such that $l(p_t) < \bar{p} < r(p_t)$, and for all v_t s.t. $b(p_t, v_t) \in [l(p_t), r(p_t)]$ we have $\mu(m \mid p_t, b_t) = 1$ for one of such m.*

The intuition behind this statement mostly mirrors those from Sections 5 and 6.2. Namely, if one of the consumers arriving after t is able to leave informative reviews, the option value of this information for social welfare makes it optimal for the t-consumer to make the consumer at $t + 1$ buy the product when it is myopically suboptimal for the latter. To make this recommendation credible, it must also be sometimes issued after private beliefs b_t that above $\bar{p}$. The part that is worth pointing out here is the qualifier on p_t: in particular, communication at p_t is noisy in the sense described above only if at least some message is available in $\mathcal{M}(p_t)$ that does not start a cascade – i.e., if at least some future consumer can issue a pivotal review. The complementary case was discussed in Proposition 2: if all messages in $\mathcal{M}(p_t)$ start a cascade then perfect comunication in state p_t is possible.

a profitable deviation. The sender cannot benefit by sending informative messages because they are ignored by the receivers regardless, and the receivers cannot benefit by following the sender's recommendation since it is uninformative.

7 Conclusion

This paper builds a theoretical model of social learning through product reviews, focusing on the issue of information provision. We look closely at the empirically observed tension between self-interest in purchasing behavior and prosocial motives when writing a review, and we investigate how this tension affects the informational content of the reviews. The conflict emerges from the reviewers' desire to deceive future consumers into buying a potentially subpar product for sake of generating information beneficial for the society.

We show that truthful communication through reviews cannot be sustained in the equilibrium of such a model. Moreover, despite the conflict only arising under specific circumstances, the noise created by it can propagate, making *all* communication noisy in equilibrium. If, however, the reviewer can commit to a particular communication strategy before experiencing the product or, equivalently, a social norm can be chosen by a welfare-maximizing principal that all consumers will have to follow, then the noise in communication, while still present, is more confined, and some experiences can be communicated truthfully.

This paper contributes to the broader literature on social learning, helping to identify the issues that can deteriorate the quality of learning via product reviews, demonstrating that even in the absence of any kind of interference from the sellers or other external actors (platforms, competing sellers), reviews may not be the perfect source of information about products with uncertain characteristics. Further research may help understand the relative magnitudes of informational noise stemming from different sources, as well as identify possible ways in which these different sources of noise interact.

References

S. N. Ali and N. Kartik. Herding with collective preferences. *Economic Theory*, 51(3):601–626, November 2012.

R. Alonso and O. Camara. Bayesian persuasion with heterogeneous priors. *Journal of Economic Theory*, 165:672–706, 2016.

A. Ambrus, E. Azevedo, and Y. Kamada. Hierarchical cheap talk. *Theoretical Economics*, 8(1): 233–261, January 2013.

J. Andreoni. Impure altruism and donations to public goods: A theory of warm-glow giving. *The Economic Journal*, 100(401):464–477, June 1990.

I. Arieli and M. Mueller-Frank. Multidimensional social learning. *Review of Economic Studies*, 86 (3):913–940, May 2019.

C. Avery, P. Resnick, and R. Zeckhauser. The market for evaluations. *American Economic Review*, 89(3):564–584, June 1999. ISSN 0002-8282. doi: 10.1257/aer.89.3.564.

G. S. Becker. A theory of social interactions. *Journal of Political Economy*, 82(6):1063–1093, 1974.

S. Bikhchandani, D. Hirshleifer, O. Tamuz, and I. Welch. Information cascades and social learning. NBER Working paper 28887, 2021.

P. Bolton and C. Harris. Strategic experimentation. *Econometrica*, 67(2):349–374, March 1999.

A. Campbell. Word-of-mouth communication and percolation in social networks. *American Economic Review*, 103(6):2466–2498, October 2013.

H. H. Cao, B. Han, and D. Hirshleifer. Taking the road less traveled by: Does conversation eradicate pernicious cascades? *Journal of Economic Theory*, 146(4):1418–1436, July 2011. ISSN 00220531. doi: 10.1016/j.jet.2011.03.010.

Y.-K. Che and J. Hörner. Recommender systems as mechanisms for social learning. *Quarterly Journal of Economics*, 133(2):871–925, May 2018.

S. Chiba. Hidden profiles and persuasion cascades in group decision-making. Kyoto U. GSE DP E-18-001, 2018.

L. Cohen and Y. Mansour. Optimal algorithm for bayesian incentive-compatible exploration. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 135–151, 2019.

Competition & Markets Authority. Online reviews and endorsements. `https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/436238/Online_revi` 2015. Retrieved January, 2020.

V. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, November 1982.

eMarketer. Surprise! most consumers look at reviews before a purchase, 2018. URL `https://www.emarketer.com/content/surprise-most-consumers-look-at-reviews-before-a-purchase.` Retrieved January, 2020.

E. Fehr and K. M. Schmidt. A theory of fairness, competition, and cooperation. *The Quarterly Journal of Economics*, 114(3):817–868, August 1999.

E. Fehr and K. M. Schmidt. Theories of fairness and reciprocity: Evidence and economic application. *M. Dewatripont, L.P. Hansen, and S.J. Turnovsky, Advances in Economics and Econometrics – 8th World Congress, Econometric Society Monographs*, pages 208–257, 2003.

M. Foerster. Dynamics of strategic information transmission in social networks. *Theoretical Economics*, 14(1):253–295, January 2019.

A. Galeotti, C. Ghiglino, and F. Squintani. Strategic information transmission networks. *Journal of Economic Theory*, 148(5):1751–1769, 2013.

S. Galperti and B. Strulovici. A theory of intergenerational altruism. *Econometrica*, 85(4):1175–1218, July 2017.

P. Heidhues, S. Rady, and P. Strack. Strategic experimentation with private payoffs. *Journal of Economic Theory*, 159:531–551, 2015.

N. Inostroza and A. Pavan. Adversarial coordination and public information design. Working paper, 2022.

G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, January 2005.

J. Konow. Which is the fairest one of all? a positive analysis of justice theories. *Journal of Economic Literature*, 41(4):1188–1239, December 2003.

I. Kremer, Y. Mansour, and M. Perry. Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012, October 2014.

T. N. Le, V. G. Subramanian, and R. A. Berry. Are imperfect reviews helpful in social learning? In *2016 IEEE International Symposium on Information Theory (ISIT)*, pages 2089–2093, Barcelona, Spain, July 2016. IEEE. ISBN 978-1-5090-1806-2. doi: 10.1109/ISIT.2016.7541667.

M. T. Le Quement and A. Patel. Communication as gift-exchange. Working paper, 2018.

A. Liang and X. Mu. Complementary information and learning traps. *The Quarterly Journal of Economics*, 135(1):389–448, 2020.

I. Lobel and E. Sadler. Information diffusion in networks through social learning. *Theoretical Economics*, 10(3):807–851, September 2015.

M. Luca and G. Zervas. Fake it till you make it: Reputation, competition, and yelp review fraud. *Management Science*, 62(12):3412–3427, December 2016.

Y. Mansour, A. Slivkins, and V. Syrgkanis. Bayesian incentive-compatible bandit exploration. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pages 565–582, 2015.

C. March and A. Ziegelmeyer. Altruistic observational learning. *Journal of Economic Theory*, 190: 105–123, 2020.

S. Meier. A survey of economic theories and field evidence on pro-social behavior. Working paper, 2006.

Mintel. Seven in 10 americans seek out opinions before making purchases, 2015. URL `https://www.mintel.com/press-centre/social-and-lifestyle/seven-in-10-americans-seek-out-opi` Retrieved January, 2020.

D. Peng, Y. Rao, X. Sun, and E. Xiao. Optional disclosure and observational learning. Working paper, SSRN 3083741, 2017.

J. Renault, E. Solan, and N. Vieille. Dynamic sender–receiver games. *Journal of Economic Theory*, 148(2):502–534, 2013.

S. Schopohl. Information transmission in hierarchies. Working paper, 2017.

A. Smirnov and E. Starkov. Bad news turned good: Reversal under censorship. *American Economic Journal: Microeconomics*, 14(2):506–560, 2022.

