

Targeted Persuasion

Gabriele Gratton Christopher Teh DJ Thornton*

August 29, 2026

Abstract

A sender wishes to persuade a group of receivers but can only communicate with one target receiver. The target's actions, in turn, may influence others. Who should she target? How should she customize her communication? The sender optimally customizes communication to the target as if other receivers did not exist and trades off popularity against influence: she may optimally target less popular receivers. We characterize a measure of a target's market influence that the sender maximizes and show that, in large markets, she strictly prefers less customizable communication if the market is more polarized or if high-popularity influencers are rare.

*Gratton: UNSW Business School, UNSW Sydney and CEPR. Email: g.gratton@unsw.edu.au. Teh: Toulouse School of Economics. Email: christopher.teh@tse-fr.eu. Thornton: UNSW Business School, UNSW Sydney. Email: d.thornton@unsw.edu.au. We thank Anton Kolotilin, Barton E. Lee, Clement Minaudier, and Nageeb Ali for useful feedback. Gratton is the recipient of an Australian Research Council Future Fellowship (FT210100176) and his research was supported under the Australian Research Council's Discovery Projects funding scheme (Project DP240103257). Teh gratefully acknowledges financial support from the ERC Grant 101088307.

1 Introduction

Startup entrepreneurs aim to persuade many investors to fund their startups. But conveying their business idea often requires complex explanations and lengthy demonstrations. As a result, they are limited to privately communicate with only a handful of potential investors they have targeted in their networking activities and with whom they have built stronger connections. Yet, if an entrepreneur successfully persuades her targets to invest, then she may reasonably hope that their example will compel others to do the same. Therefore, an entrepreneur needs to choose both *who to target* in her networking activities and *how to persuade them*. Because she communicates with her target privately, whether a target's example influences others depends on what they believe she told him. This in turn will affect the optimal choice of who to target. The entrepreneur may, for instance, choose to target fans of her ideas, since they are easier to persuade. Or, on the contrary, the entrepreneur may target famously skeptical receivers, in the hope that their example will be more compelling to others. Finally, the entrepreneur's choice of target may take into account that an example set by a more popular target is likely to be observed by—and hopefully influence—a greater number of other investors.

This dual problem of who to target and how to persuade them is common to many markets and organizations, from marketing strategies aimed at persuading “influencers” to adopt a new product or technology, to lobbying efforts aimed at popular policymakers whose example may compel others to vote for a new law or adopt a new policy. In fact, this problem may be more common in the era of social media, when many users can observe the example set by early adopters of a new product, technology, or policy position, but can only indirectly infer what these adopters knew when they made their choice.

In all these situations, an interested party (“Sender”) can choose (i) who to target in networking activities and (ii) what and how to communicate to her target. Crucially, she cares about persuading as many decision-makers (“receivers”) as possible, but non-targeted receivers can only observe actions taken by the target (and not what was communicated). In this paper we introduce a class of one-Sender, many-receivers models of these problems of targeted persuasion. We show that, in this class of models, strategic interactions are driven by two intuitive properties of all perfect Bayesian equilibria. First, in choosing how to persuade a target, Sender acts *as if* no other receiver exists: she fully customizes her communication to maximize the probability that the target buys the widget, even though Sender's objective is, in fact, to persuade as many receivers as possible. Second, in choosing who to target, Sender recognizes that a target's popularity is a double-edged sword and that, in equilibrium, the most popular receiver may not be the

most influential. Together, these two insights drive the optimal choice of target and imply a third result: that Sender may prefer *less* freedom in her ability to customize information to her target. In fact, less freedom of customization reduces Sender’s need to trade off popularity with influence.

A 3-receiver example. To see how these three insights emerge, consider the following example. Sender wishes to persuade 3 receivers—Ann, Bob, and Charlie—to buy a widget that is equally likely to be good or bad. Each receiver is willing to buy if their posterior belief that the widget is good is greater than or equal to their idiosyncratic *skepticism*. Ann is a *fan* of Sender’s products: her skepticism is $\sigma_A = .4$, so that, absent any further information, she would buy. In contrast, Bob and Charlie are *skeptics*: they would need to be persuaded to buy. Charlie is more skeptical than Bob because his skepticism is $\sigma_C = .8$ while Bob’s is $\sigma_B = .6$. Ann, Bob, and Charlie also differ in how popular they are. Bob’s popularity is $\pi_B = .9$, meaning that, if Sender targets Bob, his choice of whether to buy the widget is observed by other receivers with probability .9. Charlie is the least popular of the three: if Sender targets him, his choice is observed with probability $\pi_C = .3$. Ann is the most popular: if Sender targets her, her choice is observed with probability $\pi_A = 1$.

To see how full customization emerges, consider Sender’s choice of communication if she has targeted Bob. To fix ideas, suppose that Sender can freely and flexibly customize evidence for Bob so that communication takes the form of Bayesian persuasion à la [Kamenica and Gentzkow \(2011\)](#). Sender’s choice of information structure (a Blackwell experiment) and the signal she observes are privately communicated to Bob; Ann and Charlie observe at most Bob’s choice of whether to buy. Intuitively, Sender could prepare evidence such that, if she were to persuade Bob to buy, Bob’s positive example would also compel the more skeptical Charlie to buy. For this to be the case, in equilibrium Charlie should infer that, when Bob buys the widget, his posterior belief that the widget is good is at least equal to $\sigma_C = .8$. But this cannot be an equilibrium because communication is private and Sender cannot publicly commit to how she communicates with Bob. If Bob’s purchase would bring Charlie along, Sender would simply maximize the probability that Bob buys. Standard Bayesian persuasion logic implies that Sender optimally fully customizes evidence to the level of Bob’s skepticism: Bob’s posterior when he buys equals exactly $\sigma_B = .6 < \sigma_C$. In equilibrium, Charlie correctly anticipates Sender’s optimal choice and Bob’s example does not compel him to buy. We show that this logic is not specific to Bayesian persuasion: in equilibrium, Sender always fully customizes her communication to her chosen target, communicating with the target as she would if her only objective were to persuade him alone.

To see why popularity is a double-edged sword, consider Sender’s choice of targeting Bob. Because of our first result, if Sender persuades him to buy, his positive example cannot compel any other receiver to buy: Charlie is too skeptical to be compelled and Ann is buying the widget anyway. In contrast, if Sender fails to persuade Bob, his negative example also compels Ann to change plans and not buy the widget, since she infers that Bob has learned the widget is certainly bad. Therefore, Bob’s great popularity is a liability for Sender, because it increases the chance that, by failing to persuade Bob, Sender may lose another customer. In our terminology, Bob is very popular, but his equilibrium *influence* is negative. In fact, in this example Sender strictly prefers to target Charlie: he is harder to persuade, but his lower popularity implies that failing to persuade him is less likely to lose other customers. Targeting Charlie also strictly dominates targeting Ann or targeting no receiver at all.¹ In fact, these two choices are equivalent for Sender when she can flexibly customize her communication to the target. Because Ann is a fan, Sender can say nothing when targeting her and Ann then buys for sure. Her example is uninformative and her popularity is therefore immaterial for Sender.

While in this example Bob’s greater popularity is a liability, this is only because Sender can freely customize information for the target, so that, in equilibrium, Bob’s example cannot compel Charlie to buy. But suppose instead that Sender’s hands were tied so that, when communicating with the target, she could only induce posteriors equal to 0 and $\sigma_C = .8$. Then the target’s example carries the same information regardless of the target’s skepticism. Bob’s popularity is now an asset, and Sender would now prefer to target him rather than Charlie. Yet, now that Sender’s hands are tied, Ann’s popularity also matters, since her example is no longer uninformative, and Sender optimally targets her.² In fact, Sender is now better off *because* her hands are tied and she is unable to flexibly customize information to the target’s skepticism.

Preview of the results. In Section 2 we show that these three insights—full customization, the double-edged sword of popularity, and the suboptimality of more flexible customization—are remarkably general. We study a benchmark model of targeted persuasion in which (i) Sender chooses a target receiver and (ii) privately communicates with the target about the quality of a widget; (iii) Sender’s objective is to convince as many receivers as possible to buy the widget, but non-targeted receivers can only observe the target’s

¹The expected sales from targeting Charlie are $\frac{5}{8} + (1 - \frac{3}{10} \times \frac{3}{8}) + \frac{3}{10} \times \frac{5}{8} = \frac{17}{10} > 1$, whereas targeting Ann (or no receiver at all) yields one sale to Ann. If Sender instead targets Bob, expected sales are $\frac{5}{6} + (1 - \frac{9}{10} \times \frac{1}{6}) = \frac{101}{60} < \frac{17}{10}$.

²Targeting Ann yields $15/8$ expected sales, compared with $37/20$ from targeting Bob and $17/10$ from targeting Charlie.

positive or negative example: his choice of whether to buy or not. Our model allows for an arbitrary persuasion mechanism through which Sender communicates with the target. Furthermore, our three key results carry over to a much broader range of microfoundations of the targeting technology and of Sender’s market.

In our model, the double-edged sword of popularity leads to the outcome that, in equilibrium, Sender may choose a target *despite* his popularity—the equilibrium target’s popularity may be a liability and Sender would prefer a less popular target. In fact, were this target’s popularity to increase, Sender would optimally switch to another target, or even no target at all. This perhaps counterintuitive equilibrium feature emerges because Sender may choose a target just to persuade him. In fact, Sender hopes nobody sees his example. But the gains from a single individual sale are less important for Sender as the number of receivers grows large. Therefore, this feature disappears in large markets.

In Section 3 we study optimal targeted persuasion in large markets and deliver three sharper results. First, large markets reduce optimal targeting under Bayesian persuasion to a one-dimensional problem: Sender chooses the receiver who maximizes *market influence*, the product of popularity and influence. Second, this representation yields ratio-order comparative statics for changes in both the distribution of skepticism and the distribution of popularity. Third, we derive a tight upper bound on Sender’s payoff across all persuasion mechanisms and identify two forces that make Bayesian persuasion strictly suboptimal, so that Sender strictly prefers less flexibility in customizing information for her target. The first force operates through popularity. When popularity varies across skepticism, Sender faces a tradeoff: with more flexible customization, the most popular target—whose example reaches the largest audience—need not have the highest influence. Tying Sender’s hands can align popularity with influence: Sender can preserve influence while targeting a more popular but less skeptical receiver, potentially even a fan. The ability to preserve influence when popularity varies with skepticism may be particularly important in markets with only a handful of high-popularity superstars among many low-popularity, high-skepticism receivers. The second force operates through the distribution of skepticism, even when popularity is constant. We show that Bayesian persuasion is optimal when the market has few “hardcore” fans—i.e., individuals with very low skepticism—but strictly suboptimal when skepticism is concentrated at the extremes. Sender therefore prefers to have her hands tied in polarized markets.

Related literature. Our paper contributes to a growing literature on information design with multiple receivers. In this literature, Sender publicly commits to “how” she communicates—the information structure she chooses—and either publicly communi-

cates to all receivers (Alonso and Câmara, 2016; Goldstein and Huang, 2016; Inostroza and Pavan, 2025; Khantadze et al., 2025) or privately communicates to a subset of them (Arieli and Babichenko, 2019; Awad, 2020; Awad and Minaudier, 2026; Bardhi and Guo, 2018; Caillaud and Tirole, 2007; Chan et al., 2019; Li et al., 2023; Mathevet and Taneva, 2022; Morris et al., 2024; Schnakenberg, 2017; Taneva, 2019; Wang, 2015).³ In contrast, we study environments in which Sender privately chooses both how and what she communicates to her target. This distinction is why, in our model, Sender fully customizes communication to her target, while in these other models she optimally chooses how to communicate to persuade more receivers than just the target. In other words, in our model Sender is unable to commit to how she communicates.

Akbarpour and Li (2020); Crawford and Sobel (1982); Kreutzkamp and Lou (2025); Lin and Liu (2024); Lipnowski et al. (2022); Min (2021) also study communication with limited commitment. However, in their work, limited commitment arises because Sender can modify the message the target observes, or because she partially learns the state before choosing how to communicate. Instead, in our setting, even if Sender can privately commit fully to how she communicates with her target, she cannot publicly do so. Equilibrium communication is therefore constrained not because Sender cannot fool the target, but because she cannot fool other receivers.

Finally, our paper also relates to the literature on optimal targeting in networks (e.g., Ballester et al., 2006; Galeotti and Goyal, 2009). There, optimal targets are characterized by a network-position measure (e.g., Bonacich centrality), and the most-connected (or “popular”) agent need not be the best target. Our setting differs by decomposing targeting into two independent dimensions: popularity and influence. A popular target may be unattractive if his influence is too low.

2 A model of targeted persuasion

We now introduce a benchmark model of targeted persuasion. In this model, Sender earns a linear payoff in the number of receivers who buy the widget, Sender chooses a single target while still uninformed about the widget quality, and each non-targeted receiver knows the target’s identity and may observe the target’s choice, but not the choices of other receivers.⁴ To simplify exposition, we formally study a reduced-form version of this environment in which Sender is the only strategic player while each receiver’s choice

³There are also papers on communication in networks where receivers may observe their neighbors’ signals before taking an action (Babichenko et al., 2022; Candogan et al., 2020; Egorov and Sonin, 2019; Galperti and Perego, 2025).

⁴We briefly return to these assumptions in Section 2.5.

is mechanically determined by his posterior belief about the quality of the widget. That is, we model receivers as Bayesian decision rules: after forming correct posterior beliefs, a receiver buys if and only if his posterior belief weakly exceeds his skepticism.⁵

2.1 Setup

There is a Sender (“she”) and a market of $R \geq 2$ receivers (“he”), indexed by $r \in \mathcal{R} \equiv \{1, \dots, R\}$.

Sender. Sender’s payoff is given by the number of receivers who buy a widget of uncertain quality, $\theta \in \{G, B\}$. The widget is *good* ($\theta = G$) with probability $\mu \in (0, 1)$ and *bad* ($\theta = B$) otherwise.

Market: Receivers’ skepticism and popularity. Each receiver $r \in \mathcal{R}$ has unit demand and is characterized by two parameters: his *skepticism*, $\sigma_r \in (0, 1]$, and his *popularity*, $\pi_r \in (0, 1]$. Receiver r buys the widget if and only if he believes it is good with sufficiently high probability: his posterior belief p_r that $\theta = G$ satisfies $p_r \geq \sigma_r$. Receiver r is a *fan* if $\sigma_r \leq \mu$ and a *skeptic* otherwise. Let $\mathcal{F} \equiv \{r \in \mathcal{R} : \sigma_r \leq \mu\}$ denote the set of fans. To rule out uninteresting cases, we assume that Sender’s market includes at least one fan and at least one skeptic: $\mathcal{F} \neq \emptyset$ and $\mathcal{F} \subset \mathcal{R}$.

Targeted persuasion. A targeted persuasion game plays out as follows. At the onset, Sender chooses a single *target*, $t \in \mathcal{R} \cup \{\emptyset\}$, where $t = \emptyset$ represents Sender targeting no receiver at all, in which case all receivers choose whether to buy the widget without any further observation. Instead, if $t \in \mathcal{R}$, then (1) Sender privately communicates with t according to a publicly known persuasion mechanism (specified below); (2) t chooses whether to buy the widget, and (3) each non-targeted receiver $r \neq t$ independently observes t ’s choice with probability equal to his popularity, π_t . Importantly, non-targeted receivers do not observe how Sender communicates with the target or what she said to him. They only observe, with probability π_t , his “example”, and then decide whether to buy the widget.

We characterize the set of perfect Bayesian equilibria of this class of games. See Appendix A for a detailed description of the solution concept.

⁵A fully strategic formulation would specify receivers’ payoffs for which the same threshold rule is optimal on the equilibrium path. The results below continue to hold in that formulation for equilibria satisfying a natural monotonicity restriction on off-path beliefs: another receiver’s belief about the widget’s quality cannot be lower after observing that the target buys the widget than after observing that the target does not buy it. We state this restriction formally in Appendix A.2. An analogous off-path restriction is called belief monotonicity by Halac and Kremer (2020); see also Kartik (2009, Appendix B) for a related monotonicity restriction on receiver responses to ordered messages.

