The Informational Impact of Opinion Leaders

Alesandro Arcuri

January 2021

Abstract

I consider a model of Bayesian persuasion with observational learning. A sender designs a test to encourage adoption of a new product, which may be either good or bad. The test and its outcome are viewed by an opinion leader, who chooses whether or not to adopt the sender's product. If the opinion leader chooses to adopt, the test and its outcome are next viewed by an audience who also makes an adopt/reject decision. Both opinion leader and audience strictly prefer to adopt the product when it is good, and reject it when it is bad. The cost of adoption is heterogeneous for each individual, and each individual's adoption cost is their own private information. The sender's goal is to maximize the probability of audience adoption. Compared to a benchmark in which the audience views the sender's test directly, I show that the presence of the opinion leader weakly increases the informativeness of the sender's test. However, it may lower the probability that the audience views the sender's test, and so the overall effect on audience welfare is ambiguous. I extend this model to include two forms of collusion. Under information-sharing, the opinion leader informs the sender of her private type. In this case the audience weakly prefers the opinion leader to be present, and they may or may not prefer the opinion leader to share her type with the sender. Under cooperation, the opinion leader adopts the sender's product regardless of the test result. In this case, the audience receives the same welfare as in the benchmark where the sender communicates with them directly.

1 Introduction

In persuading individuals to adopt new innovations, carefully-designed tests often provide key evidence. A doctor deciding whether to prescribe a new drug may consult manufacturer-sponsored trials comparing it to standard treatment options. An environmental group deciding whether to endorse a political incumbent or a challenger may refer to candidate-commissioned studies comparing the effects of their policies on the climate. A game developer may release a trial demo to convince consumers to pre-order a full game. In each of these cases, persuasion can have second-order effects. A doctor switching to a new prescription might encourage colleagues to follow suit. A group's support of a candidate is only valuable insofar as it convinces its audience to vote for him. And a streamer playing a game may convince her audience to at least try the demo.

This paper examines the impact of these observational learning dynamics on the design of the optimal test. A central example is in the field of pharmaceutical marketing, where the tests being studied are manufacturer-sponsored drug trials. It is well known both by physicians and marketers that most doctors learn about new prescriptions not based on their own reading, but rather by observing the prescription habits of others. Firms may conduct extensive trials to show that their products are better than their competitors, but they also face the challenge of making doctors aware of them. To address this, firms spend resources to identify and influence "key opinion leaders" in physician networks. Efforts to identify influential doctors from network and prescription data have been published in the marketing literature (e.g. Nair, Manchanda, and Bhatia (2010), Iyengar, van den Bulte, and Valente (2011)). Such "opinion leaders" are usually well-respected physicians, and many are research-active. Firms may hire opinion leaders to give talks on the company's behalf, using pre-prepared slides (Moynihan 2008). Additionally, firms may hire opinion leaders to consult on their clinical trials (Moynihan 2008). The prominence of opinion leaders in pharmaceutical marketing is a relatively new phenomenon. One analysis finds that 'the number of articles in the industry journal Pharmaceutical Executive that mention "opinion leader" roughly tripled between 2000 and 2010" (Sismondo 2013). By 2010, as the Chronicle of Higher Education puts it, "it is an article of faith among pharmaceutical executives that KOL's [key opinion leaders] are a critical part of any marketing plan (Chronicle 2010).

Growing relationships between firms and opinion leaders have sparked some controversy in the medical community. One advocacy group known as No Free Lunch asks doctors to pledge not to accept gifts or payments from pharmaceutical firms. A paper published in the Journal of Law, Medicine, and Ethics, alleges that "partly due to the use of KOLs, a small number of companies with well-defined and narrow interests have inordinate influence over how medical knowledge is produced, circulated, and consumed" (Sismondo 2013). My paper aims to shed some light on these claims from an information design perspective.

I extend the Bayesian persuasion framework of Kamenica and Gentzkow (2011) to include three agents: a sender, an opinion leader, and an audience. The sender wants to maximize the probability that the audience adopts his product. To do so, he designs an experiment π to test his product's quality. He has full control over the outcome space S of the experiment, as well as over the conditional probability distributions $\pi_0(s)$,$\pi_1(s)$ of the outcome $s \in S$. When the state of the world is ω, s is distributed according to π_ω. The experiment and its results are viewed by an opinion leader. If this persuades the opinion leader to adopt the sender's product, then the audience is made aware of the experiment and its results, and makes their own adoption decision. Opinion leader and audience preferences are each characterized by an "adoption cost" $z_{OL}, z_{Au} \in (0,1)$. The adoption costs z_{OL}, z_{Au} serve as the private types of the opinion leader and audience, and the sender ex ante knows only their distributions f^{OL}, f^{Au}. Agent i's type $z_i \in (0,1)$ denotes the posterior belief the agent needs to reach before being willing to adopt.