L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2): 371–398, March 2000.

L. Smith, P. Sørensen, and J. Tian. Informational herding, optimal experimentation, and contrarianism. *Review of Economic Studies*, 88(5):2527–2554, 2021.

O. Swank and B. Visser. Learning from others? decision rights, strategic communication, and reputational concerns. *American Economic Journal: Microeconomics*, 7(4):109–149, 2015.

Trustpilot. Why do people write reviews? what our research revealed, 2020. URL `https://business.trustpilot.com/reviews/learn-from-customers/why-do-people-write-reviews-wha` Retrieved April, 2022.

A. Wolitzky. Learning from others' outcomes. *American Economic Review*, 108(10):2763–2801, October 2018.

Appendix

A.1 Supplementary Lemmas

Fix an arbitrary state $p_t \in [\bar{p}, 1)$.[17] Let $\Phi(b|\theta, p_t)$ denote the c.d.f. of the private posterior b_t from the point of view of time-t consumer with prior p_t given θ. We next derive the exact expression for $\Phi(b|\theta, p_t)$.

Lemma 3. *The distribution of b_t conditional on θ and public belief p_t is*

$$\Phi(b|\theta, p_t) = F_\theta\left(L^{-1}\left(\ln\frac{b}{1-b} - \ln\frac{p_t}{1-p_t}\right)\right).$$

Proof. By definition

$$\begin{aligned}\Phi(b|\theta, p_t) &= \mathbb{P}\left\{b(p_t, v_t) \leqslant b \mid \theta, p_t\right\} = \mathbb{P}\left\{\ln\frac{b(p_t, v_t)}{1-b(p_t, v_t)} \leqslant \ln\frac{b}{1-b} \mid \theta, p_t\right\} = \\ &= \mathbb{P}\left\{\ln\frac{p_t}{1-p_t} + L(v_t) \leqslant \ln\frac{b}{1-b} \mid \theta, p_t\right\} = \mathbb{P}\left\{v_t \leqslant L^{-1}\left(\ln\frac{b}{1-b} - \ln\frac{p_t}{1-p_t}\right) \mid \theta, p_t\right\} = \\ &= F_\theta\left(L^{-1}\left(\ln\frac{b}{1-b} - \ln\frac{p_t}{1-p_t}\right)\right).\end{aligned}$$

□

Given $\Phi(b|\theta, p_t)$, we can define $\Phi(b|p_t)$, – the c.d.f. of the private posterior b_t from the point of view of time-t consumer with prior p_t.

$$\Phi(b|p_t) = p_t \cdot \Phi(b|\theta = H, p_t) + (1-p_t) \cdot \Phi(b|\theta = L, p_t) \tag{7}$$

and let $\phi(b|p_t) = \frac{d\Phi(b|p_t)}{db}$ denote the respective p.d.f. Then $q(m|p_t, \mu)$ can in equilibrium be written as an expectation w.r.t. $\Phi(b_t|p_t)$ of the private posterior b_t of the agent who sent report m:

$$q(m|p_t) = \mathbb{E}[b|m, p_t] = \frac{\int_0^1 b \cdot \mu(m|p_t, b) \cdot \phi(b|p_t) db}{\int_0^1 \mu(m|p_t, b) \cdot \phi(b|p_t) db}. \tag{8}$$

[17] Remember that if $p_t < \bar{p}$ the experimentation stops, while if $p_t = 1$ the distribution becomes degenerate and reduces to a point mass at $b = 1$.

The next result provides a more convenient representation for value function (3). It shows that the continuation value $V(m \mid p_t, b_t)$ after any message $m \in \mathcal{E}(p_t)$ can be characterized by two value functions $V_H(m)$ and $V_L(m)$ that do not depend on b_t or p_t.

Lemma 4. *Fix any equilibrium. For any public belief p_t, if $m \in \mathcal{E}(p_t)$ then*

$$V(m \mid p_t, b_t) = \theta(b_t) - c + \beta \cdot \Big[b_t \cdot V_H(p_{t+1}) + (1 - b_t) \cdot V_L(p_{t+1})\Big], \tag{9}$$

where $p_{t+1} = m = q(m|p_t)$ and

$$V_\theta(p_t) := \mathbb{E}\left[\sum_{s=t+1}^{+\infty} \beta^{s-t-1} \cdot \mathbb{I}\left(p_s \geqslant \bar{p}\right) \cdot (v_s - c) \,\middle|\, \theta, p_t\right]. \tag{10}$$

Moreover, $V(m \mid p_t, b_t)$ is linear and strictly increasing in b_t for any given $m \in \mathcal{E}(p_t)$.

Proof. Assumption $m \in \mathcal{E}(p_t)$ means the $t+1$-consumer will buy the product, receiving a payoff that period-t consumer estimates at $\mathbb{E}[v_{t+1}|b_t] = \theta(b_t)$. For all $s \geqslant t+2$, v_s are independent from $\{v_1, \ldots, v_t\}$, so p_s are independent from b_t given m. Hence (3) reduces to exactly (9). Linearity follows immediately, since both $\theta(b_t)$ and $b_t \cdot V_H(p_t) + (1 - b_t) \cdot V_L(p_t)$ are linear in b_t. To show monotonicity, observe first that $\theta(b_t)$ is strictly increasing in b_t. Second, for any $s \geqslant t+1$: $\mathbb{E}\left[v_s - c \mid \theta = H, p_t\right] > 0 > \mathbb{E}\left[v_s - c \mid \theta = L, p_t\right]$, and v_s is independent of p_s (because v_s is independent of $\{v_1, \ldots, v_{s-1}\}$), meaning that $V_H(p_{t+1}) \geqslant 0 \geqslant V_L(p_{t+1})$. The result then follows. ◻

Lemma 5. *Fix an arbitrary $p_t \in [\bar{p}, 1)$ and the corresponding distribution $\Phi(b|p_t)$ given by (7). Suppose $b \sim \Phi(b|p_t)$. For any $l \in [0, \bar{p}]$, let $r(l)$ be implicitly defined by the following condition, whenever a solution exists:*

$$\mathbb{E}\left[b \mid b \in [l, r]\right] = \bar{p}. \tag{11}$$

Then $r(l)$ is well defined for all $l \in [0, \bar{p}]$, and is differentiable on the interior of this interval.

Proof. Rewriting definition (11), we get

$$\int_{\bar{p}}^{r(l)} (b - \bar{p}) \cdot \phi(b|p_t) db = \int_{l}^{\bar{p}} (\bar{p} - b) \cdot \phi(b|p_t) db. \tag{12}$$

The left-hand side of (12) is continuous, strictly increasing in r, and equals zero for $r = \bar{p}$. Since $\mathbb{E}[b|p_t] = p_t \geqslant \bar{p}$, it follows that

$$\int_{\bar{p}}^{1} (b - \bar{p})\phi(b|p_t) db \geqslant \int_{0}^{\bar{p}} (\bar{p} - b)\phi(b|p_t) db. \tag{13}$$

The right-hand side of (13) is greater than that of (12) for any $l \in [\bar{p}, 1]$. Combining the observations above, we can conclude by the intermediate value theorem that $r(l)$ exists.

