Targeting in Networks and Markets: An Information Design Approach*

Roberto Corrao

MIT

May 2021

Abstract

In many economic settings, heterogeneous information is aggregated through channels such as social networks or markets' prices. Moreover, information is often controlled and manipulated as to influence the final outcome. The goal of this paper is to introduce aggregation mechanisms in an otherwise standard information design environment and analyze their effect on the information released and on economic outcomes. First, the analysis provides a benchmark irrelevance result: when the designer can target every receiver and the aggregator is linear, it is without loss of optimality to consider public experiments that do not depend on the aggregation mechanism. Differently, if the designer can target only a subset of receivers, then the most prominent individuals are chosen. Next, comparative statics results that link the informativeness of the optimal policy to the underlying aggregation process are discussed. Finally, motivated by robustness concerns, it is showed that the main findings extend to a class of nonlinear aggregation mechanisms.

*I am particularly indebted to Glenn Ellison, Joel Flynn, Giacomo Lanzani, and Stephen Morris for several suggestions that substantially improved the paper. I also wish to thank Marc de la Barrera, Esther Duflo, Drew Fudenberg, Ying Gao, Bob Gibbons, Benjamin Golub, as well as all the participants of the Theory lunch at MIT Economics for helpful comments and suggestions. All the remaining errors are my own. The financial support of the Marco Fanno scholarship is gratefully acknowledged.

1 Introduction

In various economic settings, heterogeneous information is aggregated through several channels. For example, agents in a social network that communicate and learn from each other might aggregate their private information and reach an agreement in the long-run. Differently, agents in competitive markets use their private information about traded assets when choosing the quantity to acquire. With this, the competitive equilibrium prices might aggregate all the sparse information about the goods. At the same time, information is often controlled and manipulated by other individuals as to influence the final outcome of the aggregation process described above. It is enough to think about firms advertising their products on social networks or disclosing information about their financial situation to the market as to influence the price of their stocks. Finally, in both the scenarios described, some key agents might have a more prominent role in affecting the aggregation mechanisms. For instance, the opinions of the so-called "influencers" in social networks is considered relatively more by the constellation of "simpler" agents. Similarly, consumers with an higher willingness to pay and an higher demand volume are going to influence more the market price of an asset.

These considerations raise some natural questions: Will be the optimal information policy affected by the underlying aggregation process? Will the key agents targeted by more tailored information policies? The goal of this paper is to answer to these questions by introducing aggregation mechanisms in an otherwise standard information design environment and analyze their effect on the information released and on economic outcomes. In particular, the analysis is almost entirely focused on a benchmark model where a reduced-form *linear* aggregation mechanism captures non-strategic interactions among agents (e.g., naive social learning, competitive markets). The information designer is aware of these mechanisms when choosing the information to disclose and might take in account the relative importance of the receivers in the aggregation process. In this setting, I first show that, when completely unrestricted, it is without loss of optimality for the designer to disclose public information that does not depend on the aggregation structure. On the other hand, if the designer can persuade only a subset of agent of his choice, then he will optimally target the most prominent agent according to the aggregation mechanism and the information disclosed depends on the total influence of the chosen targets.

The model I consider an information design problem with a unique sender and multiple receivers. For example, one can interpret the receivers as nodes of a social network such as a group of consumers linked on Facebook. Alternatively, it is possible to interpret the receivers as agents in an competitive market placing orders for the quantity of an asset. Regardless of the particular interpretation, it is assumed that there is a common state of uncertainty that is of interest for the receivers. In relation to the previous examples, the state of uncertainty might represent the objective quality of a new good or the objective return of a financial asset.

Moreover, all the receivers, as well as the sender, share the same common prior about this uncertain state. In light of the examples given, the common-prior hypothesis might seem too strong and restrictive. However, the aim of this paper is to isolate the effect of the aggregation mechanism on the information choice of the sender without further sources of heterogeneity across receivers.[1]

The receivers obtain some private information about the uncertain state from the sender. For instance, a group of consumers linked on Facebook might receive information about a new product in the form of personalized adds on their pages from the advertiser of the product. However, differently form large part of the previous information-design literature, the receivers do not interact strategically and their private information is aggregated through an aggregation mechanism. Here, the goal is to understand the role of these information spillovers across the agents and how they change the persuasion problem.