Persuasion mechanism. A publicly known persuasion mechanism specifies what information Sender can acquire about θ and what she can privately communicate to the target.⁶ It consists of a pair, (Φ, c) : a set Φ of information structures, $\varphi : \{G, B\} \rightarrow \Delta S$, and a cost function, $c : M \times S \rightarrow \{0, \infty\}$, where S is a set of signals Sender can observe about θ , and M a set of messages that Sender can communicate to the target. Sender first privately communicates to the target her choice of information structure: the probability $\varphi(s \mid \theta)$ that she observes signal $s \in S$ conditional on the widget quality, θ . Then, Sender observes signal s and chooses which message m to privately communicate to the target at a cost $c(m, s)$. Throughout, we restrict attention to mechanisms under which every message (after every signal) is either free or infeasible: $c(m, s) \in \{0, \infty\}$ for every $(m, s) \in M \times S$.

It will be convenient to capture a persuasion mechanism with a reduced-form notation. We show in Appendix A that an equilibrium pins down a collection of sets $\mathcal{P} = \{P_t\}_{t \in \mathcal{R}}$, one for each potential target $t \in \mathcal{R}$, of distributions over a target's posterior, p_t , that Sender can induce under the mechanism.⁷

A canonical protocol we will return to throughout this paper is *Bayesian persuasion* (Kamenica and Gentzkow, 2011): (i) $S = M$, (ii) Φ contains all information structures, and (iii) $c(m, s) = 0$ if $m = s$, and $c(m, s) = \infty$ otherwise. In this case, for all $t \in \mathcal{R}$,⁸ P_t is simply the set of Bayes-plausible distributions over posteriors:

$$P_t = P_{BP} \equiv \left\{ \rho \in \Delta[0, 1] : \int_0^1 p \, d\rho(p) = \mu \right\}. \quad (1)$$

For comparison, under *information disclosure* (Dye, 1985; Milgrom, 1981): (i) messages are subsets of signals, $M \subseteq 2^S$, (ii) Sender is endowed with a fixed information structure, $\Phi = \{\varphi_0\}$, and (iii) $c(m, s) = 0$ if $s \in m$, and $c(m, s) = \infty$ otherwise. We illustrate the reduced-form P_t in the simple binary case in which the information structure generates a signal $s \in S = \{G, B\}$ that fully reveals quality: $\varphi_0(\theta \mid \theta) = 1$. The set of feasible messages upon observing signal s is then given by $M(s) = \{\{s\}, n\}$ for all s , where $n \equiv S$ is nondisclosure. For each signal s , let η_s be the probability that Sender discloses s (i.e., sends $m = \{s\}$ rather than n). For each target t : sending message $m = \{G\}$ induces posterior $p_t(\varphi_0, \{G\}) = 1$; message $m = \{B\}$ induces $p_t(\varphi_0, \{B\}) = 0$; message n induces

⁶This is a special case of the persuasion mechanisms introduced in the online appendix supplement to Kamenica and Gentzkow (2011), which allow for message costs to be strictly positive but not infinite.

⁷While working directly with the distributions over posteriors induced by communication is standard in the literature (e.g., Gentzkow and Kamenica, 2016; Kamenica and Gentzkow, 2011; Lipnowski and Ravid, 2020), our reduced-form objects are *equilibrium outcomes* of the underlying communication game, *not* primitives. This distinction matters because Sender's communication with the target is private and other receivers do not observe how Sender communicates with him (the chosen information structure, ϕ).

⁸Notice that under this communication protocol, Sender can choose to "say nothing" to the target.

some $p_t(\varphi_0, n)$. Let $\rho_{\eta_G, \eta_B}(p_t)$ denote the distribution of the target's posterior induced by (η_G, η_B) . Then, for all $t \in \mathcal{R}$,⁹

$$P_t = P_{\text{ID}}(p_t) \equiv \{\rho_{\eta_G, \eta_B}(p_t) : (\eta_G, \eta_B) \in [0, 1]^2\}. \quad (2)$$

2.2 How to persuade a target

We now study how Sender optimally chooses to persuade a target. That is, suppose that Sender has targeted a receiver $t \in \mathcal{R}$ —possibly not the equilibrium target. How should Sender communicate with him in order to maximize the number of receivers who buy the widget? Theorem 1 says that Sender fully customizes her communication: she communicates with t as if no other receiver existed.

Theorem 1 (Full customization). *In any equilibrium, for any receiver $t \in \mathcal{R}$, Sender's communication when targeting t induces a distribution $\rho_t^* \in P_t$ over t 's posteriors, p_t , that maximizes the probability*

$$B_t(\rho_t) \equiv \int_0^1 \mathbb{I}\{p_t \geq \sigma_t\} d\rho_t(p_t) \quad (3)$$

that t buys the widget. That is, $\rho_t^* \in \arg \max_{\rho_t \in P_t} B_t(\rho_t)$.

To gain intuition, consider how the target's example affects other receivers. In any continuation equilibrium, if t buys the widget, then t 's posterior is weakly greater than σ_t . If instead he does not buy the widget, his posterior is strictly less than σ_t . Thus, from the point of view of any other receiver, observing that t buys the widget is information that increases the probability that the widget is good, while observing that t does not buy the widget is information that decreases the probability that the widget is good. It follows that Sender's sales are increasing in the probability $B_t(\rho)$ that t buys the widget. Thus, the Sender's optimal communication strategy is to fully customize her communication to t 's skepticism—as if no other receiver existed.

An immediate consequence of Theorem 1 is that, in equilibrium, non-targeted receivers' behavior depends on Sender's private communication with t only through three summary statistics: the equilibrium probability $B_t(\rho_t^*)$ that t buys, and the two posterior beliefs induced by observing, respectively, that t buys and does not buy the widget. When

⁹Unlike under Bayesian persuasion, where P_{BP} is target- and equilibrium-independent, here the set P_{ID} depends on the equilibrium belief function p_t , illustrating why the sets P_t are equilibrium outcomes rather than primitives. We give further details in Appendix A.1, where we also discuss cheap talk (Crawford and Sobel, 1982) and constrained Bayesian persuasion Doval and Skreta (2024).

a non-targeted receiver observes that t buys the widget, he learns that t 's posterior lies in $[\sigma_t, 1]$; when he observes that t does not buy the widget, he learns that t 's posterior lies in $[0, \sigma_t]$. Hence, when $B_t(\rho_t^*) \in (0, 1)$ —so the target's choice is *informative*—these posteriors are given by,

$$p_t^1(\rho_t^*) \equiv \mathbb{E}_{\rho_t^*}[p_t \mid p_t \geq \sigma_t], \quad p_t^0(\rho_t^*) \equiv \mathbb{E}_{\rho_t^*}[p_t \mid p_t < \sigma_t]. \quad (4)$$

Notice that the triplet $(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*))$ is sufficient to pin down the expected probability that each receiver buys (and hence Sender's expected payoff) if Sender targets receiver t .

Bayesian persuasion. We now demonstrate the implications of Theorem 1 when communication takes the form of Bayesian persuasion. Suppose first that Sender targets a skeptic: $\sigma_t > \mu$. By [Kamenica and Gentzkow \(2011\)](#), the distribution of posteriors maximizing $B_t(\rho)$ has mass μ/σ_t on posterior σ_t and mass $1 - \mu/\sigma_t$ on posterior 0. The resulting triplet is

$$(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*)) = \left(\frac{\mu}{\sigma_t}, \sigma_t, 0 \right) \quad (5)$$

and the target's choice is informative: he buys with probability μ/σ_t , other receivers who observe a positive example update to σ_t , and other receivers who observe a negative example update to 0. Notice that, if Sender fails to persuade t to buy the widget, then all receivers who observe his example also do not buy. If Sender instead targets a fan ($\sigma_t \leq \mu$) then she optimally induces him to buy with certainty, so the resulting triplet is¹⁰

$$(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*)) = (1, \mu, p_t^0) \quad (6)$$

and the target's action is uninformative.

Information disclosure. Full customization is not specific to Bayesian persuasion. Return, for example, to the binary information-disclosure case and suppose that the target is a skeptic, $\sigma_t > \mu$. First, notice that his equilibrium belief following nondisclosure must satisfy $p_t(\varphi_0, n) < \sigma_t$. Otherwise, nondisclosure would induce him to buy and Sender would never disclose $s = B$, so that Bayes' rule would yield $p_t(\varphi_0, n) \leq \mu < \sigma_t$. It follows that, among the posteriors in $P_{\text{ID}}(p_t)$, only posterior 1 induces the target to buy. Theorem 1

¹⁰The off-path belief p_t^0 is not uniquely pinned down by ρ_t^* and is irrelevant to Sender's equilibrium payoff.

then implies that Sender maximizes

$$\int_0^1 \mathbb{I}\{p \geq \sigma_t\} d\rho_{\eta_G, \eta_B}(p | p_t) = \mu\eta_G.$$

That is, Sender optimally chooses $\eta_G = 1$: she discloses the good signal, no feasible message following the bad signal induces the target to buy, and the target buys with probability μ . The resulting triplet is therefore

$$(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*)) = (\mu, 1, 0)$$

and the target's choice fully reveals the quality of the widget. Contrasting this result to the case of Bayesian persuasion, changing the persuasion mechanism changes Sender's feasible set, but not the theorem's conclusion: in equilibrium, Sender never sacrifices the probability that the target buys in order to make his example more compelling to other receivers.

When instead Sender targets a fan, the case of information disclosure highlights the importance of recognizing that P_t is indeed an equilibrium outcome. I.e., Theorem 1 says that Sender communicates with her target as if no other receiver existed and maximizes the probability that the target buys. However, the communication game between Sender and her target may itself have multiple equilibria, so that the optimal communication that maximizes the probability that the target buys depends on the equilibrium target's posteriors. In the binary case, information disclosure can support two types of equilibrium outcomes. In a pooling equilibrium (or partial-pooling equilibrium with $p_t(\varphi_0, n) \geq \sigma_t$), Theorem 1 says that Sender does not disclose $s = B$ and the target buys with probability 1. His purchase is therefore uninformative. Yet there are also separating equilibria supported by $p_t(\varphi_0, n) < \sigma_t$. In these equilibria, Sender discloses $s = G$, the target buys if and only if the widget is good, and the triplet is $(\mu, 1, 0)$, just as for a skeptic.

2.3 Customization and communication unraveling

The key insight captured by Theorem 1 is that observing the target's example (his choice) *unravels* all the information that Sender has privately communicated to the target. Ideally, Sender would prefer to credibly announce that she will communicate with the target in the way that maximizes the compelling power of his example. Yet, in equilibrium, non-targeted receivers correctly anticipate that, in the privacy of her communication with the target, Sender would renege on her announcement and fully customize communication to maximize the probability that the target buys the widget. Any attempt to fool

other receivers by promising a different type of communication therefore also unravels—in equilibrium, all receivers know that the target’s posterior distribution is given by ρ_t^* .

We now make this point more precise by illustrating the *cost* of full customization in the case of Bayesian persuasion. Recall that—as in (5)—Sender optimally induces posterior $p_t^0(\rho_t^*) = 0$ among those receivers who observe that the target does not buy. That is, if Sender fails to persuade t to buy the widget, then all receivers who observe his example also do not buy.

It is instructive to compare our model of *private* and *targeted* persuasion with one in which Sender can *publicly* and *credibly* commit to how she will communicate with the target. That is, she can publicly commit to the information structure φ that generates information for the target. Notice that, like in our model, only the target (*not* other receivers) observes *what* Sender communicates: the signal generated by φ . Yet, in contrast with our model, all receivers observe, in addition to the target’s example, also *how* Sender communicates with him: the information structure, φ .

We now show that the Sender’s optimal choice in this case does not fully customize communication to the target. To see this, focus on the set of binary-support distributions over t ’s posteriors that Sender may induce with her choice of φ .¹¹ I.e., Sender chooses two probabilities, p^0 and p^1 , with mean μ and such that $p^0 < \sigma_t \leq p^1$. Fully customizing information to t implies choosing $p^0 = 0$ and $p^1 = \sigma_t$. To characterize Sender’s problem with public commitment, we introduce new notation that will be useful throughout the remainder of the paper. Let

$$\hat{F}_t(x) \equiv \frac{1}{R-1} |\{r \in \mathcal{R} \setminus \{t\} : \sigma_r \leq x\}|$$

be the *empirical distribution of skepticism among receivers different from t* . With public commitment Sender’s optimal (p^0, p^1) solves

$$B_t(\rho) + (R-1)\pi_t[B_t(\rho)\hat{F}_t(p^1) + (1-B_t(\rho))\hat{F}_t(p^0)] \quad (7)$$

In many cases, and in contrast with our model, the solution to Sender’s problem entails a non-zero p^0 . That is, in contrast to our model of private communication, Sender prefers *not* to fully customize her communication to t . This allows Sender to avoid losing all non-targeted receivers who observe a negative example: with public commitment, even when Sender fails to persuade the target, the least skeptical non-targeted receivers (those with $\sigma_r \leq p^0$) buy the widget even if they observe the target’s negative example.

¹¹By Proposition 1 of [Kamenica and Gentzkow \(2011\)](#), it is without loss to restrict attention to binary support distributions under this form of public Bayesian persuasion.

The gap between private and public persuasion highlights that Sender would benefit from the opportunity of being able to commit *not* to fully customize her communication to the target. A critical observation is that Sender can partially recover this commitment gain through a more constrained persuasion mechanism that rules out some information structures so that her hands are tied and she is unable to simply maximize the probability that the target buys the widget. There are two ways this can happen. First, the underlying persuasion mechanism may prevent Sender from inducing $p^0 = 0$, so that her negative example is less damaging to sales to non-targeted receivers. Second, the persuasion mechanism may decouple the high posterior p^1 from the target's skepticism, allowing Sender to induce a high posterior $p^1 > \sigma_t$ while targeting a more popular receiver. We discuss these possibilities and the related tradeoffs in greater detail in Section 3.

2.4 Who to target

We now turn to Sender's choice of target $t \in \mathcal{R} \cup \{\emptyset\}$. To this end, we characterize Sender's equilibrium gains and losses from a given target's example: how many other receivers may be compelled to buy or not buy the widget if they were to observe his choice. Importantly, Sender always retains the option of choosing not to target any receiver ($t = \emptyset$), in which case Sender sells to all (and only all) fans and obtains a payoff of $|\mathcal{F}|$.

Suppose instead Sender targets a receiver $t \in \mathcal{R}$. Recall that, in any equilibrium, targeting receiver t entails inducing a triplet $(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*))$. Sender's gain from t 's positive example is measured by its compelling power over skeptical non-targeted receivers whose skepticism is less than or equal to $p_t^1(\rho_t^*)$ —those who would not have bought the widget but will after observing t 's example. This gain is

$$G_t(p_t^1(\rho_t^*)) \equiv (R - 1) \left(\hat{F}_t(p_t^1(\rho_t^*)) - \hat{F}_t(\mu) \right).$$

Sender's loss from t 's negative example is measured by its negative compelling power over the non-targeted fans whose skepticism is greater than $p_t^0(\rho_t^*)$ —those who would have bought the widget but do not after observing the target's example. This loss is

$$L_t(p_t^0(\rho_t^*)) \equiv (R - 1) \left(\hat{F}_t(\mu) - \hat{F}_t(p_t^0(\rho_t^*)) \right).$$

Weighting gains and losses by the probability $B_t(\rho_t^*)$ of generating a positive example, and normalizing by the size of Sender's market, R , we can define a quantity that will be useful to characterize the optimal target in the remainder of the paper. This quantity captures a target's equilibrium average effect, or *influence*, on other receivers.

Definition 1 (Influence). *For any receiver $t \in \mathcal{R}$, in any equilibrium in which targeting t induces $(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*))$, t 's equilibrium influence $\mathcal{I}_t(\rho_t^*)$ is given by*

$$\mathcal{I}_t(\rho_t^*) \equiv \frac{1}{R} [B_t(\rho_t^*) G_t(p_t^1(\rho_t^*)) - (1 - B_t(\rho_t^*)) L_t(p_t^0(\rho_t^*))].$$

Sender's value of targeting $t \in \mathcal{R}$, V_t , is then given by her direct probability of persuading him, $B_t(\rho_t^*)$, plus his *market influence*: his influence, $\mathcal{I}_t(\rho_t^*)$, multiplied by his popularity, π_t .¹²

Proposition 1 (The value of targeting a receiver). *If there exists an equilibrium in which a receiver $t \in \mathcal{R}$ has value*

$$V_t \equiv \frac{R-1}{R} \hat{F}_t(\mu) + \frac{1}{R} B_t(\rho_t^*) + \pi_t \mathcal{I}_t(\rho_t^*) > \frac{|\mathcal{F}|}{R} \quad (8)$$

then in that equilibrium, Sender targets a receiver $t \in \mathcal{R}$ with the greatest value. Otherwise, in all equilibria, Sender attains a payoff equal to that induced by targeting no receiver at all: $|\mathcal{F}|$.