It may be helpful to consider a brief example. Suppose the sender knows the adoption thresholds z_{OL}, z_{Au} of the opinion leader and audience respectively, and suppose $z_{OL} > \mu$ and $z_{Au} > \mu$, where μ is the prior belief that the sender's product is good. In order to convince the audience to adopt his product, the sender must also convince the opinion leader to do so. Thus the sender only gains when his test convinces both parties to adopt. Then the sender can restrict his attention to signals with only two outcomes: one which recommends both parties adopt, and one which recommends both parties reject the sender product. For any signal that has more than two outcomes, the sender could design a binary signal which achieves the same payoff, by running the former signal, but instead of reporting the realization, reporting only whether or not both agents would find it incentive-compatible to adopt the sender's product after viewing the realization. An optimal binary signal for the sender will have a negative outcome s_R, which induces posterior belief $Pr(\omega = 1|s_R) = 0$, and a positive outcome s_A, which induces posterior belief $Pr(\omega = 1|s_A) = \max\{z_{OL}, z_{Au}\}$. The only constraint on the distribution of posterior beliefs that the sender can induce is Bayes-plausibility: the expectation of the posterior equals the prior, that is $\pi(s_A) * 0 + \pi(s_A) * \max\{z_{OL}, z_{Au}\} = \mu$, where $\pi(s) = \mu\pi_1(s) + (1-\mu)\pi_0(s)$. Thus ex ante the sender's optimal signal will induce posterior belief $\max\{z_{OL}, z_{Au}\}$ with probability $\frac{\mu}{\max\{z_{OL}, z_{Au}\}}$, and posterior belief 0 with probability $1 - \frac{\mu}{\max\{z_{OL}, z_{Au}\}}$. Note that the sender's optimal signal will never have $\pi_1(s_R) > 0$: the sender can strictly increase the probability of audience adoption by instead choosing $\pi_1(s_R) = 0$. Thus the sender's optimal signal must have a "false negative rate" $\pi_1(s_R)$ of 0. This pins down the conditional distributions of the sender's optimal signal: $\pi_1(s_A) = 1$, $\pi_1(s_R) = 0$, $\pi_0(s_A) = \frac{\mu(1-\max\{z_{OL}, z_{Au}\})}{\max\{z_{OL}, z_{Au}\}(1-\mu)}$, $\pi_0(s_R) = 1 - \frac{\mu(1-\max\{z_{OL}, z_{Au}\})}{\max\{z_{OL}, z_{Au}\}(1-\mu)}$. If the opinion leader were not present, and the sender could show his experiment directly to the audience, he would choose $\pi(s_A) = \frac{\mu}{z_{Au}}$, $\pi(s_R) = \frac{\mu}{z_{Au}}$. The conditional distributions implied by this are $\pi_1(s_A) = 1$, $\pi_1(s_R) = 0$, $\pi_0(s_A) = \frac{\mu(1-z_{Au})}{z_{Au}(1-\mu)}$. This signal thus has the same false negative rate, and a weakly higher false positive rate, than the signal the sender would choose with the opinion leader present.

Some features of the sender's optimal signal in this example also hold when the sender is not sure of the types of the opinion leader and audience. I show that the sender's optimal signal still needs at most two outcomes, with the sender's optimal binary signal π^* "targeting" at most one non-zero posterior $s^* \in Supp(f^{OL}) \cup Supp(f^{Au})$. This optimal signal still has a zero false negative rate. I also show that the opinion leader's presence still makes the sender's equilibrium signal weakly more informative. However, one important element is missing. In the example, the opinion leader was always willing to adopt after seeing a positive outcome, by design. This need not be the case when the sender does not know the opinion leader's type. If the presence of the opinion leader does not change the sender's choice of signal, but some opinion leader types are not willing to adopt after seeing a positive experiment result, then the audience will worse off than if the sender could communicate with them directly.

I also allow for collusion between the sender and opinion leader in two different forms. Under information sharing, when the sender learns the opinion leader's type, I show that the opinion leader's presence still makes the sender's signal weakly more informative, and the that the opinion leader will no longer block the audience from seeing a positive test result. As a result, the audience weakly prefers the opinion leader to be present under information-sharing, which is not always true without collusion. Under cooperation, when the opinion leader adopts regardless of the sender's information, the outcome is exactly the same as in the benchmark where the opinion leader is not present. Notably, transparency has no effect on audience welfare under either regime.

This paper is structured as follows. The remainder of this section reviews related literature. Section 2 describes the model. Section 3 presents the main results. Section 4 discusses the results, and Section 5 concludes.

Related Literature - This paper is one of many to study Bayesian persuasion since the seminal work of Kamenica and Gentzkow (2011). Wang (2013) studies persuasion with multiple receivers, as do Alonso and Camara (2016), and Arieli and Babichenko (2019). Their papers are concerned with persuading a majority of voters, whereas in mine the sender must unanimously convince both parties. Additionally, in my model the sender faces uncertainty over the preferences of the audience and the opinion leader. With a different form of information transmission, Ambrus, Azevedo, and Kamada (2013) investigate how the equilibria of a cheap talk game change when the sender's message can only reach the receiver via a chain of intermediaries. The observational learning structure of my model is also reminiscent of the herding models dating back to Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992).

