(Pro-)Social Learning and Strategic Disclosure*

Roland Bénabou[†] Nikhil Vellodi[‡]

Current Draft : May 14, 2024

Abstract

We study a sequential experimentation model with endogenous feedback. Agents choose between a safe and risky action, the latter generating stochastic rewards. When making this choice, each agent is selfishly motivated (myopic). However, agents can disclose their experiences to a public record, and when doing so are pro-socially motivated (forward-looking). When prior uncertainty is large, disclosure is both *polarized* (only extreme signals are disclosed) and *positively biased* (no feedback is bad news). When prior uncertainty is small, a novel form of unraveling occurs and disclosure is complete. Subsidizing disclosure costs can perversely lead to less disclosure but more experimentation.

Keywords: social learning, experimentation, dynamic disclosure, consumer reviews, time-inconsistent preferences, motivated beliefs. **JEL Classification**: D82, L11, L12.

1 Introduction

In many settings, agents face a choice between safe and risky actions, with different individuals facing these choices in sequence. Agents might benefit from the information generated by those who preceded them. For instance, consider consumers choosing whether or not to dine at a restaurant with unknown quality, or to watch a new movie. Those who do so can

*We thank Aislinn Bohren, Gabriel Carroll, Navin Kartik, Alex Wolitzky and various seminar participants for useful suggestions. Vellodi acknowledges funding from the EUR grant ANR-17-EURE-0001, as well as Princeton University, where this research was initiated in 2019.

[†]Princeton University, NBER, CEPR, IZA, BREAD, and briq. rbenabou@princeton.edu

[‡]Paris School of Economics. n.vellodi@gmail.com

then leave feedback, helping later-arriving consumers make more informed choices. Similar settings include the adoption of new products and technologies, employment choices, and sequential voting.

A well-known dynamic externality emerges in such settings, namely that agents do not internalize the benefit to future consumers of taking the risky action, leaving feedback and thereby generating socially valuable information.

To remedy this inefficient under-exploration, a planner would direct agents to take the risky action even when it is unprofitable to them provided the informational gain to future agents more than compensates; oftentimes, direct incentives are either absent or forbidden.[1]

This question has been studied extensively in economics and computer science under the label of "incentivized exploration (IE)" (Kremer et al., 2014; Che and Hörner, 2018), itself part of the broader literatures on social learning and sequential experimentation (Banerjee, 1992; Bikhchandani et al., 1992; Smith and Sørensen, 2000; Smith et al., 2021).[2] Papers in the IE literature largely take a *normative* approach to the problem. Namely, they assume the presence of a benevolent designer who can control the provision of incentives either via dynamic information provision or through direct recommendations. Furthermore, these works largely assume that, once generated, individual signals are perfectly observed by the planner or designer in charge of public information provision.

Such studies are thus silent on a particularly salient issue within the online feedback setting — why and when do people leave feedback in the first place? For instance, when leaving feedback regarding their dining experience, consumers might be driven by a desire to help future consumers make informed choices, or alternatively to reward or punish the restaurant for a positive or negative experience. Conversely, while the availability of feedback is often highly valued by consumers, the vast majority fail to provide it.[3] And those that do provide feedback generate well-known biases such as *positive selection* (Nosko and

[1] For instance, the US *Consumer Review Fairness Act* forbids eliciting reviews for payment. See https://tinyurl.com/5v3mt846.

[2] See Slivkins (2022) and Bikhchandani et al. (2022) for recent surveys on IE and social learning more generally.

[3] Recent surveys report that only around 10% of consumers regularly leave reviews. See https://tinyurl.com/mrrsf9v5, https://tinyurl.com/ux3zyuem.

Tadelis, 2015; Hui et al., 2024) — undisclosed experiences are on average negative — and *polarization* (Schoenmueller et al., 2020; Marinescu et al., 2021) — extreme reviews are more prevalent than average reviews. Understanding the costs and benefits of leaving feedback is crucial, both to better understand observed patterns of behavior and for designing policy interventions that could improve market efficiency.

We take a step towards addressing these questions by providing a *positive* theory of IE. Namely, we propose a simple model of sequential experimentation, in which we endow our agents with *pro-social* motives to leave feedback. In our model, agents arrive sequentially over three periods, and have the choice between a safe (S) and risky action (R). Action S generates a deterministic reward, whereas R yields a random reward (which we sometimes refer to as a signal) correlated to an underlying hidden state, for instance, the unknown quality of a product. If an agent plays R, they can disclose their signal. Our main analysis models disclosure via hard evidence with noisy transmission (Dye, 1985a); an agent is able to leave feedback with probability $\alpha \in [0, 1]$, but can either truthfully report their signal or not report it at all. Two crucial assumptions determine an agent's payoff. First, when choosing their *action,* we assume that they are fully *self-interested*, maximizing only their personal reward. Second, when making their subsequent *disclosure* choice, we assume agents are fully *pro-social.* Formally, they have lexicographic preferences over making optimal consumption decisions for themselves, and then transmitting useful information to help others do the same.

This simple combination of ingredients delivers a rich theory of selective disclosure, which both accords with well-documented phenomena and provides new testable predictions. In particular, our first main result (Theorem 1) demonstrates that equilibrium disclosure is both positively selected and polarized. The intuition is simple. When player 1 (P1) plays R, their disclosure choice is governed equally by the subsequent payoffs of P2 and P3. This is not the case for P2, who values their own payoff far more. Thus, P1 may be inclined to strategically disclose or conceal their experience, in order to induce P2 to take an action unfavorable to themselves but beneficial to P3. This is the case when an informed P2 would fail to experiment, even though the loss to themselves is smaller than the gain to P3 that the information generated by doing so would provide. It is precisely such experiences that P1 strategically conceals, inducing P2 to experiment against their interests for the sake of P3. In this case, we say P1 induces P2 to *experiment*. Simply put, an early adopter would rather not take responsibility for causing the untimely demise of a new product, if there is a reasonable chance the product is in fact worth a second chance, and in this case they keep quiet. On the other hand, when experiences are sufficiently negative, P1 is convinced that no further experimentation should occur and thus terminates it by posting their feedback, while for (even marginally) positive experiences, there is no downside to disclosure. Thus, strategic non-disclosure is used exclusively by P1 to foster efficient experimentation by P2. Of course, models that assume that leaving feedback is costly and done only when sufficiently informative also generate polarized feedback, but struggle to also deliver positive selection from a single behavioral foundation.[4]

Beyond these, our model delivers further predictions. For instance, we fully characterize how equilibrium non-disclosure, and thus experimentation, varies with the prior belief regarding risky payoffs (Theorem 2). We view this exercise as capturing, in a reduced-form manner, how disclosure varies with how old or well-established the product market in question is. We show that the extent of experimentation is hump-shaped in the prior. Moreoever, experimentation completely fails to occur in equilibrium for a range of sufficiently high prior beliefs. The result relies on incentives that are distinct from the standard unraveling mechanism (Milgrom, 1981; Grossman, 1981; Dye, 1985a). In particular, the classic result does not depend on the prior distribution, whereas it does in our analysis. The key distinction is that in our setting, the sender's preferences over induced posterior beliefs are type-dependent, whereas in the classic setting, each type strictly prefers to induce the highest possible posterior. This type-dependence stems from the misaligned preferences at the heart of the model: for an intermediate range of signals, P1 strictly prefers to induce P2 to hold the lowest belief at which P2 still plays R and thus experiments, as described

[4]For instance, Hui et al. (2024) allow feedback to be positively biased for unmodelled reasons, suggesting reasons such as fear of retaliation or a simple aversion to providing negative criticism.

above. Outside of this range, P1 strictly prefers to have P2 share their belief. By this token, this failure of experimentation is hence critically linked to the ex-post optimality of disclosure; we show that it fails to hold when agents can commit to a disclosure policy prior to receiving their signal (Kamenica and Gentzkow, 2011).

Finally, we show that the extent of experimentation is also hump-shaped in α, the feedback opportunity parameter. This insight has important implications for real-world interventions; practitioners argue that the lack of feedback in online markets leads to biased inference, and that making feedback less costly (e.g. by providing explicitly monetary incentives) would thus lead to more information and thus experimentation (Marinescu et al., 2021). If we take the natural interpretation that α corresponds to the fraction of agents for whom feedback is costless and $1-\alpha$ the fraction for whom it is prohibitively costly, our result suggests that making feedback less costly could perversely lead to *less* disclosure, and more generally that the effectiveness of such interventions in stimulating feedback rates will vary by products and markets.

The joint assumption of selfish consumption and pro-social disclosure is appealing on three separate fronts. First, from a positive perspective, surveys suggest that the welfare of future consumers is a key driver when leaving feedback, as is rewarding or punishing sellers for good versus bad experiences.[5] (Our analysis is flexible and can accommodate both of these motivations.) At the same time, empirical evidence suggests that incentives to provide feedback are divorced from actual consumption choices in online settings (Cabral and Li, 2015). We present a first attempt at formalizing these arguments, with a view to understanding both their theoretical foundations as well as their ability to organize empirical findings.

Second, endowing agents with benevolent preferences in this manner allows our theory to be viewed as a minimal departure from the normative analyses in the IE literature. That is, our agents are effectively mini-planners when disclosing, facing the same trade-off between *exploration* (long-run information gains) and *exploitation* (short-run consumption gains) as

[5] For example, see https://tinyurl.com/mrybw969.

in previous work, but they also face additional constraints imposed on them in equilibrium, such as ex-post optimality of disclosure rules. Our results thus demonstrate how such constraints shape the degree to which disclosure can be used to incentivize exploration.

Third, the informational externality described above derives fundamentally from the structure of intertemporal preferences, namely that agents are "present-biased" when making their consumption choices. This gives rise to an alternative, psychological interpretation of the model. Instead of a sequence of agents, consider a single decision maker with the following dynamically inconsistent preferences. When taking actions that affect current payoffs, they are myopic (completely present biased), whereas when deciding what available information to store in memory to inform future choices, they are patient. This corresponds to a limiting case of quasi-hyperbolic, or $\beta\delta$ (Laibson, 1997), preferences wherein β is arbitrarily small. By modeling the disclosure objective as altruistic, our framework permits this application to an individual who selectively encodes their experiences in order to become less "conservative" — that is, more open to trying and learning from new experiences. Our work thus demonstrates a close conceptual connection between IE and motivated reasoning (Bénabou and Tirole, 2002, 2004; Carrillo and Mariotti, 2000).

To identify how ex-post constraints shape disclosure, we also analyze the case where P1 can commit to a signaling rule. (We also study cheap-talk in our Online Appendix.) We find that communication is again polarized and positively selected, but *regardless* of the starting prior (Proposition 3). This result is in contrast to Theorem 2. In concurrent and independent work, Smirnov and Starkov (2024) analyze a very similar model, focusing on the persuasion benchmark as well as cheap talk. They also study a three-period model but also have some partial results for the infinite horizon case. We focus on disclosure, as our interest is in understanding how and why consumers might choose not to disclose their experiences. By doing so, we also uncover tight comparative-statics implications on the nature and degree of equilibrium communication; our results regarding non-monotone disclosure (Theorem 2 and Corollary 1) have no analog under alternative forms of information transmission.

The paper proceeds as follows. After introducing the model (Section 2), we perform a belief-based analysis of optimal disclosure, abstracting from the details of the communication technology. Armed with these insights, we fully characterize equilibrium disclosure in Section 4, and contrast it to the persuasion (commitment) benchmark in Section 5. We discuss our modeling choices as well as other extensions in Section 6, and conclude with thoughts on future research in Section 7. Unless otherwise mentioned, proofs are gathered in the Appendix.

