

The Limits of Social Learning*

Aleksei Smirnov[†], Egor Starkov[‡]

March 17, 2020

PRELIMINARY DRAFT

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 in this setting, regardless of whether the reviewers can commit to a communication strategy or have to resort to cheap talk. We further show that in the latter case, the communication must necessarily have the interval structure, meaning that the noise persists even when the conflict between the reviewers and future consumers vanishes.

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

JEL Codes: C73, D83, L15.

*We are grateful to Renato Gomes, Marek Pycia, Armin Schmutzler, Jakub Steiner, Bauke Visser, Jidong Zhou and seminar participants at University of Zürich for valuable feedback and helpful comments.

[†]Department of Economics, University of Zürich, Blümlisalpstrasse 10, CH-8006, Zürich, Switzerland; e-mail: aleksei.smirnov@econ.uzh.ch.

[‡]Department of Economics, University of Copenhagen, Øster Farimagsgade 5, bygning 26, 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], eMarketer [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 even in the absence of any intervention from sellers, reviews can get noisy organically.

To understand the source of this noise, which stems from *how* customers write reviews, one must first ask *why* customers write reviews. Surveys consistently produce a few modal answers to this question, with one of the most popular ones being “to help other consumers” (Trustpilot [2018]). 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*. In particular, when choosing which product to buy, the consumers appear to focus primarily on their own expected utility from consumption, rather than on their desire to provide helpful information to others.

This inconsistency in altruism, as we show, must lead to noise in product reviews. When product quality is uncertain, purchases have an informational externality, since in addition to direct consumption utility they allow informative reviews to be written, which allow future consumers to make more efficient decisions. However, when deciding on the purchase, a self-interested consumer does not internalize this information-generating effect, and so their private expected value from buying a product is always lower than social value. This discrepancy is, in turn, recognized by an altruistic reviewer, who may in some circumstances want to mislead a future consumer into buying a product when it is not individually optimal to do so.

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 expected consumption utility warrants this. The realized utility is informative about the product quality. The consumer can leave a review describing her consumption experience, and when doing so she wishes to maximize welfare of consumers that arrive at the market after her.

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 (which should be interpreted

¹See, e.g., Luca and Zervas [2016] for an exploration of the effects of fake reviews and Smirnov and Starkov [2018] for a model of censorship in product reviews.

as a shared social norm among consumers) or not. In the latter case, every consumer leaves a review they believe to be socially optimal given their experience. This scenario produces even more interesting results.

In particular, we show that if a consumer cannot commit to a communication strategy, then despite the conflict arising only in a special set of circumstances – when the product is believed to be good enough to experiment with socially, but not good enough to buy just for the sake of doing so, – we show that the effects of this conflict propagate and distort communication in other cases as well. More specifically, communication must then take the interval structure known in the cheap talk literature, when senders with similar private beliefs pool on 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 significantly 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 are 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 just fake. Contrary to those explanations, in our case inflation arises endogenously as a result of interaction between consumers, with no intervention from the firm whatsoever.

This paper contributes to the social learning literature. A lot of the existing literature has focused on non-strategic learning in local settings, such as networks. This is driven by recognition that we as consumers receive a lot of our information second-hand, so it may be distorted by other agents' perceptions and beliefs. In turn, the part of the literature that deals with *strategic* learning has mostly focused on strategic information *acquisition*, forcing the agent with a limited learning capacity to choose their information sources carefully. Our paper focuses instead on social learning with strategic information *provision*. For a detailed literature review, see below.

The paper is organized as follows. Section 2 contains a review of the relevant literature. In Section 3 we formulate the general version of the model. Section 4 presents the main result – that no perfect communication is possible in our setting – and explains the intuition behind it. Sections 5 and 6 then discuss what communication structures *can* arise in our model, with Section 5 exploring an illustrative three-period example, and Section 6 generalizing the insights to an infinite-horizon problem. Section 7 concludes. All proofs are relegated to the Appendix.

2 Literature Review

The current paper mainly contributes to two strands of literature: *social learning* and *dynamic cheap talk*.

The literature on social learning is vast. Our paper is closest to the literature on herding and cascades in sequential learning (Banerjee [1992], Bikhchandani, Hirshleifer, and Welch [1992], Smith and Sørensen [2000]). In these models the agents choose actions which are payoff-relevant for agents themselves and, at the same time, signal their private information to subsequent agents. Smith and Sørensen [2011] provide an excellent overview of the topic. The most recent general treatment of the setting is provided by Xu [2018]. Most relevant are works by Ali and Kartik [2012] and Smith,

Sørensen, and Tian [2017], who consider sequential observational learning with others-regarding preferences. However, in the observational learning framework agents receive private information and act on it; actions are the only source of information for future agents who observe neither past signals, nor past outcomes, while our paper explores learning under strategic communication. Wolitzky [2018], in contrast, considers sequential learning when players observe outcomes of previous players, but not their actions. This is closer to our paper, except outcomes in our model are observed through noisy communication rather than directly. Ali [2018] studies observational learning with costly information acquisition.

Social learning with strategic information provision was explored by Swank and Visser [2015] when the conflict arises from senders’ career concerns. Liang and Mu [2019] consider a model where agents can, similarly to our paper, be tempted by exploiting myopic benefits which prevents future generations from learning the state correctly. Au [2019] presents a model of (non-social) learning, in which experts’ recommendations to the agent are distorted even despite the seeming absence of conflict between the parties.

A separate strand of the social learning literature focuses exclusively on learning from customer reviews (e.g., Acemoglu, Makhdoumi, Malekian, and Ozdaglar [2017] and Vaccari, Maglaras, and Scarsini [2018]).

The decentralized literature of decision-making and communication in our model relates us to 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. Migrow [2018] studies how a manager should design a communication network in an organization to optimally elicit the information from employees. The conflicts explored in these papers are different from what we focus on in this paper.

Literature on the *design* of social learning considers the information structures that incentivize short-lived agent to experiment for the sake of society. Notable references include Kremer, Mansour, and Perry [2014], Che and Hörner [2018], Mansour, Slivkins, and Syrgkanis [2019], and Cohen and Mansour [2019].² We explore what effectively is a decentralized version of these models, with each agent trying to communicate in a way so as to create optimal experimentation incentives, but lacking the commitment power and memory of a single principal.³

Our paper models communication via cheap talk à la Crawford and Sobel [1982]. Other models (apart from ours) of sequential communication include Ambrus, Azevedo, and Kamada [2013], Renault, Solan, and Vieille [2013], and Chiba [2018]. Le Quement and Patel [2018] explore cheap talk with preferences for reciprocity.

Our model presents consumers as altruistic when they are writing reviews. It has been argued for a long time that the economic model of homo economicus as a self-interested agent does not fully capture real-world behavior, which often exhibits regard for others. Various classical explanations

²Optimal experimentation by a group of long-lived agents with incentives for free-riding was studied by Bolton and Harris [1999], Keller, Rady, and Cripps [2005], and Hörner, Klein, and Rady [2015]. Heidhues, Rady, and Strack [2015] move from observable to private payoffs and explore communication in this setting.

³We also consider the case when consumers can commit to a specific message structure.

are Andreoni [1990] (impure altruism), Fehr and Schmidt [1999] (inequality aversion) and Becker [1974] (pure altruism). The literature is surveyed in 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].

Within the context of social learning, experiments by March and Zieglmeyer [2016] and Peng, Rao, Sun, and Xiao [2017] find evidence of altruistic motives when testing standard models of observational learning.

3 The Model

3.1 Primitives

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

Seller. There is a single long-lived seller, who offers for sale a single product that he has in infinite supply. Product quality θ , which represents the average consumption utility of the product, can be either *low* or *high*: $\theta \in \{L, H\}$, with $0 \leq L < H$. 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 risk-neutral consumer arrives at the market. The consumer can either purchase the good at cost c or leave the market forever, receiving the reservation utility normalized to 0. In case of purchase, the consumer receives random consumption utility s , distributed according to quality-contingent cdf F^θ with mean θ and respective pdf f^θ . We assume that both F^L and F^H have full support on the same open interval $S = (\underline{s}, \bar{s}) \subseteq \mathbb{R}$.⁴ Both measures are absolutely continuous on S , and their respective densities are continuously differentiable and bounded from above. In addition, we assume that MLRP holds:

Assumption (MLRP). *Ratio $\frac{f^H(s)}{f^L(s)}$ is a strictly increasing and continuous function of s on S . Moreover, $\lim_{s \rightarrow \underline{s}} \frac{f^H(s)}{f^L(s)} = 0$, and $\lim_{s \rightarrow \bar{s}} \frac{f^H(s)}{f^L(s)} = +\infty$.⁵*

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

$$\theta(p) := Hp + L(1 - p) \geq c \Leftrightarrow p \geq \bar{p},$$

where $\bar{p} := \frac{c-L}{H-L}$.⁶ This purchasing strategy will be taken as given in what follows.

Reviews. If the good was purchased, the consumer then sends a cheap talk message $m \in \mathcal{M}$ (writes a review) to subsequent consumers, describing her experience with the product. The message set \mathcal{M} is assumed to be arbitrarily rich, with $[0, 1] \subseteq \mathcal{M}$. When leaving a review, the consumer maximizes the expected discounted sum of consumption utilities of all future consumers.

⁴Here $\underline{s} = -\infty$ and $\bar{s} = +\infty$ are both admissible values.