2 Model

State Space - The state of the world is a random variable ω distributed over $\{0, 1\}$ with $Pr(\omega = 1) = \mu \in (0, 1)$.
Players - There are three players: a sender (he), an opinion leader (she), and an audience (they).
Types - The opinion leader and audience each possess a private type z_{OL} and z_{Au} respectively, which are distributed over $(0, 1)$ independently of the state and each other, according to probability mass functions f^{OL} and f^{Au} respectively. Each party's type is its own private information.
Beliefs - It is common knowledge ex ante that $Pr(\omega = 1) = \mu$, and that z_{OL} and z_{Au} are jointly distributed according to f^{OL} and f^{Au} respectively.
Signals - A signal $\pi = (S, \pi_0, \pi_1)$ consists of an outcome space S and two probability distributions π_0, π_1 over S. The outcome s of π is a random variable distributed over S according to π_0 when the state of the world is 0, and according to π_1 when the state of the world is 1. I will refer to the universal set of signals as Π.
Actions and Timing - First, ω, z_{OL}, and z_{Au} are realized. Second, the sender commits to a choice of signal $\pi \in \Pi$, and the outcome of π is realized. Third, the

opinion leader observes the sender's choice of π and its outcome, and chooses an action $a_{OL} \in \{Adopt, Reject\}$. If the opinion leader chooses to reject, the game ends and payoffs are realized. If the opinion leader chooses to adopt, the audience is fourth made aware of the sender's choice of π and its realization, then chooses $a_{Au} \in \{Adopt, Reject\}$, after which the game ends and payoffs are realized.

Payoffs - All agents are rational expected utility maximizers. The sender's payoff $u_s(a_{Au}) = \mathbb{1}_{\{a_{Au}=Adopt\}}$ is determined solely by the audience's decision. The payoffs of opinion leader and audience, $u_{OL}(a_{OL}, \omega)$ and $u_{Au}(a_{Au}, \omega)$, are given as follows. For $i \in \{OL, Au\}$,

$$u_i(Reject, 0) = 0, \quad u_i(Reject, 1) = 0,$$

$$u_i(Adopt, 0) = -z_i \quad u_i(Adopt, 1) = 1 - z_i$$

Note that with two states and two actions, any preferences for which a decision-maker (i) strictly prefers to adopt in the good state, and (ii) strictly prefers to reject in the bad state, can be normalized to this form.[1]. The expected payoff from adopting for an agent with utility $\tilde{u}$ is $\mu'(1-\frac{x-y}{x-y+w-z})+(1-\mu')\frac{x-y}{x-y+w-z}$ A useful feature of this normalization is that agent i is indifferent between adopting and rejecting when $Pr(\omega = 1|\pi, s) = z_i$, for $i \in \{OL, Au\}$.

Equilibrium - I am interested in *sender-preferred subgame perfect equilibrium*, which requires that each agent's strategy maximizes their expected utility given the strategies of others, and furthermore that the opinion leader and audience take the sender's preferred action when they are indifferent between their two actions. This fixes the equilibrium strategies of the opinion leader and audience: adopt whenever the expected value of adopting is non-negative, and otherwise reject. In the next subsection I will examine the optimal strategy for the sender.

[1] Consider an agent with arbitrary preferences $u(Reject, 0) = x, u(Adopt, 0) = y, u(Reject, 1) = z, u(Adopt, 1) = w$. The agent will adopt if $\mu' w + (1-\mu')y \geqslant \mu' z + (1-\mu')x$, where μ' is the agent's belief that the state is 1. Then the agent is willing to adopt if and only if $\mu' \geqslant \frac{x-y}{x-y+w-z}$. Now consider the normalized utility function $\tilde{u}$, where $\tilde{u}(Reject, 0) = 0, \tilde{u}(Adopt, 0) = -\frac{x-y}{x-y+w-z}, \tilde{u}(Reject, 1) = 0, \tilde{u}(Adopt, 1) = 1 - \frac{x-y}{x-y+w-z}$

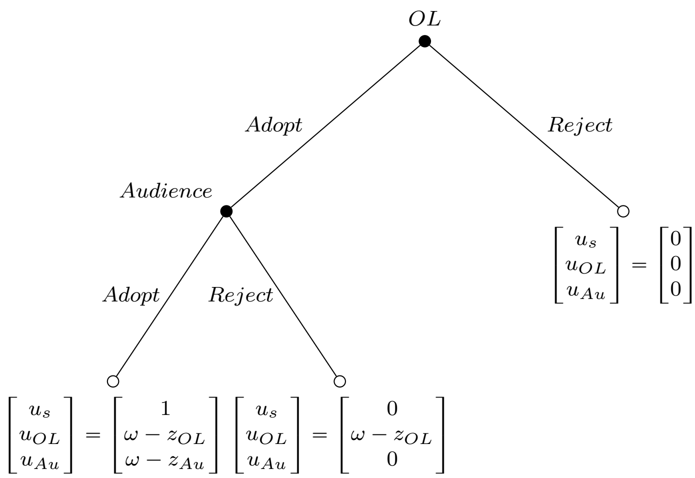

2.1 The Sender's Problem

In this section I will simplify the sender's problem by showing that there is an optimal signal for the sender which has only two outcomes. Note that for any signal $\pi = (S, \pi_0, \pi_1)$, we can without loss of generality re-name the elements s of S according to the posterior belief they induce, so that $S \subset [0,1]$, and for all $s \in S$, we have $s = Pr(\omega = 1|s) = \frac{\mu\pi_1(s)}{\mu\pi_1(s)+(1-\mu)\pi_0(s)}$. Given the earlier normalization of opinion leader and audience utility, this means that after seeing outcome s, an agent with adoption cost z_i is willing to adopt if $z_i \leqslant s$.