Applying the implicit function theorem to (12) we obtain

$$\frac{dr(l)}{dl} = -\frac{\bar{p} - l}{r(l) - \bar{p}} \cdot \frac{\phi(l|p_t)}{\phi(r(l)|p_t)}$$

The first fraction in this expression is trivially bounded for all $l \in [0, \bar{p})$, and the second fraction is finite for $r(l) \in [\bar{p}, 1)$ since $\phi(b|p)$ is continuous, strictly positive, and finite. $\square$

Lemma 6. *Fix an arbitrary $p_t \geqslant \bar{p}$ and the corresponding distribution $\Phi(\cdot|p_t)$ given by (7); let $r(l)$ be defined as in Lemma 5. Then for any $\bar{\varepsilon} \in (0, \bar{p})$ there exists $\delta > 0$ such that for all $l \in [\bar{p} - \bar{\varepsilon}, \bar{p}]$: $r(l) - \bar{p} < \delta(\bar{p} - l) = \delta\varepsilon \leqslant \delta\bar{\varepsilon}$.*

Proof. Fix some $\bar{\varepsilon}$. Let $B_l := \inf\{\phi(b_t|p_t) \mid b_t \in (\bar{p} - \bar{\varepsilon}, r(\bar{p} - \bar{\varepsilon}))\}$ and $B_h := \sup\{\phi(b_t|p_t) \mid b_t \in (\bar{p} - \bar{\varepsilon}, r(\bar{p} - \bar{\varepsilon}))\}$. These bounds exist because $\phi(b|p_t)$ is continuous, strictly positive, and finite. We can therefore obtain from (12) that

$$\frac{B_l}{2}\left(r(l) - \bar{p}\right)^2 < \frac{B_h}{2}\left(\bar{p} - l\right)^2 \qquad \Longleftrightarrow \qquad r(l) - \bar{p} < \sqrt{\frac{B_h}{B_l}}\left(\bar{p} - l\right) < \sqrt{\frac{B_h}{B_l}}\varepsilon.$$

Recall that $B_l > 0$ (since $\phi(b_t|p_t) > 0$ and $\bar{\varepsilon} < \bar{p}$, so $r(\bar{p} - \bar{\varepsilon}) < 1$), so by setting $\delta := \sqrt{\frac{B_h}{B_l}}$ we get the result. $\square$

Lemma 7. *Suppose $f(x)$ is a [weakly] convex and differentiable function on $[\bar{p}, 1]$, $a, b > 0$ and $f(\bar{p}) = a\bar{p} + b$, $f'(\bar{p}) = a$. Then $\frac{f(x)}{ax+b}$ is a [weakly] increasing function on $[\bar{p}, 1]$.*

Proof. Consider $y > x \geqslant \bar{p}$. Then

$$\frac{f(y)}{ay+b} - \frac{f(x)}{ax+b} = \frac{y-x}{(ay+b)(ax+b)}\left((ax+b)\frac{f(y)-f(x)}{y-x} - af(x)\right).$$

Because $f(x)$ is convex we have that $\frac{f(y)-f(x)}{y-x} \geqslant \frac{f(x)-f(\bar{p})}{x-\bar{p}}$. Therefore

$$(ax+b)\frac{f(y)-f(x)}{y-x} - af(x) \geqslant (ax+b)\frac{f(x)-f(\bar{p})}{x-\bar{p}} - af(x) = (a\bar{p}+b)\frac{f(x)-f(\bar{p})}{x-\bar{p}} - af(\bar{p}).$$

$f(\bar{p}) = a\bar{p} + b$ and using convexity of $f(x)$ once again we know that $\frac{f(x)-f(\bar{p})}{x-\bar{p}} \geqslant f'(\bar{p}) = a$. Therefore, the expression above is non-negative. $\square$

Lemma 8. *In any equilibrium, there exists $\delta > 0$ such that for all p_t, all $m \in \mathcal{E}(p_t)$, and all $b_t \leqslant \delta$: $V(m \mid p_t, b_t) < 0$.*

Proof. Since $m \in \mathcal{E}(p_t)$, representation (9) applies. It is then enough to recognize that $V_H(m) \leqslant \frac{H-c}{1-\beta}$ and $V_L(m) \leqslant 0$ to conclude that $V(m \mid p_t, b_t) \leqslant 0$ for all $b_t \leqslant \delta := \frac{(1-\beta)(c-L)}{H-\beta c-(1-\beta)L}$. $\square$

Lemma 9. *Fix some equilibrium. Let $V^*(b_t) := V(b_t \mid p_t, b_t)$ denote the equilibrium continuation value to consumer t of communicating her private belief b_t truthfully. Then $V^*(b)$ is strictly increasing and convex for $b \geqslant \bar{p}$, and $V^*(b) = 0$ for $b < \bar{p}$. Finally, for any $b_t, p_{t+1} \geqslant \bar{p}$, it is true that*

$$V(p_{t+1} \mid p_t, b_t) \leqslant V^*(b_t). \tag{14}$$

Proof. First note that $V^*(b_t)$ is well defined since $V(b_t \mid p_t, b_t)$ does not depend on p_t. Indeed, the path of play from $t+1$ onwards only depends on p_{t+1}, and consumer t's expectation of this play only depends on b_t and p_{t+1}. Then the claim that $V^*(b) = 0$ for $b < \bar{p}$ is immediate from the fact that $t+1$-consumer does not buy the product and receives utility zero when $p_{t+1} < \bar{p}$.

Inequality (14) can be expanded as

$$\begin{aligned} V(p_{t+1} \mid p_t, b_t) &= \mathbb{E}\left[v_{t+1} - c \mid b_t\right] + \beta \cdot \mathbb{E}\left[V(m_{t+1} \mid p_{t+1}, b_{t+1}) \mid \mu^*(p_{t+1}), b_t\right] \\ &\leqslant \mathbb{E}\left[v_{t+1} - c \mid b_t\right] + \beta \cdot \max_{\mu} \mathbb{E}\left[V(m_{t+1} \mid b_t, b_{t+1}) \mid b_t\right] = V(b_t \mid p_t, b_t) = V^*(b_t), \end{aligned}$$

where $\mu^*(p_{t+1}) := \arg\max_{\mu} \mathbb{E}\left[V(m_{t+1} \mid p_{t+1}, b_{t+1}) \mid p_{t+1}\right]$ is the optimal strategy for a $t+1$-consumer given p_{t+1}. The inequality follows from the fact that $\mu^*(p_{t+1})$ does not necessarily maximize $\mathbb{E}\left[V(m_{t+1} \mid p_{t+1}, b_{t+1}) \mid b_t\right]$, as well as from the observation that $V(m_t \mid p_t, b_t)$ does not depend on p_t.

To show that $V^*(b_t) > 0$ for $b_t > \bar{p}$, note that by (14),

$$V^*(b_t) \geqslant V(1 \mid p_t, b_t) = \mathbb{E}\left[\sum_{s=t+1}^{+\infty} \beta^{s-t-1}(v_s - c) \,\middle|\, b_t\right] = \frac{1}{\beta}\mathbb{E}\left[(v_s - c) \,\middle|\, b_t\right] > 0.$$

Finally, to show monotonicity and convexity, observe from Lemma 4 that $V(m_t \mid p_t, b_t)$ is linear and strictly increasing in b_t for all $m_t \geqslant \bar{p}$, i.e., $V(m_t \mid p_t, b_t) = C_1(m_t) + C_2(m_t)b_t$ with $C_2(m_t) > 0$. Inequality (14) then implies that for any p and $b', b'' \geqslant \bar{p}$: $V^*(b'') \geqslant C_1(b') + C_2(b')b''$. Therefore, $C_1(b') + C_2(b')b$ is a subderivative of $V^*(b)$ at $b = b'$. This holds for all b', meaning that $V^*(b)$ is convex. Further, fixing any p and $b' \geqslant \bar{p}$, for any $b'' > b'$:

$$V^*(b') = V(b' \mid p, b') = C_1(b') + C_2(b')b' < C_1(b') + C_2(b')b'' = V(b' \mid p, b'') \leqslant V(b'' \mid p, b'') = V^*(b''),$$

implying that $V^*(b'') > V^*(b')$, so $V^*(b)$ is indeed strictly increasing. This concludes the proof. □

Now we are ready to prove the main theorem.

A.2 Proof of Theorem 1

The proof of the Theorem proceeds in a number of steps.

Step 1. Without loss, we can assume that in equilibrium, at any state $p_t \geqslant \bar{p}$: if $b_t > \bar{p}$, then $p_{t+1} \geqslant \bar{p}$ (i.e., $\mu(m|p_t, b_t) > 0$ only if $m \geqslant \bar{p}$). To show this, proceed by contradiction and assume that given an equilibrium strategy μ, there exist such $p_t \geqslant \bar{p}$, $b_t > \bar{p}$, and $m < \bar{p}$ that $\mu(m|p_t, b_t) > 0$. Consider then an alternative strategy μ', which coincides with μ, with the exception that it communicates b_t in state p_t truthfully: $\mu'(b_t|p_t, b_t) = 1$. Doing so yields continuation value $V^*(b_t)$ in that state, and this value is strictly positive for any $p_t > \bar{p}$ and can only be zero at $p_t = \bar{p}$ by Lemma 9, whereas sending message $m < \bar{p}$ yields a continuation value of zero. At the same time, $q(m|p_t, \mu') \leqslant q(m|p_t, \mu)$ since m is not sent after $b_t > \bar{p}$ under μ', hence message m stops experimentation under μ' same as it does under μ. Therefore, μ' yields higher value than μ in private state (p_t, b_t) and performs equally well in all other states, hence either μ is not optimal, or μ' performs equally well.