⁵Also note that MLRP implies that F^H first order stochastically dominates F^L .

⁶We assume that the consumer purchases the product when indifferent.

We consider two regimes. Under the *commitment regime* a consumer can commit to some utility-contingent reporting strategy $(\underline{s}, \bar{s}) \rightarrow \mathcal{M}$ before a purchase. The interpretation of this regime is that there exists a welfare-maximizing social norm, which prescribes the mapping from experiences to reviews. Under the *no commitment regime*, the consumer chooses m after observing her consumption utility s . The latter regime is also referred to as the decentralized scenario.

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 $(m_1, m_2, \dots, m_{t-1})$ and forms belief p_t about the quality of the product.
2. The consumer decides whether to purchase the product at cost c or not.
3. After a purchase she receives random consumption utility $s_t \sim F^\theta$ 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: $m_t = \emptyset$.⁷

3.2 Histories and State Variables

Review history $R_t := (m_1, m_2, \dots, m_{t-1})$ 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 $p_t := \mathbb{P}(\theta = H \mid R_t)$. The prior $p_1 = \mathbb{P}(\theta = H \mid \emptyset)$ is exogenously fixed and commonly agreed upon.

The *private posterior belief* of time- t consumer in case she purchased and consumed the product is denoted by $b_t := \mathbb{P}(\theta = H \mid R_t, s_t)$. Given p_t and s_t we can compute b_t as

$$b_t = \frac{p_t f^H(s_t)}{p_t f^H(s_t) + (1 - p_t) f^L(s_t)}. \quad (1)$$

Let $\mu^\theta(b_t \mid p_t)$ denote the cdf of a distribution of b_t induced by s_t conditional on p_t and true state θ .⁸

The 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. In what follows, we will focus on a setting, in which we treat belief p_t and current time t as a sufficient statistic of the review history R_t .⁹ Because of this, we call the tuple (p_t, b_t, t) the *private state* of time- t consumer, and we refer to (p_t, t) as the time- t *public state*. We will typically omit t from the description of states, given that it can be inferred from belief indexing.

⁷For simplicity, we assume that $\emptyset \notin \mathcal{M}$, i.e., a purchasing consumer cannot stay silent and must leave a meaningful review.

⁸It can be computed explicitly: $\mu^\theta(b_t \mid p_t) = F^\theta \left(l^{-1} \left(\ln \left(\frac{b_t}{1-b_t} \right) - \ln \left(\frac{p_t}{1-p_t} \right) \right) \right)$, where l^{-1} is an inverse function to $\ln \left[\frac{f^H(s)}{f^L(s)} \right]$.

⁹This is not without loss: if two time- t review histories produce the same p_t , they will be treated as equivalent. This would preclude the possibility of having different continuation equilibria after the two histories.

Given that the consumers' purchasing decisions are myopic and described by "buy iff $p_t \geq \bar{p}$ ", from this point onward we will be focusing on consumers' *communication* strategies. The time- t consumer's behavioral strategy is r , where $r(m|p_t, b_t)$ is the probability with which the time- t consumer sends message $m \in \mathcal{M}$ in private state (p_t, b_t) . Let $\mathcal{M}(p_t) = \{m \in \mathcal{M} \mid \exists b_t : r(m|p_t, b_t) > 0\}$. Then the public belief p_{t+1} induced by message $m \in \mathcal{M}(p_t)$ is given by

$$p_{t+1} = q(p_t, m) := \frac{p_t \cdot \int_0^1 r(m|p_t, b_t) d\mu^H(b_t|p_t)}{p_t \cdot \int_0^1 r(m|p_t, b_t) d\mu^H(b_t|p_t) + (1 - p_t) \cdot \int_0^1 r(m|p_t, b_t) d\mu^L(b_t|p_t)}. \quad (2)$$

We let $\mathcal{P}(p_t) = \{q(p_t, m) \mid m \in \mathcal{M}(p_t)\}$ denote the set of all posteriors which are induced by time- t consumer in equilibrium. We partition this set into $\mathcal{S}(p_t) \cup \mathcal{E}(p_t) = \mathcal{P}(p_t)$. Here $\mathcal{E}(p_t) = \{q \in \mathcal{P}(p_t) \mid q \geq \bar{p}\}$ includes all posteriors for which the next consumer purchases the product, while $\mathcal{S}(p_t) = \{q \in \mathcal{P}(p_t) \mid q < \bar{p}\}$ contains all posteriors which deter the next consumer from the purchase. Note that if $p_t \geq \bar{p}$ then $\mathcal{E}(p_t) \neq \emptyset$, as the public belief p_t is a martingale. Further, note that if $p_t < \bar{p}$ then the market shuts down: 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 (i.e., $p_{t+1} = p_t$) and does not purchase the product either. Therefore, all $q \in \mathcal{S}(p_t)$ are equivalent in the sense of shutting the market down. Hereinafter we will without loss only consider a representative element of $\mathcal{S}(p_t)$ whenever it is nonempty.

3.3 Maximization Problem

When a consumer sends message m at private state (p_t, b_t) and induces public belief $p_{t+1} = q$, her value (the discounted sum of future consumers' utilities) from doing so is equal to

$$V(q \mid p_t, b_t) := \mathbb{E} \left[\sum_{j=t+1}^{+\infty} \beta^{j-t-1} \cdot \mathbb{I}(p_j \geq \bar{p}) \cdot s_j \mid p_{t+1} = q \right]. \quad (3)$$

The expectation is taken over all future histories that start with public belief $p_{t+1} = q$. Implicit in (3) is the correlation between future s_j and future p_j stemming from future consumers' equilibrium strategies. Maximizing (3) over all available messages, we get the consumer's optimal value in private state (p_t, b_t) :

$$V(p_t, b_t) = \max_{p \in \mathcal{P}(p_t)} V(p \mid p_t, b_t).^{10}$$

For a given equilibrium, the time- t consumer's *ex ante* expected continuation value conditional on public belief p_t is given by

$$V(p_t) := \mathbb{E}[V(p_t, b_t)] = p_t \cdot \int_0^1 V(p_t, b_t) d\mu^H(b_t|p_t) + (1 - p_t) \cdot \int_0^1 V(p_t, b_t) d\mu^L(b_t|p_t).$$

¹⁰This representation implies that the consumer chooses a message from $\mathcal{M}(p_t)$ rather than \mathcal{M} . This is a simplifying assumption: we do not allow to send out-of-equilibrium messages so that we do not have to keep track of beliefs after such messages. This restriction is without loss.

When talking about values, we will use superscripts C or D to distinguish commitment and no-commitment (decentralized) solutions.

3.4 Equilibrium Definition

We are looking for Perfect Bayesian Equilibria of the game, which consist of a strategy profile $r(m|p_t, b_t)$ and updating rules for beliefs p_t and b_t such that

Belief Consistency: (1) holds at all private histories (p_t, b_t) , and (2) holds after all p_t and $m \in \mathcal{M}(p_t)$;

C-Optimality: For commitment regime: $r(m | p_t, b_t)$ is chosen so as to maximize $V^C(p_t)$ for all p_t ;

D-Optimality: For decentralized regime: if $m \in \mathcal{M}(p_t)$ and $r(m | p_t, b_t) > 0$ then $V^D(p_t, b_t) = V^D(q(p_t, m) | p_t, b_t)$.

Belief consistency condition ensures that consumers use Bayes' rule whenever possible to update their belief. C-Optimality states that in the commitment scenario, the consumer chooses a mapping from private belief b_t to messages (conditional on p_t and subject to Belief Consistency) so as to maximize the ex ante value. In the decentralized game, D-Optimality requires that the consumer maximizes her ex post value (after learning s_t).

4 No Perfect Communication

This section demonstrates the main idea of this paper: that truth-telling is neither an optimal social “norm” for writing reviews (i.e., it’s not a commitment solution), nor it is an equilibrium in the decentralized market. This conclusion is driven by the implicitly lexicographic nature of consumers’ preferences. When buying the product, a consumer maximizes own expected utility, but when writing a review, she cares about all future generations. The consumer can thus be represented as having lexicographic preferences: the first-order preference is for own well-being, while the warm glow from social welfare is second-order. The consumer is thus unwilling to sacrifice her consumption utility for sake of the society. This creates a conflict, since time- t consumer would like the consumer at $t + 1$ to conduct socially efficient experimentation, possibly by buying the product even when it is not myopically optimal, so that more information about product quality is generated. This conflict introduces noise into communication between the two generations of consumers, i.e., into the review of time- t consumer.

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). We use it in the same context.

Definition. Message $m \in \mathcal{E}(p(R_t))$ at public history R_t starts a cascade if $p(R_s) \geq \bar{p}$ for all $R_s = (R_t, m, \dots)$.

In other words, we say that a recommendation to purchase issued at R_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 consumers' behavior. There are two things to note in relation to cascades. First, any message $m \in \mathcal{S}(p_t)$ starts a cascade as well, in the sense that no future consumers buy the product again, as discussed in Section 3.2.¹¹ Second, in the no-commitment scenario there always exists a continuation equilibrium in which any given $m \in \mathcal{E}(p(R_t))$ starts a cascade. One example is the babbling equilibrium, one in which all future reviews are uninformative and are perceived as such, and thus the public belief remains frozen at $p(R_t, m)$.¹² However, in general, a cascade need not shut down the information transmission completely: reviews may be informative and affect the public belief p_t as long as they do not affect future consumers' actual purchasing decisions.