Let $\pi(s) = \mu\pi_1(s) + (1-\mu)\pi_0(s)$ be the ex ante probability of viewing outcome s, equivalently the probability of inducing the posterior s. Let F^{OL} and F^{Au} respectively be the cumulative density functions of f^{OL} and f^{Au}. Then the probability that the opinion leader (audience) is willing to adopt after observing outcome s is given by $F^{OL}(s)$ $(F^{AU}(s))$. Recall that agent preferences are independent by assumption. The sender's problem is thus

$$\max_{\pi \in \Pi} \int_{s \in [0,1]} \pi(s) F^{OL}(s) F^{Au}(s) ds$$

First, note that the sender can choose his signal π to induce a probability distribution of posteriors if and only if the distribution's mean is equal to the prior, i.e. $\sum_{s \in S} \pi(s)s = \mu$ (see for example Kamenica and Gentzkow (2011)).Here that constraint serves as our "Bayesian budget constraint" as in Kolotilin (2018). Second, note that we can replace the integrals here with summations based on the following observation. If some outcome $s' \notin \{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})$ were realized with positive probability, the sender could replace s' with $s'' = \operatorname{argmax} \{s \in \{\{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})\} : s < s'\}$. This modified outcome s'' would persuade the same types as s' would, and could be sent with higher probability without violating Bayes-plausibility. Specifically, for $s'' > 0$, the sender can send s'' more often in the bad state than he can s', and for $s'' = 0$, the sender can send s'' less frequently in the good state than s'. Either change weakly benefits the sender. Thus an optimal signal need not induce any posterior belief from outside $\{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})$. With these two points in mind we can then rewrite the sender's problem as

$$\max_{(\pi(s))_s} \sum_{s \in \{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})} \pi(s) F^{OL}(s) F^{Au}(s)$$

subject to

$$\sum_{s \in \{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})} s\pi(s) = \mu,$$

$$\sum_{s \in \{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})} \pi(s) = 1$$

The sender's objective function and constraint are both linear in $(\pi(s))_s$. The sender's choice set is $[0,1]^{|\{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})|}$ which is compact. Since the opinion leader and audience will both take the sender's preferred action when indifferent, the sender's utility is upper-semicontinuous in π. This ensures that a solution to the sender's problem exists, and therefore that an equilibrium exists as well. The linearity of the objective function and constraint also yield the following useful proposition, which extends an observation from Kolotilin (2018) to our setting. Namely, it shows that the the sender has an optimal signal which needs no more than two realizations. His optimization problem can be solved by spending his entire "budget" on one non-zero posterior.

Proposition 2.1 *There exists an optimal signal for the sender which has at most two outcomes.*

Proof Consider a more general problem:

$$\max_{(\pi(s))_s} \sum_{s \in \{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})} \pi(s) F^{OL}(s) F^{Au}(s)$$

subject to

$$\sum_{s \in \{0\} \cup Supp(f^{Au}) \cup Supp(f^{OL})} s\pi(s) = \mu.$$

This is a problem of consumer choice, where the goods being purchased are posterior beliefs. The amount of good s in the consumer's bundle is given by $\pi(s)$. The price of good s is s, and the marginal utility of good s is $F^{OL}(s)F^{Au}(s)$. The consumer has budget μ. This a linear optimization problem, and the consumer's utility is maximized by spending his entire budget on one good, specifically $s^* = \text{argmax}_{s \in Supp(f^{Au}) \cup Supp(f^{OL})} \frac{F^{OL}(s)F^{Au}(s)}{s}$. By the budget constraint, we must have $\pi(s^*) = \frac{\mu}{s^*}$

In the case where $s^* \leqslant \mu$, both the opinion leader and audience will adopt with probability 1 without the sender's signal. Thus the sender can focus his attention on signals with $s^* > \mu$. In this case, by the budget constraint, $\pi(s^*) <$

1. As a result, when we add the additional constraint $\sum \pi(s) = 1$, the sender can choose $\pi(0) = 1 - \pi(s^*)$, and earn the same payoff as he did before the introduction of the additional constraint. Thus choosing $\pi(s^*) = \frac{\mu}{s^*}$, $\pi(0) = 1 - \frac{\mu}{s^*}$ must be optimal for this problem as well. But this is the sender's problem, and thus shows that the sender's optimal signal requires only two outcomes when $s^* > \mu$ (and only one outcome when $s^* < \mu$). $\blacksquare$