2 Model

Players and signals – At each date $t = 1, 2, 3$, a short-lived agent arrives and takes a binary decision $a_t \in \{0, 1\}$, corresponding to safe and risky actions respectively. The safe action generates a payoff 0. The risky action incurs a cost $c \in (0, 1)$ and generates a payoff $x \sim F_\theta$ supported on $[0, 1]$ that depends on a hidden state $\theta \in \{0, 1\}$. We shall often refer to $a_t = 1$ as "consuming". Conditional on receiving outcome x, the agent may then have the opportunity to provide feedback regarding their experience via direct communication. As we will vary the precise form of communication available, we defer providing further details. We make the following standard assumptions on F_θ (Smith et al., 2021):

Assumption 1. ***(1.a)*** *F_L, F_H are differentiable and mutually absolutely continuous with common, convex support $X = [0, 1]$ and densities f_H, f_L satisfying the monotone likelihood ratio property (MLRP).*

(1.b) $\inf_x \left(\frac{f_L}{f_H}\right) = 0$, $\sup_x \left(\frac{f_L}{f_H}\right) = \infty$.

(1.c) $\mathbb{E}(x \mid \theta = H) = 1$ *and* $\mathbb{E}(x \mid \theta = L) = 0$.

Assumption **(1.a)** states that higher signals are more likely in the high state, and that no perfectly revealing signal exists in either state. Assumption **(1.b)** is the "unbounded beliefs" assumption of Smith and Sørensen (2000), stating that there always exists a signal strong enough to almost completely overturn any prior belief. Assumption **(1.c)** is a normalization ensuring that beliefs and expected payoffs coincide, i.e. $\mathbb{E}(x \mid p) = p$, and is made simply for algebraic convenience. As in Smith et al. (2021), we further assume that the distributions of the log-likelihood ratio of signals are log-concave. This ensures an intuitive feature of belief updating known as "posterior monotonicity" (PM) holds under Bayesian updating.

Assumption 2. *Let $\phi_\theta(l)$ denote the state-contingent densities for the transformed variable $l = \log(x/(1-x))$. Then $\phi_\theta(\cdot)$ is log-concave for $\theta \in \{0, 1\}$.*

Let p^x denote the posterior belief formed by combining the belief p with the outcome $x \in [0, 1]$. That is,

$$p^x \equiv \frac{pf_H(x)}{f_p(x)} \equiv \frac{pf_H(x)}{pf_H(x) + (1-p)f_L(x)} \quad \text{for } x \in X. \tag{1}$$

Note that for all $p \in (0, 1)$, $p^x = p$ if and only if $f_H(x) = f_L(x)$. Let $\hat{x}$ denote the "neutral" signal that satisfies this equality, and more generally, let $x(p, q)$ solve $p^{x(p,q)} = q$, i.e. it is the signal required to achieve posterior q starting from prior p. Both $\hat{x}$ and $x(p, q)$ are guaranteed to exist and be unique for all $p, q \in (0, 1)$ by Assumption 1.

We will sometimes use a natural transformation from signal space X into belief space $[0, 1]$. Namely, we denote by G the distribution over posterior beliefs induced by the signal distribution: for each $p, q \in [0, 1]$, let $G_p(q) \equiv F_p(x(p, q))$. By the absolute continuity assumption (**1.b**), G_p is continuously differentiable and admits a density g_p.

Payoffs – Each agent values the payoffs to both themselves and future agents, but very differently. We assume a form of lexicographic preferences, in which players care infinitely more about their own consumption than that of any other consumer.[6] On the other hand, once their consumption choice has been made, they value the welfare of future consumers equally. Formally, given a prior belief p_t, agent t chooses a_t to maximize their expected consumption payoff $a_t(p_t - c)$, so that $a_t(p_t) = \mathbb{I}_{p_t \geq c}$, whereas given $a_t = 1$ and the signal x thereby obtained, they provide feedback to maximize

$$V_t(p_t^x \mid p_t) \equiv \mathbb{E}\left[\sum_{s=t+1}^{3} a_s(\theta - c) \mid p_t, x\right] = \mathbb{E}\left[\sum_{s=t+1}^{3} \mathbb{I}_{p_s \geq c}(\theta - c) \mid p_t, x\right].$$

[6]We discuss this assumption in Section 6.

3 The Value of Communication

Before beginning our analysis of equilibrium disclosure, we provide an initial, belief-based approach. Since we focus on ex-post disclosure without commitment, we will need to account for divergent private and public beliefs on the equilibrium path of play. Let $u(r \mid q) \equiv \mathbb{I}_{r \geqslant c}(q - c)$ be the utility, as judged by an agent with private belief q, that a successor with belief r will derive from their own consumption decision. Next, let $V_t(r \mid q)$ denote the present value to player t if they hold private belief q, and the public continuation belief in period $t+1$ is r.

We will restrict our attention throughout the paper to equilibria in which P2 fully reveals. This is natural for several reasons. First, it is immediate that truthful revelation is weakly dominant for P2. To see this, note that $V_2(r \mid q) = u(r \mid q) = \mathbb{I}_{r \geq c}(q - c)$ is constant over $r \geq q$, and maximized at $r = q$. Second, we will show in Online Appendix A that truthful revelation by P2 is strictly dominant in the presence of (possibly arbitrarily small) shocks to players' payoffs, and is thus uniquely selected by an argument of robustness to such perturbations.

With this restriction in mind, we can write the value function for P1 in the following succinct form:

$$V_1(r \mid q) = \begin{cases} u(r \mid q) + \alpha \Lambda(r \mid q) + (1 - \alpha) u(r \mid q) & \text{if} \quad r \geqslant c \\ 0 & \text{if} \quad r < c, \end{cases} \tag{2}$$

where

$$\Lambda(r \mid q) \equiv \mathbb{E}\left[u(r^z \mid q^z)\right] = \int_{x(r,c)}^{1} (q^z - c)\, f_q(z)\, dz$$

denotes the expected consumption value of P3 from P1's perspective, given that P1 holds private belief q and that P2 both holds belief r and consumes. Basic algebra confirms that:

$$\Lambda(r \mid q) = q(1 - F^H(x(r,c)))(1 - c) + (1 - q)(1 - F^L(x(r,c)))(-c). \tag{3}$$

When P1 holds belief q and P2 belief r, P1 believes the state is high with probability q and that P3 will consume with probability $1 - F^H(x(r, c))$, receiving a payoff $1 - c$, and similarly when the state is low.

3.1 Experimentation Versus Accuracy

The function $\Lambda(\cdot \mid \cdot)$ is crucial in determining P1's preferences for strategic disclosure. The following lemma therefore provides a complete characterization of its key properties.

Lemma 1. *1.* $r \mapsto \Lambda(r \mid q)$ *is strictly increasing on* $[0, q)$ *and strictly decreasing on* $(q, 1]$.

2. $q \mapsto \Lambda(r \mid q)$ *is strictly increasing (and in particular affine) for all* $r \geqslant c$.

3. $\Lambda(c \mid c) > 0$.

Most importantly, $r \mapsto \Lambda(r \mid q)$ is single-peaked at q. Thus, Λ measures the loss (from P1's perspective) from inducing an incorrect belief on P2 — it increases the likelihood that P3 consumes when they shouldn't, or doesn't consume when they should. That $q \mapsto \Lambda(r \mid q)$ is increasing simply reflects that, for any given profile of consumption choices, P1 is better off holding a higher belief. Finally, that $\Lambda(c \mid c) > 0$ quantifies the *option value* from consumption; note that $u(c \mid c) = 0$, so while the immediate return from P2 consuming at belief c is 0, the gain to P3 from such consumption is strictly positive, as there is a chance P2 receives a positive outcome, acquiring useful information and thus providing an expected gain to P3. Since $\Lambda(c \mid c) > 0$, by continuity $\Lambda(c \mid q) > 0$ for q in some neighborhood below c. However, since inducing a belief $r < c$ leads to non-consumption, P1 would rather suffer the loss in accuracy than terminate consumption when their belief is just below q.

We can leverage this structure to characterize V_1. To do so, it is often convenient to study the "relaxed" value function

$$W_1(r \mid q) \equiv q - c + \alpha\Lambda(r \mid q) + (1 - \alpha)(q - c), \tag{4}$$

which denotes P1's value given a continuation belief r and private belief q, *assuming* that

P2 consumes and P3 consumes if they do not see a signal. That is, $V_1(r \mid q) = \mathbb{I}_{r \geq c} W_1(r \mid q)$. From equation (4), it is straightforward to demonstrate that W_1 obtains the same properties as Λ, since u also preserves these properties.

Lemma 2. *1. $r \mapsto W_1(r \mid q)$ is strictly increasing on $[c, q)$ and strictly decreasing on $(q, 1]$.*

2. $q \mapsto W_1(r \mid q)$ is strictly increasing (and affine) on $[0, 1]$.

3. $W_1(c \mid c) > 0$.

Lemma 2 provides a comprehensive characterization of P1's preferences over which beliefs to induce, conditional upon receiving a given signal. In particular, it uncovers when the trade-off between fostering experimentation and minimizing consumption errors is active and, when so, how it is resolved. First, if $q \geq c$, then the trade-off is inactive and P1 prefers to truthfully reveal; experimentation is guaranteed and consumption errors are minimized by doing so. Second, if $q < c$ and P1 wants to foster experimentation then they would prefer to do so by inducing a belief c, as this is the closest belief to theirs that induces experimentation, thus minimizing consumption errors. Third, if P1 finds it optimal to foster experimentation when $q < c$, then they find it optimal to do so at all $q' \in [q, c)$. Intuitively, the closer is P1's posterior to c from below, the greater is the option value in having P2 experiment.

4 Missing Feedback – Hard Evidence Disclosure

The results in Section 3.1 provide merely a guide to P1's preferences. We now proceed to analyze optimal disclosure rules to uncover the extent to which the ex-post constraints imposed by strategic information transmission limit the degree to which these preferences can be respected. We adopt hard evidence (verifiable) disclosure (Dye, 1985a; Jung and Kwon, 1988), wherein P1: (i) with probability α, is able to freely disclose their signal x, and chooses whether or not to do so; (ii) with probability $1-\alpha$, has no such opportunity, due for instance to a prohibitively high disclosure cost.[7] In this context, and since P2 is assumed to fully disclose (see Section 3) the disclosure rule is a function $d : X \times [0,1] \to \{0,1\}$, where $d(x,p) = 1$ denotes disclosure of signal x at prior p. Note that we keep the dependence on the prior for ease of notation, given the comparative statics exercise we perform in Section 4.3.

Since disclosure is verifiable, the following statements are immediate. First, if a signal is disclosed, it is simply combined with the current belief according to Bayes' rule (1). Second, if a realized signal x is not disclosed, then the update rule must account for all other signals at which non-disclosure also occurs, as well as the possibility that disclosure was not feasible. For disclosure rule d, let $\mathcal{D}(p,d) = \{x \in X \mid d(x,p) = 1\}$, $D(p) \equiv \mathcal{D}(p,d)$, and $N(p) \equiv \mathcal{D}(p,d)^c$.[8] We have:

$$p^{\emptyset} \equiv \frac{\mathbb{P}(d = \emptyset \mid \theta = \theta_H)}{\mathbb{P}(d = \emptyset)} = \frac{(1-\alpha)p + \alpha \int_{N(p)} p^x f_p(x)dx}{(1-\alpha) + \alpha \int_{N(p)} f_p(x)dx}. \tag{5}$$

The relevant incentive compatibility (IC) constraint for the disclosure choice by P1 is then: for all $x \in X$, $d(x,p) = 1$ if and only if

$$V_1(p^x \mid p^x) \geqslant V_1(p^{\emptyset} \mid p^x). \tag{6}$$

An equilibrium is simply a disclosure rule d for P1 such that: 1) given the non-disclosure belief $p^{\emptyset}$, d is incentive compatible, and 2) given d, $p^{\emptyset}$ is correctly computed:

Definition 1. An equilibrium is a disclosure rule d such that (5) and (6) are satisfied.

4.1 Full Disclosure

We begin with a trivial but important observation, namely that when ($\alpha = 1$), the *only* equilibrium outcome is full revelation.[9] The logic is as follows. First, we show that for any

[7] In a slight variant of the model, the arrival of agents is random, occurring with probability α in each period.

[8] For any $p < c$, P_1 abstains from consuming ($a_1 = 0$) and thus has no signal to report, making $D(p)$ irrelevant. In what follows we will therefore focus on values $p \geq c$.