Step 2. In equilibrium, there exists $\delta > 0$ such that at any state $p_t \in [\bar{p}, 1)$: if $b_t < \delta$ then $p_{t+1} < \bar{p}$. We will show the statement for $\delta := \frac{(1-\beta)(c-L)}{H-\beta c-(1-\beta)L}$. Suppose by way of contradiction that there exist such p_t, $b_t < \delta$, and $m_t \geqslant \bar{p}$ that $\mu(m_t|p_t, b_t) > 0$ according to the equilibrium strategy μ. Let γ be determined from the condition

$$\frac{\int_0^1 b\mu(m_t|p_t, b)d\Phi(b|p_t)}{\int_0^1 \mu(m_t|p_t, b)d\Phi(b|p_t)} = \frac{\int_\delta^\gamma b\mu(m_t|p_t, b)d\Phi(b|p_t)}{\int_\delta^\gamma \mu(m_t|p_t, b)d\Phi(b|p_t)}, \tag{15}$$

where $\Phi(b|p_t)$ is defined by (7). In words, we look at the equilibrium distribution of b_t conditional on message m_t, and select γ in such a way that the expectation of b_t according to this distribution does not change if we restrict its domain from $[0,1]$ to $[\delta,\gamma]$. Note that $\gamma < 1$ is well defined, since $\mathbb{E}[b|p_t, m_t] = m_t > \bar{p} > \delta$, and the intermediate value theorem applies.

Analogous to the argument in the previous step, we construct an alternative strategy μ' that is identical to μ, except whenever μ prescribes that message m_t is sent, μ' prescribes the following:

1. if $b_t \in [0,\delta]$, then send a truthful message $m = b_t$;

2. if $b_t \in (\delta,\gamma)$, then send m_t as before;

3. if $b_t \in [\gamma,1]$, then send a truthful message $m = b_t$.

Lemma 8 states that if $b_t \leqslant \delta$, then $V(m \mid p_t, b_t) < 0$, i.e., the continuation value of t-consumer from sending any message $m \geqslant \bar{p}$ is negative. Hence part 1 of the modification above strictly (because the interval has non-zero measure) increases the ex post payoffs. If $b_t \in (\delta,\gamma)$ then the consumer receives the exact same payoff as above, since (15) ensures that m_t generates the same posterior under both μ and μ': $q(m_t|p_t,\mu') = q(m_t|p_t,\mu)$. Finally, if $b_t \in [\gamma,1]$, then from Lemma 9 (and inequality (14) in particular) we know that truthful reporting is weakly better for t-consumer than inducing any $m_t \neq b_t$. We conclude that strategy μ' weakly improves t-consumer's continuation value after all private posteriors (and strictly at some) relative to μ, hence μ was not optimal – a contradiction.

Step 3. It is never optimal to pool posteriors above the cutoff $\bar{p}$ (without also pooling them with some posteriors below $\bar{p}$). Formally, in equilibrium, at any p_t, if there exists such a message $m \geqslant \bar{p}$ that $\mu(m|p_t,b'), \mu(m|p_t,b'') > 0$ for some $b', b'' \geqslant \bar{p}$ and $\mu(m|p_t,b) = 0$ for all $b < \bar{p}$, then there exists a weak improvement. Specifically, consider an alternative strategy μ' such that at all (p_t,b_t), at which μ prescribes leaving review m, μ' prescribes instead revealing b_t truthfully with the same probability, and μ' is equivalent to μ otherwise. At any such b_t, reporting m yields continuation value $V(m|p_t,b_t)$, while a truthful report yields $V^*(b_t)$, which is weakly larger by (14), hence μ' is a weakly better strategy for t-consumer than μ.

Step 4. In equilibrium, at any $p_t \geqslant \bar{p}$, if there exists some pooling message $\bar{m}$ such that $\mu(\bar{m}|p_t,b'), \mu(\bar{m}|p_t,b'') > 0$ for some $b' < \bar{p} \leqslant b''$, then $q(\bar{m}|p_t) = \bar{p}$. Suppose by way of contradiction that there exists an equilibrium with a corresponding message strategy μ, in which a pooling message $\bar{m} > \bar{p}$ is sent with positive probability. Consider an alternative strategy μ' for period-t consumer in state p_t which is equivalent to μ, except message $\bar{m}$ is replaced by two. Whenever μ prescribes message $\bar{m}$, μ' sends message $m = \bar{p}$ if $b_t < \gamma$ for some $\gamma_k > \bar{p}$ and sends a new pooling message $\hat{m} := \mathbb{E}[b_t \mid p_t, \bar{m}, b_t \geqslant \gamma]$ whenever $b_t \geqslant \gamma$. Select γ in such a way that $q(\bar{p}|p_t,\mu') = q(\bar{m}|p_t,\mu,b_t < \gamma) = \bar{p}$.

Consider then the t-consumer's expected continuation value conditional on message $\bar{m}$ (or conditional on messages $m \in \{\bar{p},\hat{m}\}$ under μ', which is the same event) under the two strategies:

$$\begin{aligned}
V(p_t|\mu,\bar{m}) &:= \mathbb{E}\left[V(\bar{m} \mid p_t,b_t) \mid p_t,\mu,\bar{m}\right] = V^*(\bar{m}_k),\\
V(p_t|\mu', m \in \{\bar{p},\hat{m}\}) &:= \mathbb{E}\left[V(m \mid p_t,b_t) \mid p_t,\mu', m \in \{\bar{p},\hat{m}\}\right]\\
&= \mathbb{P}\left\{\bar{p} \mid p_t,\mu', m \in \{\bar{p},\hat{m}\}\right\} \cdot V^*(\bar{p}) + \mathbb{P}\left\{\hat{m} \mid p_t,\mu', m \in \{\bar{p},\hat{m}\}\right\} \cdot V^*(\hat{m}).
\end{aligned}$$

Since $V^*(b_t)$ is convex and, by construction, $\bar{m} = \mathbb{P}\left\{\bar{p} \mid p_t,\mu', m \in \{\bar{p},\hat{m}\}\right\} \cdot \bar{p} + \mathbb{P}\left\{\hat{m} \mid p_t,\mu', m \in \{\bar{p},\hat{m}\}\right\} \cdot \hat{m}$ (and the two probabilities sum up to one), Jensen's inequality implies that $V(p_t|\mu', m \in \{\bar{p},\hat{m}\}) > V(p_t|\mu,\bar{m})$. Since the two strategies are otherwise equivalent and thus yield the same continuation values after all other realizations of v_t (and respective b_t), it follows that μ' is a profitable deviation for t-consumer relative to μ. (Note that step 3 further implies that μ' is itself not optimal, and the consumer should reveal b_t truthfully instead of sending $\hat{m}$).

Step 5. It is without loss to restrict attention to Markov equilibria with at most one pooling message. This is because by Step 4, any pooling message must induce the same public posterior $p_{t+1} = \bar{p}$, and by the Markov property, the continuation play must then also be the same after any pooling message.

To summarize the steps above: we now know that in any Markov equilibrium, the t-consumer either reveals her belief truthfully so that $p_{t+1} = b_t$, or sends a pooling message $m = \bar{p}$ such that $p_{t+1} = \bar{p}$.

Step 6. The pooling region (if it is nonempty) is convex left of the cutoff. Formally, in equilibrium, at any $p_t \geqslant \bar{p}$, if there exists such $\bar{m}$ that $\mu(\bar{m}|p_t, b'), \mu(\bar{m}|p_t, b'') > 0$ for some b', b'', then there exists $l(p_t) < \bar{p}$ such that $\mu(\bar{m}|p_t, b_t) = 0$ if $b_t < l(p_t)$ and $\mu(\bar{m}|p_t, b_t) = 1$ if $b_t \in (l(p_t), \bar{p})$. Suppose not – that there exist $b' < b'' < \bar{p}$ such that $\mu(\bar{m}|p_t, b') > 0$ and $\mu(\bar{m}|p_t, b'') < 1$. By Step 4 above, we can let $\bar{m} = \bar{p}$.