In the leading examples, the aggregation mechanisms represent either the social network communication among agents or the market prices. In the first case, I model the aggregation of opinions in the network according to the simple DeGroot model. Formally, after receiving their private initial information, the agents periodically update their opinions (i.e., beliefs) by taking simple averages of their neighbors' opinions. If the social network is connected, then the agents will eventually reach a consensus (i.e., aggregated opinion) that I interpret as the outcome of the aggregation mechanism.[2] In the second case, the agents are consumers characterized by quadratic preferences over quantities of the unique good traded in competitive markets. The consumers are uncertain about the objective quality of the good and receive private information from the sender. Therefore, with a fixed supply, the equilibrium price of the good results from a linear aggregation of the posterior beliefs of the agent.

The sender's payoff depends on the outcome of the interaction of the agents and he can influence it by manipulating their private information. For example, the advertiser might be interested in the consensus opinion about the new product emerging from the communication process within the social network. Similarly, a firm disclosing information to the agents in a market is interested in the effect on the resulting market price of the asset.

Unrestricted information design I start my analysis by considering the case where the information designer is completely unrestricted in his choice of the information policy. In this case, the answer to the previous questions is that restricting on *public signals* is without loss of optimality for the sender. Moreover, the optimal public information policy does not depend on the aggregation mechanism (e.g., the network structure or the willingness to pay). The intuition of this irrelevance result is that, anticipating the aggregation process, the sender knows that

[1] See Section 5 for a further discussion on this point.

[2] Here, the inconsistency between the Bayesian updating for the sender's signals and the DeGroot updating for the other receivers' opinion might be surprising. I postpone the detailed discussion of this point to Section 5.

consensus will be eventually reached and he can replicate any such outcome by just disclosing the same information to all agents.

Next, I test the robustness of the irrelevance results with respect to both more general forms of aggregation that allow for less naiveté of the agents (i.e., beyond the DeGroot model), and uncertainty of the designer about the aggregation mechanism. In both cases, and under some mild assumptions, the irrelevance result still holds, hence downplaying the effect of the aggregation process. However, targeting policies and personalized information disclosure are pervasive and often effective in social networks (e.g., Banerjee et al 2013, 2019). Motivated, by these concerns, I next analyze a version of the model with additional restriction on the designer.

Restricted information design In the second part of the paper, I consider the case where the sender can persuade only a subset of agents of his choice. With this, the standard information design problem is augmented with an additional targeting choice. On the one hand, this modeling choice is very similar to standard models of seeding in networks (see for example Kempe et al 2003 and Akbarpour et al 2018). On the other hand, two important distinctive elements makes the analysis different. First, rather than a model of diffusion of information, I consider aggregation mechanism that can capture a broader spectrum of economic settings than just networks.[3] Second, the targeting problem is just the first of the two relevant choices of the sender, hence it needs to be paired with the design of the experiment for the targeted agents.

In this case, I find that the irrelevance result is partially overcome. Indeed, the Sender can still restrict without loss of optimality public signals for the subset of targeted agents. However, both the identity of the targeted agents and the informativeness of the chosen signal depend on the underlying aggregation structure. In particular, the target set is chosen as to maximize the total sum of the aggregation weights of the set itself. Interestingly, under the social network interpretation of my model, this simple result provides a foundation for targeting policies based on centrality measure on networks (i.e., the most influential agents are optimally targeted with information).

Next, I exploit the targeting result in order to link properties of the optimal information structure to the aggregation mechanism considered. First, I find that, as the size of the population of receivers grows, extreme information structures such as full-disclosure and no-disclosure are optimal depending on the curvature of the payoff of the sender around the prior mean. This result exploits the observation that, in the restricted case, the problem of the sender is equivalent to a single-receiver Bayesian persuasion problem with a degree of stickiness of the receiver to the prior mean. Second, I derive a comparative static result linking the informativeness of the optimal information policy to variations of the underlying aggregation structure

[3] See the survey by Breza et al (2019) for a detailed classification of models of aggregation and diffusion of information in networks.