[9] Note that were $\alpha < 1$ but reviewing opportunities observable, the result would also apply.

$\alpha \in (0,1]$, if a signal x is acquired with $p^x \geq c$, then it is disclosed. Intuitively, such types have a preference for truthful disclosure since they face no tradeoff between accuracy and experimentation, as explained following Lemma 2, and thus the hard-evidence constraint leads to a simple unraveling argument. Equation (5) then tells us that it is impossible to foster experimentation via strategic non-disclosure, as non-disclosure will necessarily induce a belief below c.[10]

Lemma 3. *(Positive Selection) If $p^x \geqslant c$, then $d(x,p) = 1$ is a strictly dominant strategy.*

Corollary 1. *If $\alpha = 1$, then all equilibria are outcome-equivalent to full disclosure.*

4.2 Positively Biased and Polarized Disclosure

Having seen that when disclosure is always feasible this leads to unraveling, we turn to the more realistic setting in which reviewing opportunities are random. We focus on a particular equilibrium, namely the *maximal experimentation equilibrium* (MEE). To justify this choice, we introduce various relevant concepts, the first of which is the *experimentation region* of an equilibrium d, denoted by $X_E(d)$:

Definition 2. For an equilibrium d, let

$$X_E(d) = \{x \in X \mid x \in N(p) \text{ and } a_2(p^\emptyset) > a_2(p^x)\}.$$

A signal $x \in X_E(d)$ if under equilibrium d, player 1 chooses not to disclose it, and by so doing *induces* player 2 to consume with strictly higher probability.[11]

Definition 3. An equilibrium d is an *experimentation equilibrium* (EE) if $X_E(d)$ has strictly positive (Lebesgue) measure. Let $\mathcal{E}$ denote the set of all such equilibria. An equilibrium d is a *maximal experimentation equilibrium* (MEE) if $d \in \mathcal{E}$ and $d' \in \mathcal{E}$ implies $X_E(d') \subset X_E(d')$.

[10] Of course, indifference allows for non-disclosure, but this would lead to identical outcomes.

[11] In the current baseline setting, this last point amounts to player 2 consuming at $p^\emptyset$ but not at p^x. In Section A of the Online Appendix, we study an extension with idiosyncratic preference shocks that will make both probabilities interior.

Thus, a MEE contains the largest experimentation region out of all equilibria. The first main result we now establish is that when P1's prior is intermediate, any MEE exhibits both *polarity bias* (only sufficiently low and sufficiently high signals are disclosed) and *positively selected disclosure* (non-disclosure is bad news). We next formalize these notions:

Definition 4. A disclosure rule d is:

1. *Polarized*if there exist $\underline{\varepsilon}, \overline{\varepsilon} > 0$ such that $d(p, x) = 1$ for all $x \in [0, \underline{\varepsilon}) \cup [1 - \overline{\varepsilon}, 1]$ and $N(p)$ has strictly positive measure.

2. *Positively selected at p* if

$$\mathbb{E}(p^z \mid z \in N(p)) \equiv \frac{\int_{N(p)} p^z \, dF_p(z)}{\int_{N(p)} dF_p(z)} \leqslant p. \tag{7}$$

Lemma 3 tells us that disclosure occurs in equilibrium for all signals such that $p^x \geq c$, i.e. for sufficiently high signals. on the other hand, Lemma 2 tells us that if non-disclosure occurs, it will do so on a single interval. Finally, it is easily shown that for sufficiently low x, P1 strictly prefers to disclose. Intuitively, after such a bad experience, P1 holds little hope that the product is good and thus prefers to cease experimentation immediately. This logic gives rise to our first main result:

Theorem 1. *(Interval non-disclosure) In any EE, player 1 adopts the disclosure strategy:*

$$d(x, p) = \begin{cases} 1 & \text{if} \quad p^x \geqslant c \\ 0 & \text{if} \quad p^x \in [\underline{q}(p), c) \\ 1 & \text{if} \quad p^x < \underline{q}(p), \end{cases}$$

for some $\underline{q}(p) \in (0, c]$.

In equilibrium, $P1$ discloses only those signals that lie on either side of the interval $[x(p, \underline{q}(p)), x(p, c)]$, thus exhibiting both polarity and positive selection (since $x(p, c) \leq x(p, p) = \hat{x}$). P1 thus thinks along the following lines. If they are sufficiently confident

Figure 1: Equilibrium Non-Disclosure as prior p varies.

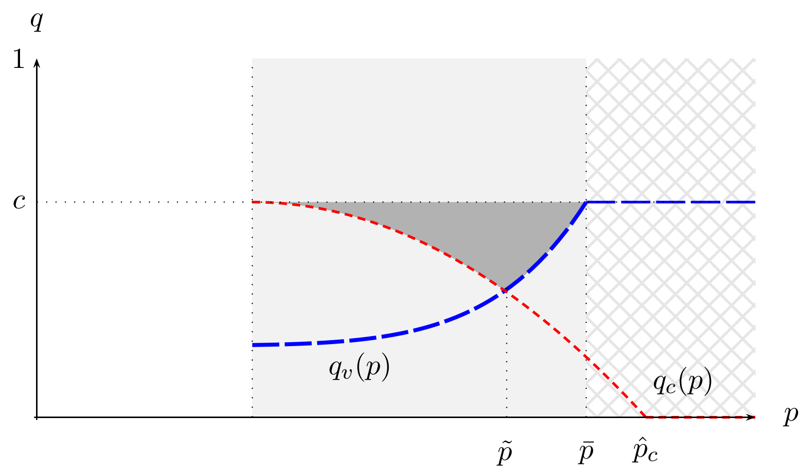

Non-disclosure in the MEE, as the prior p varies. Dark-shaded region: non-disclosed posterior beliefs; light-shaded region: disclosed posterior beliefs; hatched region: full disclosure. Long-dashed line: incentive constraint, $q_v(p)$; short-dashed line: belief constraint, $q_c(p)$.

regarding the product's quality, they should let others know. But if they are not, they would rather not take the risk of ceasing further experimentation prematurely and thus should keep their opinions to themselves.

4.3 U-shaped Disclosure with respect to prior

Theorem 1 tells us that non-disclosure occurs on an interior interval of signals. But how does this interval depend on primitives, such as P1's prior belief p, and the disclosure parameter α? In this section, we show how the equilibrium disclosure threshold $\underline{q}$ varies with the prior belief p. One can view this exercise as comparing across products that differ in how well-established they are. For instance, if p is close to 1 then the product is well-established, while for p close to c, the product is close to exit. For p in the interior of this region, products can be viewed as novel. The following result, illustrated in Figure 2, summarizes these findings and constitutes our second main result:

Theorem 2. *There exist unique thresholds $\tilde{p}, \bar{p} \in (c, 1)$ and functions $p \mapsto q_c(p), q_v(p)$ on $[c, 1]$, with $p \mapsto q_c(p)$ strictly decreasing and $p \mapsto q_v(p)$ strictly increasing on $[c, \bar{p}]$, such that:*

1. *If* $p \in [c, \tilde{p}]$ *then setting* $\underline{q}(p) = q_c(p)$ *constitutes a MEE.*

2. *If* $p \in [\tilde{p}, \bar{p}]$*, then setting* $\underline{q}(p) = q_v(p)$ *constitutes a MEE.*

3. *If* $p \in (\bar{p}, 1]$*, no EE exists and all equilibria are outcome-equivalent to full disclosure.*

As will become clear through the subsequent analysis, the functions $p \mapsto q_c(p), q_v(p)$ represent two important constraints on strategic non-disclosure. The first, $q_c(p)$, is a "belief constraint", and arises due to the classic form of unraveling; Lemma 3 shows that $d(p, x) = 1$, occurs for all posterior beliefs $p^x > c$. Non-disclosure by P1 is thus positively selected, and so in order to induce experimentation by P2, P1 cannot conceal signals that are too negative –or else the resultant posterior $p^{\emptyset}$ would drop below c. This constraint tightens as p gets closer to c; the close is p to c, the less room there is for negatively selected non-disclosure to keep $p^{\emptyset} > c$ and thus induce experimentation.

The second, $q_v(p)$, is an "incentive constraint"; it tracks the marginal posterior for P1 at which they are indifferent between disclosing and not, assuming posteriors in $[q_v(p), c)$ are concealed in equilibrium. That $q_v(p)$ is increasing constitutes yet another expression of the experimentation-accuracy trade-off; in particular, when p is close to 1, the continuation belief r is also close to 1, *regardless* of P1's disclosure strategy. As such, non-disclosure plays no role in altering P2's experimentation incentives, and in fact would only induce P3, being less well-informed, to make more mistakes in their consumption choice. Thus, despite there being a range of beliefs q just below c at which P1 would prefer to induce experimentation, they would suffer too great a loss through increased consumption errors by P3 to warrant doing so.

To establish Theorem 2, we will use a series of lemmas characterizing beliefs following non-disclosure and the functions $q_c(p), q_v(p)$. Let $\phi(p, q)$ denote the continuation public belief if signals in the range $[x(p, q), x(p, c))$ are not disclosed:

$$(q \leqslant c): \quad \phi(p, q) = \frac{(1-\alpha)p + \alpha \int_{x(p,q)}^{x(p,c)} p^z dF_p(z)}{(1-\alpha) + \alpha \int_{x(p,q)}^{x(p,c)} dF_p(z)}, \qquad (q > c): \quad \phi(p, q) = p. \qquad (8)$$

We showed in Theorem 1 that, if non-disclosure happens, it is over an interval of exactly this type, so $\phi(p, q)$ is indeed the relevant computation for the equilibrium belief $p^{\emptyset}$ following non-disclosure.

Lemma 4. *1. For $p \geqslant c$, $q \mapsto \phi(p, q)$ is strictly increasing and differentiable on $[0, c]$, with $\phi(p, 0) \in (0, p)$ and $\phi(p, c) = p$.*

2. For $q \leq c$, $p \mapsto \phi(p, q)$ is strictly increasing on $[c, 1]$.

Intuitively, Lemma 4 part i) tell us that, fixing a prior p, if the lowest marginal concealed signal increases, so too will the average concealed belief (since the upper marginal belief is fixed at $x(p, c)$), and thus so will the continuation belief. The proof of part ii) relies on the "posterior monotonicity" property that Smith et al. (2021) show is equivalent to Assumption 2. Intuitively, fixing the interval non-disclosure rule, increasing P1's prior leads to P2 holding a higher belief, conditional on non-disclosure.

The Belief Constraint – We can now use Lemma 4 to draw out a key constraint bearing on strategic non-disclosure: as P1's prior declines toward c, the maximal interval of signals they could conceal and still induce experimentation *shrinks*. Indeed, when P1 is more pessimistic to begin with, P2 is more easily dissuaded from consuming. To formalize this intuition let, for $p \geqslant c$,

$$q_c(p) = \inf \{q \in [0, 1] \mid \phi(p, q) \geqslant c\}. \tag{9}$$

Lemma 4 and Corollary 4 tell us that $p \mapsto q_c(p)$ is well-defined on $[c, 1]$. We now establish key properties of the function $q_c(p)$, illustrated by the dashed line in Figure 2:

Lemma 5. *The map $p \mapsto q_c(p)$ is everywhere continuous, with $q_c(c) = c$. Furthermore, there exists $\hat{p}_c \in (c, 1)$ such that: (i) on $[c, \hat{p}_c]$, $q_c(p)$ is strictly decreasing, differentiable and solves $\phi(p, q_c(p)) = c$; (ii) on $[\hat{p}_c, 1]$ $q_c(p) = 0$.*

The Incentive Constraint — Having characterized the "belief constraint" $q_c(p)$ bearing on P1's disclosure rule, we next turn to the second key constraint involved: an "incentive constraint" $q_v(p)$ that identifies the marginal signal at which P1 is indifferent between disclosing or not, assuming that experiences between that level and $x(p,c)$ are also concealed (Theorem 1 demonstrated that this is necessarily the form of non-disclosure in an EE).