Consider an alternative strategy μ' for period-t consumer that is equivalent to μ after all $b_t \geqslant \bar{p}$, sends message $\bar{p}$ for all $b_t \in [l(p_t), \bar{p})$ for some $l(p_t)$, and sends message $m \in \mathcal{S}(p_t)$ if $b_t < l(p_t)$. Select $l(p_t)$ in such a way that $q(\bar{p} \mid p_t, \mu') = \mathbb{E}[b_t \mid p_t, \mu', m = \bar{p}] = \bar{p}$. Then

$$\frac{\int_0^1 b_t \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)}{\int_0^1 \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)} = \bar{p} = \frac{\int_0^1 b_t \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)}{\int_0^1 \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)} \tag{16}$$

We want to show that μ' yields a higher payoff to the t-consumer, i.e., that

$$\mathbb{E}\left[V(m|p_t, b_t) \mid p_t, \mu\right] \leqslant \mathbb{E}\left[V(m|p_t, b_t) \mid p_t, \mu'\right] \tag{17}$$

$$\Leftarrow \int_0^1 V(\bar{p} \mid p_t, b_t) \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \leqslant \int_0^1 V(\bar{p} \mid p_t, b_t) \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \tag{18}$$

where (18) is sufficient for (17) because $\mathbb{E}\left[V(m|p_t, b_t) \mid p_t, \mu, m > \bar{p}\right] = \mathbb{E}\left[V(m|p_t, b_t) \mid p_t, \mu', m > \bar{p}\right]$ (since μ and μ' are equivalent for $m > \bar{p}$) and $V(m \mid p_t, b_t) = 0$ for all $m < \bar{p}$.

Recall that $V(\bar{p} \mid p_t, b_t)$ is a linear function of b_t, i.e., given p_t, $V(\bar{p} \mid p_t, b_t) = \alpha b_t + \beta$ for some $\alpha, \beta \in \mathbb{R}$ (under either messaging strategy). This means that (18) can be written as

$$\alpha \int_0^1 b_t \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) + \beta \int_0^1 \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \leqslant$$
$$\alpha \int_0^1 b_t \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) + \beta \int_0^1 \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t).$$

$$\iff \left[\alpha \frac{\int_0^1 b_t \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)}{\int_0^1 \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)} + \beta\right] \cdot \int_0^1 \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \leqslant$$
$$\left[\alpha \frac{\int_0^1 b_t \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)}{\int_0^1 \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)} + \beta\right] \cdot \int_0^1 \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \tag{19}$$

Note that (16) implies that the square bracket terms on both sides of (19) are equal. Further, since the distribution of b_t conditional on p_t and $m = \bar{p}$ under μ' first-order stochastically dominates that under μ, it follows that $\int_0^1 \mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \leqslant \int_0^1 \mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t)$, and hence the inequality does indeed hold.

Therefore, we conclude that μ' is indeed weakly better for the consumer, hence the original statement is true.

Step 7. The pooling region (if it is nonempty) is convex right of the cutoff. Formally, in equilibrium, at any $p_t \geqslant \bar{p}$, if there exists such $\bar{m}$ that $\mu(\bar{m}|p_t, b'), \mu(\bar{m}|p_t, b'') > 0$ for some b', b'', then there exists $r(p_t) > \bar{p}$ such that $\mu(\bar{m}|p_t, b_t) = 0$ if $b_t > r(p_t)$ and $\mu(\bar{m}|p_t, b_t) = 1$ if $b_t \in (\bar{p}, r(p_t))$. Suppose not – that there exist $\bar{p} < b' < b''$ such that $\mu(\bar{m}|p_t, b') < 1$ and $\mu(\bar{m}|p_t, b'') > 0$. By Step 4 above, we can let $\bar{m} = \bar{p}$.

Consider an alternative strategy μ' for period-t consumer that is equivalent to μ for $b_t < \bar{p}$, sends message $\bar{p}$ for all $b_t \in [\bar{p}, r(p_t))$ for some $r(p_t)$, and reveals b_t truthfully for all $b_t \geqslant r(p_t)$. Select $r(p_t)$ in such a way that $q(\bar{p} \mid p_t, \mu') = \mathbb{E}[b_t \mid p_t, \mu', m = \bar{p}] = \bar{p}$. Then (16) holds for μ and μ' defined in this step, implying that

$$\begin{aligned}
&\int_0^1 (b_t - \bar{p})\mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) = 0 = \int_0^1 (b_t - \bar{p})\mu'(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \\
\Rightarrow &\int_{\bar{p}}^1 (b_t - \bar{p})\mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) = \int_{\bar{p}}^{r(p_t)} (b_t - \bar{p}) d\Phi(b_t|p_t) \\
\iff &\int_{r(p_t)}^1 (b_t - \bar{p})\mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) = \int_{\bar{p}}^{r(p_t)} (b_t - \bar{p}) \left(1 - \mu(\bar{p} \mid p_t, b_t)\right) d\Phi(b_t|p_t). \qquad (20)
\end{aligned}$$

We want to show that μ' yields a higher payoff to the t-consumer, i.e., that

$$\mathbb{E}\left[V(m|p_t, b_t) \mid p_t, \mu\right] \leqslant \mathbb{E}\left[V(m|p_t, b_t) \mid p_t, \mu'\right]. \qquad (21)$$

The equivalence of μ and μ' for $b_t < \bar{p}$ implies that we only need to look at $b_t \geqslant \bar{p}$. From Steps 1-5, for every b_t, μ either sends a pooling message $\bar{p}$, or reveals b_t truthfully. The same holds by construction for μ'. Hence (21) is equivalent to

$$\begin{aligned}
&\int_{\bar{p}}^1 \Big[V(\bar{p} \mid p_t, b_t)\mu(\bar{p} \mid p_t, b_t) + V^*(b_t)\left(1 - \mu(\bar{p} \mid p_t, b_t)\right)\Big] d\Phi(b_t|p_t) \leqslant \\
&\qquad \int_{\bar{p}}^{r(p_t)} V(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) + \int_{r(p_t)}^1 V^*(b_t) d\Phi(b_t|p_t)
\end{aligned}$$

$$\iff \int_{\bar{p}}^1 \Big[V^*(b_t) - V(\bar{p} \mid p_t, b_t)\Big]\mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \geqslant \int_{\bar{p}}^{r(p_t)} \Big[V^*(b_t) - V(\bar{p} \mid p_t, b_t)\Big] d\Phi(b_t|p_t)$$

$$\begin{aligned}
\iff &\int_{r(p_t)}^1 \Big[V^*(b_t) - V(\bar{p} \mid p_t, b_t)\Big]\mu(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) \geqslant \\
&\qquad \int_{\bar{p}}^{r(p_t)} \Big[V^*(b_t) - V(\bar{p} \mid p_t, b_t)\Big] \left(1 - \mu(\bar{p} \mid p_t, b_t)\right) d\Phi(b_t|p_t), \qquad (22)
\end{aligned}$$

where the second inequality is obtained by flipping the signs and then adding $\int_{\bar{p}}^{r(p_t)} V^*(b_t) d\Phi(b_t|p_t)$ to both sides.