(e.g., additional links in a social network). Again, the equivalent representation of the sender's problem allows me to adapt results in the information design literature to obtain the desired comparative statics. Finally, I extend the main finding of the section for a subclass of nonlinear aggregation mechanism. Also in this case, the sender optimally targets the set of agents that maximize a (nonlinear) measure of influence of the agents in the aggregation.

Literature review This paper lies at the intersection of several literatures, the most relevant ones being information design and targeting in networks.

From the seminal contributions of Kamenica and Gentzkow (2011) and Bergemann and Morris (2013, 2016), to the most recent contributions such as Kolotilin (2018), Dworczak and Martini (2019), and Mathevet et al (2020), the works analyzing information disclosure in economic settings have been exponentially growing in the recent years.[4] Following this new tradition, the current paper analyzes the problem of a sender that designs information for multiple receivers and is able to commit ex-ante.[5] Differently from the mentioned papers, and more in general from majority of the literature, the model proposed here considers non-strategic interactions among multiple receivers in the form, for example, of communication in networks and competitive markets. This hypothesis greatly simplifies the analysis as it avoids dealing with the induced distributions over higher-order beliefs of the agents that do not always admit an explicit characterization (cf. Mathevet et al 2020). The fruitfulness of this approach has also been recently showed in Arieli et al (2020) where the authors characterize, in a binary-state environment, the feasible joint distributions over first-order posterior beliefs that can be induced by some information structure. In particular, Arieli et al (2020) define the *social Bayesian persuasion* problem where, as in this paper, the payoff of the sender depends on the actions of agents that do not interact strategically, and solve for some simple examples exploiting their characterization. My work goes a step further in this direction by avoiding the restriction on binary-state environments and by characterizing the sender's optimal value for a broad class of problems. On the hand, this allows me to establish the optimality of public experiments that can be computed by standard tools analyzed in single-receiver Bayesian persuasion problems as in, for example, Kolotilin (2018) and Dworczak and Martini (2019). On the other hand, the value characterization substantially simplifies the targeting problem that I analyze in the second part of paper. In this sense, the recent contribution by Galperti and Perego (2019) also analyzes the problem of a sender that faces multiple receivers playing a game on a network and designs the information only for an exogenously determined subset of them. Differently, in this paper, I let the sender endogenously select the target set and I provide explicit properties both for the optimal target set and the optimal information policies.

There is a vast literature that analyzes, both theoretically and empirically, targeting inter-

[4]See Bergemann and Morris (2019) for a comprehensive survey on information design.

[5]See Section 5 for a discussion on how to relax the commitment assumption.

ventions in networks.[6] In particular, it is possible to categorize the theoretical contributions in two broad categories: models of diffusion of information in networks as in Kempe et al (2013), Banerjee et al (2013, 2019), and Akbarpour et al (2018), and models of aggregation of information in networks as in DeMarzo et al (2003), Golub and Jackson (2010, 2012), Cerreia-Vioglio et al (2020), and Chandrasekhar et al (2020). In the first case, the sequential diffusion of a binary piece of information or action (like adoption of a new product) is analyzed. In the second case, as in this paper, the agents is a social network aggregate their information through direct observation of neighbors' continuous opinions. While the first group of contributions directly deal with optimality concerns for targeting intervention in networks, the second group of contributions focus on a description of the opinion dynamics without directly addressing the intervention problem. In this sense, this paper is a step forward as it directly analyzes the choice of optimal interventions, in the form of information disclosure, within a model of information aggregation in social networks. Moreover, coherently with both the theoretical literature (for example, Ballester et al 2006 and Galeotti et al 2020) and the empirical literature (Banerjee et al 2013, 2019), the findings of this paper confirm the optimality of targeting interventions in networks based on centrality measures.

More in general, to the best of my knowledge, this is the first work that analyzes targeted information policies in social networks and markets under the lenses of the recent information design literature. On the one hand, this approach allows me to derive novel results that link the underlying aggregation mechanism to structural properties of both optimal information and targeting policies. On the other hand, differently from the literature on information disclosure to financial markets, this approach does not rely on ad hoc parametric restrictions on the feasible information structures. Finally, differently from the rest of the works mentioned, my analysis also deals with robustness concerns both from the perspective of the designer within the model and the analyst outside the model.