To begin, we demonstrate that disclosure is strictly optimal after sufficiently extreme signal realizations. We do so by proving a property of the relaxed value function $W_1(r \mid q)$ defined in Section 3.1.

Lemma 6. *For all* $p \in [c, 1)$,

$$\lim_{q \to 0,1} \left[W_1(\phi(p,q) \mid q) - V_1(q \mid q)\right] < 0.$$

Away from these limits, we will show that there can exist signals that render P1 indifferent between disclosing and not. This indifference condition defines the incentive constraint: for $p \geqslant c$, let

$$\hat{q}(p) \equiv \inf\left\{q \in [0,1] \mid W_1\left(\phi(p,q), q\right) = V_1(q \mid q)\right\} \tag{10}$$

$$q_v(p) = \min\{c, \hat{q}(p)\}. \tag{11}$$

The minimization defining $\hat{q}(p)$ is well defined, as the constraint is always satisfied at $q = p$. Thus, $\hat{q}(p)$ defines the lowest belief at which the above indifference is satisfied, while $q_v(p)$ simply stores the value c whenever this belief is greater than c. To show that $\hat{q}(p)$ is strictly increasing in p, as illustrated by the solid line in Figure 2, we first show that as p increases, *all* solutions (in q) to $W_1\left(\phi(p,q), q\right) = V_1(q \mid q)$ decrease (and thus so does the smallest one, $\hat{q}(p)$).

Lemma 7. *The map* $p \mapsto \hat{q}(p)$ *is strictly increasing, with* $\hat{q}(c) < c$ *and* $\hat{q}(p) > 0$ *for all* $p \geqslant c$*. Furthermore, there exists* $\bar{p} \in (c,1)$ *such that* $\hat{q}(p) > c$ *if and only if* $p > \bar{p}$.

Combining the proven properties of $q_c(\cdot)$ and $\hat{q}(\cdot)$ immediately shows that the two loci cross at a unique interior point:

Lemma 8. *There exists* $\tilde{p} \in (c,1)$ *such that* $\hat{q}(p) \leqslant q_c(p)$ *if and only if* $p \leqslant \tilde{p}$.

Figure 2: Equilibrium Non-Disclosure

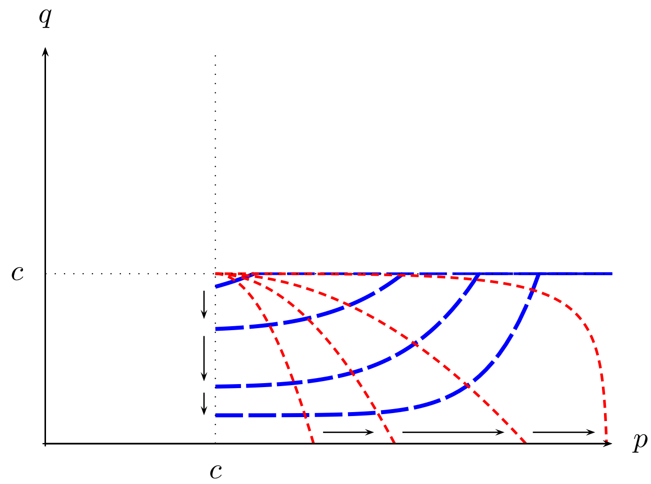

Non-disclosure as α varies. As in Figure 1, long-dashed lines are incentive constraints, $q_v(p)$, short-dashed line are belief constraints, $q_c(p)$. The former move up with α, while the latter move down with α. Black arrows: direction in which constraints move as α increases.

The proof of Theorem 2 concludes by setting $\underline{q}(p) \equiv \max\{q_c(p), q_v(p)\}$ for all $p \geq c$.

4.4 Subsidizing Disclosure

Theorem 2 uncovers the complex relation between the degree of non-disclosure and the prior belief p, due to the two key constraints $q_c(p), q_v(p)$ working against each other. But how are these constraints themselves determined by the ability to disclose, α? We first show that the belief constraint q_c is increasing in α, converging pointwise to $q_c(p) = c$ for all $p < 1$ as $\alpha \to 1$; intuitively, a higher α leaves less room for strategic (non)disclosure. We then show that the incentive constraint q_v is decreasing in α, converging pointwise to $q_v(p) = c$ for all $p < 1$ as $\alpha \to 0$. Intuitively, a higher α leads to fewer mistakes by P3 simply by virtue of disclosure opportunities being more probable, and thus P1 is more willing to strategically non-disclose. (See Figure 2.) As a result, for a fixed prior p, the experimentation region is non-monotone in α, converging to the empty set when α tends to 0 or 1. (See Figure 3)

Lemma 9. *1. For fixed $p \in (c, 1)$, there exists $\hat{\alpha}(p) \in (0, 1)$ such that $q_c(p)$ is strictly increasing in α for $\alpha \in [\alpha(\hat{p}), 1]$ and $q_c(p) = 0$ otherwise. Furthermore, $\lim_{\alpha\to 1} q_c(p) = c$.*

Figure 3: Non-Disclosure as α varies

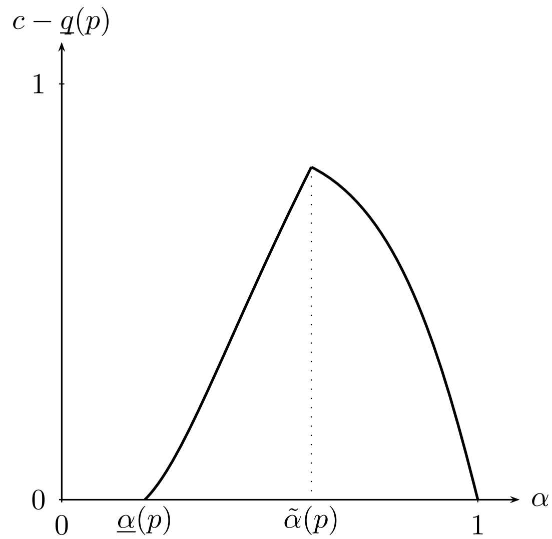

Non-disclosure as α varies for fixed p.

2. For fixed $p \in (c, 1)$, $\hat{q}(p)$ is strictly decreasing in α. Furthermore, there exists $\underline{\alpha}(p) > 0$ such that for all $\alpha \in [0, \underline{\alpha}]$, $\hat{q}(p) = p$.

Corollary 2. *For fixed $p \in (c, 1)$, there exists $\underline{\alpha}(p) > 0$ such that for all $\alpha \in [\underline{\alpha}(p), 1]$, $q_v(p)$ is strictly decreasing in α and for all $\alpha \in [0, \underline{\alpha})$, $q_v(p) = c$.*

Finally, we can put these results together to demonstrate that the region of experimentation $[\bar{q}, c]$ is non-monotone in α.

Proposition 1. *Fix $p \in (c, 1)$. Then there exist $0 < \underline{\alpha}(p) < \tilde{\alpha}(p) < 1$ such that $c - \hat{q}$ is: (1) equal to 0 for all $\alpha \in [0, \underline{\alpha})$; (2) strictly increasing for all $\alpha \in [\underline{\alpha}, \tilde{\alpha}(p))$; (3) strictly decreasing for all $\alpha \in [\tilde{\alpha}, 1)$.*

5 Optimal Feedback – Persuasion

We now turn to the benchmark wherein P1 can commit to an arbitrary messaging rule prior to receiving their private signal x (Kamenica and Gentzkow, 2011). Formally, P1 chooses an *information structure*, consisting of a message space $\mathcal{S}$ along with a collection of conditional probabilities $(\pi(\cdot \mid x))_{x \in [0,1]}$, where $\pi(s \mid x)$ denotes the likelihood of P1 sending the message

s given that they received signal x. Let $\mathcal{M} = [0,1] \cup \{\emptyset\}$ denote the (rich) message space that naturally associates messages with outcomes, as well as a privileged message $\emptyset$ that denotes no signal reported. We may take $\mathcal{S} = \mathcal{M}$. Since communication is no longer constrained to be verifiable, we can set $\alpha = 1$ without loss of generality. Contrasting this case with that of hard-evidence disclosure will thus shed light on how ex-post IC constraints shape optimal feedback. Recently developed techniques in the persuasion literature allow us to completely characterize the solution (Dworczak and Martini, 2019). Denote $V_1(q \mid q)$ by $V_1(q)$ for simplicity.

Proposition 2. *There exist $\underline{q}^*(p) < c < \bar{q}^*(p)$ such that the solution to the persuasion problem takes the following form: reveal x if either $p^x < \underline{q}^*(p)$ or $p^x \geqslant \bar{q}^*(p)$, and pool all x such that $p^x \in [\underline{q}^*(p), \bar{q}^*(p))$. Furthermore, $\underline{q}^*(p), \bar{q}^*(p)$ solve*

$$\mathbb{E}_p(q \mid q \in [\underline{q}^*(p), \bar{q}^*(p))) \equiv \frac{\int_{\underline{q}^*(p)}^{\bar{q}^*(p)} q \, dG_p(q)}{\int_{\underline{q}^*(p)}^{\bar{q}^*(p)} dG_p(q)} = c, \tag{12}$$

and

$$\frac{V_1(\bar{q}^*(p))}{V_1(c)} = \frac{\bar{q}^*(p) - \underline{q}^*(p)}{c - \underline{q}^*(p)}. \tag{13}$$

Communication under persuasion is also both polarized (pooling takes place on an interior interval) and positively selected (the average belief conditional on pooling is c, which is less than p). However, under persuasion, posteriors above c can be non-disclosed ($\bar{q}^*(p) > c$), whereas under ex-post disclosure, unraveling above c prevents this from happening (Lemma 3). This contrast clearly highlights the role of ex-post incentives. When p is close to either c or 1, the constraints $q_c(p)$ and $q_v(p)$ bind and completely undo P1's ability to foster experimentation (Theorem 2). Without the ex-post IC constraints, P1 is able to non-disclose above c, thereby relaxing these constraints and enabling P1 to credibly non-disclose and foster experimentation. The following result, analogous to Theorem 2, formalizes this logic.

Corollary 3. *Both $\underline{q}^*(p)$ and $\bar{q}^*(p)$ are strictly decreasing in p. Furthermore, $lim_{p \to c,1} \underline{q}^*(p) < c < \lim_{p \to c,1} \bar{q}^*(p)$.*

Figure 4: Persuasion Solution

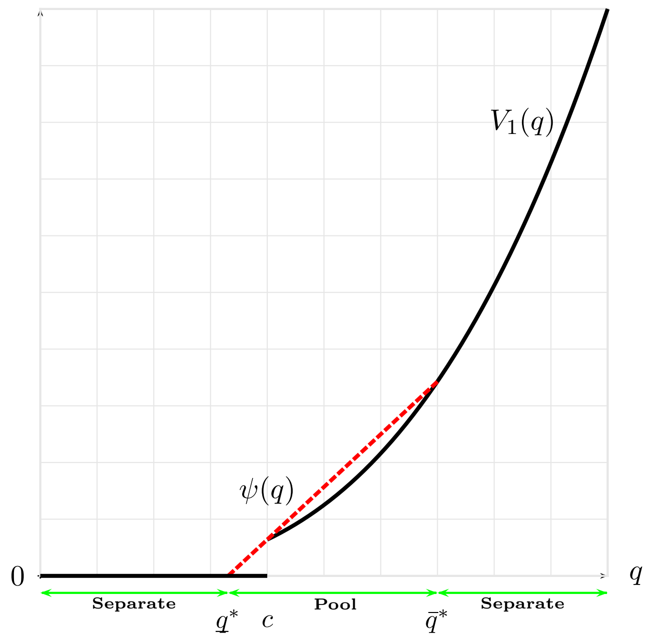

Disclosure under commitment. Value function $V_1(q) \equiv V_1(q \mid q)$ solid black lines. $\underline{q}^*, \bar{q}^*$ are determined by both $\mathbb{E}(q \mid q \in [\underline{q}^*, \bar{q}^*)) = c$ and lying on a straight-line segment $\psi(q)$ (dotted red) intersecting $V_1(q)$ at $\underline{q}^*, \bar{q}^*$ and c.