With this in mind, for the remainder of this paper I will restrict the sender's attention to binary signals when he chooses π. Here I am using the term "binary signal" to refer to signals with at most two outcomes, so that the signal with only one realization (i.e. $S = \{\mu\}$) also counts as a binary signal. An optimal binary signal π^* can be identified with the non-zero posterior s^* it induces, which I will also refer to as the posterior "targeted" by π^*. Given s^*, π^* must satisfy $\pi^*(s^*) = \frac{\mu}{s^*}$, $\pi^*(0) = 1 - \frac{\mu}{s^*}$. In terms of conditional probabilities, this requires that $\pi_1(s^*) = 1$, $\pi_1(0) = 0$, $\pi_0(s^*) = \frac{\mu(1-s^*)}{s^*(1-\mu)}$, and $\pi_0(0) = 1 - \frac{\mu(1-s^*)}{s^*(1-\mu)}$.

2.2 Benchmark

In what follows I will compare the audience's welfare to their welfare when the opinion leader is not present, and the sender can show the audience his signal and its outcome directly. In this benchmark case, the sender's problem is to maximize $\sum_{s\in\{0\}\cup Supp(f^{Au})} \pi(s)F^{Au}(s)$ subject to the constraint $\sum_{s\in Supp(f^{Au})} s\pi(s) = \mu$. As before, the sender's problem has a linear objective function and linear constraint and so he has an optimal binary signal $\hat{\pi}$, which targets the posterior $\hat{s} = \text{argmax}_{s\in Supp(f^{Au})} \frac{F^{Au}(s)}{s}$.

2.3 Collusion

I model collusion here in two different ways. Under *information sharing*, the sender is made aware of the type z_{OL} of the opinion leader. In this case the sender has an optimal binary signal which targets the posterior $s^* = \text{argmax}_{s\in Supp(f^{OL})} \frac{f^{Au}}{s} \mathbb{1}_{\{s\geqslant z_{OL}\}}$. Under *cooperation*, the opinion leader always adopts regardless of the sender's signal or its outcome. In this case, the sender's problem is identical to his problem in the benchmark where the audience is always made aware of his signal. Thus the sender's optimal signal under cooperation is the same as his optimal signal in the benchmark, $\hat{\pi}$.

3 Analysis

For any two posteriors $s < s' \in (0, 1)$, the binary signal that targets s is a garbling of the binary signal that targets s'.[2] The following proposition uses this observation.

[2]Specifically, a garbling that truthfully reports any realizations of s', but misreports some realizations of 0 as realizations of s', when applied to the binary signal that targets posterior s', yields a signal equivalent to the binary signal that targets posterior s.

Proposition 3.1 *The sender's equilibrium signal π^* is weakly more informative than the benchmark signal $\hat{s}$.*

Proof When the opinion leader is present, the sender's signal will target the posterior s^* which maximizes $\frac{F^{OL}(s)F^{Au}(s)}{s}$. In the benchmark where the opinion leader is not present, the sender targets the posterior $\hat{s}$ which maximizes $\frac{F^{Au}(s)}{s}$. Assume towards a contradiction that $s^* < \hat{s}$. Since $\hat{s}$ maximizes $\frac{F^{Au}(s)}{s}$, it must be true that $F^{OL}(s^*) > F^{OL}(\hat{s})$. But F^{OL} is a CDF, so this is impossible. Thus it must be true that $s^* \geqslant \hat{s}$.

As a result, the welfare of the opinion leader is weakly increasing in the posterior s^* targeted by the sender's equilibrium choice of π. The relationship between s^* and audience welfare is slightly more subtle. While the audience always prefers to view a more informative signal, the audience's probability of viewing the sender's signal need not be continuous in s^*, due to the filtering effect of the opinion leader.

Say that the opinion leader *blocks* the audience whenever $z_{OL} > s^* > z_{Au}$ and the realization of π^* is s^*. The probability of such a block occurring is $(1 - F^{OL}(s^*))\pi^*(s^*)$. The disutility due to such a block, when it occurs, is the expected gain to audience of seeing the sender's signal.

Formally, we can write the audience's expected utility in the benchmark regime as

$$\hat{u}_{Au} = \hat{\pi}(\hat{s}) \sum_{z_{Au} \leqslant \hat{s}} f^{Au}(z_{Au})(\hat{s} - z_{Au})$$

and the audience's expected utility in the opinion leader regime as

$$u^*_{Au} = \pi(s^*)F^{OL}(s^*) \sum_{z_{Au} \leqslant s^*} f^{Au}(z_{Au})(s^* - z_{Au}).$$

The audience prefers the opinion leader to mediate their communication with the sender if $\pi(s^*)F^{OL}(s^*) \sum_{z_{Au} \leqslant s^*} f^{Au}(z_{Au})(s^* - z_{Au}) > \hat{\pi}(\hat{s}) \sum_{z_{Au} \leqslant \hat{s}} f^{Au}(z_{Au})(\hat{s} - z_{Au})$, and prefers the sender to communicate with them directly if $\pi(s^*)F^{OL}(s^*) \sum_{z_{Au} \leqslant s^*} f^{Au}(z_{Au})(s^* - z_{Au}) < \hat{\pi}(\hat{s}) \sum_{z_{Au} \leqslant \hat{s}} f^{Au}(z_{Au})(\hat{s} - z_{Au})$.