Structure of the paper Section 2 presents the model and illustrates it through two examples. Section 3 analyzes the unrestricted information design case and derives the main irrelevance results. Section 4 analyzes the case where the sender is restricted to target a subset of the receivers, and shows that the most prominent receivers are optimally targeted. Finally, Section 5 summarizes the entire analysis and discusses the more relevant extensions of the main model by highlighting further links with the literature.

[6] For a textbook exposition on these topics see Jackson (2008), while for a recent survey with applications to development economics, see Breza et al (2019).

2 Model

This section introduces a simple model of *targeting and persuasion* of a group of agents, called receivers, that interact non-strategically with each other. The initial information of the agents is designed by an external agent, called Sender, who wants to influence the final outcome of their interaction. In particular, the final outcome is generated by an *aggregation mechanism* (or aggregators) that maps the receiver's posterior beliefs given their information to an aggregated variable. On the one hand, most of the analysis will focus on affine aggregators of agents' beliefs. On the other hand, this class is flexible enough to capture in a simple and reduced form several examples of economic non-strategic interactions. Specifically, the two main examples, presented at the end of this section, describe with stylized models the process of information aggregation in social networks and in competitive markets through prices. We first start with some notation and the formalization of the primitives of the model.

Preliminary notation For all subsets Ω of an Euclidean space, let $\mathcal{B}_\Omega$ denote the class of Borel subsets of Ω, let $\Delta(\Omega)$ denote the set of Borel probability measures over Ω and endow it with the topology of weak convergence, and let $C(\Omega)$ denote the set of continuous functions over Ω.[7] For all $\mu, \nu \in \Delta([0,1])$, μ dominates ν in the *concave order*, written $\mu \succsim_{cv} \nu$, if

$$\int \phi(z)\, d\mu(z) \geq \int \phi(z)\, d\nu(z)$$

for all concave and continuous functions $\phi : [0,1] \to \mathbb{R}$. In this case, it is possible to obtain ν as a *mean-preserving spread* of μ.

2.1 Targeting, information design and aggregators

Agents and common prior Consider a finite set $N_0 = \{0, 1, ..., n\}$ of agents where 0 is the sender or (information) designer and each $i \in N = \{1, ..., n\}$ is a receiver. I will use the male pronoun he and the female pronoun she to refer respectively to the sender and an arbitrary receiver. There is a common state of uncertainty $\omega \in \Omega = [0,1]$, hence the space of first-order beliefs for every agent is $\Delta(\Omega)$.[8] Before any information is revealed, all the agents, included the sender, share the same common prior $\pi \in \Delta(\Omega)$ with corresponding common expectation $\bar{\pi} \in [0,1]$. This is going to be an important assumption for the derivation of the main results.

[7] Recall that a sequence of probability measures $\{\mu_n\}_{n\in\mathbb{N}} \subseteq \Delta(\Omega)$ weakly converges to $\mu \in \Delta(\Omega)$ if

$$\int_\Omega f(\omega)\, d\mu_n(\omega) \to \int_\Omega f(\omega)\, d\mu(\omega)$$

for all bounded and continuous functions $f : \Omega \to \mathbb{R}$. For details, see, for example, Aliprantis and Border (2006, Chapter 15).

[8] I consider $[0,1]$ for the sake of expositional simplicity but all the results would remain true for any closed and bounded interval of the real line.

The discussion concerning how to extend the model for the heterogenous case is postponed to Section 5. As detailed above, the receivers will receive (potentially private) information from the sender and update their beliefs according to Bayes rule. In particular, for all $i \in N$, let $p_i \in \Delta(\Omega)$ denote an arbitrary posterior belief for i. Given a posterior $p_i \in \Delta(\Omega)$, let $x_i = \mathbb{E}_{p_i}[\omega] \in [0,1]$ denote the corresponding conditional expectation or *opinion*. Similarly, let $x = (x_i)_{i \in N} \in X = [0,1]^n$ denote an arbitrary profile of opinions for all the receivers.[9]

Experiments The sender can design a public or private experiment whose realizations are observed by the receivers. Formally, an experiment is a pair (Y, σ) of *signal space* $Y = \prod_{i \in N} Y_i$, where each *personal signal space* Y_i is an arbitrary measurable space, and $\sigma : \Omega \to \Delta(Y)$ is Markov kernel mapping states to distributions over observables. The interpretation is that each receiver i observes her realization $y_i \in Y_i$ in her personal signal space and updates her prior expectation $\bar{\pi}$ to the opinion

$$x_i(y_i) = \mathbb{E}_{\pi,\sigma}[\omega | y_i]. \tag{1}$$

With this, an arbitrary profile of signal realizations $y = (y_i)_{i \in N} \in Y$ induces a profile of opinions $x(y) \in X$ as defined in eq. (1).