Finally, notice that the persuasion outcome — which did not assume information to be verifiable – can be implemented via commitment to the verifiable disclosure rule

$$d(x,p) = \begin{cases} 1 & \text{if} \quad p^x \geqslant \bar{q}^*(p) \\ \emptyset & \text{if} \quad p^x \in [\underline{q}^*(p), \bar{q}^*(p)) \\ 1 & \text{if} \quad p^x < \underline{q}^*(p). \end{cases}$$

This is due to the simple structure of optimal persuasion; it is not only monotone partitional (Dworczak and Martini, 2019), but includes only one pooling region (see Figure 4). Thus, the pooling region can be interpreted as non-disclosure and the separating regions as disclosure, satisfying the verifiability assumption. In this sense, the benefit of persuasion over (ex-post) verifiable disclosure comes directly from which posteriors (signals) are able to be non-disclosed.

6 Discussion

6.1 Model Discussion

Our model differs from standard social learning settings in some important ways. First, whereas typically agents' private signals are hidden while their consumption choices are public (Banerjee, 1992; Bikhchandani et al., 1992), here it is the reverse, in the sense that (some) private signals are publicly disclosed.[12] Second, whereas typically agents receive a private signal prior to making a consumption choice, here our agents can only receive their signal if they consume. In our setting, the link between actions and private signal acquisition is crucial; were future agents to receive (and then disclose) private signals regardless of their action choices, current agents would never seek to distort these choices by withholding information, and thus full revelation would be (weakly) dominant. We thus view our paper as belonging more to the literature on experimentation.

The three-period horizon on which we focus provides the simplest, most transparent setting in which strategic (non-)disclosure for purposes of inducing experimentation will arise. With an infinite horizon, each agent's disclosure choice would need to internalize not only its impact on future experimentation decisions but also its impact on future disclosure choices, each affecting all subsequent ones, etc. This "induced chain of strategic disclosures" aspect renders the problem analytically intractable, and it also does not seem so empirically relevant to natural applications such as consumer product reviews, employees rating their firm, etc. The three-period setup, in contrast, allows for sharp and distinctive predictions on what signals will be disclosed or hidden, and how these regions vary with prior beliefs.

In Section 4 we modeled hard-evidence disclosure through a binary choice of whether or not to disclose the exact signal received. It is almost immediate that broadening the feasible messages to allow agents to report *any interval* that contains their true signal does not affect these results. To see this, note first that for $q \geq c$, $r \mapsto V_1(r \mid q)$ is single-peaked

[12] Wolitzky (2018) also studies a social learning model with unobservable actions, but without strategic disclosure.

at q, and thus Lemma 3 continues to hold since each type has a strict incentive to separate. Next, note that for $q < c$, deviating to an off-path message results in a payoff of 0, since all types above c fully separate and are thus uniquely identified in equilibrium.

6.2 Rewarding and Punishing Sellers

We now demonstrate yet another interpretation of our framework and results. Consider a firm selling a product of unknown quality θ, that is free to produce. Consumers arrive in sequence and can either purchase the product (risky action) or not (safe action), with consumption payoffs $\theta - c$ for some $c > 0$ and 0 respectively.

The firm posts a take-it-or-leave-it price offer, but is constrained by limited liability, so that it cannot set a negative price. That is, if the prevailing public belief is p, then if $p \geq c$ the seller's price is $p - c$ and the consumer purchases, whereas if $p < c$, trade is infeasible as it would violate limited liability, and so the consumer does not purchase. Now suppose that at the point of disclosure, the first consumer's objective is *total firm profit*

$$\Pi_t = \mathbb{E}\left[\sum_{s=2}^{3} a_s(p_s - c) \mid p, x\right].$$

Notice that this is precisely the same objective as given in Section 2.

In short, limited liability prevents the firm from internalizing the full social surplus when setting prices, and therefore the fundamental externality present in the baseline analysis remains here: consumers with beliefs just below c do not consume, despite there being positive social value in doing so due to informational externalities.

The same analysis remains applicable, but can now be reinterpreted in terms of consumers dictating the firm's future revenue. That is, after a positive enough experience ($q > c$) they are happy to disclose, thus "rewarding" the firm with continued demand. After a sufficiently bad experience ($q < \underline{q}(p)$) they disclose, thereby "punishing" it with zero future revenue. For mildly negative experiences, the consumer is not sufficiently unhappy to kill the firm's prospects, and thus strategically non-discloses to encourage further experimentation. Of course, one could formulate the notion of "reward and punishment" in ways that go beyond our framework. For instance, a consumer's desire to affect the fate of the firm may be driven by reference-dependence — the greater is the difference $|p^x - p|$, the more inclined they are to leave feedback, as in Hui et al. (2024) — or involve asymmetries such as loss aversion, if they are more incensed by bad outcomes than grateful for good ones.

7 Conclusion

We studied a model of strategic information transmission, driven by a tension between selfish consumption and pro-social disclosure. Our analysis sheds light on an important question: when might consumers choose not to leave feedback in order to improve overall welfare? We showed that equilibrium disclosure is necessarily *extreme* and *positively biased*, two well-established empirical regularities found in consumer reviewing behavior. We further showed that disclosure is hump-shaped with respect to both agents' prior and their opportunities for leaving reviews. Taking the prior as a proxy for the age of the product (a higher prior implies an older product, due to learning and ex-post selection), the first result implies full disclosure for well-established products, while the second cautions that making feedback less costly could potentially reduce learning. There are undoubtedly many economic forces at play that govern consumers' incentives to leave reviews, so we view our results as taking an important first step toward understanding a very natural one — that consumers might be pro-socially motivated when leaving feedback.

References

Banerjee, A. V., "A simple model of herd behavior," *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Bénabou, R. and J. Tirole, "Self-Confidence and Personal Motivation," *Quarterly Journal of Economics*, 2002, *117* (3), 871–915.

— **and** — , “Willpower and Personal Rules,” *Journal of Political Economy*, 2004, *112* (4), 848–886.

Bikhchandani, S., D. Hirshleifer, and I. Welch, “A theory of fads, fashion, custom, and cultural change as informational cascades,” *Journal of Political Economy*, 1992, *100* (5), 992–1026.

— **,** — **, O. Tamuz, and I. Welch**, “Information Cascades and Social Learning,” *Working Paper*, 2022.

Cabral, L. and L.I. Li, “A Dollar for Your Thoughts: Feedback-Conditional Rebates on eBay,” *Management Science*, 2015, *61* (9).

Carrillo, J. D. and T. Mariotti, “Strategic Ignorance as a Self-Disciplining Device,” *Review of Economic Studies*, 2000, *67* (3), 529–544.

Che, Y-K. and J. Hörner, “Recommender Systems as Mechanisms for Social Learning,” *Quarterly Journal of Economics*, 2018, *133*, 871–925.

Dworczak, P. and G. Martini, “The Simple Economics of Optimal Persuasion,” *Journal of Political Economy*, 2019, *127* (5).

Dye, R.A., “Disclosure of Nonproprietary Information,” *Journal of Accounting Research*, 1985a, *23*, 123–145.

Grossman, S. J., “The Informational Role of Warranties and Private Disclosure about Product Quality,” *The Journal of Law and Economics*, 1981, *24* (3), 461–483.

Hui, X., T.J. Klein, and K. Stahl, “Learning from Online Ratings,” *working paper*, 2024.

Jung, W. and Y. Kwon, “Disclosure When the Market is Unsure of Information Endowment of Managers,” *Journal of Accounting Research*, 1988, *26*, 146–153.

Kamenica, E. and M. Gentzkow, “Bayesian Persuasion,” *American Economic Review*, 2011, *101*, 2590–2615.

Kremer, I., Y. Mansour, and M. Perry, “Implementing the Wisdom of the Crowd,” *Journal of Political Economy*, 2014, *122* (5), 988–1012.

Laibson, D., “Golden Eggs and Hyperbolic Discounting,” *The Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

Marinescu, Ioana, Andrew Chamberlain, Morgan Smart, and Nadav Klein, "Incentives can reduce bias in online employer reviews," *Journal of Experimental Psychology: Applied*, 2021.

Milgrom, P. R., "Good News and Bad News: Representation Theorems and Applications," *The Bell Journal of Economics*, 1981, *12* (2), 380–391.

Nosko, C. and S. Tadelis, "The Limits of Reputation in Platform Markets: An Empirical Analysis and Field Experiment," *NBER working paper*, 2015.

Rockafellar, R.T., *Convex Analysis*, New Jersey: Princeton University Press, 1997.

Schoenmueller, V., O. Netzer, and F. Stahl, "The Polarity of Online Reviews: Prevalence, Drivers and Implications," *Journal of Marketing Research*, 2020, *57* (5), 853–877.

Slivkins, A., "Exploration and Persuasion," in V.V. Vazirani F. Echenique, N. Immorlica, ed., *Online and Matching-Based Market Design*, Cambridge University Press, 2022.

Smirnov, A. and E. Starkov, "Designing Social Learning," *working paper*, 2024.

Smith, L. and P.N. Sørensen, "Pathological Outcomes of Observational Learning," *Econometrica*, 2000, *68* (2), 371–398.

_ , _ , and J. Tian, "Informational Herding, Optimal Experimentation, and Contrarianism," *Review of Economics Studies*, 2021, *8* (5), 2527–2554.

Wolitzky, A., "Learning from Others' Outcomes," *American Economic Review*, 2018, *108*, 2763–2801.

A Proofs

A.1 Proof of Lemma 1

To prove the first part of the lemma, it suffices to note that $\frac{\partial \Lambda(r|q)}{\partial r}$ is proportional to $q^{x(r,c)} - c$, which has the sign of $q - r$, by the MLRP. The second part follows immediately from equation (3). $(\partial F/\partial q)(x) = (F_H - F_L)(x) < 0$, which is implied by MLRP. The third part follows from

$$\Lambda(c \mid c) = \int_{\hat{x}}^{1} (c^z - c) \, f_c(z) \, dz > 0,$$

since $c^z > c$ for $z > \hat{x}$.

A.2 Proof of Lemma 3

The proof proceeds in two cases. Let $q = p^x$. First, suppose that $p^{\emptyset} < c < q$, so that non-disclosure causes consumption to stop. Then, using (2),

$$V_1(p^{\emptyset} \mid q) = 0 < u(q \mid q) + \alpha\Lambda(q \mid q) + (1 - \alpha)u(q \mid q) = V_1(q \mid q),$$

since by Lemma 1,

$$\Lambda(q \mid q) \geq \Lambda(c \mid q) \geq \Lambda(c \mid c) = \int_{\hat{x}}^{1} (c^z - c) \, f_c(z) \, dz > 0.$$

Next, suppose that $p^{\emptyset} \geqslant c$, so that non-disclosure leads to consumption (and subsequent disclosure by P2) in spite of a lower belief. In this case,

$$V_1(q \mid q) - V_1(p^{\emptyset} \mid q) = \alpha \left[\Lambda(q \mid q) - \Lambda(p^{\emptyset} \mid q) \right] \geq 0,$$

since the first part of Lemma 1 showed that that $r \mapsto \Lambda(r \mid q)$ is maximized at q, for $q \geqslant c$.

A.3 Proof of Lemma 4

For part i), differentiability is clear, and $\partial\phi(p,q)/\partial q$ has the sign of

$$-\frac{\partial x(p,q)}{\partial q} p^{x(p,q)} \left[(1-\alpha) + \alpha \int_{x(p,q)}^{x(p,c)} dF_p(z)\right] + \frac{\partial x(p,q)}{\partial y} \left[(1-\alpha)p + \alpha \int_{x(p,q)}^{x(p,c)} p^z dF_p(z)\right].$$

Given that $q \mapsto x(p,q)$ is increasing by the MLRP and $p^{x(p,q)} \equiv q$, that sign is also that of

$$(1-\alpha)(p-q) + \alpha \int_{x(p,q)}^{x(p,c)} (p^z - q)\, dF_p(z) > 0,$$

since $q \leq c \leq p$ and $p^z \geq q$ for $z > x(p,q)$.The bounds on $q \mapsto \phi(p,q)$ follow immediately.