One immediate lower bound for $F^{OL}(s^*)$ is μ: if the sender's supposedly optimal signal does not convince the opinion leader with probability at least μ, he could deviate to a signal which fully reveals the state. This deviation would give him expected utility μ, a strict improvement.

A slightly better lower bound for $F^{OL}(s^*)$ can be derived if we instead think of the sender deviating to a signal which targets $\bar{z} = \max_{Supp(f^{OL}) \cup Supp(f^{Au})} z_i$. Such a signal would never be blocked, and $F^{OL}(\bar{z}) = 1$. This results in opinion leader adoption with probability $\mu + (1 - \mu)\frac{\mu(1-\bar{z})}{(1-\mu)\bar{z}} = \mu(1 + \frac{(1-\bar{z})}{\bar{z}}) > \mu$. From this we can conclude that for any equilibrium s^*, we must have $F^{OL}(s^*) \geqslant \mu(1 + \frac{(1-\bar{z})}{\bar{z}})$

In the remainder this section I extend the environment to consider potential collusion between the firm and the opinion leader. I model such collusion in two ways. Under *information sharing*, I assume the firm can learn the opinion leader's private type. Under *cooperation*, the opinion leader will choose to always adopt, regardless of the sender's trial or its result. I first show that while information-sharing may make the audience worse off, the audience weakly prefers the opinion leader to be present. I then show that under cooperation the audience receives their benchmark utility whether or not the opinion leader is present. These results are somewhat surprising - the only cases in which opinion leader presence strictly lowers audience welfare are cases in which there is no collusion.

3.1 Information Sharing

Under information sharing, the sender is made aware of the opinion leader's type z_{OL} before making his choice of π. Thus the sender's problem is the same as if the opinion leader's type were deterministically z_{OL}. The signal that the sender chooses as a result, by Proposition 3.1, will be no less informative than the signal that the audience would view in the benchmark without the opinion leader. Note that under information sharing, the opinion leader can never block the audience from viewing the firm's signal. This is because the sender receives a payoff of 0 if he targets some posterior $s' < z_{OL}$. Thus under information sharing the sender will always choose to target a posterior $s^{IS} \geqslant z_{OL}$. This yields the following proposition.

Proposition 3.2 *The audience weakly prefers the opinion leader to be present under information sharing for any* f^{OL}, f^{Au}, μ.

The following example shows how information sharing can hurt audience welfare (compared to the regime where the opinion leader is present but there is no information sharing). Without information sharing, the sender's optimal signal will target posterior $s^* = .6$ with probability 1. With information sharing, the sender's optimal signal will target posterior $s^* = .6$ when $z_{OL} = .6$ and $s^* = .55$ when $z_{OL} = .4$. This results in the audience having weaker information in expectation under information sharing.

Example 1 Suppose that $f^{OL}(.4) = .5 = f^{OL}(.6)$ and $f^{Au}(.45) = .5 = f^{Au}(.55)$. When the sender does not know the opinion leader's type, his optimal binary signal targets the posterior $s^* = \text{argmax}_{s \in \{.4,.45,.55,.6\}} \frac{F^{OL}(s) F^{Au}(s)}{s} = .6$.

When the sender can learn the opinion leader's type, if $z_{OL} = .6$ the sender's signal will target $s^* = \text{argmax}_{s \in \{.4,.45,.55,.6\}} \frac{F^{Au}(s)}{s} \mathbb{1}_{\{s \geqslant .6\}} = .6$ once more, but if $z_{OL} = .4$, the sender will instead target the posterior $s^* = \text{argmax}_{s \in \{.4,.45,.55,.6\}} \frac{F^{Au}(s)}{s} \mathbb{1}_{\{s \geqslant .4\}} = .55$. Thus the audience's ex ante expected utility is strictly higher when the sender cannot learn the opinion leader's type.

In contrast, the next example shows that information sharing can increase audience welfare in expectation. Without information sharing, the sender's optimal signal will target the posterior $s^* = .55$ with probability 1. With information sharing, the sender's optimal signal will target the posterior $s^* = .55$ when $z_{OL} = .1$ and will target the posterior $s^* = .9$ when $z_{OL} = .9$. In this case, the audience is more informed in expectation under information sharing.

Example 2 Now suppose that $f^{OL}(.9) = .1$, $f^{OL}(.1) = .9$, and $f^{Au}(.45) = .5 = f^{Au}(.55)$. When the sender does not know the opinion leader's type, his optimal binary signal targets the posterior $s^* = \text{argmax}_{s\in\{.1,.45,.55,.9\}} \frac{F^{OL}(s)F^{Au}(s)}{s} = .55$.

When the sender can learn the opinion leader's type, if $z_{OL} = .1$ the sender's signal will be target $s^* = \text{argmax}_{s\in\{.1,.45,.55,.9\}} \frac{F^{Au}(s)}{s} \mathbb{1}_{\{s\geqslant .1\}} = .55$, but if $z_{OL} = .9$, the sender's signal will target posterior $s^* = \text{argmax}_{s\in\{.1,.45,.55,.9\}} \frac{F^{Au}(s)}{s} \mathbb{1}_{\{s\geqslant .9\}} = .9$. Thus the audience's ex ante expected utility is strictly higher when the sender can learn the opinion leader's type.