The following proposition presents the main result of this section, which motivates further discussion. This result shows that truth-telling (fully revealing communication) is neither a welfare-maximizing social norm, nor it is an equilibrium in a decentralized game.

Proposition 1. *In the commitment regime, $\mathcal{P}(p_t) = [0, 1]$ for all p_t does not deliver a maximum for $V^C(p_t)$.*

In the no-commitment regime, $\mathcal{P}(p_t) = [0, 1]$ if and only if any message available at p_t starts a cascade.

Proposition 1 shows that the conflict between the sender and the receiver of a review precludes perfect communication. If the sender'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 sender thus wants to misrepresent her posterior as if it was barely above \bar{p} . As we show in the following sections, in the absence of commitment this noisiness of communication unravels – even though the sender-receiver conflict only exists for some $b_t < \bar{p}$ the noise propagates to all $b_t > \bar{p}$.

The statement for the no-commitment regime also illustrates that the noise arises exactly from the sender'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.

In case of commitment, the reasoning behind the result is more involved. On the one hand, it is still desirable for the sender to pool the states just below \bar{p} with those just above it to induce more experimentation at $t + 1$. On the other hand, however, distorting communication of states $b_t > \bar{p}$ is costly, since this is not only concealing some information about the state from the future consumers, but it may also decrease the amount of experimentation from $t + 2$ onwards. The latter statement holds because pooling depresses beliefs p_{t+1} induced after $b_t > \bar{p}$, as compared to

¹¹The definition above relates to positive cascades, while $m \in \mathcal{S}(p_t)$ starts a negative cascade.

¹²Babbling is prominent in cheap talk literature; to see that it is an equilibrium note that neither player has 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.

truth-telling. However, the gains from pooling states $b_t \in (\bar{p} - \varepsilon, \bar{p} + \varepsilon)$ are approximately equal to $V^C(\bar{p}) \cdot \varepsilon > 0$, while the losses are of order $\mathcal{O}(\varepsilon^2)$ because $V^C(p_t)$ is continuous in p_t . Truth-telling is thus not optimal under commitment either. A more detailed exposition of this logic is presented in Subsection 5.4.

In this section we explored how equilibria cannot look. The following sections provide more insight into how they do look, with and without commitment. We begin with exploring the most basic version of the model.

5 Three-Period Example

This section demonstrates the main insights in the most basic three-period setting. For sake of this example, assume that consumption utilities s_t are normally distributed with mean θ and variance σ^2 . Suppose further that there is no discounting. We shall denote the three consumers as C1, C2, and C3 respectively. We solve the example by backward induction.

C3 purchases the product if and only if $\theta(p_3) = Hp_3 + L(1 - p_3) \geq \bar{p}$, and her messaging strategy is irrelevant, since no consumers arrive at the market after her.

5.1 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. All payoffs starting from $t = 2$ are zero in this case. Conversely, if $p_2 \geq \bar{p}$ then C2's continuation value equals C3's expected consumption utility: $V(p_3|p_2, b_2) = \theta(p_3) - c$. Therefore, there is no conflict between C2 and C3, and truthful communication, where C2 reports $m_2 = p_2$, is possible in equilibrium (in both regimes – with and without commitment).

Note, however, that the only 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, all continuation equilibria from p_2 with $\mathcal{S}(p_2) \neq \emptyset$ are payoff-equivalent to the one where only two messages are used: $\mathcal{M} = \{\text{“buy”}, \text{“do not buy”}\}$. In this case “buy” is sent by C2 whenever $b_2 \geq \bar{p}$, and “do not buy” is sent when $b_2 < \bar{p}$. Then

$$V(p_2, b_2) = \max\{\theta(b_2) - c, 0\}.$$

If $\mathcal{S}(p_2) = \emptyset$ then any message at p_2 starts a cascade, so $V(p_2, b_2) = \theta(b_2) - c$ in this case.

5.2 First Period

In this section we analyze $V(p_2|p_1, b_1)$, C1's continuation value from inducing prior belief p_2 for C2, when her own private belief is b_1 . We again look at states $p_1 \geq \bar{p}$ (otherwise all values are zero). For simplicity we assume that all time-2 continuation equilibria are informative, i.e., $\mathcal{S}(p_2) \neq \emptyset$ for all $p_2 \geq \bar{p}$.¹³ Note that since truth-telling is both a first-best and an equilibrium outcome at $t = 2$ – and any informative continuation is equivalent to truth-telling, as shown above – the value

¹³The case when a cascade is started at $t = 2$ (after any message at $t = 1$) is trivial, since then truth-telling is an equilibrium by the same argument as in the second period. The case when a cascade is started by some but not all messages $m \in \mathcal{E}(p_1)$ is non-trivial, but we do not deem it worthy of careful consideration.

$V(p_2 | p_1, b_1)$ for given p_2 and b_1 is the same in both commitment and no-commitment regimes. The difference between the two only lies in what communication strategies can be optimal for C1 given $V(p_2 | p_1, b_1)$. These strategies are explored in sections 5.3 and 5.4.

With $p_1 \geq \bar{p}$, C1 buys the good and receives utility s_1 . If she sends a message $m \in \mathcal{S}(p_1)$ that induces no further purchases, her continuation value equals zero. When she sends $m \in \mathcal{E}(p_1)$, C2 purchases the product and obtains utility s_2 . Following that, C3 purchases the product if and only if $p_3 = b_2(s_2) \geq \bar{p}$. Signal \bar{s}_2 which induces $b_2 = \bar{p}$ can be found from

$$\frac{b_2}{1 - b_2} \equiv \frac{p_2}{1 - p_2} \cdot \frac{f^H(\bar{s}_2)}{f^L(\bar{s}_2)} = \frac{\bar{p}}{1 - \bar{p}}.$$

Given that $\frac{f^H(\bar{s}_2)}{f^L(\bar{s}_2)} = e^{\frac{H-L}{\sigma^2}(\bar{s}_2 - \frac{H+L}{2})}$, we have

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

Therefore, if C2 buys the product, C1's continuation value from inducing some belief p_2 is given by

$$V^*(p_2 | p_1, b_1) = \mathbb{E} [s_2 + s_3 \cdot \mathbb{I}\{s_2 \geq \bar{s}_2(p_2)\} | b_1]. \quad (5)$$

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

$$V(p_2 | p_1, b_1) = \begin{cases} V^*(p_2 | p_1, b_1) & \text{if } p_2 \geq \bar{p}, \\ 0 & \text{if } p_2 < \bar{p} \end{cases}$$

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 $1 - F^H(\bar{s}_2)$ 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 $1 - F^L(\bar{s}_2)$. In the end, expression (5) can be rewritten as

$$V^*(p_2 | p_1, b_1) = \theta(b_1) - c + b_1 \cdot (1 - F^H(\bar{s}_2(p_2))) (H - c) + (1 - b_1) \cdot (1 - F^L(\bar{s}_2(p_2))) (L - c), \quad (6)$$

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

Analyzing (6), we can identify several important properties of $V^*(p_2 | p_1, b_1)$. Firstly, it is strictly positive at $b_1 = \bar{p}$ for all p_2 . This follows from the fact that $F^H(\bar{s}_2(p_2)) < F^L(\bar{s}_2(p_2))$. The function 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 $p_2 \geq \bar{p}$ for at least some values of $b_1 < \bar{p}$: she wants C2 to purchase the product despite believing that this is not myopically optimal. This is due to the social value of experimentation (i.e., of information generated by the purchase at $t = 2$), which is internalized by C1 in her review strategy, but not by C2 in her purchasing strategy. There is thus a conflict 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 $b_1 \geq \bar{p}$, C1 would




Figure 1: $V^*(p_2 | p_1, b_1)$ as a function of p_2 .

Note: the parameter values are $H = 3, L = 0, c = 1, \sigma = 4, b_1 = \frac{1}{5}, \frac{1}{3}, \frac{1}{2}$.

prefer to tell the truth to C2 and induce the correct belief $p_2 = b_1$. To see this, observe that

$$\frac{\partial V(p_2 | p_1, b_1)}{\partial p_2} = (1 - b_1) \bar{p} f^L(\bar{s}_2(p_2)) \cdot \frac{\sigma^2}{(1 - \bar{p})p_2(1 - p_2)} \left(\frac{b_1}{1 - b_1} \cdot \frac{1 - \bar{p}}{\bar{p}} \cdot \frac{f^H(\bar{s}_2(p_2))}{f^L(\bar{s}_2(p_2))} - 1 \right).$$

Since $\frac{f^H(\bar{s}_2(p_2))}{f^L(\bar{s}_2(p_2))}$ is strictly decreasing in p_2 from $+\infty$ to 0, and the fraction multiplying the bracket is positive, we get that $V(p_2 | p_1, b_1)$ is single-peaked. We can find the peak by equating $\frac{\partial V(p_2 | p_1, b_1)}{\partial p_2}$ to zero, which yields

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

which together with (4) gives condition $p_2 = b_1$.

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

The two properties of $V^*(p_2 | p_1, b_1)$ 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.

5.3 Interval Structure of Equilibrium Communication

We now show that in the no commitment regime, communication in the first period must have interval structure. In other words, there exists a partition $0 = \Delta_0 \leq \Delta_1 < \Delta_2 < \dots = 1$ and messages m_1, m_2, \dots such that if $b_1 \in (\Delta_{j-1}, \Delta_j)$ then $r_{m_j}(p_1, b_1) = 1$.