For part ii), let us first re-write $\phi(p,q)$ as

$$\phi(p,q) = \frac{(1-\alpha)p + \alpha \int_{x(p,q)}^{x(p,c)} p^z\, dF_p(z)}{(1-\alpha) + \alpha \int_{x(p,q)}^{x(p,c)} dF_p(z)} = \frac{(1-\alpha)p + \alpha \int_q^c r\, dG_p(r)}{(1-\alpha) + \alpha \int_q^c dG_p(r)}.$$

Let $a(p) \equiv \int_q^c r\, dG_p(r)$ and $b(p) \equiv \int_q^c dG_p(r)$. By Proposition 4 in Smith et al. (2021), Assumption 2 implies that

$$\frac{d}{dp}\left(\frac{a(p)}{b(p)}\right) > 0.$$

In our case, P_2's not having received a signal may also be due to $P1$ not having had the opportunity to leave feedback, which occurs with probability $1-\alpha$. As a result, $\partial\phi(p,q)/\partial p$ has the sign of

$$\begin{aligned}
&[(1-\alpha) + \alpha b(p)][(1-\alpha) + \alpha a'(p)] - [(1-\alpha)p + \alpha a(p)][\alpha b'(p)] \\
&= (1-\alpha)^2 + (1-\alpha)\alpha a'(p) + \alpha b(p)(1-\alpha) - \alpha(1-\alpha)p b'(p) + \alpha^2 \underbrace{(b'(p)a(p) - b(p)a'(p))}_{>0} \\
&\geqslant \alpha(1-\alpha)(a'(p) - p b'(p) + b(p)).
\end{aligned}$$

But

$$\begin{aligned}
a'(p) - p b'(p) + b(p) &= \frac{\partial}{\partial p}\int_q^c r\, dG_p(r) - p\frac{\partial}{\partial p}\int_q^c dG_p(r) + \int_q^c dG_p(r) \\
&= \int_q^c \left[r\frac{\partial g_p(r)}{\partial p} - p\frac{\partial g_p(r)}{\partial p} + g_p(r)\right] dr = \int_q^c r(g_1(q) - g_0(q)) + g_0(q)\, dr \\
&= \int_q^c r g_1(q) + (1-r) g_0(q)\, dr = \int_q^c g_r(r)\, dr > 0,
\end{aligned}$$

thus proving the claim.

A.4 Proof of Lemma 5

Note that $p \mapsto q_c(p)$ is defined as the minimum of a continuous function, $(q \mapsto \phi(p,q))$, on a compact set. Therefore, the infimum is attained by Weierstrass' Theorem, and continuity follows from Berge's Theorem (note that the constraint $\phi(p,q) \geqslant$ defines an upper-hemicontinuous correspondence, since $\phi(p,q)$ is continuous). That $q_c(c) = c$ follows from Corollary 4, part 2.

Next, that there exists a $\hat{p}_c \in (c,1)$ such that $q_c(p) = 0$ for all $p \in [\hat{p}_c, 1]$ follows from the definition of $\phi(p,q)$, since as $p \to 1$,

$$\phi(p,0) \to \frac{(1-\alpha)\cdot 1 + \alpha \cdot 1}{1-\alpha+0} = 1.$$

Finally, that $p \mapsto q_c(p)$ is strictly decreasing on $[c, \hat{p}_c]$ follows directly from Lemma 4.

A.5 Proof of Lemma 6

The lower limit follows immediately since $\Lambda(r \mid q) < 0$ and $q - c < 0$ for sufficiently small q. The upper limit obtains by noting that as $q \to 1$, $V_1(q \mid q)$ achieves the upper bound on V_1.

A.6 Proof of Lemma 7

For $p \geqslant c$, let $q(p) < c$ be any solution to the equation $W_1(\phi(p,q), q) = V_1(q \mid q)$. From Lemmas 2 and 4 and the chain rule, it follows that if $q(p) < c$, then $q'(p) > 0$. Therefore, $p \mapsto \hat{q}(p)$ must be strictly increasing. For sufficiently high p, $W_1(\phi(p,q) \mid q) = V_1(\phi(p,q) \mid q)$, so the existence of $\bar{p}$ follows from Lemma 3. To see that $\hat{q}(c) < c$, note that $\lim_{q\to c}[W_1(\phi(p,q) \mid c) - V_1(q \mid c)] = W_1(p \mid c) > 0$. Lemma 6 combined with the Intermediate Value Theorem then implies there exists $q' \in (0,c)$ such that $W_1(\phi(p,q') \mid c) = V_1(q' \mid c)$, which is the constraint defining $\hat{q}(c)$, and thus $\hat{q}(c) \leq q' < c$. That $\hat{q}(p) > 0$ follows from the fact that $W_1(\phi(p,0) \mid 0) < 0 = V_1(0 \mid 0)$ and the continuity of $q \mapsto W_1(\phi(p,q) \mid q)$.

A.7 Proof of Lemma 8

Follows from Lemma 5 and Lemma 7 part ii).

A.8 Proof of Theorem 1

To establish the result, note first that following a signal leading P1 to hold a posterior q, their disclosure decision hinges on the sign of $V_1(q \mid q) - V_1(p^{\emptyset} \mid q)$. First, disclosure for $q \geq c$ follows from Lemma 3. Lemma 2 then implies that if non-disclosure occurs in equilibrium, it must take a interval form; $D(p) = [\underline{q}, c]$. As such, if it exists, the MEE under any prior p is (essentially) unique, as it is fully determined by its corresponding $\underline{q}(p)$. Finally, that $\underline{q} > 0$ follows from the fact that $V_1(r \mid 0) < 0 = V_1(0 \mid 0)$ for $r \geq c$, and thus by continuity revealing is strictly preferred to inducing experimentation for sufficiently low q.

A.9 Proof of Theorem 2

We proceed in three cases:

1. If $\hat{q}(p) \in [0, q_c(p))$, then setting $\underline{q} = q_c(p)$ defines the MEE. To see this, note first that the equilibrium belief condition (5) is satisfied by definition. Next, we will verify the IC condition (6), which in this case amounts to $V_1(c \mid q) \geq V_1(q \mid q)$ for all $q \in [q_c(p), c)$. But if $\hat{q}(p) \leqslant q_c(p)$ then $\phi(p, \hat{q}(p)) \leqslant \phi(p, q_c(p)) = c$ by (4), and so for all $q \in [q_c(p), c)$,

   $$V_1(c \mid q) \geq V_1(\phi(p, q_c(p)) \mid q) = W_1(\phi(p, q_c(p)) \mid q) \geq W_1(\phi(p, \hat{q}(p)) \mid q) = 0 = V_1(q \mid q),$$

   with the first equality holding since $V_1(r \mid q) = W_1(r \mid q)$ for all $r \geq c$, and the second one holding by Lemma 2. This verifies incentive compatibility. That $\underline{q}$ defines an EE is then immediate. To verify that this is a MEE, note that were $\underline{q} < q_c(p)$, then one would have $\phi(p, \underline{q}) < c$ and thus no experimentation by P2 could be supported.

2. If $\hat{q}(p) \in [q_c(p), c)$, then set $\underline{q} = \hat{q}(p)$. Again, (5) is satisfied immediately since $\hat{q}(p) \geq q_c(p)$. Next, note that $\underline{q} = \hat{q}(p) \geq q_c(p)$ implies that $\phi(p, \underline{q}) \geq \phi(p, q_c(p))$, and so $W_1(\phi(p, \bar{q}) \mid q) = V_1(\phi(p, \bar{q}) \mid q) \geq 0$ for all $q \in [\underline{q}, c)$. Thus, (6) is verified. Furthermore, since (6) is binding, this must also be a MEE (setting $\underline{q} < \hat{q}(p)$ would violate (6)).

3. If $\hat{q}(p) \in [c, 1]$, then set $\underline{q} = c$. In this case, full revelation is the only equilibrium: since $\hat{q}(p) \geq c$, it must be that $V_1(q \mid q) \geq V_1(\phi(p, q) \mid q)$ for all $q < c$.

A.10 Proof of Lemma 9

For part i), note that by Lemma 4, $q \mapsto \phi(p,q)$ is strictly increasing. Furthermore, from equation (8), $\phi(p,q)$ is strictly decreasing in α. Since $q_c(p)$ solves $\phi(p, q_c(p)) = c$, this proves the first claim. For the second part, note that from equation (8), $\lim_{\alpha\to 1}\phi(p,q) = \mathbb{E}(r \mid r \in [q,c])$, and hence $\lim_{\alpha\to 1}\phi(p,c) = c$, while $\lim_{\alpha\to 0}\phi(p,q) = p$, and hence $\lim_{\alpha\to 0}\phi(p,c) = 0$.

For part ii), note that $\partial W_1(\phi(p,q) \mid q)/\partial q > 0$, as asserted in Lemma 7. Next, by Lemma 4, $q \mapsto \phi(p,q)$ is strictly increasing. Thus, the first part of the claim obtains provided that $\partial W_1(\phi(p,q) \mid q)/\partial\alpha > 0$. To see that such is the case, note that $\partial W_1(r \mid q)/\partial r > 0$ as argued in Lemma 7, and from equation (8), $\phi(p,q)$ is strictly decreasing in α. Finally,

$$\frac{\partial W_1(r \mid q)}{\partial \alpha} = \Lambda(r \mid q) - (q - c) = \int_{x(r,c)}^{1} (q^z - c) f_r(z)\, dz \geq 0,$$

since $q^z > c$ for $z > x(r,c)$. Hence

$$\frac{\partial W_1(\phi(p,q) \mid q)}{\partial \alpha} = \frac{\partial W_1(r \mid q)}{\partial \alpha} + \frac{\partial W_1(r \mid q)}{\partial r}\frac{\partial \phi(r,q)}{\partial \alpha} > 0.$$

To prove the second claim, suppose that for all $p \in (c,1)$, $\hat{q}(p) < c$ for all $\alpha \geq 0$. For α small enough, $\phi(p,\hat{q}(p)) > c$ and thus $W_1(\phi(p,\hat{q}(p)) \mid q) \approx 2(\hat{q}(p) - c) < 0$, violating the indifference condition that $\hat{q}(p)$ must satisfy. We can of course then conclude that $q_v(p)$ is also strictly increasing whenever $\hat{q}(p) \leq c$, and is equal to c otherwise.

A.11 Proof of Proposition 1

For $p \in (c,1)$, let $\tilde{\alpha}(p)$ be that value of α such that $q_v(p) = q_c(p)$. Such a value exists and lies in $(0,1)$ by Lemmas 9 and 2; we have that $q_v(p) = c > q_c(p)$ at $\underline{\alpha}(p)$, and thus by continuity, $\tilde{\alpha}(p) > \underline{\alpha}(p)$. Finally, $\tilde{\alpha}(p) < 1$ since $\tilde{p} < 1$ for all $\alpha \in (0,1)$.

A.12 Proof of Proposition 2

Since $q \mapsto V_1(r \mid q)$ is affine, standard arguments imply the problem faced by P1 under commitment is to solve

$$v^*(p) = \max_{H\in\Delta([0,1])} \int_0^1 V_1(q)\, dH(q), \tag{A.1}$$

subject to the constraint that H is a mean-preserving contraction of G_p (Kamenica and Gentzkow, 2011). First, we prove that $V_1(q)$ is convex on $[c, 1]$. To see this, note that Lemma 1 implies that $V_1(q \mid q) = \sup_{r \in [0,1]} V_1(r \mid q)$ for $q \in [c, 1]$, and that $q \mapsto V_1(r \mid q)$ is affine. The convexity of $V_1(q)$ then follows from standard results in convex duality (Rockafellar, 1997, Theorem 13.2).