Notice that we normalize the value V_t by the size of Sender's market, R . The first term in V_t is the share of non-targeted fans who buy the widget if Sender reveals to t no information about the widget quality at all. The second term captures the direct sale to the target. The third captures the average informational effect of the target's example, his influence, scaled by the target's popularity, π_t .

We now study Proposition 1's implications for how a receiver's popularity affects his value as a potential target. The next result formalizes the paper's second main insight.

Proposition 2 (The double-edged sword of popularity). *For any receiver $t \in \mathcal{R}$, in any equilibrium in which targeting t induces $(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*))$, Sender's payoff from targeting $t \in \mathcal{R}$ is increasing in π_t if $\mathcal{I}_t(\rho_t^*) > 0$, decreasing in π_t if $\mathcal{I}_t(\rho_t^*) < 0$, and independent of π_t if $\mathcal{I}_t(\rho_t^*) = 0$.*

A more popular target has greater reach, but this amplifies both the upside of a positive example and the downside of a negative one. Whether greater popularity of a target helps Sender therefore depends on the sign of his influence, while popularity itself determines only the reach of his influence.

In our model, the double-edged sword of popularity may also have bite in equilibrium: if the equilibrium target's popularity increases, Sender may prefer to switch to

¹²Notice that because a target's equilibrium influence can be negative, Sender's value from targeting $t \in \mathcal{R}$ can be equal to $|\mathcal{F}|/R$, so the expected number of sales is equal to that when Sender targets no receiver, even if Sender's equilibrium communication to t is informative, so $p_t^1(\rho_t^*) \neq p_t^0(\rho_t^*)$.

a different target (as we show in the example below), including perhaps a fan or no receiver at all. This happens when the equilibrium target's influence is negative, but Sender prefers to target him anyway because she hopes at least to generate a direct sale to the target. However, as the number of receivers R increases, the direct-sale term, $B_t(\rho_t^*)/R$, becomes negligible, and Sender's comparative value of targeting a skeptic is captured entirely by his market influence, $\pi_t \mathcal{I}_t(\rho_t^*)$. In the limit as $R \rightarrow \infty$, if in equilibrium Sender targets a skeptic, then the target's influence (and therefore his market influence) must be non-negative; otherwise, Sender would strictly prefer to target a fan or no receiver at all. In large markets, popularity therefore ceases to be a double-edged sword in equilibrium: only the positive edge can cut. We investigate this limiting case with a continuum of receivers in Section 3. Before doing so, we illustrate how Propositions 1 and 2 precisely pin down the intuition behind our motivating example in Section 1.

Bayesian persuasion. Recall that when communication takes the form of Bayesian persuasion, Sender's equilibrium distribution over posteriors from targeting a skeptic t is captured by the triplet $(B_t(\rho_t^*), p_t^1(\rho_t^*), p_t^0(\rho_t^*)) = (\mu/\sigma_t, \sigma_t, 0)$. Therefore,

$$\begin{aligned} \mathcal{I}_t(\rho_t^*) &= \frac{R-1}{R} \left[\frac{\mu}{\sigma_t} \left(\hat{F}_t(\sigma_t) - \hat{F}_t(\mu) \right) - \left(1 - \frac{\mu}{\sigma_t} \right) \hat{F}_t(\mu) \right] \\ &= \frac{R-1}{R} \left[\frac{\mu}{\sigma_t} \hat{F}_t(\sigma_t) - \hat{F}_t(\mu) \right]. \end{aligned} \quad (9)$$

Sender's value of targeting a skeptic t is therefore

$$V_t = \frac{1}{R} \left[(R-1) \hat{F}_t(\mu) + \frac{\mu}{\sigma_t} + \pi_t (R-1) \left(\frac{\mu}{\sigma_t} \hat{F}_t(\sigma_t) - \hat{F}_t(\mu) \right) \right]. \quad (10)$$

If instead Sender targets a fan t , $\mathcal{I}_t(\rho_t^*) = 0$ and $V_t = |\mathcal{F}|/R$, because—when the communication protocol is Bayesian persuasion—Sender can simply “say nothing” to him. I.e., targeting a fan induces the same payoff as targeting no receiver at all. Hence, in equilibrium, Sender targets the skeptic t that maximizes (10) whenever there is a skeptic whose value exceeds $|\mathcal{F}|/R$, and targets a fan or no receiver at all otherwise.

Intuitively, a more skeptical target is harder to persuade: μ/σ_t decreases with σ_t . Yet, his positive example compels more receivers to buy the widget: $\hat{F}_t(\sigma_t) - \hat{F}_t(\mu)$ increases with σ_t . Furthermore, notice that $\mathcal{I}_t(\rho_t^*) > 0$ if and only if

$$\frac{\sigma_t}{\mu} < \frac{\hat{F}_t(\sigma_t)}{\hat{F}_t(\mu)}.$$

Proposition 2 therefore specializes to a simple condition on the distribution of skepticism among non-targeted receivers.

An example. To see the double-edged sword of popularity at work, return to the introductory example with three receivers: Ann ($\sigma_A = .4$), Bob ($\sigma_B = .6$), and Charlie ($\sigma_C = .8$), with $\mu = .5$. Ann is a fan; Bob and Charlie are skeptics. Charlie’s skepticism is greater than Bob’s—so that he’s harder to persuade. How does the popularity of Bob and Charlie (i.e., π_B and π_C) affect Sender’s choice between the two?

Under Bayesian persuasion, targeting Bob yields $B_B(\rho_B^*) = 5/6$ and $p_B^1(\rho_B^*) = \sigma_B = .6 < \sigma_C$, so Bob’s positive example does not persuade Charlie. However, $p_B^0(\rho_B^*) = 0 < \sigma_A$, so Bob’s negative example causes Ann to stop buying. Because Bob’s example can only hurt Sender, $\mathcal{I}_B(\rho_B^*) < 0$, and hence his popularity is a liability: a higher π_B makes targeting him strictly worse. Sender’s value of targeting Bob equals

$$\begin{aligned} V_B &= \frac{1}{3} \left[1 + \frac{5}{6} + \pi_B \cdot 2 \left(\frac{5}{6} \cdot \frac{1}{2} - \frac{1}{2} \right) \right] \\ &= \frac{11}{18} - \pi_B \frac{1}{18}. \end{aligned}$$

In contrast, targeting Charlie yields $B_C(\rho_C^*) = 5/8$ and $p_C^1(\rho_C^*) = \sigma_C = .8 > \sigma_B$, so Charlie’s positive example persuades Bob. Like for Bob’s case, $p_C^0(\rho_C^*) = 0 < \sigma_A$, so a negative example causes Ann to stop buying. Even if Charlie is very hard to persuade, the value of his example is positive, and hence his popularity is an asset for Sender. Sender’s value of targeting Charlie equals

$$\begin{aligned} V_C &= \frac{1}{3} \left[1 + \frac{5}{8} + \pi_C \cdot 2 \left(\frac{5}{8} - \frac{1}{2} \right) \right] \\ &= \frac{13}{24} + \pi_C \frac{1}{12}. \end{aligned}$$

In our initial example, Bob is much more popular than Charlie: $\pi_B = .9 > .3 = \pi_C$. Hence, Sender optimally targets Charlie: $V_B < V_C$. Yet, were Bob to be less popular—e.g., $\pi_B = .7$ —Sender would in fact prefer to target him.

2.5 Robustness and extensions

Readers may be interested in knowing whether the key insights from this section extend to other market microfoundations that better fit specific applications. In [Gratton et al. \(2025\)](#), we show how our key results carry over (and what additional forces arise) if we

allow for: more general utility functions of Sender; the receivers' choices to be strategic complements; Sender to choose multiple targets at once; non-targeted receivers not to observe Sender's choice of target but only if he buys the widget; non-targeted receivers to also observe other non-targeted receivers' choices. Finally, we show that our central full customization result in Theorem 1 does not change if Sender is informed about the quality of the widget before choosing her target.

3 Targeted persuasion in large markets

To obtain sharper comparative static results, we now study a version of our targeted persuasion model in which (i) there is a unit mass of receivers (i.e., Sender's market is large) and (ii) the persuasion mechanism is Bayesian persuasion. Since Bayesian persuasion is the persuasion mechanism that affords Sender maximum flexibility in customizing her communication to the target's skepticism, this is a natural benchmark to understand when and why Sender may benefit from less flexibility.

We write F for the cumulative distribution of skepticism among receivers and assume that F is differentiable with full support on $[0, 1]$. We allow popularity to depend on skepticism: each receiver with skepticism σ has popularity $\pi(\sigma) \in (0, 1]$,¹³ where π is continuous. Notice that a triple (π, F, μ) completely characterizes Sender's market.

In this limit model,¹⁴ the same full-customization logic of Theorem 1 applies: if Sender chooses target t , then in equilibrium Sender maximizes the probability that t buys the widget. Proposition 1 also holds with the additional feature that both a target's influence and his value for Sender can be expressed as functions of his skepticism using the market parameters (π, F, μ) :

$$\mathcal{I}(\sigma \mid F, \mu) \equiv \mathcal{I}_t(\rho_t^*)|_{\sigma_t=\sigma} = \begin{cases} \frac{\mu}{\sigma} F(\sigma) - F(\mu), & \sigma > \mu, \\ 0, & \sigma \leq \mu, \end{cases} \quad (11)$$

$$V(\sigma \mid \pi, F, \mu) \equiv V_t|_{\sigma_t=\sigma} = F(\mu) + \pi(\sigma)\mathcal{I}(\sigma \mid F, \mu). \quad (12)$$

¹³As we will make clear below, when multiple receivers share the same skepticism σ but differ in popularity, and Sender finds it optimal to target a receiver with popularity σ , then Sender is weakly better off from targeting the most popular among them. Therefore, writing $\pi(\sigma)$ as a function rather than a correspondence is without loss, with $\pi(\sigma) = \max_t \{\pi_t : \sigma_t = \sigma\}$.

¹⁴Appendix A.4 formalizes this interpretation. Given any large market (π, F, μ) , we construct finite markets for which, for any targeted receiver, the distribution over other receivers' skepticism converges to F . Sender's optimal finite-market payoff converges to the payoff characterized in Proposition 3 below, and every limit point of equilibrium finite-market targets is optimal under (π, F, μ) .

As we noted above, when the persuasion mechanism is Bayesian persuasion, Sender's choice to target a fan induces the same payoff as targeting no receiver at all: $F(\mu)$. Furthermore, as discussed in Section 2.4, in the limit case as $R \rightarrow \infty$, the comparative value of targeting a skeptic is entirely captured by his market influence, $\pi(\sigma)\mathcal{I}(\sigma \mid F, \mu)$. Most importantly, when the persuasion mechanism is Bayesian persuasion, a receiver's influence, and therefore his market influence, is uniquely determined by the market parameters (π, F, μ) and his skepticism, σ . Therefore, the equilibrium is essentially unique.

Proposition 3 (Optimal targeted Bayesian persuasion). *In equilibrium, Sender's choice of target maximizes the target's market influence. I.e., if there exists*

$$\sigma^* \in \arg \max_{\sigma \in (0,1]} \pi(\sigma)\mathcal{I}(\sigma \mid F, \mu)$$

such that $\pi(\sigma^)\mathcal{I}(\sigma^* \mid F, \mu) > 0$, then Sender targets a receiver with skepticism σ^* and her payoff equals $V(\sigma^* \mid \pi, F, \mu)$. Otherwise, she either targets a fan or no receiver at all and her payoff equals $F(\mu)$.*

An immediate implication of this result is that Sender may need to trade off popularity against influence and that the optimal target—the one with maximum market influence—may neither be the most popular, nor the most influential.

These observations drive the results that follow. Section 3.1 studies how the equilibrium target depends on both the correlation between skepticism and popularity, π , and the distribution of skepticism itself, as captured by how the influence of a target changes with his skepticism, σ_t . Section 3.2 studies how the distribution of skepticism in Sender's market determines when Bayesian persuasion itself is suboptimal—that is, when Sender would prefer a persuasion protocol that ties her hands so that she can credibly induce posteriors different from 0 and σ_t .

3.1 Optimal targeting

By Proposition 1, Sender targets a skeptic only if some skeptic has value at least $F(\mu)$ —the payoff she can guarantee by targeting a fan and saying nothing. Corollary 1 is then immediate from (11) and (12).

Corollary 1 (Equilibria when influence is negative). *If $\mathcal{I}(\sigma \mid F, \mu) \leq 0$ for all $\sigma > \mu$, then:*

1. *There exists an equilibrium in which Sender targets a fan and the target buys the widget with probability 1.*

2. In all equilibria, Sender's payoff equals $F(\mu)$.

If instead the condition in Corollary 1 is not satisfied, Sender's optimal targeting choice depends on how a target's skepticism affects both his influence and his popularity. Consider the simplest case: $\mathcal{I}$ and π are both nondecreasing on $[\mu, 1]$. Then V is nondecreasing in σ , so that targeting the most skeptical receiver is optimal. But when σ affects $\mathcal{I}$ and π in opposite directions, finding the optimal target involves a nontrivial tradeoff between a target's influence and his popularity.

Let

$$T^*(\pi, F, \mu) = \arg \max_{\sigma \in [0,1]} V(\sigma \mid \pi, F, \mu)$$

be the set of receivers who are targeted in at least one equilibrium in market (π, F, μ) . We say that Sender *targets more skeptical receivers* in market (π, F, μ) than in (π', F', μ') if $T^*(\pi, F, \mu)$ dominates $T^*(\pi', F', \mu')$ in the strong set order.¹⁵

Proposition 4 (Influence ratio ordering). *Let (π, F, μ) and (π, F', μ) be two markets differing only by the distribution of skepticism, and suppose that $\mathcal{I}(\sigma \mid F, \mu) > 0$ and $\mathcal{I}(\sigma \mid F', \mu) > 0$ for all $\sigma \in (\mu, 1]$, such that in all equilibria Sender targets a skeptic. If*

$$\frac{\mathcal{I}(\sigma \mid F, \mu)}{\mathcal{I}(\sigma \mid F', \mu)} \tag{13}$$

is increasing in σ on $(\mu, 1]$, then Sender targets more skeptical receivers in (π, F, μ) than in (π, F', μ) .

The key step to derive this result is to show that $V(\sigma \mid \pi, F, \mu)$ single-crossing dominates $V(\sigma \mid \pi, F', \mu)$; the result then follows from Milgrom and Shannon (1994).

Condition (13) is an *influence ratio ordering*: it requires the influence under F to grow faster with skepticism than under F' . Intuitively, F concentrates relatively more mass at higher skepticism levels, after adjusting for the prior μ . Although related to standard distributional orderings, this condition is distinct: it operates on the influences $\mathcal{I}$ rather than directly on the cumulative distributions of skepticism, since Sender's value of targeting a given level of skepticism depends on these distributions only through $\mathcal{I}$.¹⁶ The following example illustrates Proposition 4 for a well-known parametric family.

¹⁵For subsets $A, B \subseteq [0, 1]$, A dominates B in the strong set order if, for every $a \in A$ and $b \in B$, $\max\{a, b\} \in A$ and $\min\{a, b\} \in B$ (see Milgrom and Shannon, 1994).

¹⁶Condition (13) can be interpreted via likelihood ratio dominance by defining a *gain distribution* $H(\sigma) \equiv \frac{F(\sigma)/\sigma - F(\mu)/\mu}{1 - F(\mu)/\mu}$. If F and F' both satisfy *increasing average density*, i.e., $F(\sigma)/\sigma$ is nondecreasing on $[\mu, 1]$, then H is a valid distribution function on $[\mu, 1]$ with $\mathcal{I}(\sigma) = (\mu - F(\mu))H(\sigma)$. Condition (13) then reduces to H/H' increasing—the cumulative distribution function ratio order on gain distributions—which is implied by likelihood ratio dominance of H over H' .