First note that if $\mathcal{S}(p_1)$ is nonempty – i.e., if there is a review that will prevent C2 from buying the product, – this review will be used by C1 after at least some b_1 low enough.¹⁴ Assume that this is the case in what follows.

Consider now the smallest posterior belief among those available in equilibrium that leads C2

¹⁴For $b_1 \approx 0$, expression (6) reduces to $V^*(p_2 | p_1, b_1) = (L - c) \cdot (1 + 1 - F^L(\bar{s}_2(p_2)))$, which is negative because $L < c$.

to purchase the product, $e_1 = \min \mathcal{E}(p_1)$. Let Δ_1 denote the level of posterior belief b_1 with which C1 is indifferent between leaving a review in $\mathcal{S}(p_1)$ and review e_1 . From the fact that $V^*(p_2 | p_1, b_1)$ is positive at $b_1 = \bar{p}$ it is immediate that $\Delta_1 < e_1$. The fact that $V^*(p_2 | p_1, b_1)$ has a peak at $p_2 = b_1$ implies, in turn, that all types b_1 of C1 also prefer to leave review e_1 rather than any review $p_2 > e_1$.¹⁵

In other words, if there exists a way to make the “most cautious recommendation to buy”, then C1 would like to adopt that phrasing for a wide range of posteriors b_1 . 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. These two goals conflict with each other, since C1 only has one stone – her review – to hit both birds.

Recall, however, that C2 is rational and Bayesian – in particular, when forming her belief p_2 she takes C1’s incentives into account. Therefore, it must be the case that the prior belief $p_2 = e_1$ of C2 must incorporate the information contained in the posteriors b_1 of the versions of C1 who write reviews that induce e_1 . We have argued above that there are types $b_1 < e_1$ that induce $p_2 = e_1$, so there must also be types $b_1 > e_1$ that do the same. 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 V^*) be indifferent between inducing e_1 and the next-lowest available posterior e_2 . However, we know that $V^*(p_2 | p_1, b_1)$ is single peaked in p_2 with a peak at $p_2 = b_1$, hence the indifference condition $V^*(e_1 | p_1, \Delta_2) = V^*(e_2 | p_1, \Delta_2)$ implies that $e_2 > \Delta_2$. By iterating the argument, we get that

$$\dots < \Delta_j < e_j < \Delta_{j+1} < e_{j+1} < \dots$$

Plainly speaking, the fact that the aforementioned “most cautious recommendation to buy” is noisy and not perfectly revealing of b_1 implies eventually 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.

Figure 2 plots the continuation payoff of C1 in an interval equilibrium with three messages: $e_1 = \bar{p}$, e_2 , and “stop experimentation”. This payoff coincides with the unconstrained maximum (when C1 can choose any p_2 and force C2 to purchase the item) whenever $b_1 = e_1$, 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 this is compensated by the more efficient purchasing decision of C2 as compared to the case when C1 can choose any p_2 but cannot force C2 to buy.

5.4 Commitment Solution

By committing to truthful communication, C1 can achieve value

$$V(b_1 | p_1, b_1) = \begin{cases} V^*(b_1 | p_1, b_1) & \text{if } b_1 \geq \bar{p}, \\ 0 & \text{if } b_1 < \bar{p} \end{cases}$$

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). The source of improvement lies in the discontinuity of $V(b_1 | p_1, b_1)$ at $b_1 = \bar{p}$.

¹⁵This also implies that $\mathcal{E}(p_1)$ must be closed at the bottom in any equilibrium, i.e., e_1 does actually exist.




Figure 2: $V^*(b_1 | p_1, b_1)$ and value $V^D(p_1, b_1)$ in an interval 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 (one can think that 9 consumers arrive in the third period).

To evaluate the trade-offs introduced by imperfect communication of posteriors b_1 , it is useful to understand how $V(p_2 | p_1, b_1)$ depends on the induced prior p_2 for a given posterior b_1 . To do this, we ask the dual question and visualize the dependence of $V(p_2 | p_1, b_1)$ on b_1 for a given p_2 . In particular, (6) is a linear function of b_1 , meaning that once we fix p_1 and p_2 , value $V(p_2 | p_1, b_1)$ as a function of b_1 is given by a tangent to the maximal value $V^*(b_1 | p_1, b_1)$ at $b_1 = p_2$. This means that $V^*(b_1 | p_1, b_1)$, as well as $V(p_1, b_1)$ in any equilibrium, are convex in b_1 , since all of these are upper envelopes of respective sets of linear functions.

This convexity implies that C1 cannot benefit from sending messages that pool different posteriors b_1 above \bar{p} . On the other hand, it is also strictly optimal to stop experimentation for all private beliefs below \bar{p} , where \bar{p} is determined from condition $V(\bar{p} | p_1, \bar{p}) = 0$. Therefore, benefits can only arise from pooling posteriors b_1 in the neighborhood of \bar{p} , and perfect communication is optimal for posteriors b_1 above the pooling region. The benefits of pooling come from inducing experimentation after $b_1 < \bar{p}$, which means that the posterior induced by pooling must exactly equal \bar{p} . Indeed, otherwise C1 can lower the cutoff above which perfect communication occurs, forcing C2 to buy the item for all the same $b_1 < \bar{p}$, and conveying better information for some $b_1 > \bar{p}$, which is an improvement.

Finally, it is always optimal to pool at least some private beliefs to the left and to the right of \bar{p} . Indeed, suppose that a consumer sends the same message m for all private beliefs ε -below and $C \cdot \varepsilon$ -above \bar{p} such that resulting $q(p_t, m) = \bar{p}$.¹⁶ This is equivalent to substituting value from truth-telling in this interval with a line tangent to $V(b_1 | p_1, b_1)$ at $b_1 = \bar{p}$. The gains from it are approximately equal to $V(\bar{p} | p_1, \bar{p}) \cdot \varepsilon > 0$, because $V(\bar{p} | p_1, \bar{p}) > 0$. Losses associated with pooling above the cutoff are less than $\frac{C(H-L)}{2} \varepsilon^2 = \mathcal{O}(\varepsilon^2)$. Therefore, there always exists $\varepsilon > 0$ such that it is optimal to pool at least some small neighborhood of private beliefs b_t around \bar{p} . Figure 3 plots the value attained by C1 under this communication mechanism. Although this argument shows that it is always beneficial to pool private beliefs below and above \bar{p} , it does not however show *how*

¹⁶As distributions $\mu^\theta(b_t | p_t)$ are arbitrary the weights between two sides from \bar{p} do not have to be equal.




Figure 3: $V^*(b_1 | p_1, b_1)$ and the commitment value $V^C(p_1, b_1)$.

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. The left panel plots functions for all $b_1 \in [0, 1]$; the right panel focuses on the neighborhood of \bar{p} .

it is optimal to pool beliefs. As we formally show in the next chapter this region is indeed always a convex interval around \bar{p} .

6 General Case

This section describes the equilibrium outcomes and the commitment solution arising in the infinite-horizon version of the game. We show that neither of the two feature perfect communication, with the exception that it may be an equilibrium outcome in some period as long as future communication is uninformative.

6.1 Commitment Solution

We begin by looking at the commitment solution this time around. The main result and the argument behind it mirror those that we have discovered in the three-period example: the reviewer's desire to inflate the review of a marginally-bad item for sake of social experimentation results in him pooling moderately bad experiences with some good experiences so as to push the posterior of the next consumer up to \bar{p} . After all other experiences the reviewer reports truthfully. This idea is formalized in Theorem 1 below.

Theorem 1. *At every public state p_t the optimal commitment solution is characterized by cutoffs $l^C(p_t) < \bar{p} < r^C(p_t)$ such that:*

1. *For all $b_t \in [0, l^C(p_t)]$ the consumer sends message $m \in \mathcal{S}(p_t)$, i.e., the experimentation stops.*
2. *For any $b_t \in [r^C(p_t), 1]$ the consumer truthfully transmits her private belief b_t (or what is the same his private payoff) to the public, that is $p_{t+1} = b_t$.*
3. *For all $b_t \in (l^C(p_t), r^C(p_t))$ the consumer sends message m such that $q(p_t, m) = \bar{p}$.*

The proof of the theorem in the appendix proceeds in three main steps. First, we show that $V(b_t | p_t, b_t)$ is continuously differentiable and [weakly] convex in b_t above \bar{p} . This implies that

pooling is only beneficial around the cutoff. Since $V(p_t, \bar{p}) > 0$, there are potential gains from pooling posteriors $b_t < \bar{p}$ with posteriors above the cutoff. This, however, would at the same time decrease the quality of information transmitted after those $b_t > \bar{p}$ that are pooled with posteriors below the cutoff, which is socially costly. The second major step of the proof is thus in showing that the gains from pooling over an arbitrarily small interval of posteriors will be of first order, while the losses will be of second order. Finally, we show that the optimal commitment strategy exists within the class of strategies restricted to such combinations of pooling around the cutoff and truth-telling otherwise. This existence is proved with the help of Arzela-Ascoli Theorem. Furthermore, we show that the optimal strategy is Markovian, i.e., depends only on public belief p_t but does not explicitly depend on t .