We may now apply (Dworczak and Martini, 2019, Theorem 1). In particular, consider the function ψ defined by

$$\psi(q) = \begin{cases} V_1(q) & \text{if} \quad p^x \geqslant \bar{q}^*(p) \\ V_1(c)\left(\frac{q - \underline{q}^*(p)}{c - \underline{q}^*(p)}\right) & \text{if} \quad p^x \in [\underline{q}^*(p), \bar{q}^*(p)) \\ V_1(q) & \text{if} \quad p^x < \underline{q}^*(p), \end{cases}$$

and the distribution $H_p : [0, 1] \to [0, 1]$ defined by

$$H_p(q) = \begin{cases} G_p(q) & \text{if} \quad p^x \geqslant \bar{q}^*(p) \\ G_p(c) + \mathbb{I}_{q \geq c}[G_p(\bar{q}^*(p)) - G_p(\underline{q}^*(p))] & \text{if} \quad p^x \in [\underline{q}^*(p), \bar{q}^*(p)) \\ G_p(q) & \text{if} \quad p^x < \underline{q}^*(p). \end{cases}$$

which reveals q when either $q \geq \bar{q}^*(p)$ or $q \leq \underline{q}^*(p)$ and pools otherwise. It is readily verified that ψ and H together satisfy conditions 3.1-3.3 of (Dworczak and Martini, 2019, Theorem 1), and thus constitute a solution to the commitment problem. Finally note that since $q \mapsto G_p(q)$ is continuous and strictly increasing, so too are $\underline{q}^*(p), \bar{q}^*(p)$.

A.13 Proof of Corollary 3

Note that the constraint (13) is independent of p, whereas a simple application of the posterior monotonicity property (Proposition 4 in Smith et al. (2021)) implies that for fixed $\underline{q}, \bar{q}, \mathbb{E}_p(q \mid q \in [\underline{q}, \bar{q}))$ is strictly increasing in p. Thus, to keep $\mathbb{E}_p(q \mid q \in [\underline{q}, \bar{q}))$ fixed, we must lower both $\underline{q}$ and $\bar{q}$. The final part of the corollary follows by noting that $V(q)$ is strictly increasing and convex for $q \geq c$ and strictly positive at c, and thus for all $p \in [c, 1]$ the line segment intersecting the three points $(\underline{q}^*(p), 0), (c, V_1(c)), (\bar{q}^*(p)$ and $V_1(\bar{q}^*(p)))$ can only exist if $\underline{q}^*(p) \neq \bar{q}^*(p)$, while the constraint that $\mathbb{E}_p(q \mid q \in [\underline{q}^*(p), \bar{q}^*(p))) = c$ further implies that $\underline{q}^*(p) < c < \bar{q}^*(p)$.

Online Appendix

"(Pro)-Social Learning and Selective Disclosure"

Roland Bénabou, Nikhil Vellodi

A Heterogeneous Payoffs

We now extend the analysis to allow for heterogeneous payoffs, by introducing an idiosyncratic component to utility. Besides adding realism this will serve to show that, under general conditions on the form of this heterogeneity, disclosure is still polarized and positively biased, and that all equilibria are *necessarily* EE's in which P2 strictly prefers to disclose.

Let the payoff to agent t from receiving signal x now be $x\epsilon_t$, where each ϵ_t is drawn from a distribution H, independently from x. Without loss of generality, we assume that $\mathbb{E}(\epsilon) = 1$ and H has full support on $[0, \infty)$, with a density h that is everywhere positive. We further assume that the realization of their own ϵ_t is observable to an agent prior to their consumption decision –e.g., it represents the intensity of their need for such a product– whereas the value $x\epsilon_t$ (or, equivalently, x itself) is revealed only when consumption occurs. Thus ϵ_t guides the experimentation decision a_t, but when $a_t = 1$ the relevant information for the disclosure decision d_t remains x itself, since ϵ_t is irrelevant to any successor. Formally, consumption rules now map both from beliefs and shocks, i.e. $a_t : [0, 1] \times [0, \infty) \to \{0, 1\}$, while disclosure rules remain as before.

The expected values, from P1's perspective, of subsequent players' consumptions are now:

$$u(r \mid q) \equiv \int_{c/r}^{\infty} (q\epsilon - c)\, dH(\epsilon),$$

$$\Lambda(r \mid q) \equiv \mathbb{E}_{\epsilon,z}(u\,(r^z\epsilon \mid q^z\epsilon)) = \int_0^1 \int_{\frac{c}{r^z}}^{\infty} (q^z\epsilon - c)\, dH(\epsilon)dF_q(z).$$

We start with some basic properties of u and Λ.

Lemma A.1. *1. Both maps $r \mapsto u(r \mid q)$ and $\Lambda(r \mid q)$ are strictly maximized at q.*

2. $\Lambda(r \mid q) \geq (>) u(r \mid q)$ for all $r \leq (<) q$.

Proof. Direct calculation verifies that $\frac{\partial u}{\partial r} = -\frac{c^2}{r^2}\left(\frac{q}{r} - 1\right) h\left(\frac{c}{r}\right)$,

which is equal to zero if and only if $q = r$. Since $\Lambda(r \mid q) = \mathbb{E}_{\epsilon,z}(u(r^z\epsilon \mid q^z\epsilon))$, point 1 is verified. To verify point 2, note that

$$\begin{aligned}\Lambda(r \mid q) - u(r \mid q) &= \int_0^1 \int_{\frac{c}{r^z}}^{\infty} (q^z\epsilon - c)\, dH(\epsilon) dF_q(z) - u(r \mid q) \\ &= \int_0^{\hat{x}} \int_{\frac{c}{r^z}}^{\infty} (q^z\epsilon - c)\, dH(\epsilon) dF_q(z) + \underbrace{\int_{\hat{x}}^1 \int_{\frac{c}{r^z}}^{\infty} (q^z\epsilon - c)\, dH(\epsilon) dF_q(z)}_{\geq u(r|q)} - u(r \mid q) \geq 0,\end{aligned}$$

where the last inequality holds because $z \mapsto q^z\epsilon - c$ is positive on the range $[c/r^z, \infty)$ by the MLRP, since by assumption $r \leq q$. $\square$

Note that since $V_2(r \mid q) = u(r \mid q)$, Lemma A.1 implies that full disclosure by P2 is a *strictly* dominant strategy. In Section 2, P2 was indifferent over posterior beliefs that induce the same action by P3. Now, greater accuracy leads to a strictly lower chance of erroneous consumption choices by P3 due to idiosyncratic shocks.

As before, this allows us to simplify player 1's value function,

$$V_1(r \mid q) = u(r \mid q) + \Big(\alpha C(r)\Lambda(r \mid q) + (1 - \alpha + \alpha(1 - C(r)))\, u(r \mid q)\Big), \tag{A.1}$$

where

$$C(r) \equiv \int_{c/r}^{\infty} dH(\epsilon)$$

is the probability of consumption given a prior belief r, prior to the realization of ϵ.

A.0.1 Selected Disclosure

In order to draw comparison to the results in Sections 4.2 and 4.3, we first adapt the definition of experimentation equilibria in the most natural manner. Now, let

$$X_E(\sigma) = \{x \in N_1(p) \mid a_2(p^{\emptyset}, \epsilon) > a_2(p^x, \epsilon) \quad \forall \epsilon \in [0, \infty)\}$$

denote the experimentation set for an equilibrium σ. First, we recover the result of polarized disclosure.

Lemma A.2. *(Polarized disclosure) Fix $r \in (0,1)$. Then, $\lim_{q \to 0,1} [V_1(q \mid q) - V_1(r \mid q)] > 0$.*

Proof. For the lower limit, note that

$$u(r \mid 0) = \int_{c/r}^{\infty} -c \, dH(\epsilon) < 0, \quad \Lambda(r \mid 0) = \int_0^1 \int_{c/r^z}^{\infty} -c \, dH(\epsilon) dF_q(z) < 0,$$

whereas $u(0 \mid 0) = \Lambda(0 \mid 0) = 0$. Thus, by the expression for $V_1(r \mid q)$ given in (A.1), $V_1(r \mid q) < 0 = V_1(q \mid q)$. For the upper limit, note that

$$\Lambda(r \mid 1) = \int_0^1 \int_{\frac{c}{r^z}}^{\infty} (\epsilon - c) \, dH(\epsilon) dF_q(z),$$

which is strictly increasing in r by the MLRP, since the integrand is strictly positive. Similarly, $r \mapsto u(r \mid 1)$ is strictly increasing. Finally, $r \mapsto C(r)$ is also strictly increasing, and thus so is $r \mapsto V_1(r \mid 1)$. Therefore, the claim is verified. □

Next, we demonstrate that for any prior $p \in (0,1)$, any posterior $q \geqslant p$ (i.e. any signal $x \geqslant \hat{x}$) is disclosed by P1. Note that whereas in the baseline model (Lemma 3) it was dominant for all posteriors $q \geq c$ to be disclosed, here this is no longer necessarily the case.

Lemma A.3. *If $p^x \geqslant p$, then $d_1(x, p) = 1$ is a strictly dominant strategy.*

Proof. Suppose not, so that there exists an $x > \hat{x}$ such that $d_1(x, p) = 0$. Take the largest such x and let $q = p^x$. By construction, to satisfy the equilibrium belief condition (5) it must be that $p^\emptyset < q$. But then by Lemma A.1, $u(p^\emptyset \mid q) < u(q \mid q)$ and $\Lambda(p^\emptyset \mid q) < \Lambda(q \mid q)$, while we also have $C(p^\emptyset) < C(q)$ as $r \mapsto C(r)$ is strictly increasing. Combining, we have that $V_1(p^\emptyset \mid q) < V_1(q \mid q)$. □

Finally, we prove that non-disclosure of signals that convey marginally bad news (namely, such that the posterior p^x is just below the prior p) is optimal. This result has no direct analogue in the baseline model, insofar as non-disclosure now occurs at signals the revelation of which would have induced consumption with strictly positive probability ($c < p^x < p$). Furthermore, the result holds at *all* $p \in (0, 1)$, thus demonstrating the existence of an EE at *all* priors. This is again distinct from Theorem 2, which showed that for high enough prior p, no EE existed.

Lemma A.4. *(Positive selection) Let* $\tilde{V}_1(q) \equiv V_1(r \mid q)$. *Then* $\tilde{V}_1'(q) > \frac{\partial V_1}{\partial q}|_{r=q}$.

Proof. Since $\tilde{V_1}'(q) = \partial V_1(r \mid q)/\partial r|_{r=q} + \partial V_1(r \mid q)/\partial q|_{r=q}$, the claim is equivalent to proving that $\partial V_1(r \mid q)/\partial r|_{p=q} > 0$. But

$$\frac{\partial V_1|(r \mid q)}{\partial r}|_{r=q} = \underbrace{\frac{\partial u}{\partial r}_{r=q}}_{=0} + \frac{c}{q^2} h\left(\frac{c}{q}\right)\left[\Lambda(q \mid q) + (1 - C(q))u(q \mid q)\right]$$

$$+ C\left(\frac{c}{q}\right)\left[\underbrace{\frac{\partial \Lambda}{\partial r}|_{r=q}}_{=0} + (1 - C(q))\underbrace{\frac{\partial u}{\partial r}|_{r=q}}_{=0} - C'(q)u(q \mid q)\right]$$

$$= \Lambda(q \mid q) - C\left(\frac{c}{q}\right) u(q \mid q) \frac{c}{q^2} h\left(\frac{c}{q}\right) > 0,$$

where the last inequality holds because $C(q \mid q) < 1$ and $\Lambda(q \mid q) > u(q \mid q)$. □

In particular, for $x = \hat{x} - \varepsilon$ where ε is small, non-disclosure is optimal. Combining Lemmas A.3 and A.4 with a continuity argument yields that non-disclosure takes place in (at least) some interval $[\hat{x} - \varepsilon, \hat{x})$, and thus disclosure is positively biased. Furthermore, we have:

Lemma A.5. *Any equilibrium is an EE.*

Proof. Since $d_2 \equiv 1$ is now a strictly dominant strategy, all equilibria are DE's. To see that all admit a non-empty experimentation region, note that Lemmas A.2 and A.3 imply that in any equilibrium σ, for each p there exists a minimal posterior $\underline{q}(p) < p$ that is concealed. Continuity of $r \mapsto V_1(r \mid q)$ then ensures the existence of a $\delta > 0$ such that posteriors in the interval $[\underline{q}(p), \underline{q}(p) + \delta)$ are concealed. But for δ sufficiently small, $\underline{q}(p) + \delta < p$, and so $[\underline{q}(p), \underline{q}(p) + \delta) \subset X_E(\sigma)$. □

B Biased Feedback – Cheap Talk

We now consider a natural variant on our baseline model by relaxing the requirement of hard evidence disclosure and instead permitting arbitrary message reporting (cheap talk). Such a variant is important for several reasons. First, in many applied settings, it might not only be feasible but strategically optimal for consumers to mis-report their experiences. The hard-evidence baseline abstracts from this possibility, thus providing a useful benchmark; even when fake reviews are impossible, might their be scope for strategic disclosure? In this section, we explore the extent of strategic information transmission when lying is both feasible and costless. Second, by studying an alternative, well-established form of equilibrium information transmission, we make clear the features of strategic disclosure that are invariant to the information sharing technology available to agents.