Example (Power family). Consider $F_\alpha(\sigma) = \sigma^\alpha$ on $[0, 1]$, with $\alpha > 1$. This is a $\text{Beta}(\alpha, 1)$ distribution where the mean skepticism $\alpha/(\alpha + 1)$ is increasing in α , i.e., higher α concentrates more mass at higher skepticism levels. The influence is

$$\mathcal{I}(\sigma \mid F_\alpha, \mu) = \mu(\sigma^{\alpha-1} - \mu^{\alpha-1}),$$

which is positive for all $\sigma > \mu$ when $\alpha > 1$, so Sender targets a skeptic. The example-value ratio is

$$\frac{\mathcal{I}(\sigma \mid F_\alpha, \mu)}{\mathcal{I}(\sigma \mid F_{\alpha'}, \mu)} = \frac{\sigma^{\alpha-1} - \mu^{\alpha-1}}{\sigma^{\alpha'-1} - \mu^{\alpha'-1}}.$$

This ratio is increasing in σ on $(\mu, 1)$ whenever $\alpha > \alpha' > 1$, by the monotone L'Hôpital rule: the pointwise derivative ratio $\frac{(\alpha-1)\sigma^{\alpha-2}}{(\alpha'-1)\sigma^{\alpha'-2}} = \frac{\alpha-1}{\alpha'-1} \sigma^{\alpha-\alpha'}$ is increasing in σ . Hence, by Proposition 4, Sender targets more skeptical receivers under F_α than under $F_{\alpha'}$, holding π and μ fixed.

We now turn to the role of popularity. Fixing the other two market parameters, F and μ , it is immediate that Sender's payoff is weakly greater in market (π, F, μ) than in (π', F, μ) whenever $\pi(\sigma) > \pi'(\sigma)$ for all σ . In this sense, *greater market popularity is unambiguously better for Sender*. However, popularity may be concentrated among more or less skeptical targets, thus affecting the relative reach of their example. The following proposition says that greater concentration of popularity among more skeptical receivers induces Sender to target a more skeptical receiver.¹⁷

Proposition 5 (Popularity ratio ordering). Let (π, F, μ) and (π', F, μ) be two markets differing only in popularity, and suppose that $\mathcal{I}(\sigma \mid F, \mu) > 0$ for every $\sigma \in (\mu, 1]$ so that in all equilibria Sender targets a skeptic. If $\pi(\sigma)/\pi'(\sigma)$ is increasing in σ on $(\mu, 1]$, then Sender targets more skeptical receivers in (π, F, μ) than in (π', F, μ) .

3.2 Suboptimality of flexible communication

Under Bayesian persuasion, when Sender chooses a target with skepticism σ , she induces equilibrium posteriors:

¹⁷The comparative statics in Propositions 4 and 5 do not generally reverse when influence is negative. Because targeting a fan or no receiver yields $F(\mu)$, a skeptic with $\mathcal{I}(\sigma \mid F, \mu) < 0$ cannot be optimal. Moreover, the same single-crossing argument can be applied on any region of skepticism that contains the optimal targets of both markets, provided influence is strictly positive and the relevant ratio condition holds on that region. In Proposition 5 the two markets share F , so the positive-influence region is common to both and this local restriction can always be done. In Proposition 4, however, a change in F can change the positive-influence region itself. In that case, the proposition does not directly apply: the optimal target may jump to another positive-influence region or to targeting a fan or no receiver at all.

(if $\sigma > \mu$): with support $\{0, \sigma\}$ among receivers who observe the target's example, and equal to μ among receivers who do not;

(if $\sigma \leq \mu$): equal to μ for all receivers.

This simple structure highlights two features of equilibrium. First, by Theorem 1, Sender always fully customizes communication within the available persuasion mechanism. Second, Bayesian persuasion places the fewest constraints on that customization. Together, these features mean that Bayesian persuasion delivers the highest possible probability of persuading the target.

As discussed in Section 2.3, Sender may prefer to have her hands tied so that she can credibly communicate to the target in a way that maximizes her real objective: the expected number of receivers who buy the widget. Such a mechanism can, by restricting Sender's access to different information structures, induce a distribution over posteriors under which the target buys with lower probability but expected sales to other receivers rise. Or it may allow Sender to induce the same distribution over posteriors 0 and σ , but by targeting a more popular receiver, increasing how many receivers observe the target's positive example.¹⁸

We now identify the characteristics of Sender's market that determine whether Bayesian persuasion is suboptimal. To do so, we first characterize the supremum payoff achievable across all persuasion mechanisms. Recall that, because other receivers only observe the target's positive or negative example, Sender induces a distribution over posteriors with at most two values among those receivers who observe the target's example. Specifically, if she chooses a target t with skepticism σ , then non-targeted receivers who observe the target's positive and negative example hold posteriors, respectively,

$$p^1 = p_t^1(\rho_t^*) = \mathbb{E}_{\rho_t^*}[p_t \mid p_t \geq \sigma_t], \quad p^0 = p_t^0(\rho_t^*) = \mathbb{E}_{\rho_t^*}[p_t \mid p_t < \sigma_t].$$

Therefore, the share of receivers who buy the widget conditional on a positive and negative example is, respectively,

$$\begin{aligned} v(p^1 \mid \pi(\sigma)) &\equiv \pi(\sigma)F(p^1) + (1 - \pi(\sigma))F(\mu) \\ v(p^0 \mid \pi(\sigma)) &\equiv \pi(\sigma)F(p^0) + (1 - \pi(\sigma))F(\mu) \end{aligned}$$

¹⁸One example of a hands-tying protocol is information disclosure (Dye, 1985; Milgrom, 1981). There, Sender is committed to a fixed information structure and may only truthfully report a subset of what she observes; she cannot tailor the high posterior to the target's skepticism. See Gratton et al. (2025) for a 3-receiver example in which Sender strictly prefers information disclosure to Bayesian persuasion.

Furthermore, Bayes-plausibility implies $p^0 \leq \mu \leq p^1$ and, whenever both inequalities are strict, pins down the probability with which the target sets a positive example,

$$B(p^1, p^0) \equiv \frac{\mu - p^0}{p^1 - p^0}.$$

Equivalently, $B(p^1, p^0)$ is the probability that a non-targeted receiver's posterior is p^1 when he observes the target's example.

An upper bound on Sender's payoff across all mechanisms is the supremum of expected sales over all Bayes-plausible posterior pairs (p^0, p^1) , and all targets who set different examples under the two posteriors. That is, $V^*(\pi, F, \mu)$ solves

$$\begin{aligned} V^*(\pi, F, \mu) &\equiv \sup_{(\sigma, p^0, p^1) \in [0,1]^3} \{B(p^1, p^0)v(p^1 | \pi(\sigma)) + (1 - B(p^1, p^0))v(p^0 | \pi(\sigma))\} \\ \text{subject to } &p^0 \leq \mu \leq p^1, \\ \text{and } &\sigma \in (p^0, p^1] \text{ whenever } p^0 < p^1. \end{aligned}$$

Recall from Section 3.1 that Sender is never worse off from an increase in the target's popularity. Therefore,

$$V^*(\pi, F, \mu) = \sup_{p^0 \leq \mu \leq p^1} \{B(p^1, p^0)v(p^1 | \bar{\pi}(p^1, p^0)) + (1 - B(p^1, p^0))v(p^0 | \bar{\pi}(p^1, p^0))\}, \quad (14)$$

where

$$\bar{\pi}(p^1, p^0) \equiv \sup_{\sigma \in (p^0, p^1]} \pi(\sigma)$$

is the highest popularity among targets who set different examples under posteriors p^0 and p^1 .¹⁹

The quantity $V^*(\pi, F, \mu)$ has a natural interpretation when $\bar{\pi}(p^1, p^0)$ is attained: it equals Sender's payoff in a hypothetical environment in which she can publicly and credibly commit to induce posteriors $\{p^0, p^1\}$ and targets the most popular receiver compatible with such posteriors. Then a fraction $\bar{\pi}(p^1, p^0)$ of receivers observe the target's example, while the remaining receivers retain the prior μ .²⁰ The following Lemma says that $V^*(\pi, F, \mu)$ is also a *tight* upper bound on the Sender's equilibrium payoff across all persuasion mechanisms.

¹⁹When $p^0 = p^1$, Bayes-plausibility forces $p^0 = p^1 = \mu$. We adopt the conventions $B(\mu, \mu) \equiv 0$ and $\bar{\pi}(\mu, \mu)$ equal to any value in $[0, 1]$. These conventions do not affect V^* : at $p^0 = p^1 = \mu$, both terms in the displayed objective equal $v(\mu | \bar{\pi}(\mu, \mu)) = F(\mu)$ regardless.

²⁰If $\bar{\pi}(p^1, p^0)$ is not attained, take a sequence $\{\sigma_n\} \subset (p^0, p^1]$ such that $\pi(\sigma_n) \rightarrow \bar{\pi}(p^1, p^0)$ and interpret this description in the limit.

Lemma 1 (Maximum payoff across persuasion mechanisms). *For every $\varepsilon > 0$, there exists a persuasion mechanism and an equilibrium under it in which Sender’s payoff is at least $V^*(\pi, F, \mu) - \varepsilon$.*

We identify two forces that make flexible communication (Bayesian persuasion) suboptimal for Sender. The first force arises only if popularity is not (weakly) monotonically increasing in skepticism. In this case, Sender must trade off popularity and influence.

Proposition 6 (The popularity-influence tradeoff). *If there exists $\sigma \in T^*(\pi, F, \mu)$ with $\mathcal{I}(\sigma \mid F, \mu) > 0$, and $\sigma' \in (0, \sigma)$ such that $\pi(\sigma') > \pi(\sigma)$, then Sender’s equilibrium payoff when the persuasion mechanism is Bayesian persuasion is strictly less than $V^*(\pi, F, \mu)$.*

Intuitively, when communication is fully flexible, Sender chooses to target a skeptic only if his influence is positive and maximizes market influence, $\pi(\sigma)\mathcal{I}(\sigma \mid F, \mu)$. However, Sender benefits both from more popularity and more influence and she may be forced to trade off one for the other because the skeptic with the most influence is not the skeptic with most popularity. But if Sender’s hands are tied so that she can only choose to induce posteriors $p^0 = 0$ and $p^1 = \sigma$, then every receiver with skepticism in $(0, \sigma]$ takes a different action at the two posteriors. His example therefore carries the same information and has influence $\mathcal{I}(\sigma \mid F, \mu)$, whether he is a skeptic or a fan. So Sender can optimally choose, among them, the target with maximum popularity. In particular, unlike under Bayesian persuasion, a sufficiently popular fan may be the optimal target under a hands-tied mechanism.²¹

We remark that this result is perhaps of greatest importance in markets in which popularity is concentrated at a few levels of skepticism that are rare in the population. For example, it implies that flexible communication is suboptimal in a large social network where most users are very skeptical and have very few followers, while the handful of “superstar” influencers with millions of followers are much less skeptical. In such a market, a persuasion mechanism inducing fixed posteriors $\{0, \sigma\}$ enables Sender to freely target superstars to leverage their popularity, even if all of them are less skeptical than σ , while fine-tuning influence to the distribution of skepticism in the market. Thus, restricting customization can increase market influence by allowing Sender to target a more popular receiver without sacrificing the compelling power of that receiver’s example.

²¹While Proposition 6 is defined over equilibrium objects, we can give sufficient conditions on the primitives of the model for its conclusion to be valid. Suppose there exists $\bar{\sigma} > \mu$ such that $\frac{\mu}{\bar{\sigma}}F(\bar{\sigma}) > F(\mu)$, and π is strictly decreasing on $(\mu, 1]$. Then some skeptic has positive influence, so Bayesian persuasion optimally targets a skeptic $\sigma^* > \mu$ with positive influence. Yet every skeptic with skepticism $\sigma' \in (\mu, \sigma^*)$ is more popular than the target. Therefore, Proposition 6 implies that Bayesian persuasion is strictly suboptimal.

This last result does *not* imply that more flexible communication is optimal if popularity is constant. We now turn to a second force by which flexible communication is suboptimal that arises even if popularity is constant. The following proposition says that whether flexible communication is optimal depends on the distribution of skepticism.

Proposition 7 (Flexible communication is suboptimal in polarized markets). *Suppose that $\pi(\sigma) = \pi(\sigma')$ for all $\sigma, \sigma' \in [0, 1]$.*

1. *If F is convex on $[0, \mu]$, then Sender's equilibrium payoff when the persuasion mechanism is Bayesian persuasion is equal to $V^*(\pi, F, \mu)$.*
2. *If there exists $\hat{\sigma} \in (0, \mu)$ such that $F(\hat{\sigma}) > \hat{\sigma}$ and F is strictly convex on $[\mu, 1]$, then Sender's equilibrium payoff when the persuasion mechanism is Bayesian persuasion is strictly less than $V^*(\pi, F, \mu)$.*

The key to understanding Proposition 7 is that fully flexible communication implies large losses from a negative example: the target's posterior when he does not buy is $p^0 = 0$, so that all receivers who observe his example do not buy the widget. Part 1 then simply says that this excessive cost of a negative example is not too damaging when there are sufficiently few hardcore fans whose skepticism is very low compared to the prior μ . In this case, raising the posterior p^0 does not increase Sender's expected sales significantly to overcome the cost of reducing the gains from a positive example (the reduction in p^1 implied by Bayes' plausibility).

Part 2, conversely, says that fully flexible communication is suboptimal if skepticism is sufficiently polarized, with a large enough mass of hardcore fans with skepticism below a threshold $\hat{\sigma}$ and many skeptics concentrated among the highest levels of skepticism—but not many moderate skeptics. Then the extra sales induced by increasing p^0 more than compensate the losses induced by decreasing p^1 , so that Sender would strictly benefit from having her hands tied.

4 Conclusions

When choosing whom to target in networking activities, salespersons, entrepreneurs, or lobbyists need to anticipate how persuading a target will affect other potential customers, investors, or politicians. Our [Full customization](#) theorem allows us to characterize optimal targeted persuasion independently of the specific assumptions about how the communication with a target takes place. Two implications of this result capture the intuitive ideas that the popularity of a target is a double-edged sword and that the optimal targeted

persuasion may benefit from less flexibility in customizing communication to persuade the target. Limited flexibility reduces the need to trade off influence with popularity. Less flexibility is optimal in markets that feature only a few high-popularity superstars and when customers' or investors' tastes are more polarized.

References

- Akbarpour, M. and S. Li (2020). Credible auctions: A trilemma. *Econometrica* 88(2), 425–467.
- Alonso, R. and O. Câmara (2016). Persuading Voters. *American Economic Review* 106(11), 3590–3605.
- Arieli, I. and Y. Babichenko (2019). Private bayesian persuasion. *Journal of Economic Theory* 182, 185–217.
- Awad, E. (2020). Persuasive lobbying with allied legislators. *American Journal of Political Science* 64(4), 938–951.
- Awad, E. and C. Minaudier (2026). Persuasive lobbying and the value of connections. *American Journal of Political Science* 70(2), 539–554.
- Babichenko, Y., I. Talgam-Cohen, H. Xu, and K. Zabarnyi (2022). Multi-channel bayesian persuasion. In *13th Innovations in Theoretical Computer Science Conference*, Volume 215 of *Leibniz International Proceedings in Informatics*, pp. 11:1–11:20. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.
- Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006). Who's who in networks. wanted: the key player. *Econometrica* 74(5), 1403–1417.
- Bardhi, A. and Y. Guo (2018). Modes of persuasion toward unanimous consent. *Theoretical Economics* 13(3), 1111–1149.
- Caillaud, B. and J. Tirole (2007). Consensus Building: How to Persuade a Group. *American Economic Review* 97(5), 1877–1900.
- Candogan, O., Y. Guo, and H. Xu (2020). On information design with spillovers. Ssrn working paper.
- Chan, J., S. Gupta, F. Li, and Y. Wang (2019). Pivotal persuasion. *Journal of Economic Theory* 180, 178–202.
- Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica* 50(6), 1431–1451.
- Doval, L. and V. Skreta (2024). Constrained information design. *Mathematics of Operations Research* 49(1), 78–106.