3.2 Cooperation

Another way to model collusion is to assume the sender can compensate the opinion leader for choosing to adopt his product. In this case, if the opinion leader is willing to always adopt the sender's product, the audience will always be made aware of the sender's signal and its outcome. The sender's problem is then the same as it is in the benchmark where there is no opinion leader and the audience always sees the sender's signal and its outcome. The optimal signal for the sender to choose in this case is his benchmark optimum $\hat{\pi}$, so the audience sees the same information as they are would if the the opinion leader were present. This logic yields the following proposition.

Proposition 3.3 *Under cooperation, the audience is unaffected by the opinion leader's presence for any f^{OL}, f^{Au}, μ.*

3.3 Summary of Audience Welfare Results

Define regimes R, R^B, R^{IS}, R^C as follows:

- R: Model with opinion leader present, and no collusion.
- R^B: Benchmark model where opinion leader is not present, and sender communicates with audience directly.
- R^{IS}: Model with opinion leader present, and information sharing between opinion leader and sender.
- R^C: Model with opinion leader present, and cooperation between opinion leader and sender.

Define the relation $\geqslant_{Au}$ between two regimes R and R' such that $R \geqslant_{Au} R'$ if for all f^{OL}, f^{Au}, μ, the audience's expected utility is weakly higher under regime R than under regime R' . Then we can summarize the results of this section in the following theorem.

Theorem 3.4 *(i)* $R \not\geqslant_{Au} R^B$ *and* $R^B \not\geqslant_{Au} R$*, (ii)* $R^{IS} \geqslant_{Au} R^B$*, (iii)* $R^C = R^B$*,*[3] *(iv)* $R \not\geqslant_{Au} R^{IS}$ *and* $R^{IS} \not\geqslant_{Au} R$*.*

3.4 Transparency

In this model, the only information that is payoff-relevant to the audience is the design of the sender's signal and its outcome. Learning whether or not there is collusion between the opinion leader and audience, in either form, would not have any effect on the audience's optimal decision. This is an interesting observation, because the most prominent regulatory efforts affecting relationships between firms and opinion leaders have sought to increase transparency. For example, the 2010 Physician Payment Sunshine Act requires pharmaceutical firms to disclose any financial relationships they have with physicians totalling over $100. This observation that transparency has no effect is in line with Sismondo (2013)'s argument that "transparency, while it is ethically desirable and provides useful data for analysis, is of limited immediate value in changing the situation." It may be interesting to see if the same result holds in a model where the opinion leader and audience have heterogeneous information, instead of heterogeneous preferences.

4 Discussion

In the setting considered here, we have seen that the presence of the opinion leader makes the sender's optimal signal more informative. The analysis here is in part facilitated by the binary state space. With a larger state space, however, optimal signals, as well as the definition of a "more informative" signal, become more complicated, and so it is more difficult to compare equilibrium signals across regimes. However, it would be interesting to see if the results discussed here apply to a model with a general state space.

In constructing the binary-state model, I implicitly assumed that both the opinion leader and audience strictly preferred to adopt in state 1 and reject in state 0. That assumption, or at least one like it, is in fact necessary for these results to hold. If the opinion leader instead preferred to reject in state 1 and adopt in state 0, while the audience prefers adopt in state 1 and reject in state 0, the opinion leader's presence can instead make the audience weakly worse off. For example, if the opinion leader isn't willing to adopt without seeing some information that the state is 0, and the audience isn't willing to adopt without seeing some information that the state is 1, no communication between the sender and audience is possible. This makes the audience weakly worse off than they are in the benchmark where the sender can communicate with them directly. In moving to three dimensions, this implicit preference alignment is lost, and it is not clear if it is possible to replace it.

Finally, all of the binary-state results in this paper still hold if we relax the assumption that all agents share a common prior belief μ. If the audience

[3] i.e. $R^C \geqslant_{Au} R^B$ and $R^B \geqslant_{Au} R^C$

and opinion leader are allowed to vary in their beliefs μ^{Au}, μ^{OL}, as well as their preferences z_{OL}, z_{Au}, the agent's preferences will be identical to a model in which all agents share a common prior. I prove this claim in the following proposition.