It is straightforward that the commitment solution induces underexperimentation relative to the first best (in which the consumer can transmit the information perfectly while also having perfect control over the future consumers' actions). This is because 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 $b_t \geq \bar{p}$. In particular, this distortion is downwards, meaning it makes all future consumers more pessimistic and so exacerbates the problem of underexperimentation after those histories. In other words, the reviewer in period t has to trade off underexperimentation at $t + 1$ against underexperimentation from $t + 2$ onwards.

It is also worth pointing out the differences between our model of committed sender and the model of Che and Hörner [2018]. One obvious difference lies in the signal structure, where we allow for a wide class of private signals compared to binary (good news Poisson) signals in their case. This allows us to give a richer characterization of within-period outcomes at the cost of the tractability of the dynamics in the model. However, a more substantial difference between the two models lies in the technologies: in the model of Che and Hörner [2018] the principal is long-lived, and has memory of old information even if it was not publicly disclosed at the time, so this information may be disclosed at a later time. In our model, in contrast, all consumers are short-lived, hence the public record of product reviews is the only past information available today. This constraint to a public memory limits the principal designing a reviewing mechanism to “now-or-never” revelation schemes, eliminating the opportunity to delay.

6.2 Decentralized Equilibrium

We now move on to exploring the equilibria of the decentralized game. The analysis is complicated by the fact that an equilibrium at t is determined by the continuation equilibria after various induced priors p_{t+1} , which in turn depend on continuation equilibria after p_{t+2} and so on. Backwards induction is not available in an infinite-horizon game, and even restricting ourselves to Markov setting, where strategies only depend on the public prior p but not calendar time t , does not render the problem tractable enough to provide a full characterization of the set of equilibria.

We are, however, able to provide a partial characterization of equilibria. In particular, Theorem 2 below provides two statements pertaining to such characterization. First, it claims that commitment is always valuable in the sense that no equilibrium of the decentralized game can generate a higher level of social welfare. Second, it shows that experiences b_t in some neighborhood of the myopic cutoff \bar{p} are always pooled together into a single review – this applies to any equilibrium of the

game.

Theorem 2. 1. For any equilibrium in the decentralized game, $V^D(p_t) \leq V^C(p_t)$ for any p_t .

2. For any p_t for which there exists a $p \in \mathcal{P}(p_t)$ that does not start a cascade, there exist $l^D(p_t)$, $r^D(p_t) > l^D(p_t)$ and $m \in \mathcal{M}$ such that $\bar{p} \in [l(p_t), r(p_t)]$, and for all $b_t \in [l(p_t), r(p_t)]$ we have $r(m | p_t, b_t) = 1$.

The first statement is relatively straightforward, since the principal in the commitment scenario has access to any communication structure that can arise in the equilibrium of a decentralized game.

The second statement mostly mirrors the intuition from the three-period example. The part that is worth pointing out is the qualifier on p_t : in particular, communication at p_t is noisy only if at least some message is available in $\mathcal{M}(p_t)$ that does not start a cascade. The complementary case was discussed in Proposition 1: if all messages in $\mathcal{M}(p_t)$ start a cascade then perfect communication is possible.

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 propagates, 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 is more confined, and perfect communication is possible when the reviewer sees the product as very good.

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 firm, reviews may not be the perfect source of information about products with uncertain characteristics.

References

- D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Fast and slow learning from reviews. working paper, 2017.
- S. N. Ali. Herding with costly information. *Journal of Economic Theory*, 175:713–729, May 2018.
- S. N. Ali and N. Kartik. Herding with collective preferences. *Economic Theory*, 51(3):601–626, November 2012.

- 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.
- P. H. Au. The loser’s curse in the search for advice. *Theoretical Economics*, 14(3):1015–1061, December 2019.
- A. V. Banerjee. A simple model of herd behavior. *Quarterly Journal of Economics*, 107(3):797–817, August 1992.
- G. S. Becker. A theory of social interactions. *Journal of Political Economy*, 82(6):1063–1093, 1974.
- S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of Political Economy*, 100(5):992–1026, October 1992.
- 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.
- 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. working paper, 2018.
- L. Cohen and Y. Mansour. Optimal algorithm for bayesian incentive-compatible exploration. working paper, 2019.
- Competition & Markets Authority. Online reviews and endorsements. https://www.gov.uk/government/uploads/system/uploads/attachment_data/file/436238/Online_reviews_and_endorsements.pdf, 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.
- J. Hörner, N. Klein, and S. Rady. Strongly symmetric equilibria in bandit games. working paper, 2015.
- G. Keller, S. Rady, and M. Cripps. Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68, January 2005.
- A. Kolmogorov and S. Fomin. *Elements of the Theory of Functions and Functional Analysis*. Dover Publications, 1957.
- 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.
- M. T. Le Quement and A. Patel. Communication as gift-exchange. working paper, 2018.
- A. Liang and X. Mu. Complementary information and learning traps. working paper, 2019.
- 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. working paper, 2019.
- C. March and A. Ziegelmeyer. Altruistic observational learning. working paper, 2016.
- S. Meier. A survey of economic theories and field evidence on pro-social behavior. working paper, 2006.
- D. Migrow. Designing communication hierarchies. working paper, 2018.
- 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-opinions-before-making-purchases>. Retrieved January, 2020.

- D. Peng, Y. Rao, X. Sun, and E. Xiao. Optional disclosure and observational learning. working paper, 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. working paper, 2018.
- L. Smith and P. Sørensen. Pathological outcomes of observational learning. *Econometrica*, 68(2): 371–398, March 2000.
- L. Smith and P. Sørensen. Observational learning. *New Palgrave Dictionary of Economics Online Edition*, pages 29–52, 2011.
- L. Smith, P. Sørensen, and J. Tian. Informational herding, optimal experimentation, and contrarianism. working paper, 2017.
- 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, 2018. URL <https://business.trustpilot.com/reviews/learn-from-customers/why-do-people-write-reviews-what-our-research-revealed>. Retrieved January, 2020.
- S. Vaccari, C. Maglaras, and M. Scarsini. Social learning from online reviews with product choice. working paper, 2018.
- A. Wolitzky. Learning from others’ outcomes. *American Economic Review*, 108(10):2763–2801, October 2018.
- W. Xu. Social learning in general information environments. working paper, 2018.

Appendix

When talking about current strategy in some public state p_t we always assume that equilibrium is not babbling in this state.

We also introduce the following notation

$$G_p(b) = p \cdot \mu^H(b|p) + (1 - p) \cdot \mu^L(b|p).$$

This is a cdf of the distribution of private posterior b as perceived by the consumer with the prior given by public belief p .

We will also use extensively the following representation of messaging strategies.

Definition. A messaging partition $\Sigma(p_t) = \{I_j\}$ in public state p_t given strategy profile r is a (possibly uncountable) collection of disjoint sets I_j such that

1. $\bigcup_j I_j = [0, 1]$,

2. for each I_j there exists $m_j \in \mathcal{M}$ such that $r(m_j | p_t, b_t) = 1$ if and only if $b_t \in I_j$.

Strategy profile r admits representation with a partition at state p_t if there exists a respective messaging partition $\Sigma(p_t)$.

Definition. Consider two public beliefs p'_t and $p''_t > p'_t$ and corresponding messaging partitions $\Sigma(p'_t) = \{I'_j\}$ and $\Sigma(p''_t) = \{I''_j\}$. Then we call messaging partition $\Sigma(p''_t)$ a parallel shift of $\Sigma(p'_t)$ if for any $b'_t \in I'_j$ we have

$$b''_t := b'_t + \ln\left(\frac{p''_t}{1-p''_t}\right) - \ln\left(\frac{p'_t}{1-p'_t}\right) \in I''_j.$$

We call $\Sigma(p''_t)$ a consistent parallel shift of $\Sigma(p'_t)$ if

$$\mathbb{E}[b_t | b_t \in I'_j] < \bar{p} \Rightarrow \mathbb{E}[b_t | b_t \in I''_j] < \bar{p}.$$

The first result provides a more convenient representation for value function (3). It shows that any posterior $p \in \mathcal{E}(p_t)$ is fully characterized by two numbers $V^H(p)$ and $V^L(p)$.

Lemma 2. If $p \in \mathcal{E}(p_t)$ then

$$V(p | p_t, b_t) = \theta(b_t) - c + \beta (b_t V^H(p) + (1 - b_t) V^L(p)), \quad (7)$$

where

$$V^i(p) = \mathbb{E} \left[\sum_{j=t+2}^{+\infty} \beta^{j-t-2} \cdot \mathbb{I}(p_j \geq \bar{p}) \cdot s_j \mid p, \theta = i \right], \quad i = \{H, L\}.$$

Moreover, $V(p | p_t, b_t)$ is linear and strictly increasing in b_t for any given p .

Proof. Because $p \in \mathcal{E}(p_t)$ then the next consumer buys the product. Therefore (3) reduces to

$$V(p | p_t, b_t) = b_t (H - c) + (1 - b_t) (L - c) + \beta (b_t V^H(p) + (1 - b_t) V^L(p)),$$

which is (7). At the same time if $\theta = H$ then $\mathbb{I}(p_j \geq \bar{p}) \cdot s_j \geq 0$ for any j , while if $\theta = L$ then $\mathbb{I}(p_j \geq \bar{p}) \cdot s_j \leq 0$. Therefore $(H - c + \beta V^H(p)) > (L - c + \beta V^L(p))$, and due to (7) $V(p_t, b_t | p)$ is thus strictly increasing in b_t . \square

Because $V(p | p_t, b_t)$ is linear in b_t , the higher b_t is, the more weight is assigned to $V^H(p)$ and less to $V^L(p)$. Therefore at a higher b_t a consumer would prefer to induce a posterior with higher $V^H(p)$ and lower $V^L(p)$. The next lemma shows it formally.