- Dye, R. A. (1985). Disclosure of nonproprietary information. *Journal of Accounting Research* 23(1), 123–145.
- Egorov, G. and K. Sonin (2019). Persuasion on networks. Technical report, SSRN.
- Fudenberg, D. and J. Tirole (1991). Perfect bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory* 53(2), 236–260.
- Galeotti, A. and S. Goyal (2009). Influencing the influencers: A theory of strategic diffusion. *The RAND Journal of Economics* 40(3), 509–532.
- Galperti, S. and J. Perego (2025). Games with information constraints: Seeds and spillovers. *Theoretical Economics* 20(2), 667–711.
- Gentzkow, M. and E. Kamenica (2016). A Rothschild–Stiglitz approach to Bayesian persuasion. *American Economic Review Papers and Proceedings* 106(5), 597–601.
- Goldstein, I. and C. Huang (2016, May). Bayesian persuasion in coordination games. *American Economic Review* 106(5), 592–596.
- Gratton, G., C. Teh, and D. J. Thornton (2025, September). Targeted persuasion. SSRN Scholarly Paper 6769998. Available at SSRN: <https://ssrn.com/abstract=6769998>.
- Halac, M. and I. Kremer (2020). Experimenting with career concerns. *American Economic Journal: Microeconomics* 12(1), 260–288.
- Inostroza, N. A. and A. Pavan (2025). Adversarial coordination and public information design. *Theoretical Economics* 20(2), 763–813.
- Kamenica, E. and M. Gentzkow (2011). Bayesian Persuasion. *American Economic Review* 101(6), 2590–2615.
- Kartik, N. (2009). Strategic communication with lying costs. *Review of Economic Studies* 76(4), 1359–1395.
- Khantadze, D., I. Kremer, and A. Skrzypacz (2025). Persuasion with multiple actions. *Journal of Political Economy* 133(5), 1497–1526.
- Kreutzkamp, S. and Y. Lou (2025). Persuasion without ex-post commitment. *Journal of Economic Theory* 228, 106058.
- Li, F., Y. Song, and M. Zhao (2023). Global manipulation by local obfuscation. *Journal of Economic Theory* 207, 105575.
- Lin, X. and C. Liu (2024). Credible persuasion. *Journal of Political Economy* 132(7), 2228–2273.
- Lipnowski, E. and D. Ravid (2020). Cheap talk with transparent motives. *Econometrica* 88(4), 1631–1660.
- Lipnowski, E., D. Ravid, and D. Shishkin (2022). Persuasion via weak institutions. *Journal of Political Economy* 130(10), 2705–2730.

- Mathevet, L. and I. Taneva (2022). Organized information transmission. Discussion Paper 16959, CEPR.
- Milgrom, P. and C. Shannon (1994). Monotone comparative statics. *Econometrica* 62(1), 157–180.
- Milgrom, P. R. (1981). Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics* 12(2), 380–391.
- Min, D. (2021). Bayesian persuasion under partial commitment. *Economic Theory* 72(3), 743–764.
- Morris, S., D. Oyama, and S. Takahashi (2024). Implementation via information design in binary-action supermodular games. *Econometrica* 92(3), 775–813.
- Schnakenberg, K. E. (2017). Informational lobbying and legislative voting. *American Journal of Political Science* 61(1), 129–145.
- Taneva, I. (2019). Information design. *American Economic Journal: Microeconomics* 11(4), 151–185.
- Wang, Y. (2015). Bayesian persuasion with multiple receivers. Working paper, Xiamen University.

Appendix

A Additional Details

This section provides additional details that were omitted from the main text. Section A.1 discusses how to map persuasion mechanisms to other prominent examples from the literature. Section A.2 formally defines perfect Bayesian equilibria in our setting. Section A.3 shows that every PBE can be uniquely identified (up to payoff-equivalence) by the collection of sets of distributions over posteriors that Sender can induce for the target in equilibrium and the system of beliefs for non-targets. Section A.4 shows that our analysis in Section 3 can be viewed as the limit of finite market analyses as the number of receivers grows large.

A.1 Examples of Persuasion Mechanisms

Information disclosure (Dye, 1985; Milgrom, 1981). Here, $\Phi = \{\varphi_0\}$, i.e., Sender is endowed with some (partially) informative information structure φ_0 . Furthermore, the message space is $M \subseteq 2^S$, so each message is a subset of S that the target interprets as “the

realized signal lies in this set", while the cost function satisfies $c(m, s) = 0$ if and only if $s \in m$, so Sender must tell the truth by reporting a subset that contains the realized signal but need not tell the full truth (reporting the singleton subset containing the realized signal). The case where φ_0 is fully informative, so $\varphi_0(s_G | G) = \varphi_0(s_B | B) = 1$ for some $s_G \neq s_B$, coincides with the model of [Milgrom \(1981\)](#). The case where φ_0 is only partially informative coincides with the model of [Dye \(1985\)](#).

Cheap talk ([Crawford and Sobel, 1982](#)). Here, $S = \{G, B\}$ and $\Phi = \{\bar{\varphi}\}$ where $\bar{\varphi}(\theta | \theta) = 1$, i.e., Sender fully learns the state. Furthermore, $c(m, s) = 0$ for all $s \in S$ and $m \in M$, so Sender can send any message at zero cost regardless of what she observes.

Constrained Bayesian persuasion ([Doval and Skreta, 2024](#)). Here, $S = M$, Φ is a strict subset of the set of all information structures and $c(m, s) = 0$ if $m = s$ and $c(m, s) = \infty$ otherwise. That is, Sender is restricted to choosing among a small set of information structures and must truthfully communicate the realized signal.

A.2 Perfect Bayesian Equilibrium

A *strategy* for Sender is a tuple $(t, \{(\varphi_t, \eta_t)\}_{t \in \mathcal{R}})$, where $t \in \mathcal{R} \cup \{\emptyset\}$ is the chosen target and, for each targeted receiver $t \in \mathcal{R}$,

- $\varphi_t \in \Phi$ is Sender's choice of information structure; and
- $\eta_t : S \rightarrow \Delta M$ is the distribution over messages Sender sends to the target t .

A *system of beliefs* for the receivers is a collection $p \equiv \{(p_t, \mathbf{p}_{-t})\}_{t \in \mathcal{R} \cup \{\emptyset\}}$, where $\mathbf{p}_{-t} \equiv (p_{-t,r})_{r \neq t}$. Conditional on targeting receiver $t \in \mathcal{R}$, his posterior $p_t(\varphi_t, m)$ is a function of Sender's information structure φ_t and the message m he receives from Sender. Meanwhile, for each non-targeted receiver $r \neq t$, Sender's posterior $p_{-t,r}(o_t)$ is a function of his observation $o_t \in \{1, 0, \emptyset\}$ of the target t 's action, where $o_t = 1$ and $o_t = 0$ represent the non-target observing t buy and not buy respectively, and $o_t = \emptyset$ represents the non-target failing to observe t 's action. In the special case with $t = \emptyset$, there is a singleton information set in which all receivers hold a belief of $p_{-\emptyset,r}(\emptyset)$.

Given $\mathbf{p}_{-t}$, let $\tilde{F}_t(\cdot | \mathbf{p}_{-t}) : \{0, 1, \emptyset\} \rightarrow [0, R - 1]$ denote the number of non-targeted receivers who buy the widget given observation o_t , i.e.,

$$\tilde{F}_t(o_t | \mathbf{p}_{-t}) \equiv |\{r \in \mathcal{R} \setminus \{t\} : \sigma_r \leq p_{-t,r}(o_t)\}|.$$

Given a target t , information structure φ_t and system of beliefs p , let $V_t(m \mid \varphi_t, p_t, \mathbf{p}_{-t})$ denote Sender's ex-post payoff from sending message m , i.e.,

$$V_t(m \mid \varphi_t, p_t, \mathbf{p}_{-t}) = \mathbb{I}[\sigma_t \leq p_t(\varphi_t, m)] + \pi_t \tilde{F}_t(\mathbb{I}[\sigma_t \leq p_t(\varphi_t, m)] \mid \mathbf{p}_{-t}) + (1 - \pi_t) \tilde{F}_t(\emptyset \mid \mathbf{p}_{-t}).$$

An *assessment* is a collection of strategies for Sender and a system of beliefs for receivers, $(t, \{(\varphi_t, \eta_t)\}_{t \in \mathcal{R}}, \{p_t, \mathbf{p}_{-t}\}_{t \in \mathcal{R}})$. An assessment is a *perfect Bayesian equilibrium (PBE)* (Fudenberg and Tirole, 1991) if for each target, receivers' beliefs are formed by Bayes' rule whenever possible and, fixing receivers' beliefs, Sender's communication and her choice of target are sequentially optimal.

To simplify exposition, we first explain the belief restrictions implied by the structure of our reduced-form model and, where applicable, by the "no signaling what you don't know" condition of PBE. First, when Sender targets no receiver at all ($t = \emptyset$), all receivers hold belief $p_{-\emptyset, r}(\emptyset) = \mu$.²² Second, when a receiver only observes the identity of the target but *not* his choice of whether to buy the widget, he holds belief $p_{-t, r}(\emptyset) = \mu$.²³ Finally, when a target $t \in \mathcal{R}$ (whether the equilibrium target or not) receives a message $m \in M$ that is communicated by Sender with probability 0 (i.e., $\eta_t(m \mid s) = 0$ for all s), then t knows that Sender has observed a signal s that allows communicating m .²⁴ Importantly, when target t receives a message m , he can rule out any signal after which m would be infeasible. Suppose the set of signals observed with positive probability under Sender's chosen information structure $\varphi \in \Phi$ and after which m is feasible

$$\left\{ s \in S : \sum_{\theta \in \{G, B\}} \Pr(\theta) \varphi(s \mid \theta) > 0 \text{ and } c(m, s) = 0 \right\}$$

is a singleton, say $\{s^*\}$. Then observing m reveals to the target that Sender observed s^* . Hence the target's posterior must equal Sender's posterior after observing s^* :

$$p_t(\varphi, m) = \frac{\mu \varphi(s^* \mid G)}{\sum_{\theta \in \{G, B\}} \Pr(\theta) \varphi(s^* \mid \theta)}.$$

²²To see this, notice that our model can be represented in extensive form as Sender choosing t at the initial node and $p_{-\emptyset, r}(\emptyset)$ being beliefs at a singleton information set (with respect to θ) that immediately follows the initial node.

²³To see this, notice that whether a receiver observes the choice of the target is entirely determined by the choice of Nature, whose strategy π is independent of the quality of the good.

²⁴To see this, notice that, in our model, Sender cannot misrepresent the persuasion mechanism or her choice of information structure, φ_t . Thus, in any extensive form representation of our model, when the target observes the message m , his information set only contains histories in which the realizations of θ and s have $\varphi_t(s \mid \theta) > 0$ and $c(m, s) = 0$.

A final detail concerns a non-targeted receiver's belief after observing a target action that has zero probability in equilibrium. In our reduced-form model, the target's action is mechanically determined by his posterior. Thus, a receiver who observes that t buys *knows* that t 's posterior lies in $[\sigma_t, 1]$; conversely, a receiver who observes that t does not buy *knows* that t 's posterior lies in $[0, \sigma_t)$. Conditions 1(a) and 1(b) impose this action-consistency both on and off the equilibrium path. In a fully strategic formulation of our model, the proof of Theorem 1 only requires that observing that the target buys does not induce a lower posterior belief than observing that he does not buy, including when one of these actions is off the equilibrium path: $p_{-t,r}(1) \geq p_{-t,r}(0)$ for every $r \neq t$. For a closely related forward-induction argument supporting this type of belief monotonicity, see Halac and Kremer (2020).

We can now complete the description of our solution concept. An assessment $(t, \{(\varphi_t, \eta_t)\}_{t \in \mathcal{R}}, \{p_t, p_{-t}\}_{t \in \mathcal{R}})$ is a PBE if and only if:

1. **Non-targeted receivers update by Bayes' rule.** For each target $t \in \mathcal{R}$ and non-target $r \neq t$,

(a) **Belief upon observing target buy:**

$$p_{-t,r}(1) = \frac{\mu \sum_{s \in S} \varphi_t(s | G) \sum_{m: p_t(\varphi_t, m) \geq \sigma_t} \eta_t(m | s)}{\sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi_t(s | \theta) \sum_{m: p_t(\varphi_t, m) \geq \sigma_t} \eta_t(m | s)}, \quad (15)$$

if $\sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi_t(s | \theta) \sum_{m: p_t(\varphi_t, m) \geq \sigma_t} \eta_t(m | s) > 0$, and $p_{-t,r}(1) \in [\sigma_t, 1]$ otherwise;

(b) **Belief upon observing target not buy:**

$$p_{-t,r}(0) = \frac{\mu \sum_{s \in S} \varphi_t(s | G) \sum_{m: p_t(\varphi_t, m) < \sigma_t} \eta_t(m | s)}{\sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi_t(s | \theta) \sum_{m: p_t(\varphi_t, m) < \sigma_t} \eta_t(m | s)}, \quad (16)$$

if $\sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi_t(s | \theta) \sum_{m: p_t(\varphi_t, m) < \sigma_t} \eta_t(m | s) > 0$, and $p_{-t,r}(0) \in [0, \sigma_t)$ otherwise;

(c) **Belief upon observing nothing:** $p_{-t,r}(\emptyset) = \mu$.

2. **The target updates by Bayes' rule.** For each target t and each message $m \in M$

communicated by Sender with strictly positive probability:

$$p_t(\varphi_t, m) = \frac{\mu \sum_{s \in S} \varphi_t(s | G) \eta_t(m | s)}{\sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi_t(s | \theta) \eta_t(m | s)}. \quad (17)$$

In addition, for any pair $(\varphi, m) \in \Phi \times M$, if there exists a unique $s^* \in S$ such that $\sum_{\theta \in \{G, B\}} \Pr(\theta) \varphi(s^* | \theta) > 0$ and $c(m, s^*) = 0$, then

$$p_t(\varphi, m) = \frac{\mu \varphi(s^* | G)}{\sum_{\theta \in \{G, B\}} \Pr(\theta) \varphi(s^* | \theta)}.$$

3. Sender's communication is sequentially optimal. For each target t ,

$$(\varphi_t, \eta_t) \in \arg \max_{\substack{\varphi \in \Phi, \\ \eta = \{\eta(\cdot | s)\}_{s \in S}}} \sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi(s | \theta) \sum_{m \in M} \eta(m | s) [V_t(m | \varphi, p_t, \mathbf{p}_{-t}) - c(m, s)]. \quad (18)$$

4. Sender's target choice is sequentially optimal.

$$t \in \arg \max_{t' \in \mathcal{R} \cup \{\emptyset\}} \mathcal{V}_{t'}, \quad (19)$$

where

$$\mathcal{V}_{t'} \equiv \begin{cases} \sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi_{t'}(s | \theta) \sum_{m \in M} \eta_{t'}(m | s) [V_{t'}(m | \varphi_{t'}, p_{t'}, \mathbf{p}_{-t'}) - c(m, s)], & t' \in \mathcal{R} \\ |\mathcal{F}|, & t' = \emptyset \end{cases}.$$

A.3 Representing equilibria by sets of distributions over posteriors

We now show that a PBE can be summarized by two objects. The first is, for each possible target t , the set P_t of distributions over posteriors that Sender can induce for that target. The second is the system of beliefs $\mathbf{p}_{-t}$ held by non-targeted receivers after observing the target's action. The first object pins down what Sender can do to the target's posterior. The second object pins down how other receivers respond to each possible observation of the target's action. Once both objects are fixed, the payoff-relevant content of the PBE is pinned down.