Since $V(\bar{p} \mid p_t, b_t)$ is linear in b_t by Lemma 4, and $V^*(b_t)$ is convex by Lemma 9, their difference $V^*(b_t) - V(\bar{p} \mid p_t, b_t)$ is convex in b_t. Hence by Lemma 7, $\frac{V^*(b_t) - V(\bar{p}|p_t, b_t)}{b_t - \bar{p}}$ is increasing for $b_t > \bar{p}$. This monotonicity implies that

$$\int_{r(p_t)}^{1}(b_t-\bar p)\cdot\frac{V^*(r(p_t))-V(\bar p\mid p_t,r(p_t))}{r(p_t)-\bar p}\mu(\bar p\mid p_t,b_t)d\Phi(b_t|p_t)\leqslant$$
$$\int_{r(p_t)}^{1}(b_t-\bar p)\cdot\frac{V^*(b_t)-V(\bar p\mid p_t,b_t)}{b_t-\bar p}\mu(\bar p\mid p_t,b_t)d\Phi(b_t|p_t)=$$
$$\int_{r(p_t)}^{1}\Big[V^*(b_t)-V(\bar p\mid p_t,b_t)\Big]\mu(\bar p\mid p_t,b_t)d\Phi(b_t|p_t),$$

where the top line is the LHS of (20) multiplied by a constant $\frac{V^*(r(p_t))-V(\bar p|p_t,r(p_t))}{r(p_t)-\bar p}$, and the bottom line is the LHS of (22). For the RHS parts of the two expressions, we similarly get

$$\int_{\bar p}^{r(p_t)}(b_t-\bar p)\cdot\frac{V^*(r(p_t))-V(\bar p\mid p_t,r(p_t))}{r(p_t)-\bar p}\left(1-\mu(\bar p\mid p_t,b_t)\right)d\Phi(b_t|p_t)\geqslant$$
$$\int_{\bar p}^{r(p_t)}(b_t-\bar p)\cdot\frac{V^*(b_t)-V(\bar p\mid p_t,b_t)}{b_t-\bar p}\left(1-\mu(\bar p\mid p_t,b_t)\right)d\Phi(b_t|p_t)=$$
$$\int_{\bar p}^{r(p_t)}V^*(r(p_t))-V(\bar p\mid p_t,r(p_t))\left(1-\mu(\bar p\mid p_t,b_t)\right)d\Phi(b_t|p_t).$$

The inequalities above together with (20) imply that (22) holds, meaning that μ' is a profitable deviation from μ, and hence μ cannot be an equilibrium strategy.

Step 8. In equilibrium, $V^*(\bar p)>0$. To see this, invoke the definitions:

$$V^*(\bar p)=V(\bar p\mid p_t,\bar p)=\mathbb{E}\left[v_{t+1}-c\mid\bar p\right]+\beta\max_{\mu}\mathbb{E}\left[V(m_{t+1}\mid\bar p,b_{t+1})\mid\mu,\bar p\right]$$

By definition of $\bar p$, $\mathbb{E}\left[v_{t+1}-c\mid\bar p\right]=0$. The latter expectation is strictly positive, since a truthful reporting strategy μ^* (which is admissible but not necessarily optimal) yields a positive payoff:

$$\mathbb{E}\left[V(m_{t+1}\mid\bar p,b_{t+1})\mid\mu^*,\bar p\right]=\mathbb{E}\left[V^*(b_{t+1})\mid\bar p\right]>0.$$

Here the latter expectation is positive by the properties of $V^*(b)$ outlined in Lemma 9, and the fact that b_{t+1} has full support on $[0,1]$, which follows from our MLRP assumption.

Step 9. In equilibrium, at any $p_t\geqslant\bar p$, the t-consumer's pooling interval is nonempty: $l(p_t)<\bar p<r(p_t)$. Suppose the contrary – which, given the steps above, means that the equilibrium communication strategy μ is truthful: $\mu(b_t\mid p_t,b_t)=1$ for all b_t. Take some $\bar\varepsilon\in(0,\bar p)$. Define $l(p_t,\varepsilon):=\bar p-\varepsilon$ for all $\varepsilon\in(0,\bar\varepsilon]$ and set $r(p_t,\varepsilon)$ in such a way that $\mathbb{E}[b_t\mid p_t,b_t\in(l(p_t,\varepsilon),r(p_t,\varepsilon))]=\bar p$, which is possible by Lemma 5. Consider then a family of strategies $\{\mu'_\varepsilon\}_{\varepsilon\in(0,\bar\varepsilon]}$ that send a pooling message $\bar m$ for all $b_t\in(l(p_t,\varepsilon),r(p_t,\varepsilon))$ and are truthful otherwise. We now show that there exists $\varepsilon\in(0,\bar\varepsilon]$ such that μ'_ε yields a higher expected payoff for the consumer in state p_t than μ.

The difference in the t-consumer's expected payoff under μ'_ε relative to that under μ is given by

$$\int_{l(p_t,\varepsilon)}^{r(p_t,\varepsilon)}V(\bar p\mid p_t,b_t)d\Phi(b_t|p_t)-\int_{\bar p}^{r(p_t,\varepsilon)}V^*(b_t)d\Phi(b_t|p_t).$$

Our goal is to show that for some ε, the expression above is strictly positive, i.e.,

$$\int_{l(p_t,\varepsilon)}^{\bar{p}} V(\bar{p} \mid p_t, b_t) d\Phi(b_t|p_t) > \int_{\bar{p}}^{r(p_t,\varepsilon)} \left[V^*(b_t) - V(\bar{p} \mid p_t, b_t)\right] d\Phi(b_t|p_t). \tag{23}$$

Since $V(m \mid p_t, b_t)$ is linear in b_t, we can use (12) to claim that

$$\int\limits_{l(p_t,\varepsilon)}^{\bar{p}} \left[V^*(\bar{p}) - V(\bar{p} \mid p_t, b_t)\right] d\Phi(b|p_t) = \int\limits_{\bar{p}}^{r(p_t,\varepsilon)} \left[V(\bar{p} \mid p_t, b_t) - V^*(\bar{p})\right] d\Phi(b_t|p_t).$$

Adding the equality above to (23) we get

$$\int_{l(p_t,\varepsilon)}^{\bar{p}} V^*(\bar{p}) d\Phi(b_t|p_t) > \int_{\bar{p}}^{r(p_t,\varepsilon)} \left[V^*(b_t) - V^*(\bar{p})\right] d\Phi(b_t|p_t) \tag{24}$$

$$\Leftarrow V^*(\bar{p}) \cdot \left[\Phi(\bar{p}) - \Phi(l(p_t,\varepsilon))\right] > \left[V^*(r(p_t,\varepsilon)) - V^*(\bar{p})\right] \cdot \left[\Phi(r(p_t,\varepsilon)) - \Phi(\bar{p})\right], \tag{25}$$

where (25) is sufficient for (24) because monotonicity of $V^*(b_t)$ implies that $V^*(r(p_t,\varepsilon)) \geq V^*(b_t)$ for all $b_t \in (\bar{p}, r(p_t,\varepsilon))$.

Recall from Lemma 6 that there exists $\delta > 0$ such that $r(p_t,\varepsilon) - \bar{p} < \delta\varepsilon$ for all $\varepsilon \in (0, \bar{\varepsilon}]$. Let $B_l := \min\{\phi(b_t|p_t) \mid b_t \in [\bar{p} - \bar{\varepsilon}, \bar{p} + \delta\bar{\varepsilon}]\}$ and $B_h := \max\{\phi(b_t|p_t) \mid b_t \in [\bar{p} - \bar{\varepsilon}, \bar{p} + \delta\bar{\varepsilon}]\}$ denote the bounds on density $\phi(b_t|p_t)$ on the relevant interval of b_t. Since $\phi(b|p_t)$ is continuous, strictly positive, and bounded (see the proof of Lemma 6), these bounds exist and $0 < B_l < B_h < \infty$. These bounds then imply that

$$\Phi(\bar{p}) - \Phi(l(p_t,\varepsilon)) \geqslant B_l\varepsilon, \qquad\qquad \Phi(r(p_t,\varepsilon)) - \Phi(\bar{p}) \leqslant B_h\delta\varepsilon. \tag{26}$$

On the other hand, $V^*(r(p_t,\varepsilon)) - V^*(\bar{p})$ is also bounded from above by a factor of ε. To see this, note that by the convexity of $V^*(b_t)$, for all $b_t \in [\bar{p}, 1]$,

$$V^*(b_t) = \frac{b_t - \bar{p}}{1 - \bar{p}} \cdot V^*(1) + \frac{1 - b_t}{1 - \bar{p}} \cdot V^*(\bar{p}).$$