Lemma 3. Suppose that $m', m'' \in \mathcal{M}(p_t)$ and $\bar{p} \leq q(p_t, m') < q(p_t, m'')$. Then either $V^H(q(p_t, m')) = V^H(q(p_t, m''))$ and $V^L(q(p_t, m')) = V^L(q(p_t, m''))$, or $V^H(q(p_t, m')) < V^H(q(p_t, m''))$ and $V^L(q(p_t, m')) > V^L(q(p_t, m''))$.

Proof. If $V^H(q(p_t, m')) < V^H(q(p_t, m''))$ and $V^L(q(p_t, m')) < V^L(q(p_t, m''))$ then $V(q(p_t, m') | p_t, b_t) < V(q(p_t, m'') | p_t, b_t)$ and therefore $m' \notin \mathcal{M}(p_t)$, – a contradiction. Analogously, it can not be that $V^H(q(p_t, m')) > V^H(q(p_t, m''))$ and $V^L(q(p_t, m')) > V^L(q(p_t, m''))$. Finally, if for some p' and p'' we have $V^H(p') > V^H(p'')$ and $V^L(p') < V^L(p'')$ it implies that there exists $\bar{b} \in [0, 1]$ such that $V(p' | p_t, b_t) < V(p'' | p_t, b_t)$ for $b_t \in [0, \bar{b})$ and $V(p' | p_t, b_t) > V(p'' | p_t, b_t)$ for $b_t \in (\bar{b}, 1]$. Therefore we must have $p'' < p'$. \square

Corollary 4. *In any equilibrium for any p_t*

1. $V^H(p)$ is a [weakly] increasing function of p on $\mathcal{E}(p_t)$,
2. $V^L(p)$ is a [weakly] decreasing function of p on $\mathcal{E}(p_t)$.

Proof. Directly follows from Lemma 3. □

Lemma 5. *At any public state p_t there exists $\bar{b}(p_t) \geq 0$ such that*

1. *If $r(m|p_t, b_t) > 0$ for some $b_t < \bar{b}_t$ then $q(p_t, m) \in \mathcal{S}(p_t)$;*
2. *If $r(m|p_t, b_t) > 0$ for some $b_t \geq \bar{b}_t$ then $q(p_t, m) \in \mathcal{E}(p_t)$;*
3. *If public state $p_\tau < 1$ with $\mathcal{S}(p_\tau) \neq \emptyset$ occurs with strictly positive probabilities $A^H > 0$ if $\theta = H$ and $A^L \in [0, A^H]$ if $\theta = L$ from public state p_t then $V(p_t, \bar{p}) > 0$;*
4. *$\bar{b}_t \leq \bar{p}$ and the equality is attained if and only if any chosen $p \in \mathcal{E}(p_t)$ leads to a continuation equilibrium where experimentation never stops, i.e. to an equilibrium which is payoff-equivalent to a babbling equilibria played from period $t + 1$ onward.*

Proof. If $\mathcal{S}(p_t) = \emptyset$ then we can take $\bar{b}_t = 0$. Therefore we hereafter assume that $\mathcal{S}(p_t) \neq \emptyset$. To prove the first two claims assume the contrary. That is there exist $b'_t < b''_t$, m' , m'' such that $q(p_t, m') \in \mathcal{E}(p_t)$, $q(p_t, m'') \in \mathcal{S}(p_t)$ and $r(m'|p_t, b'_t) > 0$, $r(m''|p_t, b''_t) > 0$. Then $V(p_t, b'_t) = 0$, and $V(p_t, b'_t) \geq 0$ as $\mathcal{S}(p_t) \neq \emptyset$. At the same time due to Lemma 2 we have

$$0 = V(p_t, b'_t) \geq V(q(p_t, m') | p_t, b'_t) > V(q(p_t, m') | p_t, b'_t) \geq 0,$$

which leads to a contradiction. This argument proves first two parts of the lemma.

By the first two parts it follows that if $\mathcal{S}(p_\tau) \neq \emptyset$ then $\bar{b}_\tau > 0$. Next, note that from point of view of consumer at p_t who holds private belief $b_t = \bar{p}$ at every future history stage payoff is not less than zero. Indeed, if there is no experimentation, then payoff is equal to 0, while if there is it is also equal to 0. Therefore

$$V(p_t, \bar{p}) \geq \beta^{\tau-t} [A^H \bar{p} (1 - F^H(\bar{s}_\tau)) (H - c) + A^L (1 - \bar{p}) (1 - F^L(\bar{s}_\tau)) (L - c)] > 0,$$

where $\bar{s}_\tau = l^{-1} \left(\ln \left(\frac{\bar{b}_\tau}{1 - \bar{b}_\tau} \right) - \ln \left(\frac{p_\tau}{1 - p_\tau} \right) \right)$, and l^{-1} is an inverse function to $\ln \left[\frac{f^H(s)}{f^L(s)} \right]$.

Finally, we prove the last part. If any posterior $p \in \mathcal{P}(p_t)$ induces babbling from $(t+1)$, then the current consumer decides whether all future consumers will buy the product or avoid it. Their expected utility from her point of view is then $\frac{1}{1-\beta} (\theta(b_t) - c)$ or zero in the two respective scenarios. Therefore as equilibrium is not babbling the current consumer makes all subsequent consumers buy the product if and only if $b_t \geq \bar{p}$.

To show the reverse statement it suffices to show that $V(p_t, \bar{p}) > 0$. Then by continuity it would imply $\bar{b}(p_t) < \bar{p}$. If at any public state p_τ which is a part of a history originating from p_t we have $\mathcal{S}(p_\tau) = \emptyset$, then this equilibrium is payoff-equivalent to a one where babbling equilibrium is played from period $t + 1$ onwards. Therefore suppose there exists a public belief p_τ such that $\mathcal{S}(p_\tau) \neq \emptyset$. Denote the path of public beliefs that lead to this public belief as $p_{\tau-1}, p_{\tau-2}, \dots$. Without loss also take minimal τ with such property. Consider public state $p_{\tau-1}$ such that $\mathcal{E}(p_{\tau-1}) = p_\tau$. Then by part 3 there exists \tilde{p}_τ such that $V(\tilde{p}_\tau | p_{\tau-1}, \bar{p}) = \bar{p} V^H(\tilde{p}_\tau) + (1 - \bar{p}) V^L(\tilde{p}_\tau) > 0$. Because $\mathcal{S}(p_{\tau-1}) = \emptyset$ as τ was taken to be minimal and by Corollary 4 and Lemma 3 we have that either \tilde{p}_τ or any other public belief with the same V^H and V^L is induced for all private beliefs $b_{\tau-1} \leq \bar{p}$ and by belief consistency in some neighborhood above \bar{p} . With respect to $G_{p_{\tau-1}}$ this region has a strictly positive measure. Therefore probability A^L

that private belief $b_{\tau-1}$ is within this region if $\theta = L$ is positive and is lower than probability A^H that private belief $b_{\tau-1}$ is within this region if $\theta = H$. Therefore at $p_{\tau-2}$ there exists $\tilde{p}_{\tau-1} \in \mathcal{E}(p_{\tau-2})$ such that $V(\tilde{p}_{\tau-1} | p_{\tau-2}, \bar{p}) = \bar{p}V^H(\tilde{p}_{\tau-1}) + (1 - \bar{p})V^L(\tilde{p}_{\tau-1}) \geq \beta A^H \bar{p}V^H(p_{\tau-1}) + \beta A^L(1 - \bar{p})V^L(p_{\tau-1}) > 0$. As τ is finite going backward through the history of public states we get the result for p_t which concludes the argument. \square

Proofs of the Main Results

Proof of Proposition 1. For decentralized case the proof directly follows from part 4 of Lemma 5. The proof for commitment case follows from Theorem 1. \square

Lemma 6. *Suppose there exist $0 < l < \bar{p} < r < 1$, $\varepsilon > 0$ and $r(\varepsilon) > r$ such that*

$$\mathbb{E}_{x \sim G_p} [x | x \in [l, r]] = \mathbb{E}_{x \sim G_p} [x | x \in [l - \varepsilon, r(\varepsilon)]] = \bar{p}.$$

Then for any ε there exist $\underline{\gamma}, \bar{\gamma} > 0$ such that $\underline{\gamma}\varepsilon < r(\varepsilon) - r < \bar{\gamma}\varepsilon$.