Let $\{P_t\}_{t \in \mathcal{R}}$ denote a collection of subsets of distribution over posteriors. Further recall that $\{\mathbf{p}_{-t}\}_{t \in \mathcal{R}}$ denotes a collection of systems of beliefs for non-targeted receivers. Say that $\{(P_t, \mathbf{p}_{-t})\}_{t \in \mathcal{R}}$ is *implementable* if there exists a PBE $(t, \{(\varphi_t, \eta_t)\}_{t \in \mathcal{R}}, \{p_t, \mathbf{p}_{-t}\}_{t \in \mathcal{R}})$ such that, for every target $t \in \mathcal{R}$, P_t is the set of distributions over posteriors that Sender can induce

for target t under some feasible strategy. That is $\rho \in P_t$ if and only if there exists a strategy for Sender (φ'_t, η'_t) that is feasible (i.e., $\eta'_t(m | s) = 0$ whenever $c(m, s) = \infty$) and such that

$$\rho(K) = \sum_{\theta \in \{G, B\}} \Pr(\theta) \sum_{s \in S} \varphi'_t(s | \theta) \sum_{m \in M} \eta'_t(m | s) \mathbb{1}\{p_t(\varphi'_t, m) \in K\} \quad \forall \text{ measurable } K \subseteq [0, 1]. \quad (20)$$

That is, holding fixed the system of beliefs for targets $\{p_t\}_{t \in \mathcal{R}}$ in the PBE, P_t contains all distributions over posteriors that Sender can induce for target t under some feasible strategy. If the above holds under a given PBE, then say that the PBE *implements* $\{(P_t, \mathbf{p}_{-t})\}_{t \in \mathcal{R}}$.

The next result says that the payoff-relevant information of every PBE can be fully summarized by the pair $\{(P_t, \mathbf{p}_{-t})\}_{t \in \mathcal{R}}$ that it implements.

Proposition A.1 (PBE characterization). *Take any PBE which implements $\{(P_t, \mathbf{p}_{-t})\}_{t \in \mathcal{R}}$ and, for each t , let ρ_t^* denote the distribution over posteriors induced by Sender's equilibrium strategy (φ_t, η_t) . In this PBE, if Sender targets receiver $t \in \mathcal{R}$, then*

1. *The induced distribution satisfies*

$$\rho_t^* \in \arg \max_{\rho_t \in P_t} \int_0^1 \mathbb{1}\{\sigma_t \leq p\} d\rho_t(p);$$

2. *The expected sales to the target equal*

$$\max_{\rho_t \in P_t} \int_0^1 \mathbb{1}\{\sigma_t \leq p\} d\rho_t(p); \quad (21)$$

3. *The expected sales to other receivers equal*

$$\pi_t \left(\int_0^1 \tilde{F}_t(\mathbb{1}\{\sigma_t \leq p\} | \mathbf{p}_{-t}) d\rho_t^*(p) \right) + (1 - \pi_t) \tilde{F}_t(\emptyset | \mathbf{p}_{-t}). \quad (22)$$

Therefore, every PBE which implements the same pair induces, for each targeted receiver $t \in \mathcal{R}$, the same expected sales to the target and to non-targeted receivers.

Proof. Fix a PBE $(t, \{(\varphi_t, \eta_t)\}_{t \in \mathcal{R}}, \{p_t, \mathbf{p}_{-t}\}_{t \in \mathcal{R}})$ which implements $\{(P_t, \mathbf{p}_{-t})\}_{t \in \mathcal{R}}$. For each target $t \in \mathcal{R}$, recall that ρ_t^* denotes the distribution over posteriors induced by Sender's equilibrium strategy (φ_t, η_t) , i.e., the distribution which satisfies (20).

We first show that Conditions 1 and 2 hold. In the PBE, Sender only optimizes over distributions over posteriors ρ_t which can be induced by a feasible strategy. Since feasible

strategies send messages which have zero cost, Sender's expected payoff from choosing any such distribution over posteriors can be written as

$$\begin{aligned}
& \int_0^1 \left[\mathbb{1}\{\sigma_t \leq p\} + \pi_t \tilde{F}_t(\mathbb{1}\{\sigma_t \leq p\} \mid \mathbf{p}_{-t}) + (1 - \pi_t) \tilde{F}_t(\emptyset \mid \mathbf{p}_{-t}) \right] d\rho_t(p) \\
&= \pi_t \tilde{F}_t(0 \mid \mathbf{p}_{-t}) + (1 - \pi_t) \tilde{F}_t(\emptyset \mid \mathbf{p}_{-t}) \\
&+ \left[1 + \pi_t \left(\tilde{F}_t(1 \mid \mathbf{p}_{-t}) - \tilde{F}_t(0 \mid \mathbf{p}_{-t}) \right) \right] \int_0^1 \mathbb{1}\{\sigma_t \leq p\} d\rho_t(p). \tag{23}
\end{aligned}$$

Condition 1 implies, whether each target action is on or off the equilibrium path, that

$$p_{-t,r}(1) \geq \sigma_t > p_{-t,r}(0) \quad \text{for every } r \neq t.$$

Hence $\tilde{F}_t(1 \mid \mathbf{p}_{-t}) \geq \tilde{F}_t(0 \mid \mathbf{p}_{-t})$, and therefore

$$1 + \pi_t \left(\tilde{F}_t(1 \mid \mathbf{p}_{-t}) - \tilde{F}_t(0 \mid \mathbf{p}_{-t}) \right) > 0.$$

Since, by Condition 3 in the definition of PBE, Sender's equilibrium strategy must maximize his expected payoff from targeting t , it follows that the induced distribution over posteriors must satisfy

$$\rho_t^* \in \arg \max_{\rho_t \in P_t} \int_0^1 \mathbb{1}\{\sigma_t \leq p\} d\rho_t(p),$$

so Condition 1 holds. In turn, the probability the target buys is exactly given by (21), so Condition 2 holds.

From here, note that conditional on inducing posterior p for the target, the target's action is given by $\mathbb{1}\{\sigma_t \leq p\}$. Thus, the expected sales to non-targets under posterior p is

$$\pi_t \tilde{F}_t(\mathbb{1}\{\sigma_t \leq p\} \mid \mathbf{p}_{-t}) + (1 - \pi_t) \tilde{F}_t(\emptyset \mid \mathbf{p}_{-t}).$$

Integrating this over all posteriors, which yields the unconditional expected sales to non-targets, then yields Condition 3. $\square$

Proposition A.1 implies that, as discussed in the main text, every PBE can be associated to a collection of sets of distributions over posteriors. Furthermore, Condition 1 implies that in any equilibrium, for each target t , Sender maximizes the probability the target buys among all distributions over posteriors in P_t . This implies Theorem 1.

A.4 Approximating large markets with finite markets

We extend the finite targeted-persuasion game in Section 2 to the large-market model in Section 3 in a natural way: the timing, persuasion mechanisms, and PBE requirements are unchanged, but the finite set of receivers is replaced by a unit mass of receivers indexed by skepticism, with popularity given by the function π . Under Bayesian persuasion, the same arguments as for the finite market model imply that targeting a receiver with skepticism σ yields the large-market value $V(\sigma \mid \pi, F, \mu)$ defined in (12). Hence the equilibrium targeting problem in the large market is exactly the one characterized in Proposition 3.

This subsection shows that this large-market targeting problem is the limit of the corresponding finite-market targeting problems. We begin with several definitions. We represent a finite market of size R by a tuple (π^R, σ^R, μ) , where σ_r^R and π_r^R are the levels of skepticism and popularity of receiver $r \in \mathcal{R}$. The set of fans is given by $\mathcal{F}(\sigma^R, \mu) \equiv \{r \in \mathcal{R} : \sigma_r^R \leq \mu\}$. Note then that conditional on targeting receiver t , the normalized distribution over non-targets' skepticism is given by

$$\hat{F}_t^R(\sigma \mid \sigma^R) \equiv \frac{|\{r \in \mathcal{R} \setminus \{t\} : \sigma_r^R \leq \sigma\}|}{R-1}.$$

Let $F^R \equiv \{\hat{F}_t^R(\cdot \mid \sigma^R)\}_{t \in \mathcal{R}}$ denote the collection of target-specific empirical distributions. Then under Bayesian persuasion, (10) implies that the normalized value of targeting receiver t in the finite market is

$$V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \equiv \begin{cases} \frac{1}{R} \left[(R-1) \hat{F}_t^R(\mu \mid \sigma^R) + \frac{\mu}{\sigma_t^R} + \pi_t^R (R-1) \left(\frac{\mu}{\sigma_t^R} \hat{F}_t^R(\sigma_t^R \mid \sigma^R) - \hat{F}_t^R(\mu \mid \sigma^R) \right) \right], & \sigma_t^R > \mu \\ \frac{|\mathcal{F}(\sigma^R, \mu)|}{R}, & \sigma_t^R \leq \mu \end{cases}$$

Thus, the set of optimal targets is

$$T_R^*(\pi^R, \sigma^R, \mu) \equiv \arg \max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu).$$

Finally, define the distance between a point $x \in [0, 1]$ and a set $A \subseteq [0, 1]$ by

$$d(x, A) \equiv \inf_{a \in A} |x - a|.$$

The next result offers three main insights. First, given any large market (π, F, μ) , there exists a sequence of finite markets for which the conditional empirical distributions of non-targets' skepticism, no matter which receiver is targeted, converge uniformly to the

distribution over skepticism in the large market F . Second, along this sequence of finite markets, Sender's optimal normalized payoff converges to the large-market payoff characterized in Proposition 3. Finally, any limit (point) of a sequence of optimal targets in these finite markets is an optimal target in the large market. Thus, the large-market analysis in Section 3 is the limit of finite-market analyses as the number of receivers grows large.

Proposition A.2 (Finite-market approximation). *Fix a large market (π, F, μ) . There exists a sequence of finite markets $(\pi^R, \sigma^R, \mu)_{R=2}^\infty$ such that:*

1. *the distributions over non-targets' skepticism from targeting any receiver converge uniformly to F :*

$$\sup_{t \in \mathcal{R}} \sup_{\sigma \in [0,1]} \left| \widehat{F}_t^R(\sigma \mid \sigma^R) - F(\sigma) \right| \rightarrow 0;$$

2. *Sender's normalized equilibrium payoff in each finite market converges to her equilibrium payoff in the large market:*

$$\max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \rightarrow \max_{\sigma \in [0,1]} V(\sigma \mid \pi, F, \mu);$$

3. *if $(t_R)_{R=2}^\infty$ is any sequence of finite-market optimal targets, so $t_R \in T_R^*(\pi^R, \sigma^R, \mu)$ for all $R \geq 2$, then*

$$d(\sigma_{t_R}^R, T^*(\pi, F, \mu)) \rightarrow 0.$$

In particular, if the large market has a unique optimal target, so $T^(\pi, F, \mu) = \{\sigma^*\}$, then every sequence of finite-market optimal targets $(t_R)_{R=2}^\infty$ satisfies $\sigma_{t_R}^R \rightarrow \sigma^*$.*

Proof. Let

$$F^{-1}(q) \equiv \inf\{\sigma \in [0, 1] : F(\sigma) \geq q\}$$

denote the generalized inverse of F . For each $R \geq 2$, let (π^R, σ^R, μ) be the finite market of R receivers defined as follows:²⁵

$$\sigma_r^R \equiv F^{-1}\left(\frac{r}{R}\right) \quad \text{and} \quad \pi_r^R \equiv \pi(\sigma_r^R), \quad \forall r = 1, \dots, R.$$

We show that the sequence of finite markets $(\pi^R, \sigma^R, \mu)_{R=2}^\infty$ satisfies all three conditions.

²⁵In Section 2, we assume for expositional simplicity that every finite market contains at least one fan and one skeptic. The quantile construction below satisfies this assumption for every R such that $\lfloor RF(\mu) \rfloor \geq 1$: receiver r is a fan iff $r/R \leq F(\mu)$, while $F(\mu) < 1$ guarantees a skeptic. Since $0 < F(\mu) < 1$, the assumption therefore holds along the tail of the constructed sequence, which is all that matters for our limit statements.

Condition 1. Notice that for every $R \geq 2$, $\sigma \in [0, 1]$, and $t \in \mathcal{R}$,

$$\begin{aligned}
\left| F(\sigma) - \widehat{F}_t^R(\sigma \mid \sigma^R) \right| &= \left| F(\sigma) - \frac{|\{r \in \mathcal{R} \setminus \{t\} : \sigma_r^R \leq \sigma\}|}{R-1} \right| \\
&\leq \left| F(\sigma) - \frac{|\{r \in \mathcal{R} : \sigma_r^R \leq \sigma\}|}{R} \right| \\
&\quad + \left| \frac{|\{r \in \mathcal{R} : \sigma_r^R \leq \sigma\}|}{R} - \frac{|\{r \in \mathcal{R} \setminus \{t\} : \sigma_r^R \leq \sigma\}|}{R-1} \right| \\
&\leq \frac{1}{R} + \left| \frac{(R-1)|\{r \in \mathcal{R} : \sigma_r^R \leq \sigma\}| - R|\{r \in \mathcal{R} \setminus \{t\} : \sigma_r^R \leq \sigma\}|}{R(R-1)} \right| \\
&\leq \frac{1}{R} + \frac{R}{R(R-1)} \\
&= \frac{1}{R} + \frac{1}{R-1}.
\end{aligned}$$

Since $\frac{1}{R} + \frac{1}{R-1} \rightarrow 0$, Condition 1 holds.

Condition 2. We first show that, evaluating the large market at the corresponding finite target's skepticism, Sender's normalized equilibrium payoff in each finite market converges to her equilibrium payoff in the large market:

$$\max_{t \in \mathcal{R}} \left| V^R(\sigma_t^R \mid \pi^R, F^R, \mu) - V(\sigma_t^R \mid \pi, F, \mu) \right| \rightarrow 0. \quad (24)$$

To see this, take any $R \geq 2$ and $t \in \mathcal{R}$. If $\sigma_t^R > \mu$, then

$$\begin{aligned}
V^R(\sigma_t^R \mid \pi^R, F^R, \mu) &= \frac{1}{R} \frac{\mu}{\sigma_t^R} + \frac{R-1}{R} \left[\widehat{F}_t^R(\mu \mid \sigma^R) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} \widehat{F}_t^R(\sigma_t^R \mid \sigma^R) - \widehat{F}_t^R(\mu \mid \sigma^R) \right) \right], \\
V(\sigma_t^R \mid \pi, F, \mu) &= F(\mu) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} F(\sigma_t^R) - F(\mu) \right).
\end{aligned}$$

Therefore,

$$\begin{aligned}
& |V^R(\sigma_t^R \mid \pi^R, F^R, \mu) - V(\sigma_t^R \mid \pi, F, \mu)| \\
&= \left| \frac{1}{R} \frac{\mu}{\sigma_t^R} + \frac{R-1}{R} \left[\widehat{F}_t^R(\mu \mid \sigma^R) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} \widehat{F}_t^R(\sigma_t^R \mid \sigma^R) - \widehat{F}_t^R(\mu \mid \sigma^R) \right) \right] \right. \\
&\quad \left. - \left[F(\mu) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} F(\sigma_t^R) - F(\mu) \right) \right] \right| \\
&\leq \frac{1}{R} \frac{\mu}{\sigma_t^R} + \left| \frac{R-1}{R} \left[\widehat{F}_t^R(\mu \mid \sigma^R) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} \widehat{F}_t^R(\sigma_t^R \mid \sigma^R) - \widehat{F}_t^R(\mu \mid \sigma^R) \right) \right] \right. \\
&\quad \left. - \left[F(\mu) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} F(\sigma_t^R) - F(\mu) \right) \right] \right| \\
&\leq \frac{1}{R} + \left| \left[\widehat{F}_t^R(\mu \mid \sigma^R) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} \widehat{F}_t^R(\sigma_t^R \mid \sigma^R) - \widehat{F}_t^R(\mu \mid \sigma^R) \right) \right] \right. \\
&\quad \left. - \left[F(\mu) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} F(\sigma_t^R) - F(\mu) \right) \right] \right| \\
&\quad + \frac{1}{R} \left| \widehat{F}_t^R(\mu \mid \sigma^R) + \pi(\sigma_t^R) \left(\frac{\mu}{\sigma_t^R} \widehat{F}_t^R(\sigma_t^R \mid \sigma^R) - \widehat{F}_t^R(\mu \mid \sigma^R) \right) \right| \\
&\leq \frac{2}{R} + \left| \widehat{F}_t^R(\mu \mid \sigma^R) - F(\mu) \right| \\
&\quad + \pi(\sigma_t^R) \left| \frac{\mu}{\sigma_t^R} \left[\widehat{F}_t^R(\sigma_t^R \mid \sigma^R) - F(\sigma_t^R) \right] - \left[\widehat{F}_t^R(\mu \mid \sigma^R) - F(\mu) \right] \right| \\
&\leq \frac{2}{R} + \left| \widehat{F}_t^R(\mu \mid \sigma^R) - F(\mu) \right| + \frac{\mu}{\sigma_t^R} \left| \widehat{F}_t^R(\sigma_t^R \mid \sigma^R) - F(\sigma_t^R) \right| + \left| \widehat{F}_t^R(\mu \mid \sigma^R) - F(\mu) \right| \\
&\leq \frac{2}{R} + 3\Delta_R,
\end{aligned}$$

where

$$\Delta_R \equiv \sup_{t \in \mathcal{R}} \sup_{\sigma \in [0,1]} \left| \widehat{F}_t^R(\sigma \mid \sigma^R) - F(\sigma) \right|.$$

If instead $\sigma_t^R \leq \mu$, then

$$V^R(\sigma_t^R \mid \pi^R, F^R, \mu) = \frac{|\mathcal{F}(\sigma_t^R, \mu)|}{R} \quad \text{and} \quad V(\sigma_t^R \mid \pi, F, \mu) = F(\mu).$$

Since $\sigma_r^R = F^{-1}(r/R)$, receiver r is a fan if and only if $r/R \leq F(\mu)$, so $|\mathcal{F}(\sigma^R, \mu)| = \lfloor RF(\mu) \rfloor$. Therefore,

$$|V^R(\sigma_t^R \mid \pi^R, F^R, \mu) - V(\sigma_t^R \mid \pi, F, \mu)| = \left| \frac{|\mathcal{F}(\sigma_t^R, \mu)|}{R} - F(\mu) \right| \leq \frac{1}{R}.$$

Combining the two cases gives

$$\max_{t \in \mathcal{R}} \left| V^R(\sigma_t^R \mid \pi^R, F^R, \mu) - V(\sigma_t^R \mid \pi, F, \mu) \right| \leq \varepsilon_R \equiv \max \left\{ \frac{1}{R}, \frac{2}{R} + 3\Delta_R \right\}.$$

Since $\Delta_R \rightarrow 0$ by Condition 1, (24) holds.