Proposition 4.1 *For any agent type z_i, μ^i, the type's preferences are identical to an agent with type z'_i, μ for some $z'_i \in (0,1)$, for all $z_i \in (0,1), i \in \{OL, Au\}, \mu' \in (0,1), \mu \in (0,1)$.*

Proof When the agent has prior μ^i and switching cost z_i, they will adopt after seeing a trial outcome s^* if only if $\frac{\pi_1(s^*)\mu^i}{\pi_1(s^*)\mu^i + \pi_0(s^*)(1-\mu^i)} \geqslant z_i$. We can rewrite this as $\pi_1(s^*)\mu^i \geqslant z_i(\pi_1(s^*)\mu^i + \pi_0(s^*)(1-\mu^i))$, and again as $(1-z_i)\pi_1(s^*)\mu^i \geqslant z_i\pi_0(s^*)(1-\mu^i)$, and again as $\frac{(1-z_i)}{z_i}\frac{\mu^i}{1-\mu^i} \geqslant \frac{\pi_0(s^*)}{\pi_1(s^*)}$. Thus the agent's behavior in equilibrium will be to adopt if and only if $\frac{(1-z_i)}{z_i}\frac{\mu^i}{1-\mu^i} \geqslant \frac{\pi_0(s^*)}{\pi_1(s^*)}$. Choose $z'_i = \frac{\frac{\mu}{1-\mu}}{\frac{(1-z_i)}{z_i}\frac{\mu^i}{1-\mu^i} + \frac{\mu}{1-\mu}}$, note that by construction neither z'_i nor $(1-z'_i)$ are negative, nor are they equal to 0 or 1, so we have $z'_i \in (0,1)$. Rewriting their utility function as before, we can see that an agent with type z'_i, μ will adopt in equilibrium if and only if $\frac{(1-z'_i)}{z'_i}\frac{\mu}{1-\mu} \geqslant \frac{\pi_0(s^*)}{\pi_1(s^*)}$. Finally observe that $\frac{(1-z_i)}{z_i}\frac{\mu^i}{1-\mu^i} = \frac{(1-z'_i)}{z'_i}\frac{\mu}{1-\mu}$ by construction of z'_i. Thus the agent with type z'_i, μ has the same equilibrium behavior as with type z_i, μ^i.

Since the proposition holds for an arbitrary choice of parameters, we lose no generality in assuming that agents differ only in their adoption cost z_i, and that all agents share the prior μ.

5 Conclusion

Oftentimes, the evidence that consumers base their decisions on is passed on through observational learning. This describes many instances of word-of-mouth marketing, and is notably common in pharmaceutical marketing. In the medical community, there have been concerns raised that the presence of opinion leaders, or their relationships with pharmaceutical firms, may harm the knowledge that their audiences have access to. Here I have laid out a simple model of information production and dissemination, and shown that the presence of the opinion leader cannot lower the quality of information produced in equilibrium. However, the regime where the audience gets their information solely through the opinion leader can still be worse for the audience than the regime where they are shown the sender's test directly. Furthermore, I show that under collusion between the firm and opinion leader, the opinion leader regime is weakly better for the audience than the benchmark regime. Notably, it is possible for collusion between the firm and opinion leader to harm audience welfare, but not below the audience's benchmark welfare. It is also possible for information-sharing between the opinion leader and firm to make the audience better off in expectation. In the models of collusion presented here, transparency does nothing to improve audience welfare. This suggests that regulations aimed at improving transparency may not result in more informative trials being conducted. I also show that all of the above results hold even if agents do not share a common prior belief.

6 References

Alonso, Ricardo and Câmara, Odilon (2016) Bayesian persuasion with heterogeneous priors. Journal of Economic Theory, 165 . pp. 672-706. ISSN 1095-7235

Alonso, Ricardo, and Odilon Câmara. "Persuading voters." American Economic Review 106.11 (2016): 3590-3605.

Ambrus, Attila, Eduardo M. Azevedo, and Yuichiro Kamada. "Hierarchical cheap talk." Theoretical Economics 8.1 (2013): 233-261.

Arieli, Itai, and Yakov Babichenko. "Private bayesian persuasion." Journal of Economic Theory 182 (2019): 185-217.

Banerjee, A. V. "A Simple Model of Herding Behaviour." Quarterly Journal of Economics 107.3 (1992).

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch. "A theory of fads, fashion, custom, and cultural change as informational cascades." Journal of political Economy 100.5 (1992): 992-1026.

Elliot, Carl. "The Secret Lives of Big Pharma's 'Thought Leaders'", The Chronicle of Higher Education (2010): https://www.chronicle.com/article/The-Secret-Lives-of-Big/124335

Kamenica, E., Gentzkow, M.: Bayesian persuasion. Am. Econ. Rev. 101, 2590–2615 (2011)

Moynihan, Ray. "Key opinion leaders: independent experts or drug representatives in disguise?." bmj 336.7658 (2008): 1402-1403.

Nair, Harikesh S., Puneet Manchanda, and Tulikaa Bhatia. "Asymmetric social interactions in physician prescription behavior: The role of opinion leaders." Journal of Marketing Research 47.5 (2010): 883-895.

Raghuram Iyengar, Christophe Van den Bulte, Thomas W. Valente, (2011) Opinion Leadership and Social Contagion in NewProduct Diffusion. Marketing Science 30(2):195-212

Sismondo, S. (2013). Key Opinion Leaders and the Corruption of Medical

Knowledge: What the Sunshine Act Will and Won't Cast Light on. The Journal of Law, Medicine & Ethics, 41(3), 635–643. https://doi.org/10.1111/jlme.12073

Wang, Yun. "Bayesian persuasion with multiple receivers." Available at SSRN 2625399 (2013).