The previous step of this proof has established that $V^*(\bar{p}) > 0$, and Lemma 10 implies that $V^*(1) \leqslant \frac{H-c}{1-\beta}$, so

$$\begin{aligned} V^*(b_t) - V^*(\bar{p}) &\leqslant \frac{b_t - \bar{p}}{1 - b_t} \cdot \left(\frac{H - c}{1 - \beta} - V^*(\bar{p})\right) < \frac{b_t - \bar{p}}{1 - b_t} \cdot \frac{H - c}{1 - \beta} \\ \Rightarrow V^*(r(p_t,\varepsilon)) - V^*(\bar{p}) &< \frac{\delta\varepsilon}{1 - \bar{p} - \delta\bar{\varepsilon}} \cdot \frac{H - c}{1 - \beta}, \end{aligned} \tag{27}$$

where the latter follows from $r(p_t,\varepsilon) - \bar{p} < \delta\varepsilon \leqslant \delta\bar{\varepsilon}$. Combining (27) with (26), we have that (25) holds if the following holds:

$$V^*(\bar{p}) B_l \varepsilon > \frac{H - c}{(1 - \beta)(1 - \bar{p} + \delta\bar{\varepsilon})} B_h (\delta\varepsilon)^2. \tag{28}$$

Since the LHS is proportional to ε and the RHS to ε^2, there exists a small enough ε such that (28) holds, meaning (23) holds as well, meaning that strategy μ'_ε is strictly better for period-t consumer than μ.

This concludes the proof of the theorem. $\square$

A.3 Proof of Proposition 1

Say that $(\Omega, \mathscr{B}, \mathbb{P})$ is the probability space we're living in. Every $\omega \in \Omega$ defines a sequence of utility realizations $\{v_t\}_{t=1}^{\infty}$.

Consider the planner's optimal communication strategy μ_P. It solves

$$\max_{\mu} \mathbb{E}_{b_0}\left[V(m \mid p_0, b_0) \mid p_0\right] = \max_{\mu}\left\{\mathbb{E}\left[\sum_{s=1}^{+\infty} \beta^{s-1} \cdot \mathbb{I}\left(p_s \geqslant \bar{p}\right) \cdot (v_s - c) \,\middle|\, p_0, \mu\right]\right\}$$

Fix an arbitrary belief history $\mathbf{p}_t = (p_0, ..., p_t)$. What condition should $\mu_P(\mathbf{p}_t)$ satisfy? Since $\mu_P(\mathbf{p}_t)$ does not affect anything in periods before t and contingencies $\omega \in \Omega$ that generate belief histories different from $\mathbf{p}_t$, it is without loss to choose $\mu_P(\mathbf{p}_t)$ that maximizes

$$\mathbb{E}\left[\sum_{s=t}^{+\infty} \beta^{s-1} \cdot \mathbb{I}\left(p_s \geqslant \bar{p}\right) \cdot (v_s - c) \,\middle|\, \mathbf{p}_t, \mu_P\right].$$

Note, however, that $\mathbb{E}[\theta|\mathbf{p}_t] = \mathbb{E}[\theta|p_t]$, so p_t contains all information relevant for the distribution of $\{v_s\}_{s\geqslant t}$, and the future path of public beliefs $\{p_s\}_{s\geqslant t}$ is fully determined by the utility process $\{v_s\}$ and the communication strategies $\mu_P(p)$, with the latter restricted to be Markov. Therefore, the problem above is equivalent to

$$\max_{\mu}\left\{\mathbb{E}\left[\sum_{s=t}^{+\infty} \beta^{s-1} \cdot \mathbb{I}\left(p_s \geqslant \bar{p}\right) \cdot (v_s - c) \,\middle|\, p_t, \mu_P\right]\right\},$$

which exactly coincides with the consumer's problem in state p if all future consumers follow strategy μ_P. The planner's "best response" to future μ_P thus coincides with the consumer's best response to future consumers' strategies μ_P. From this we conclude that μ_P can be sustained in equilibrium. □

A.4 Proofs for Section 6

Lemma 10. *Fix any cheap talk equilibrium. Suppose that for some p_t, there exist $m', m'' \in \mathcal{E}(p_t)$ such that $p'_{t+1} < p''_{t+1}$, where $p'_{t+1} := q(m'|p_t)$ and $p''_{t+1} := q(m''|p_t)$. Then one of the following must hold:*

1. *$V_H(p'_{t+1}) = V_H(p''_{t+1})$ and $V_L(p'_{t+1}) = V_L(p''_{t+1})$;*
2. *$V_H(p'_{t+1}) < V_H(p''_{t+1})$ and $V_L(p'_{t+1}) > V_L(p''_{t+1})$.*

Proof. By contradiction, if $V_\theta(p'_{t+1}) < V_\theta(p''_{t+1})$ for both θ, then for any v_t: $V(m' \mid p_t, v_t) < V(m'' \mid p_t, v_t)$, so sending message m' is suboptimal. Hence $m' \notin \mathcal{M}(p_t)$ – a contradiction. Analogously, it cannot be that $V_\theta(p'_{t+1}) > V_\theta(p''_{t+1})$ for both θ. To complete the proof, we are only left to rule out the case $V_H(p'_{t+1}) > V_H(p''_{t+1})$ and $V_L(p'_{t+1}) < V_L(p''_{t+1})$. Again, assume, by way of contradiction, that this is the case. Since $m', m'' \in \mathcal{M}(p_t)$, there must exist such types $b', b'' \in [0,1]$, for whom sending m' and m'' respectively is optimal: $V(m' \mid p_t, b') \geqslant V(m'' \mid p_t, b')$ and $V(m' \mid p_t, b'') \leqslant V(m'' \mid p_t, b'')$. From Lemma 4 it is then immediate that there must exist $\bar{b} \in (0,1)$ such that $V(m' \mid p_t, b) \leqslant V(m'' \mid p_t, b)$ for $b \in [0, \bar{b})$ and $V(m' \mid p_t, b) \geqslant V(m'' \mid p_t, b)$ for $b \in (\bar{b}, 1]$. But then (2) implies that $m' > \bar{b} > m''$, which contradicts the assumption that $m' < m''$. This concludes the proof. □

Corollary 11. *In any cheap talk equilibrium, for any p_t:*

1. *$V_H(q(m|p_t))$ is a weakly increasing function of m on $\mathcal{E}(p_t)$,*

2. *$V_L(q(m|p_t))$ is a weakly decreasing function of m on $\mathcal{E}(p_t)$.*

Proof. Follows directly from Lemma 10. $\square$

To remind, we focus without loss of generality on direct communication strategies, where $m = q(m|p)$, and use $\mu(q \mid p, b)$ to denote the probability with which a consumer in private state (p, b) sends message m Further, we will be using the following shorthand notation for the consumers' optimal continuation value as a function of their private state (p_t, b_t):

$$V(p_t, b_t) = \max_{m \in \mathcal{P}(p_t)} V(m \mid p_t, b_t).$$

Lemma 12. *In any cheap talk equilibrium, at any public state $p \in (0, 1)$ there exists $\bar{b}(p) \in [0, 1]$ such that*

1. *For all $b < \bar{b}(p)$, $\mu(m \mid p, b) > 0$ only if $m \in \mathcal{S}(p)$;*
2. *For all $b \geqslant \bar{b}(p)$, $\mu(m \mid p, b) > 0$ only if $m \in \mathcal{E}(p)$.*

Proof. Assume the contrary: that there exist $b' < b''$ and $q' \in \mathcal{E}(p)$, $q'' \in \mathcal{S}(p)$ such that $\mu(q'|p, b') > 0$ and $\mu(q''|p, b'') > 0$. Then $V(q' \mid p, b') = V(p, b') \geqslant 0$ and $V(q'' \mid p, b'') = V(p, b'') = 0$. At the same time, $b' < b''$ and representation (9) from Lemma 4 imply $V(q' \mid p, b'') > V(q' \mid p, b')$, meaning

$$0 = V(p, b'') \geqslant V(q' \mid p, b'') > V(q' \mid p, b') = V(p, b') \geqslant 0,$$

which is a contradiction. This argument proves the two first parts of the lemma. $\square$