We now prove convergence of optimal payoffs. (24) implies

$$\max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \leq \max_{t \in \mathcal{R}} V(\sigma_t^R \mid \pi, F, \mu) + \varepsilon_R \leq \max_{\sigma \in [0,1]} V(\sigma \mid \pi, F, \mu) + \varepsilon_R.$$

Therefore,

$$\limsup_{R \rightarrow \infty} \max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \leq \max_{\sigma \in [0,1]} V(\sigma \mid \pi, F, \mu).$$

To show the reverse inequality, take any $\sigma^* \in T^*(\pi, F, \mu)$, so $V(\sigma^* \mid \pi, F, \mu) = \max_{\sigma \in [0,1]} V(\sigma \mid \pi, F, \mu)$. By the definition of σ^R , for all $\delta > 0$ and large enough R , there exists $r_R \in \mathcal{R}$ with

$$|\sigma_{r_R}^R - \sigma^*| < \delta.$$

Since $V(\sigma \mid \pi, F, \mu)$ is continuous in σ ,²⁶ for every $\eta > 0$ and large enough R ,

$$V(\sigma_{r_R}^R \mid \pi, F, \mu) \geq V(\sigma^* \mid \pi, F, \mu) - \eta.$$

But then (24) implies

$$V^R(\sigma_{r_R}^R \mid \pi^R, F^R, \mu) \geq V(\sigma_{r_R}^R \mid \pi, F, \mu) - \varepsilon_R \geq V(\sigma^* \mid \pi, F, \mu) - \eta - \varepsilon_R.$$

Thus,

$$\max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \geq V(\sigma^* \mid \pi, F, \mu) - \eta - \varepsilon_R.$$

Therefore,

$$\liminf_{R \rightarrow \infty} \max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \geq V(\sigma^* \mid \pi, F, \mu) - \eta.$$

Since $\eta > 0$ was arbitrary,

$$\liminf_{R \rightarrow \infty} \max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \geq \max_{\sigma \in [0,1]} V(\sigma \mid \pi, F, \mu).$$

This proves Condition 2.

²⁶Because F and π are continuous, and because $\mathcal{I}(\sigma \mid F, \mu)$ is defined to be zero for all $\sigma \leq \mu$, the large-market value function $V(\sigma \mid \pi, F, \mu)$ is continuous on $[0, 1]$.

Condition 3. Let $(t_R)_{R=2}^\infty$ be any sequence of finite-market optimal targets, so

$$V^R(\sigma_{t_R}^R \mid \pi^R, F^R, \mu) = \max_{t \in \mathcal{R}} V^R(\sigma_t^R \mid \pi^R, F^R, \mu) \quad \forall R \geq 2.$$

Suppose, by contradiction, that $d(\sigma_{t_R}^R, T^*(\pi, F, \mu)) \not\rightarrow 0$. Then there exists $\delta > 0$ and a subsequence $(t_{R_k})_{k=1}^\infty$ such that

$$d(\sigma_{t_{R_k}}^{R_k}, T^*(\pi, F, \mu)) \geq \delta \quad \forall k \geq 1.$$

Define

$$K_\delta \equiv \{\sigma \in [0, 1] : d(\sigma, T^*(\pi, F, \mu)) \geq \delta\}.$$

This set contains the sequence $(\sigma_{t_{R_k}}^{R_k})_{k=1}^\infty$ and contains no maximizer of $V(\sigma \mid \pi, F, \mu)$. Since K_δ is compact and $V(\sigma \mid \pi, F, \mu)$ is continuous in σ , there exists $\eta > 0$ such that

$$V(\sigma_{t_{R_k}}^{R_k} \mid \pi, F, \mu) \leq \max_{\sigma \in [0, 1]} V(\sigma \mid \pi, F, \mu) - \eta \quad \forall k \geq 1.$$

By (24), for every k ,

$$V^{R_k}(\sigma_{t_{R_k}}^{R_k} \mid \pi^{R_k}, F^{R_k}, \mu) \leq V(\sigma_{t_{R_k}}^{R_k} \mid \pi, F, \mu) + \varepsilon_{R_k}.$$

Therefore,

$$V^{R_k}(\sigma_{t_{R_k}}^{R_k} \mid \pi^{R_k}, F^{R_k}, \mu) \leq \max_{\sigma \in [0, 1]} V(\sigma \mid \pi, F, \mu) - \eta + \varepsilon_{R_k}.$$

Since $\varepsilon_{R_k} \rightarrow 0$, for all sufficiently large k , $\varepsilon_{R_k} \leq \frac{\eta}{2}$, so

$$V^{R_k}(\sigma_{t_{R_k}}^{R_k} \mid \pi^{R_k}, F^{R_k}, \mu) \leq \max_{\sigma \in [0, 1]} V(\sigma \mid \pi, F, \mu) - \frac{\eta}{2}.$$

But since t_{R_k} is an optimal finite-market target, this means that, for all sufficiently large k ,

$$\max_{t \in \mathcal{R}} V^{R_k}(\sigma_t^{R_k} \mid \pi^{R_k}, F^{R_k}, \mu) \leq \max_{\sigma \in [0, 1]} V(\sigma \mid \pi, F, \mu) - \frac{\eta}{2},$$

which contradicts Condition 2. □

B Proofs

Proof of Theorem 1. Follows from Condition 1 of Proposition A.1 in Appendix A. □

Proof of Proposition 1. Fix a PBE and a target $t \in \mathcal{R}$, and let ρ_t^* be the equilibrium distribution over the target's posterior chosen by Sender. If a non-targeted receiver observes the target buys (on the equilibrium path), observes the target does not buy (on the equilibrium path), and does not observe the target's action, then Properties 1a), 1b) and 1c) of a PBE respectively imply that the non-target's posterior is $p_t^1(\rho_t^*)$, $p_t^0(\rho_t^*)$ and μ respectively. Thus, the masses of non-targets who buy are, respectively, $(R-1)\hat{F}_t(p_t^1(\rho_t^*))$, $(R-1)\hat{F}_t(p_t^0(\rho_t^*))$ and $(R-1)\hat{F}_t(\mu)$. Furthermore, we recall from Theorem 1 that the target buys with probability $B_t(\rho_t^*)$. Thus, expected sales from targeting t in the PBE is

$$B_t + \pi_t(R-1) \left[B_t \hat{F}_t(p_t^1(\rho_t^*)) + (1-B_t) \hat{F}_t(p_t^0(\rho_t^*)) \right] + (1-\pi_t)(R-1)\hat{F}_t(\mu).$$

Dividing by R and adding and subtracting the no-observation benchmark gives

$$\frac{R-1}{R} \hat{F}_t(\mu) + \frac{1}{R} B_t + \frac{\pi_t}{R} \left[B_t(R-1)(\hat{F}_t(p_t^1(\rho_t^*)) - \hat{F}_t(\mu)) - (1-B_t)(R-1)(\hat{F}_t(\mu) - \hat{F}_t(p_t^0(\rho_t^*))) \right].$$

By the definitions of G_t , L_t , and $\mathcal{I}_t$, this is exactly

$$V_t = \frac{R-1}{R} \hat{F}_t(\mu) + \frac{1}{R} B_t(\rho_t^*) + \pi_t \mathcal{I}_t(\rho_t^*).$$

Finally, note that Sender can always choose to not target a receiver, $t = \emptyset$, and achieve a (normalized) payoff of $|\mathcal{F}|/R$. Hence, by Property 4) of the definition of a PBE, Sender targets a receiver if and only if there exists $t \in \mathcal{R}$ such that $V_t > |\mathcal{F}|/R$, and targets no receiver otherwise. $\square$

Proof of Proposition 2. Fix a PBE and a target $t \in \mathcal{R}$, and let ρ_t^* be the equilibrium distribution over the target's posterior chosen by Sender. By Proposition 1, Sender's normalized payoff from targeting t is affine in π_t , with slope $\mathcal{I}_t(\rho_t^*)$. Hence the payoff is increasing in π_t if $\mathcal{I}_t(\rho_t^*) > 0$, decreasing in π_t if $\mathcal{I}_t(\rho_t^*) < 0$, and independent of π_t if $\mathcal{I}_t(\rho_t^*) = 0$. $\square$

Proof of Proposition 3. Follows from the in-text discussion. $\square$

Proof of Corollary 1. If $\mathcal{I}(\sigma \mid F, \mu) \leq 0$ for every skeptic $\sigma > \mu$, then every skeptic gives Sender payoff at most $F(\mu)$, while targeting a fan and saying nothing gives exactly $F(\mu)$. Therefore there is an equilibrium in which Sender targets a fan, and no equilibrium can give Sender more than $F(\mu)$. Conversely, if a skeptic has strictly positive example value, then targeting that skeptic gives payoff strictly above $F(\mu)$ because $\pi(\sigma) > 0$. $\square$

Proof of Proposition 4. Because Sender targets only skeptics under both markets, the set of optimal targets for Sender in either market are, respectively,

$$\begin{aligned}\Sigma^*(\pi, F, \mu) &= \arg \max_{\sigma \in (\mu, 1]} V(\sigma \mid \pi, F, \mu), \\ \Sigma^*(\pi, F', \mu) &= \arg \max_{\sigma \in (\mu, 1]} V(\sigma \mid \pi, F', \mu).\end{aligned}$$

By the monotonicity theorem of [Milgrom and Shannon \(1994\)](#), a sufficient condition for Sender to target more skeptical receivers under (π, F, μ) than (π, F', μ) is for $V(\cdot \mid \pi, F, \mu)$ to single-crossing dominate $V(\cdot \mid \pi, F', \mu)$. To see that this is true, take any $\mu < \sigma' < \sigma \leq 1$ and suppose $V(\sigma \mid \pi, F', \mu) \geq (>) V(\sigma' \mid \pi, F', \mu)$. This implies

$$\frac{\pi(\sigma)}{\pi(\sigma')} \geq (>) \frac{\mathcal{I}(\sigma' \mid F', \mu)}{\mathcal{I}(\sigma \mid F', \mu)}.$$

By the assumptions on $\mathcal{I}$,

$$\frac{\mathcal{I}(\sigma' \mid F, \mu)}{\mathcal{I}(\sigma \mid F, \mu)} \leq \frac{\mathcal{I}(\sigma' \mid F', \mu)}{\mathcal{I}(\sigma \mid F', \mu)}$$

Therefore,

$$\frac{\pi(\sigma)}{\pi(\sigma')} \geq (>) \frac{\mathcal{I}(\sigma' \mid F, \mu)}{\mathcal{I}(\sigma \mid F, \mu)},$$

which implies $V(\sigma \mid \pi, F, \mu) \geq (>) V(\sigma' \mid \pi, F, \mu)$. Thus, $V(\cdot \mid \pi, F, \mu)$ single-crossing dominates $V(\cdot \mid \pi, F', \mu)$. $\square$

Proof of Proposition 5. By a similar logic to the proof of Proposition 4, it suffices to show that $V(\cdot \mid \pi_1, F, \mu)$ single-crossing dominates $V(\cdot \mid \pi_0, F, \mu)$. Take any $\mu < \sigma' < \sigma \leq 1$ and suppose $V(\sigma \mid \pi_0, F, \mu) \geq (>) V(\sigma' \mid \pi_0, F, \mu)$, so

$$\frac{\pi_0(\sigma)}{\pi_0(\sigma')} \geq (>) \frac{\mathcal{I}(\sigma' \mid F, \mu)}{\mathcal{I}(\sigma \mid F, \mu)}.$$

By the assumptions on π_1 and π_0 ,

$$\frac{\pi_1(\sigma)}{\pi_1(\sigma')} \geq \frac{\pi_0(\sigma)}{\pi_0(\sigma')}$$

Therefore,

$$\frac{\pi_1(\sigma)}{\pi_1(\sigma')} \geq (>) \frac{\mathcal{I}(\sigma' \mid F, \mu)}{\mathcal{I}(\sigma \mid F, \mu)}.$$

which implies $V(\sigma \mid \pi_1, F, \mu) \geq (>)V(\sigma' \mid \pi_1, F, \mu)$. Thus, $V(\cdot \mid \pi_1, F, \mu)$ single-crossing dominates $V(\cdot \mid \pi_0, F, \mu)$. $\square$

Proof of Lemma 1. We split the proof into two parts. Part 1 shows that, as claimed in the text, $V^*(\pi, F, \mu)$ is an upper bound on Sender's payoff under any equilibrium of any persuasion mechanism. Part 2 shows that this upper bound is tight. Throughout, we let

$$\Delta(p^1, p^0) \equiv B(p^1, p^0)F(p^1) + (1 - B(p^1, p^0))F(p^0) - F(\mu).$$

Part 1. Consider any persuasion mechanism and any equilibrium under it. If the target's example is uninformative, then all non-targeted receivers hold posterior μ whether or not they observe the target's action, and Sender's payoff is $F(\mu) \leq V^*(\pi, F, \mu)$.

Now suppose the target's example is informative. Let p^1 and p^0 be the posterior beliefs of non-targets after observing, respectively, a positive and a negative example. By Bayes plausibility, $p^1 \geq \mu \geq p^0$. Moreover, if the target has skepticism σ , then observing that he buys implies that his posterior is weakly above σ , while observing that he does not buy implies that his posterior is strictly below σ . Hence $\sigma \in (p^0, p^1]$, so Sender's payoff in this equilibrium is

$$B(p^1, p^0)v(p^1 \mid \pi(\sigma)) + (1 - B(p^1, p^0))v(p^0 \mid \pi(\sigma)) = F(\mu) + \pi(\sigma)\Delta(p^1, p^0).$$

If $\Delta(p^1, p^0) \leq 0$, then this payoff is at most $F(\mu) \leq V^*(\pi, F, \mu)$. If $\Delta(p^1, p^0) > 0$, then the payoff is increasing in the target's popularity. Therefore, using the definition

$$\bar{\pi}(p^1, p^0) \equiv \sup_{\sigma \in (p^0, p^1]} \pi(\sigma),$$

we have

$$F(\mu) + \pi(\sigma)\Delta(p^1, p^0) \leq F(\mu) + \bar{\pi}(p^1, p^0)\Delta(p^1, p^0) \leq V^*(\pi, F, \mu).$$

Thus $V^*(\pi, F, \mu)$ is an upper bound on Sender's payoff.