Proof. Because both $f^L(x), f^H(x)$ are continuously differentiable there exist $\delta > 0$ and constants $0 < B_l < B_h$ such that $G'_p(x) \in (B_l, B_h)$ for all $x \in (\delta, 1 - \delta)$. In what follows we consider this neighborhood. By definition

$$(G_p(r) - G_p(l)) \cdot \int_{l-\varepsilon}^{r(\varepsilon)} x dG_p(x) = (G_p(r(\varepsilon)) - G_p(l - \varepsilon)) \cdot \int_l^r x dG_p(x).$$

Rearranging terms we get

$$\int_{l-\varepsilon}^l (\bar{p} - x) dG_p(x) = \int_r^{r(\varepsilon)} (x - \bar{p}) dG_p(x).$$

Using boundedness of $G'_p(x)$ we can therefore obtain

$$\begin{aligned} B_l \left(\bar{p} - l + \frac{\varepsilon}{2} \right) \cdot \varepsilon &< B_h \left(r - \bar{p} + \frac{r(\varepsilon) - r}{2} \right) \cdot (r(\varepsilon) - r), \\ B_h \left(\bar{p} - l + \frac{\varepsilon}{2} \right) \cdot \varepsilon &> B_l \left(r - \bar{p} + \frac{r(\varepsilon) - r}{2} \right) \cdot (r(\varepsilon) - r). \end{aligned}$$

Solving these inequalities in $r(\varepsilon) - r$ we get

$$< r(\varepsilon) - r <$$

Both ... are continuous in ε and therefore attain their lowest and highest values respectively on $\varepsilon \in [0, 1]$. \square

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

Proof. \square

Lemma 8. *There exists $\delta > 0$ such that for any p_t and any $p \in \mathcal{E}(p_t)$ we have $V^C(p | p_t, b_t) < 0$ for all $b_t \in [0, \delta]$.*

Proof. Because $p \geq \bar{p}$ at time $t + 1$ the next consumer certainly purchases the product and therefore by Lemma 2 we get

$$V^C(p | p_t, b_t) = \theta(b_t) - c + \beta (b_t V^{C,H}(p) + (1 - b_t) V^{C,L}(p)).$$

Because $V^{C,H}(p) \leq \frac{H-c}{1-\beta}$ and $V^{C,L}(p) \leq 0$ we have that $V^C(p | p_t, b_t) \leq 0$ for all $b_t \leq \delta := \frac{(1-\beta)(c-L)}{H-\beta c-(1-\beta)L}$. \square

Lemma 9. Consider public state p_t and $p', p'' \in \mathcal{E}(p_t)$. Consider all public beliefs p'_τ and p''_τ that can be on paths originating from p' and p'' respectively and their associated messaging partitions $\Sigma(p'_\tau)$ and $\Sigma(p''_\tau)$. Then if $\Sigma(p''_\tau)$ is a consistent parallel shift of $\Sigma(p'_\tau)$ for each pair p'_τ, p''_τ by the same increment then $V^\theta(p') = V^\theta(p'')$ for $\theta \in \{H, L\}$.

Proof. Note that probability that private belief b_t is lower than some threshold conditional on some public belief p_t depends only on $\ln\left(\frac{b_t}{1-b_t}\right) - \ln\left(\frac{p_t}{1-p_t}\right)$. If one shifts all the bounds for all sets in a given partition by the same increment, then expectations over all sets in the new partition shift by the same increment. Moreover, if the new partition is consistent with the previous one it implies therefore that conditional on θ under $\Sigma(p''_\tau)$ in any public belief p''_τ the experimentation stops with the same probability as it does under $\Sigma(p'_\tau)$ in p'_τ . Therefore $V^\theta(p') = V^\theta(p'')$. \square

Proof of Theorem 1. We initially assume that at time t all consumers in the queue (current and future) can commit to a particular message structure. As the very initial step note that for every p_t an optimal $V^C(p_t)$ is achieved by choosing optimal $\Sigma(p_\tau)$ for every p_τ that can originate from p_t . Next, this partitions without loss can be assumed to be Markovian. That is $\Sigma(p_\tau)$ for all p_τ originating from p_t should not depend on t explicitly. Indeed, suppose there exists p and $\tau > t$ such that $p_t = p_\tau = p$, but optimal partitions originating from p_t and p_τ do not coincide. Then if $V^C(p_t) = V^C(p_\tau)$ without loss we can prescribe any of these partitions to p . If however, say, $V^C(p_t) < V^C(p_\tau)$ then we can prescribe partition corresponding to p_τ to p_t which then would strictly increase $V^C(p_t)$. Because partition for p_t was assumed to be optimal we get a contradiction. The case $V^C(p_t) > V^C(p_\tau)$ is analogous. Finally, in what follows we understand by $\hat{\Sigma}(p_t)$ a collection of optimal partitions for all p_τ with $\tau \geq t$ that deliver $V^C(p_t)$ (if it exists).

We next divide the proof into several steps outlined in the text.

Step 1. Note that ex ante value when public belief is p is equal to the value a consumer with *private* belief p gets when she induces public belief p . Formally, $V^C(p) = V^C(p | x, p)$ for any time- t public belief x . Indeed, public beliefs about the state coincide in these two cases and in the second case consumer's own belief is also equal to p and therefore she has the same expectations about future histories as a consumer who face public belief p .

Step 2. Take some p_t . There exists a set of public beliefs p_τ (including p_t) which can be reached on path originating from p_t . For each of them including p_t there exists an optimal partition $\Sigma(p_\tau)$. Note that for any $b_t > \bar{p}$ at any p_τ it is strictly optimal to induce experimentation. Indeed, suppose it is not, and at such b_t experimentation stops. Then we can exclude this belief from the set of private beliefs that induce no experimentation and induce babbling continuation in b_t forever after. In that case posterior for the pooled region where the experimentation stops will stay below \bar{p} , while for b_t we now get strictly positive continuation value.

We next show that for any p_t in an optimal partition $\Sigma(p_t)$ there exists $\delta > 0$ such that a consumer with private belief $b_t \in [0, \delta]$ sends a message that stops experimentation. Lemma 8 implies that if the consumer with private belief $b_t \leq \delta := \frac{(1-\beta)(c-L)}{H-\beta c-(1-\beta)L}$ induces public belief $p \in \mathcal{E}(p_t)$ then $V^C(p | p_t, b_t) < 0$. Suppose that in $\Sigma(p_t)$ there exists some set of points within $[0, \delta]$ that is pooled with private beliefs above \bar{p} and the

resulting posterior is above \bar{p} . Denote this posterior as p_{pool} . Then we can construct an alternative $\hat{\Sigma}(p_t)$ that will deliver at least the same ex ante value and where for all $b_t \in [0, \delta]$ the experimentation stops.

First, cut from the pooling region private beliefs below δ that induce experimentation and substitute continuation with no experimentation afterwards. If this set was of measure zero then we are done.¹⁷ If it was of positive measure then we can cut the right end of this interval such that the resulting posterior stays on the level of p_{pool} . Now for all points below δ by Lemma 8 this adjustment provides an improvement. For all points that were not cut the value stays the same as induced posterior stayed the same. For the points that were cut from the right end we now need to prescribe continuation partitions that also provide an improvement. We do that by inducing truthful information transmission at all such points at p_t . In all resulting public states $p_{t+1} = b_t$ afterwards we use a parallel shift of $\hat{\Sigma}(p_t)$ by the increment of $\ln\left(\frac{b_t}{1-b_t}\right) - \ln\left(\frac{p_{pool}}{1-p_{pool}}\right)$.

We then without loss we can pool together all messages that induce no further experimentation into one message. We now know that resulting posterior of this region will be *uniformly* separated from \bar{p} for any p_t .

Step 3. We show that $V^C(b_t | p_t, b_t)$ is a continuous, strictly increasing and [weakly] convex function of b_t for any p_t . Consider some b_t and denote as $\Sigma(b_t)$ the corresponding optimal partitions. Then for any point \tilde{b}_t in the *right* neighborhood of b_t consider partitions $\hat{\Sigma}(\tilde{b}_t)$ that consist of shifted partitions $\hat{\Sigma}(b_t)$.¹⁸ Because the posterior for the no experimentation region is uniformly separated from \bar{p} , in some neighborhood of b_t we can always do that in such a way that all the posteriors on path of play that were above \bar{p} stay above it, and the ones that were below \bar{p} stay below it. Then with $\hat{\Sigma}(\tilde{b}_t)$ we have $V^H(\tilde{b}_t) = V^H(b_t)$ and $V^L(\tilde{b}_t) = V^L(b_t)$.

Now suppose $V^C(b_t | p_t, b_t)$ is discontinuous at some $b_t \geq \bar{p}$. Then there exists $\varepsilon > 0$ such that for any $\delta > 0$ we can find b'_t and b''_t with $|b''_t - b'_t| < \delta$ such that

$$V^C(b''_t | p_t, b''_t) - V^C(b'_t | p_t, b'_t) > \varepsilon.$$

For any p we then can take parallel shift of partitions $\hat{\Sigma}(b'_t)$ by the increment of $\ln\left(\frac{b'_t}{1-b'_t}\right) - \ln\left(\frac{b''_t}{1-b''_t}\right)$. There also exist such δ that this shift will be consistent, which by Lemma 9 implies that $V^\theta(b''_t) = V^\theta(b'_t)$. Because $V^H(p) - V^L(p) \leq H - L$ for any p we then have that taking $\delta < \frac{\varepsilon}{H-L}$ we get

$$V^C(b''_t | p_t, b''_t) - V^C(b'_t | p_t, b'_t) < (b''_t - b'_t) \cdot (V^H(p) - V^L(p)) = \varepsilon,$$

which gives us a contradiction.

Similarly, using parallel shifts and the fact that in some right neighborhood of any private belief b_t they will be consistent, one can show strict monotonicity and convexity of $V(b_t | p_t, b_t)$.