Specifically, we endow each agent with a rich messaging space $\mathcal{M} = [0,1] \times \{\emptyset\}$ that allows not only for full separation but also for agents to send a privileged message that pools with non-arriving consumers, so that messaging rules (previously, disclosure rules) are now mappings $d_t : X \times [0,1] \to \mathcal{M}$.[13] Again, full transparency is dominant for P2, so we focus on P1's messaging strategy. Let $r^*(m)$ denote P2's equilibrium belief upon observing message m. Then the IC constraint (6) is replaced with the condition

$$d_1^*(x,p) \in \underset{m \in \mathrm{supp}(d_1^*)}{\arg\min} \; V_1(r^*(m) \mid p^x). \tag{B.1}$$

[13] We focus on pure-strategy equilibria for simplicity, noting the usual implementation via uniform randomization in cheap-talk games

We focus on the case when $\alpha = 1$. Combining various insights learned through the baseline analysis, we can immediately conclude that all equilibria must admit a partitional structure.

Proposition B.1. *All equilibria are partitional. That is, for all $r \in [0,1]$ induced in equilibrium, the set of q in which r is induced forms an interval in $[0,1]$. Furthermore, there must be at most finitely many such intervals on $[c,1]$.*

Proof. We proceed with a series of lemmas.

Lemma B.1. *All equilibria are partitional. Furthermore, there must be at most countably infinitely many such intervals on $[c,1]$.*

Proof. Lemma 1 tells us that $r \mapsto V_1(r \mid q)$ is maximized at $r = q$, and that $q \mapsto V_1(r \mid q)$ is strictly increasing, so that $\arg\max_{r\in[0,1]} V_1(r \mid q)$ is strictly increasing in q. To prove the final assertion, we argue that there can be no interval in $[c,1]$ on which separation can occur. Suppose there were, and take the lowest such interval $[q_1, q_2], q_1 \leq q_2$. If $q_2 < 1$, then we claim that types $q \in (q_2 + \varepsilon]$ have an incentive to pool with q_2. For since this was the lowest separating interval, it must be that types $q \in (q_2 + \varepsilon]$ induce a belief $\hat{q} = q_2 + \delta, \delta > 0$. By Lemma 2, $V_1(\hat{q}_2 \mid q_2 + \varepsilon) < V_1(q_2 \mid q_2 + \varepsilon) \approx V_1(q_2 \mid q_2) + \varepsilon V_1'(q_2 \mid q_2)$ for small enough $\varepsilon > 0$. If $q_2 = 1$, then we claim that $q \in (q_1 - \varepsilon]$ have an incentive to pool with q_1 by analogous reasoning. □

Lemma B.2. *There exists $\underline{q}_{min} < c$ such that full revelation is weakly dominant for all types $q \in [0, \underline{q}_{min})$.*

Proof. $\underline{q}_{min}$ is the unique root of $q \mapsto V_1(c \mid q)$ on $[0,c]$, which is well-defined since the map is continuous, strictly increasing with $V_1(c \mid 0) < 0 < V_1(c \mid c)$. □

It is thus without loss to associate an equilibrium with a lowest type $\underline{q} > 0$ that forms part of a pooling interval that itself induces experimentation. More specifically, combining with Lemma B.1, an equilibrium can be described by a (possibly infinite) sequence ($\underline{q} \equiv$

$)q_0 < q_1 < q_2 < \dots$ such that types in $[q_i, q_{i+1})$ pool and $\hat{\underline{q}} \equiv \mathbb{E}(q \mid q \in [\underline{q}, q_1)) \geq c$. More generally, we denote $\hat{q}_{i+1} = \mathbb{E}(q \mid q \in [q_i, q_{i+1}))$.

We next prove that the two first intervals $[\underline{q}, q_1), [q_1, q_2)$ cannot be "too small" as then types just below $\underline{q}$ would profitably deviate by pooling with $[q_1, q_2)$ to induce $\hat{q}_1$.

Lemma B.3. *For all $\underline{q} \in [\underline{q}_{min}, c]$ there exists $\hat{q}_{1,min} > c$ such that in any equilibrium, $\hat{q}_1 \geq \hat{q}_{1,min}$.*

Proof. If not, then for any $\varepsilon > 0$ there exists an equilibrium with $\hat{q}_1 \leq c + \varepsilon$. But since by definition $\hat{\underline{q}} \geq c$, it must be that $\underline{q} > \underline{q}_{min}$ for sufficiently small ε and by the sandwich theorem, $V_1(\hat{\underline{q}} \mid \underline{q}) > 0$, violating the IC constraint at $\underline{q}$. □

Lemma B.4. *All equilibrium partitions essentially admit at most finitely many intervals covering* $[c, 1]$.

Proof. We proceed constructively, via the following algorithm:

1. Fix a $\underline{q} \geq \underline{q}_{min}$. Compute $\hat{q}_{max} \equiv \mathbb{E}_p(q \mid q \in [\underline{q}, 1])$.
   (a) If $V_1(\hat{q}_{max} \mid \underline{q}) > 0$, then $N^*(\underline{q}) = 0$ and $\underline{q}$ cannot be implemented in equilibrium.
   (b) If not, then there exists a unique $q_1 > c$ such that $V_1(\hat{q}_1 \mid \underline{q}) = 0$, where $\hat{q}_1 \equiv \mathbb{E}(q \mid q \in [\underline{q}, q_1])$. (Such a value exists by continuity and strict monotonicity of $r \mapsto V_1(r \mid q)$ on $[q, 1]$, the IVT and because $V_1(c \mid \underline{q}) > V_1(c \mid \underline{q}_{min}) = 0$ by Lemma B.2).
2. Compute $V_1(\hat{q}_1 \mid q_1)$.
   (a) If $V_1(1 \mid q_1) \geq V_1(\hat{q}_1 \mid q_1)$, then $N^*(\underline{q}) = 1$.
   (b) If not, then there exists a unique $q_2 > q_1$ such that $V_1(\hat{q}_2 \mid q_1) = V_1(\hat{q}_1 \mid q_1)$, where $\hat{q}_2$ is analogously defined, and q_2 exists by the same reasoning as q_1.
3. Repeat from step 2.

Finally, we argue that this algorithm terminates in finitely many steps. Suppose not. Then for all $\varepsilon > 0$, there exists an equilibrium and an interval $[q_i, q_{i+1}) \subset [c, 1]$ such that $q_{i+1} - q_i \leq \varepsilon$. Without loss, assume equality. Let $\hat{q}_{i+1} = \mathbb{E}(q \mid q \in [q_i, q_{i+1}))$. Then there exists $\delta(\varepsilon) < \varepsilon$ such that $\hat{q}_{i+1} - q_i = \delta(\varepsilon)$. The Mean Value Theorem implies that

$$V_1(\hat{q}_i \mid q_i) - V_1(q_i \mid q_i) = V_1'(\varphi_1 \mid q_i)(\hat{q}_1 - q_i),$$

for some $\varphi_i \in (\hat{q}_i, q_i)$. But since $r \mapsto V_1(r \mid q)$ has a global maximum at q, we know that

$$V_1(q_i + \delta(\varepsilon) \mid q_i) - V_1(q_i \mid q_i) \approx \frac{\partial^2 V_1}{\partial r^2}(q_i \mid q_i)\delta(\varepsilon)^2.$$

Combining these terms implies that $q_i - \hat{q}_i = \kappa\delta(\varepsilon)$, for some $\kappa > 0$, and so $\hat{q}_{i+1} - \hat{q}_i = (\hat{q}_{i+1} - q_i) + (q_i - \hat{q}_i) = \kappa_i\delta(\varepsilon)$, for some $\kappa_i > 0$. Now, since $\varepsilon > 0$, there exists a finite $I > 0$ such that $q_{i-I} = \underline{q}$ (if not, then Lemma B.3) and thus a simple inductive argument implies that $\hat{q}_1 - \hat{\underline{q}} = \kappa_{i-I}\delta(\varepsilon)$, for some $\kappa_{i-I} > 0$. Taking ε (and thus $\delta(\varepsilon)$) sufficiently small violates Lemma B.3. □

□

This characterization is in stark contrast to Proposition 1, as well as the subsequent characterization under commitment (Proposition 2), both of which exhibit a "separate-pool-separate" reporting structure. The result derives from the stark nature of preference misalignment between P1 and P2; for $q \geq c$, there is no misalignment at all and thus such types have a preference to reveal themselves fully, whereas in an interval below c, all types have a preference to induce c. This latter preference toward biasing "upward" leads to a ripple effect for higher types, whereby to preserve incentives, information transmission must necessarily be coarse. Put differently, we see here an alternative manifestation of the accuracy-experimentation trade-off identified in Section 4; in order to foster experimentation by P2, equilibrium must necessarily involve consumption errors by P3.

The proof of Proposition B.1 is constructive. First, we identify a lower-bound on the

Figure 5: Cheap-Talk Equilibria: Construction

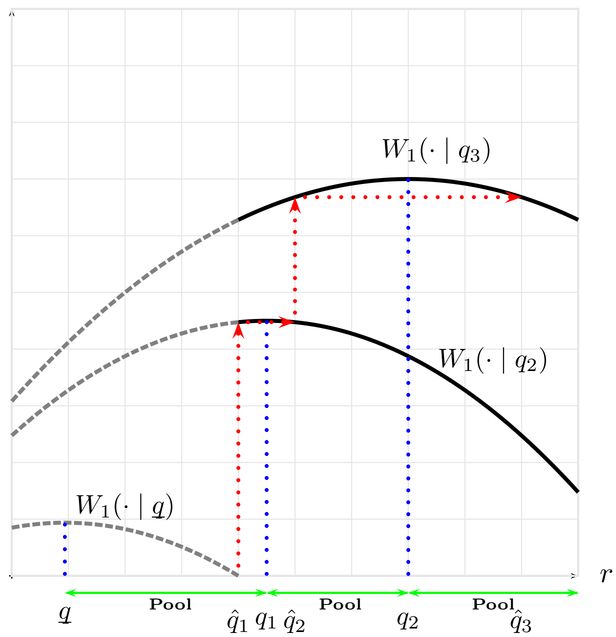

An equilibrium with three pooling intervals covering $[c, 1]$. Virtual value function $W_1(r \mid q)$. For $r \geq c$, $W_1(\cdot \mid q) = V_1(\cdot \mid q)$ (solid black lines). For $r \geq c$, $V_1(\cdot \mid q) = 0$.

degree of experimentation possible; there exists a $\underline{q}_{min}$ such that $V_1(c \mid \underline{q}_{min})$, thus any type lower prefers to terminate experimentation, regardless of the continuation belief r. Each equilibrium is essentially determined by its associated $\underline{q} \in [\underline{q}_{min}, c]$, that is the lowest type whose message induces experimentation.

Corollary B.1. *For each $\underline{q} \in [\underline{q}_{min}, c]$ there exists a unique equilibrium partition on $[\underline{q}, 1]$ induced by d_1^*.*