Let Σ denote the set of all the *conceivable* experiments for the sender. In particular, it will not be always the case that the set Σ coincides with the set of *feasible* experiments for the sender. In the most general case, the sender will have restrictions on the maximum number of receivers whose experiment is informative about the state ω. Formally, for all $k \in \{1, ..., n\}$ let Σ_k be the set of experiments such that for each receiver i in a subset $I \subseteq N$ of size k the personal signal space is a singleton $Y_i = \{\bar{y}_i\}$. It follows that, for each $i \in I$, the unique possible realization $\bar{y}_i$ does not convey any information about ω and does not alter the agent's belief. Note that, for all $k \in \{1, ..., n\}$, choosing an experiment $(Y, \sigma) \in \Sigma_k$ is equivalent to first choosing a *target* set $I \subseteq N$ of "active" receivers and then an unrestricted experiment for all the agents in I.[10] In this case, the agents in I observe signal realizations $y_I = (y_i)_{i \in I} \in \prod_{i \in I} Y_i$ and update their opinions, while the agents in I^c stick to their prior opinion $\bar{\pi}$. The interpretation here is that the sender has a maximum amount of k *seeds* to persuade all the receivers in N and $k = n$ corresponds to the standard information design problem. The dual case without seeds capacity but with a constant marginal cost of seeding for the sender has an analogous analysis.

The class of *public experiments* $\Sigma^* \subseteq \Sigma$ is going to be relevant in the analysis. Formally, an experiment $(Y, \sigma) \in \Sigma$ is public if, for all $\omega \in \Omega$, the probability measure $\sigma(\cdot|\omega) \in \Delta(Y)$ is perfectly correlated across all the dimensions of Y. In words, an experiment is public if it discloses the same information to every agent, hence inducing the same posterior beliefs.

[9] As clarified below, all the final outcomes solely depend on conditional expectations x. However, all the results would still be true if I had to consider an arbitrary linear function of the posterior beliefs $(p_i)_{i \in N}$.

[10] In this case, the personal signal spaces of the agents in $I^c = N \backslash I$ are singletons.

Indeed, if (Y, σ) is public, then $x(y)$ belongs to the diagonal of X with probability 1. Finally, for all $k \in \{1, .., n\}$, let $\Sigma_k^* \subseteq \Sigma_k$ denote the set of public experiments restricted to sets of maximum k receivers.

Affine aggregation mechanism Let $x \in X$ denote an arbitrary profile of opinions formed after that the receivers have privately observed their signal realizations. An *aggregation mechanism* (or opinion aggregator) is a *continuous* function $T : X \to [0,1]$ that satisfy:

- *Monotonicity*: For all $x, x' \in X$,

$$x' \geq x \implies T(x') \geq T(x)$$

- *Normalization*: For all $z \in [0,1]$.

$$T(z, ..., z) = z.$$

On the one hand, monotonicity has a straightforward interpretation. On the other hand, normalization is formalizing the fact that if the agents share the same opinion, then the aggregation outcome is equal to that opinion. Each aggregation mechanism T represents an unmodeled interaction among agents whose outcome $T(x)$ depends on the opinions that the agent have formed after observing the signals form the sender.

For most of the analysis, I will consider *affine* aggregation mechanisms T, that is, for all $x, x' \in X$ and $\lambda \in [0,1]$,

$$T(\lambda x + (1-\lambda) x') = \lambda T(x) + (1-\lambda) T(x') \qquad \forall x, x' \in X, \forall \lambda \in [0,1].$$

By a simple application of Riesz Theorem, we have that each linear aggregation mechanism T can be represented by a vector of weights $s \in \Delta(N)$, that is,

$$T(x) = s \cdot x = \sum_{i \in N} s_i x_i \qquad \forall x \in X.$$

The interpretation is that each $s_i \in [0,1]$ is the relative weight of receiver i in the aggregation procedure. Even if it is a quite restrictive class, affine aggregation mechanisms naturally arise from standard stylized economic settings describing information aggregation in social networks and markets (see Section 2.2). Finally, in Section 3.3, I will show how the main results can be extended to a more general class of aggregator that also capture robustness concerns of the sender.