Step 4. Because $V^C(p_t)$ is weakly convex so is $V^C(b_t | p_t, b_t)$ as a function of b_t . Therefore, first, one does not pool any private beliefs b_t above \bar{p} , and, second, if one pools beliefs below and above \bar{p} induced posterior should be \bar{p} . Indeed, if one pools some private beliefs b_t below and above \bar{p} and the resulting posterior is strictly above \bar{p} , then one can cut the right end of this region to lower induced public belief to \bar{p} . This weakly increases $V^C(p_t)$ by Jensen's inequality. Finally, the pooling region has to be convex below \bar{p} . For any $p \geq \bar{p}$, in particular \bar{p} itself, we have

$$\left. \frac{dV^C(b_t | p_t, b_t)}{db_t} \right|_{b_t=\bar{p}} > 0$$

¹⁷Measure here is understood in terms of conditional measure on this pooling region.

¹⁸All by the same increment of $\ln\left(\frac{\tilde{b}_t}{1-\tilde{b}_t}\right) - \ln\left(\frac{b_t}{1-b_t}\right)$.

Therefore gains from pooling private beliefs below \bar{p} decrease in distance from \bar{p} . At the same time due to convexity of $V^C(b_t | p_t, b_t)$ in b_t losses are increasing in distance from \bar{p} . Now suppose that the set of pooled private beliefs is not convex below \bar{p} . Then, first, we can cut private beliefs in the left end of the pooling region and attach the same measure of private beliefs directly to the left of the maximum pooling region around \bar{p} . This will weakly increase the posterior belief for the pooling region, which we can further decrease by cutting points from the right end of it. This will again improve $V^C(b_t | p_t, b_t)$ point-wise.

Therefore without loss we can assume that there exist $0 \leq l^C(p_t) \leq \bar{p}$ and set $I^C(p_t)$ with all its points above \bar{p} such that

1. for all $b_t \leq l(p_t)$ consumer sends message $m \in \mathcal{S}(p_t)$, i.e., experimentation stops.
2. for any $b_t \notin I^C(p_t)$ consumer truthfully transmits his private belief b_t (or what is the same his private payoff) to the public, that is $p_{t+1} = b_t$.
3. for all $b_t \in (l(p_t), p_t) \cup I^C(p_t)$ consumer sends message m such that $q(p_t, m) = \bar{p}$.

Step 5. We now show that it is always optimal to pool at least some private beliefs b_t around the cutoff for any p_t if $V^C(\bar{p} | p_t, \bar{p}) > 0$. In other words it has to be that $l(p_t) < \bar{p} < r(p_t)$. From the previous step we know that $V^C(\bar{p} | p_t, \bar{p}) > 0$, therefore pooling $\varepsilon \in (0, \delta)$ below \bar{p} with some above \bar{p} gives a benefit of $V^C(\bar{p} | p_t, \bar{p}) \cdot \varepsilon + \mathcal{O}(\varepsilon^2)$. At the same time because $V^C(b | p_t, b)$ is continuous in b , its slope is less than $H - L$, and by Lemma 6 losses associated with pooling beliefs above \bar{p} do not exceed $B\varepsilon^2$ for some $B > 0$. Therefore there always exists such $\varepsilon > 0$ that $V^C(\bar{p} | p_t, \bar{p}) \cdot \varepsilon + \mathcal{O}(\varepsilon^2) > B\varepsilon^2$, and therefore pooling at least some beliefs around \bar{p} is always optimal. If $V^C(\bar{p} | p_t, \bar{p}) = 0$ then it can only be the case if $V^C(b_t | p_t, b_t) = \theta(b_t) - c$ for $b_t \geq \bar{p}$. Therefore for any p_t it is not optimal to pool any private beliefs around \bar{p} , and therefore in any p_t we have that no experimentation is induced for all $b_t \in [0, \bar{p}]$. Then if $A^\theta = \Pr(b_{t+1} \geq \bar{p} | \theta)$ we have

$$V^C(\bar{p} | p_t, \bar{p}) > \bar{p}A^H \cdot (H - c) + (1 - \bar{p})A^L \cdot (L - c) > 0,$$

because $A^H > A^L$. This contradicts $V^C(\bar{p} | p_t, \bar{p}) = 0$. Therefore $V^C(\bar{p} | p_t, \bar{p}) > 0$ and it is always optimal to pool at least some private beliefs around \bar{p} .

Step 6. We show that optimal $V^C(p_t)$ exists for every p_t . We have seen that for *any* p_t optimal partition can be reduced just to one scalar $l(p_t)$ (the right end $r(p_t)$ of the pooling region is then uniquely identified). Therefore picking a *function* $l(p)$ for $p \in [\bar{p}, 1]$ we get a particular value $V_{l(p)}^C(p_t)$. Therefore to get an *optimal* value for $V^C(p_t)$ we need to maximize it with respect to $l(p)$. For any p_t we have that $l(p_t) \in [0, \bar{p}]$, and therefore $l(p)$ is uniformly bounded on $p \in [\bar{p}, 1]$. Call the space of functions bounded from above by \bar{p} and by 0 from below by $\mathcal{B}[\bar{p}, 1]$ and endow it with a standard sup-norm. That is for each pair of functions $l_1(p)$ and $l_2(p)$ define

$$\rho(l_1, l_2) := \sup_{p \in [\bar{p}, 1]} |l_1(p) - l_2(p)|.$$

Continuous functions defined on compact metric spaces attain their maximum values (see Kolmogorov and Fomin [1957], chapter 2, §19). Therefore to show that $V^C(p_t)$ attains its maximum on $\mathcal{B}[\bar{p}, 1]$ we need to establish that

1. $V^C(p_t)$ is continuous in $l(p)$,
2. there exists a compact subspace of $\mathcal{B}[\bar{p}, 1]$.

For any $\varepsilon > 0$ there exists t such that $\frac{\beta^t}{1-\beta} \cdot \max\{|H - c|, |L - c|\} < \frac{\varepsilon}{t}$.

To show the second point we use Arzela-Ascoli theorem (see Kolmogorov and Fomin [1957], chapter 2, §17). To apply it we need to establish that optimal $l(p)$ is within some family of *equicontinuous* functions.¹⁹ For that we next show that $|l'(p)| < B$ for some $B > 0$. In this case it will imply that we can restrict ourselves to a family of functions with their derivative being bounded by B , which is clearly equicontinuous. Finally, we then take closure of this set to obtain a compact subspace.²⁰

Step 7. Do the maximization procedure outlined in the previous step for each $p_t \in [bar{p}, 1]$. This delivers optimal values $V^C(b_t | p_t, b_t)$ for all b_t . Then for each p_t we can identify optimal $l(p_t)$ given $V^C(b_t | p_t, b_t)$. Then constructed $l(p)$ delivers maximum for $V(p_t)$ for any p_t .

Note that because gains from pooling are decreasing in distance from \bar{p} and losses associated with pooling private beliefs from the right [weakly] increase in the distance such $l(p_t)$ is unique for any p_t . We know that $l(p)$ is a continuous function for any $V^C(b_t | p_t, b_t)$. By Lemma 6 $r(p)$ is then continuous as well. Therefore it attains its lowest value. Denote it as \underline{r} . Consider such p_t that $r(p_t) = \underline{r}$. Because all $p_\tau > r$ appear on path from this p_t then

we can find $l(p)$ defined on $p \in [\underline{r}, 1]$ that delivers a maximum to $V(p_t)$. All $p \in [\bar{p}, \underline{r}]$ then do not appear in any future history originating from $p_t = p$. In such public states, and therefore for such values we can prescribe value for $l(p_t)$ that we obtain from a static maximization problem. That finishes construction of optimal $l(p)$.

Finally, note that if such $l(p)$ maximizes $V^C(p_t)$ for some p_t then it maximizes $V^C(p_t)$ for any p_t . Indeed, if some p_τ does not appear on path from p_t then the claim is true automatically. If however p_τ can appear on path originating from p_t then it has to be that $l(p)$ delivers a maximum (possibly not unique) to p_τ as well, as otherwise it would contradict optimality of $l(p)$ for p_t .

This finalizes the proof of the whole claim. □

Proof of Theorem 2. The first part of the Theorem follows from the fact that any decentralized solution induces partitions of private beliefs b_t for every p_t . The maximal value which can be achieved for any partitions in public state p_t is delivered by $V^C(p_t)$, and therefore $V^D(p_t) \leq V^C(p_t)$.

The second part follows from part 4 of Lemma 5. It implies that $\bar{b}_t < \bar{p}$. Therefore we can take $l^D(p_t) = \bar{b}_t$. By belief consistency and Lemma 3 there should exist $r^D(p_t) > \bar{p}$ and message $m \in \mathcal{M}$ such that m is sent for all $b_t \in [l^D(p_t), r^D(p_t)]$.

¹⁹A family \mathcal{D} of continuous functions on a closed interval $[a, b]$ is called equicontinuous if for any $\varepsilon > 0$ there exists $\delta > 0$ such that for any $f \in \mathcal{D}$ and any $x_1, x_2 \in [a, b]$ condition $|x_1 - x_2| < \delta$ implies $|f(x_1) - f(x_2)| < \varepsilon$.

²⁰Note that when taking the closure all limiting functions' values will still be within $[0, \bar{p}]$ and therefore all functions within this set are valid candidates for an optimal $l(p)$.