Part 2. Fix $\varepsilon > 0$. If $V^*(\pi, F, \mu) = F(\mu)$, then the bound is achieved by a trivial persuasion mechanism that reveals no information, together with an equilibrium in which Sender targets no receiver. Hence, for the remainder of the proof, we suppose $V^*(\pi, F, \mu) > F(\mu)$. Notice then that, defining

$$\xi \equiv \min \left\{ \frac{\varepsilon}{2}, \frac{V^*(\pi, F, \mu) - F(\mu)}{2} \right\} > 0,$$

since V^* is a supremum there exists $p^0 < \mu < p^1$ and $\hat{\sigma} \in (p^0, p^1]$ such that

$$V_0 \equiv B(p^1, p^0)v(p^1 \mid \pi(\hat{\sigma})) + (1 - B(p^1, p^0))v(p^0 \mid \pi(\hat{\sigma})) \geq V^*(\pi, F, \mu) - \xi.$$

Further note that because $V^*(\pi, F, \mu) > F(\mu)$, we must have $\Delta(p^1, p^0) > 0$. In what follows, we fix p^0, p^1 and $\hat{\sigma}$ satisfying the above.

The goal is to construct a persuasion mechanism and an equilibrium under it where Sender's payoff is at least $V^*(\pi, F, \mu) - \varepsilon$. We divide the proof into three steps. Step 1 constructs a distribution over posteriors which randomizes between the high belief p^1 and an interval of beliefs close to p^0 . Step 2 constructs a constrained Bayesian persuasion mechanism that induces this distribution and forces Sender to truthfully report the signal observed. Finally, Step 3 constructs an equilibrium under this persuasion mechanism and shows that Sender's payoff is at least $V^*(\pi, F, \mu) - \varepsilon$, which proves the claim.

Step 1. For each $\eta > 0$, let

$$q_\eta \equiv p^0 + \frac{\eta}{2},$$

and define

$$\Delta_\eta \equiv B(p^1, q_\eta)F(p^1) + (1 - B(p^1, q_\eta))F(q_\eta) - F(\mu).$$

Also define

$$V_\eta \equiv B(p^1, q_\eta)v(p^1 \mid \pi(\hat{\sigma})) + (1 - B(p^1, q_\eta))v(q_\eta \mid \pi(\hat{\sigma})).$$

Since $q_\eta \rightarrow p^0$ as $\eta \rightarrow 0$, we have $\Delta_\eta \rightarrow \Delta(p^1, p^0) > 0$ and $V_\eta \rightarrow V_0$. Therefore, by continuity we can choose $\eta > 0$ small enough such that

$$p^0 + \eta < \hat{\sigma}, \quad p^0 + \frac{\eta}{2} < \mu, \quad \Delta_\eta > 0, \quad V_\eta \geq V_0 - \frac{\varepsilon}{2}.$$

Let λ_η denote the uniform distribution over $[p^0, p^0 + \eta]$, and let ρ_η denote the distribution over posteriors defined by

$$\rho_\eta \equiv (1 - B(p^1, q_\eta))\lambda_\eta + B(p^1, q_\eta)\delta_{p^1},$$

where δ_x denotes the Dirac measure concentrated at x . Note that by the definition of q_η , ρ_η is Bayes-plausible.

Step 2. We now construct a persuasion mechanism that induces ρ_η . Let $S = M = [p^0, p^0 + \eta] \cup \{p^1\}$, $\Phi = \{\phi_\eta\}$, where ϕ_η is defined by

$$\phi_\eta(A \mid G) = \int_A \frac{s}{\mu} d\rho_\eta(s), \quad \phi_\eta(A \mid B) = \int_A \frac{1-s}{1-\mu} d\rho_\eta(s) \quad \forall \text{ Borel } A \subseteq S,$$

and define the cost function c via

$$c(m, s) = \begin{cases} 0, & m = s, \\ \infty, & m \neq s. \end{cases}$$

By construction, this information structure is *direct*, so $\Pr(\theta = G \mid s) = s$ for all $s \in S$. Furthermore, this persuasion mechanism is a constrained Bayesian-persuasion mechanism: it forces Sender to truthfully report her message to the target. It follows that the distribution over targets' posteriors is exactly ρ_η .

Step 3. We begin by computing Sender's continuation equilibrium payoff from targeting a receiver with skepticism $\sigma \in [0, 1]$, which we denote by $\widehat{V}(\sigma)$

1. **Case 1:** Suppose $\sigma > p^1$ or $\sigma \leq p^0$. Then, the target always buys or does not buy the widget, so his action is uninformative. This means that all non-targets hold belief μ regardless of observing the target buy and not buy, or observing nothing. Hence,

$$\widehat{V}(\sigma) = F(\mu)$$

2. **Case 2:** Suppose $\sigma \in (p^0, p^0 + \eta)$. Then the target does not buy after posteriors in $[p^0, \sigma)$ and buys after posteriors in $[\sigma, p^0 + \eta] \cup \{p^1\}$. Letting $\alpha_\eta(\sigma)$ and $\beta_\eta(\sigma)$ be, respectively, the probability of a negative and positive example, so

$$\alpha_\eta(\sigma) \equiv (1 - B(p^1, q_\eta)) \frac{\sigma - p^0}{\eta} \quad \text{and} \quad \beta_\eta(\sigma) \equiv 1 - \alpha_\eta(\sigma),$$

it follows that non-targets' posteriors from observing a negative and positive example from the target are, respectively,

$$p_\eta^0(\sigma) = \frac{p^0 + \sigma}{2} \quad \text{and} \quad p_\eta^1(\sigma) = \frac{\mu - \alpha_\eta(\sigma)p_\eta^0(\sigma)}{\beta_\eta(\sigma)}.$$

Therefore, Sender's continuation payoff from targeting a receiver with skepticism $\sigma \in (p^0, p^0 + \eta)$ is

$$\widehat{V}(\sigma) = \beta_\eta(\sigma)v(p_\eta^1(\sigma) \mid \pi(\sigma)) + \alpha_\eta(\sigma)v(p_\eta^0(\sigma) \mid \pi(\sigma)).$$

3. **Case 3:** Finally, suppose $\sigma \in [p^0 + \eta, p^1]$. Then the target buys if and only if he observes the high posterior p^1 . This means that non-targets' posteriors from observing a negative and positive example are, respectively, q_η and p^1 . Since the probability of

a positive example is $B(p^1, q_\eta)$,

$$\widehat{V}(\sigma) = B(p^1, q_\eta)v(p^1 \mid \pi(\sigma)) + (1 - B(p^1, q_\eta))v(q_\eta \mid \pi(\sigma)).$$

We now show that $\widehat{V}$ is upper semicontinuous. Since $\widehat{V}$ is continuous on each of the regions described above, it remains only to check what happens at the boundary points p^0 , $p^0 + \eta$, and p^1 . At p^0 ,

$$\lim_{\sigma \downarrow p^0} \widehat{V}(\sigma) = F(\mu) = \widehat{V}(p^0).$$

so $\widehat{V}$ is continuous at p^0 . At $p^0 + \eta$, the formulas in Cases 2 and 3 coincide. Indeed,

$$\alpha_\eta(p^0 + \eta) = 1 - B(p^1, q_\eta), \quad \beta_\eta(p^0 + \eta) = B(p^1, q_\eta),$$

and

$$p_\eta^0(p^0 + \eta) = q_\eta, \quad p_\eta^1(p^0 + \eta) = p^1,$$

so $\widehat{V}$ is continuous at $p^0 + \eta$. Finally, at p^1 , notice that

$$\widehat{V}(p^1) = F(\mu) + \pi(p^1)\Delta_\eta > F(\mu) = \lim_{\sigma \rightarrow p^1+} \widehat{V}(\sigma).$$

So $\widehat{V}$ has an upwards jump at p^1 .

Because $\widehat{V}$ is upper semicontinuous, $\arg \max_{\sigma \in [0,1]} \widehat{V}(\sigma)$ is non-empty. This means that there exists a PBE under this persuasion mechanism in which Sender targets $\sigma^* \in \arg \max_{\sigma \in [0,1]} \widehat{V}(\sigma)$ and achieves a payoff of $\widehat{V}(\sigma^*)$. Noticing that $\hat{\sigma} \in [p^0 + \eta, p^1]$, so

$$\widehat{V}(\hat{\sigma}) = B(p^1, q_\eta)v(p^1 \mid \pi(\hat{\sigma})) + (1 - B(p^1, q_\eta))v(q_\eta \mid \pi(\hat{\sigma})) = V_\eta.$$

It follows that

$$\widehat{V}(\sigma^*) \geq \widehat{V}(\hat{\sigma}) \geq V_0 - \frac{\varepsilon}{2} \geq V^*(\pi, F, \mu) - \xi - \frac{\varepsilon}{2} \geq V^*(\pi, F, \mu) - \varepsilon,$$

where the last inequality holds as $\xi \leq \varepsilon/2$. Thus, Sender's equilibrium payoff under this persuasion mechanism is at least equal to $V^*(\pi, F, \mu) - \varepsilon$. $\square$

Proof of Proposition 6. Notice that

$$\begin{aligned}
V^*(\pi, F, \mu) &\geq B(\sigma, 0)v(\sigma \mid \bar{\pi}(\sigma, 0)) + (1 - B(\sigma, 0))v(0 \mid \bar{\pi}(\sigma, 0)) \\
&= \bar{\pi}(\sigma, 0)\mathcal{I}(\sigma \mid F, \mu) + F(\mu) \\
&\geq \pi(\sigma')\mathcal{I}(\sigma \mid F, \mu) + F(\mu) \\
&> \pi(\sigma)\mathcal{I}(\sigma \mid F, \mu) + F(\mu).
\end{aligned}$$

The first inequality holds as $0 \leq \mu \leq \sigma$, the second inequality holds as $\sigma' \in (0, \sigma]$, and the strict inequality holds as $\mathcal{I}(\sigma \mid F, \mu) > 0$ and $\pi(\sigma') > \pi(\sigma)$. Since $F(\mu) + \pi(\sigma)\mathcal{I}(\sigma \mid F, \mu)$ is the Sender's value under Bayesian persuasion, this proves the claim. $\square$

Proof of Proposition 7. Since $\pi(\sigma)$ is constant in σ , we let $\pi(\sigma) = \pi$ throughout. We split the proof into two parts.

Part 1. Suppose that F is convex on $[0, \mu]$. Then for all $0 < p^0 \leq \mu$, $\frac{F(p^0)}{p^0} \leq \frac{F(\mu)}{\mu}$. Therefore,

$$\begin{aligned}
&B(p^1, p^0)F(p^1) + (1 - B(p^1, p^0))F(p^0) \\
&\leq B(p^1, p^0)p^1 \max\left\{\frac{F(\mu)}{\mu}, \frac{F(p^1)}{p^1}\right\} + (1 - B(p^1, p^0))p^0 \max\left\{\frac{F(\mu)}{\mu}, \frac{F(p^1)}{p^1}\right\} \\
&= \mu \max\left\{\frac{F(\mu)}{\mu}, \frac{F(p^1)}{p^1}\right\} \\
&\leq \max_{\sigma \in [\mu, 1]} \frac{\mu}{\sigma} F(\sigma).
\end{aligned}$$

This means

$$\begin{aligned}
V^*(\pi, F, \mu) &= (1 - \pi)F(\mu) + \pi \sup_{0 \leq p^0 \leq \mu \leq p^1 \leq 1} \{B(p^1, p^0)F(p^1) + (1 - B(p^1, p^0))F(p^0)\} \\
&\leq (1 - \pi)F(\mu) + \pi \max_{\sigma \in [\mu, 1]} \frac{\mu}{\sigma} F(\sigma) \\
&= \max_{\sigma \in [0, 1]} \{\pi(\sigma)\mathcal{I}(\sigma \mid F, \mu) + F(\mu)\}.
\end{aligned}$$

where the first equality holds as popularity is constant in skepticism. Since $\max_{\sigma \in [0, 1]} \{\pi(\sigma)\mathcal{I}(\sigma \mid F, \mu) + F(\mu)\}$ is the payoff under Bayesian persuasion, the claim follows.

Part 2. Suppose that there exists $\hat{\sigma} \in (0, \mu)$ such that $F(\hat{\sigma}) > \hat{\sigma}$, and that F is strictly convex on $[\mu, 1]$. The latter implies

$$\max_{\sigma \in [\mu, 1]} \frac{\mu}{\sigma} F(\sigma) = \max\{F(\mu), \mu\},$$

and so

$$\frac{F(\sigma)}{\sigma} \leq \max \left\{ \frac{F(\mu)}{\mu}, 1 \right\}.$$

We now consider two cases. If $F(\mu) \leq \mu$, then

$$\max_{\sigma \in [\mu, 1]} \frac{\mu}{\sigma} F(\sigma) = \mu,$$

and so Sender's payoff under Bayesian persuasion is

$$\max_{\sigma \in [0, 1]} \{ \pi(\sigma) \mathcal{I}(\sigma \mid F, \mu) + F(\mu) \} = (1 - \pi)F(\mu) + \pi\mu.$$

Meanwhile, observe

$$\begin{aligned} V^*(\pi, F, \mu) &\geq B(1, \hat{\sigma})v(1 \mid \bar{\pi}(1, \hat{\sigma})) + (1 - B(1, \hat{\sigma}))v(\hat{\sigma} \mid \bar{\pi}(1, \hat{\sigma})) \\ &= B(1, \hat{\sigma})[\pi F(1) + (1 - \pi)F(\mu)] + (1 - B(1, \hat{\sigma}))[\pi F(\hat{\sigma}) + (1 - \pi)F(\mu)] \\ &= (1 - \pi)F(\mu) + \pi [B(1, \hat{\sigma})F(1) + (1 - B(1, \hat{\sigma}))F(\hat{\sigma})] \\ &= (1 - \pi)F(\mu) + \pi \left[\frac{\mu - \hat{\sigma}}{1 - \hat{\sigma}} + \frac{1 - \mu}{1 - \hat{\sigma}} F(\hat{\sigma}) \right] \\ &= (1 - \pi)F(\mu) + \pi \left[\mu + \frac{1 - \mu}{1 - \hat{\sigma}} (F(\hat{\sigma}) - \hat{\sigma}) \right] \\ &> (1 - \pi)F(\mu) + \pi\mu, \end{aligned}$$

where the strict inequality follows from the fact that $F(\hat{\sigma}) > \hat{\sigma}$. Thus, Sender's payoff under Bayesian persuasion is strictly less than $V^*(\pi, F, \mu)$.

If $F(\mu) > \mu$, then

$$\max_{\sigma \in [\mu, 1]} \frac{\mu}{\sigma} F(\sigma) = F(\mu),$$

and so Sender's payoff under Bayesian persuasion is

$$\max_{\sigma \in [0, 1]} \{ \pi(\sigma) \mathcal{I}(\sigma \mid F, \mu) + F(\mu) \} = F(\mu).$$

Meanwhile, define

$$G(\varepsilon) \equiv \frac{\varepsilon}{1 - \mu + \varepsilon} + \frac{1 - \mu}{1 - \mu + \varepsilon} F(\mu - \varepsilon).$$

Notice that $G(0) = F(\mu)$, while

$$G'(0) = \frac{1 - F(\mu)}{1 - \mu} - F'(\mu) > 0,$$

where the strict inequality holds as F is strictly convex on $[\mu, 1]$. Therefore, for $\varepsilon > 0$ sufficiently small, $G(\varepsilon) > F(\mu)$ holds. In turn,

$$\begin{aligned}
V^*(\pi, F, \mu) &\geq B(1, \mu - \varepsilon)v(1 \mid \bar{\pi}(1, \mu - \varepsilon)) + (1 - B(1, \mu - \varepsilon))v(\mu - \varepsilon \mid \bar{\pi}(1, \mu - \varepsilon)) \\
&= (1 - \pi)F(\mu) + \pi [B(1, \mu - \varepsilon)F(1) + (1 - B(1, \mu - \varepsilon))F(\mu - \varepsilon)] \\
&= (1 - \pi)F(\mu) + \pi \left[\frac{\varepsilon}{1 - \mu + \varepsilon} + \frac{1 - \mu}{1 - \mu + \varepsilon} F(\mu - \varepsilon) \right] \\
&> (1 - \pi)F(\mu) + \pi F(\mu) \\
&= F(\mu).
\end{aligned}$$

Thus, Sender's payoff under Bayesian persuasion is strictly less than $V^*(\pi, F, \mu)$. $\square$