Lemma 13. *In any cheap talk equilibrium, at any public state $p \in (0, 1)$,*

1. *If there exists $m \in \mathcal{E}(p)$ s.t. $\mathcal{S}(q(m|p)) \neq \emptyset$, then $V(p, \bar{p}) > 0$;*
2. *Threshold $\bar{b}(p)$ defined in Lemma 12 satisfies $\bar{b}(p) \in [0, \bar{p}]$. Further, $\bar{b}(p) = \bar{p}$ if and only if the experimentation never stops after any $m \in \mathcal{E}(p)$.*[18]

Proof. By Lemma 12 it follows that if $\mathcal{S}(q(m|p)) \neq \emptyset$ then $\bar{b}(q(m|p)) > 0$. Next, we claim that from point of view of consumer at p who holds private belief $b = \bar{p}$, the flow payoff at every future history is weakly positive. Indeed, let $\tau(p, b) := \inf\{t | p_t < \bar{b}(p_t), p_0 = p, b_0 = b\}$ be a stopping time adapted to the filtration generated by v_t that indicates the first period when experimentation stops, depending on the path of realized utilities, from the point of view of a consumer with belief b in state p. Then

$$V(p, \bar{p}) = \mathbb{E}\left[\sum_{t=0}^{\tau(p,\bar{p})} \beta^t (v_t - c) \mid b_0 = \bar{p}\right]$$

(where the conditioning on current public belief p in the expectation is subsumed by $\tau(p, \bar{p})$). If τ is not correlated with θ, then $V(p, \bar{p}) = 0$, since $\mathbb{E}[v_t - c | b_0 = \bar{p}] = 0$. However, Lemma 12 implies that τ is positively correlated with the state: for any $t > 0$ and $p_t \in [\bar{p}, 1)$, the state-contingent continuation probability

$$\begin{aligned}\mathbb{P}[\tau > t + 1 \mid \tau > t, b_0 = b; p_t, \theta] &= \mathbb{P}[b_t \geqslant \bar{b}(p_t) \mid p_t, \theta] \\ &= \mathbb{P}\left[v_t \geqslant L^{-1}\left(\ln\left(\frac{\bar{b}(p_t)}{1 - \bar{b}(p_t)}\right) - \ln\left(\frac{p_t}{1 - p_t}\right)\right) \mid p_t, \theta\right],\end{aligned}$$

[18] "Experimentation never stops after p" means $\mathcal{S}(m) = \emptyset$ for all $m \in \cup_{t \geqslant 1} \mathcal{E}^t(p)$, where $\mathcal{E}^1(p) := \mathcal{E}(p)$, and $\mathcal{E}^t(p) := \cup_{m \in \mathcal{E}^{t-1}(p)} \mathcal{E}(m)$ for all $t \geqslant 2$.

is weakly increasing in θ due to MLRP and strictly increasing if $\bar{b}(p_t) \notin \{0,1\}$, with $L^{-1}(x)$ being an inverse function of $L(v) = \ln\left[\frac{f_H(v)}{f_L(v)}\right]$. Hence for any t and $p_t \in [\bar{p}, 1)$: $\mathbb{E}\left[v_t \mid \tau > t, b_0 = \bar{p}; p_t\right] \geqslant \mathbb{E}\left[v_t \mid b_0 = \bar{p}\right] = 0$, and $\mathbb{E}\left[v_1 \mid \tau > 1, b_0 = \bar{p}; p_1 = q(m|p)\right] > \mathbb{E}\left[v_t \mid b_0 = \bar{p}\right] = 0$ because $\mathcal{S}(q(m|p)) \neq \emptyset$. Therefore, $V(p, \bar{p}) > 0$.

Finally, we prove the last part. If experimentation never stops after any $m \in \mathcal{E}(p)$, the current consumer decides whether all future consumers will buy the product or avoid it. Their discounted expected utilities in the two cases from her point of view is then $\frac{1}{1-\beta}(\theta(b) - c)$ and zero, respectively. Therefore, the current consumer sends $m \in \mathcal{E}(p)$ if and only if $\theta(b) \geqslant c$ – equivalently, $b \geqslant \bar{p}$, – and sends $m \in \mathcal{S}(p)$ otherwise.

It is left to show the converse – that if experimentation can be stopped by some future consumer then $\bar{b}(p) < \bar{p}$. Let $\tau'(p) := \inf\{t \geqslant 1 \mid \exists q \in \mathcal{E}^t(p) : \mathcal{S}(q) \neq \emptyset\}$ denote the number of periods that must pass before one of the following consumers has the opportunity to stop experimentation, where $\mathcal{E}^t(p)$ is as in footnote 18. We will show that if $\tau'(p) < \infty$, then $V(p, \bar{p}) > 0$, which is equivalent to $\bar{b}(p) < \bar{p}$. The latter follows from $V(p, b)$ being continuous in b (due to piecewise-linearity, see Lemma 4), and the fact that $V(p, \bar{b}(p)) = 0$, because stopping the experimentation is weakly optimal at $\bar{b}(p)$. If $\tau'(p) = 1$, the implication in the lemma requires that there exists $q \in \mathcal{E}(p)$ such that $\mathcal{S}(q) \neq \emptyset$. Part 1 of this lemma then immediately yields that $V(p, \bar{p}) > 0$, which proves the statement for $\tau'(p) = 1$. We will proceed by induction on τ'.

Suppose now that an induction statement "if $\tau'(p) = t - 1$ then $V(p, \bar{p}) > 0$" is true for any p for some $t \geqslant 2$. To see that it then also holds for $\tau'(p) = t$, consider the following argument. Let q' denote the public posterior such that $q' \in \mathcal{E}(p)$ and $\tau'(q') = t - 1$ (in case of multiple such posteriors select arbitrarily). Then

$$\begin{aligned} V(p, \bar{p}) &\geqslant V(q' \mid p, \bar{p}) \\ &= \mathbb{E}[v_1 - c \mid b_0 = \bar{p}] + \beta \mathbb{E}\left[V(q', b_1) \mid q'; b_0 = \bar{p}\right], \end{aligned}$$

where $\mathbb{E}[v_1 - c \mid b_0 = \bar{p}] = 0$ by definition of $\bar{p}$. Further, $V(q', b_1) = \max_q V(q \mid q', b_1)$ for all b_1, and Lemma 4 states that $V(q \mid q', b_1)$ is linear in b_1 for all q. The two together imply that $V(q', b_1)$ is, for a given q', an upper envelope of a collection of functions that are linear in b_1. Therefore, $V(q', b_1)$ is weakly convex in b_1. By Jensen's inequality, $\mathbb{E}[V(q', b_1)] \geqslant V(q', \mathbb{E}[b_1])$, and we thus have that

$$V(p, \bar{p}) \geqslant \beta V(q', \mathbb{E}[b_1 \mid q'; b_0 = \bar{p}]).$$

By the martingale property of beliefs, $\mathbb{E}[b_1 \mid q'] = q'$. We know that $q' \geqslant \bar{p}$, since $q' \in \mathcal{E}(p)$. Splitting this into two cases: if $q' = \bar{p}$, then $\mathbb{E}[b_1 \mid q' = b_0 = \bar{p}] = \bar{p}$. If $q' > \bar{p}$ then, by MLRP and the Bayes' rule (2): $\mathbb{E}[b_1 \mid q'; b_0 = \bar{p}] > \mathbb{E}[b_1 \mid q = b_0 = \bar{p}] = \bar{p}$. Together with the fact that $V(p, b)$ is weakly increasing in b (see Lemma 4), this implies that

$$V(p, \bar{p}) \geqslant \beta V(q', \bar{p}).$$

Therefore, if $V(q', \bar{p}) > 0$ then $V(p, \bar{p}) > 0$. The former is the implication in the induction statement, and the latter is the statement we seek to prove. The proof is thus complete. □

Proof of Proposition 2. Statement of the proposition follows directly from part 2 of Lemma 13. □

Proof of Theorem 2. The statement follows from part 2 of Lemma 13, which implies that $\bar{b}(p_t) < \bar{p}$. Therefore, we can take $l(p_t) = \bar{b}(p_t)$. By belief consistency and Lemma 10, there should exist $r(p_t) > \bar{p}$ and message $m \in \mathcal{M}(p_t)$ such that m is sent for all $b_t \in [l(p_t), r(p_t)]$. □