The problem of the sender The payoff function of the sender $u : [0,1] \times \Omega \to \mathbb{R}$ depends only on the aggregation outcome $z = T(x) \in [0,1]$ and, in the general case, on the state $\omega \in \Omega$. For all the analysis, I will assume that u is *jointly continuous*.[11] Next, note that each u can be interpreted as the payoff function of a sender facing a single receiver whose final opinion is given by $z = T(x)$. Note that each u (together with aggregator T) induces the *opinions-dependent* payoff function $v : X \times \Omega \to \mathbb{R}$ defined as

$$v(x,\omega) = u(T(x),\omega) \qquad \forall (x,\omega) \in X \times \Omega.$$

This second observation clarifies that I could have considered v as the primitive of a multi-receiver information design problem and then assumed the corresponding properties directly on v. Even if this approach is strictly more general, it completely ignores aggregation mechanisms that are the main ingredients of the applications proposed. Therefore, from a pedagogical point of view, it seems reasonable to mention explicitly the aggregation channel.

For most of the analysis, I am going to focus on *state-independent* payoff functions u, that is, for all $z \in [0,1]$ and $\omega, \omega' \in \Omega$,

$$u(z,\omega) = u(z,\omega').$$

Under this restriction, the sender is purely interested in persuading the agents to obtain a certain outcome from their interaction, independently of the state.

For the rest of the analysis, I am going to make the standard commitment assumption that has characterized the recent literature in information design. In words, the sender is able to commit, before the state has been realized, to privately reveal to the receivers the outcome of a certain experiment $(Y,\sigma) \in \Sigma$. Clearly, the restrictiveness of this assumption depends on the economic context one desires to model. In particular, it depends on whether or not the sender actually observes the state before generating the stochastic message. In both the application I present, it is not unambiguous that the commitment assumption is valid. Therefore, in Section 5, I discuss how to extend the results of the analysis to the case where the sender is not able to commit ex-ante.

Before formally stating the problem of the sender, that will be the main object of analysis, I am going to summarize the timeline of the model:

1. Given the capacity $k \in \{1,...,n\}$, the sender commits to disclose the realizations of an experiment $(Y,\sigma) \in \Sigma_k$;

2. Both the state ω and the signal realizations y are drawn according to π and σ;

[11] In some applications this assumption might be too restrictive. For most of the stated results, assuming that the payoff function u is upper-semi continuous in the aggregation outcome $z \in [0,1]$ and measurable in the state ω would be sufficient.

3. The receivers observe their private realizations y and update their opinions to $x(y)$;

4. The opinions of the agents are aggregated to $T(x(y))$;

5. The sender gets his payoff $u(T(x(y)), \omega)$.

With this, given a prior belief $\pi \in \Delta(\Omega)$, a capacity $k \in \{1, ..., n\}$, an affine aggregator $s \in \Delta(N)$, and a payoff function u, the k-restricted *problem* of the sender is

$$V_n(u, \pi, s) = \sup_{(Y,\sigma)\in\Sigma_k} \left\{ \int_{\Omega\times Y} u(s \cdot x(y), \omega)\, d\sigma(y|\omega)\, d\pi(\omega) \right\}.$$

In words, the sender chooses a feasible experiment to maximize his expected payoff induced by the distribution over states and opinions.

2.2 Examples

I next illustrate the setting just described by introducing the two main examples of the paper.

Example 1: DeGroot learning in networks In this first example, the receivers represent consumers that receive some information about a payoff-relevant uncertain state ω and then start a process of opinion aggregation with each other. I consider the DeGroot model of naive learning in networks (DeGroot 1974, DeMarzo, Vayanos, Zwiebel 2003, Golub and Jackson 2010). Formally, the receivers are the nodes of an *undirected*, *connected*, and *aperiodic* network $G \in \{0, 1\}^{n\times n}$.[12] The interpretation is that the consumers are part of a social network (either a virtual one like Facebook or a real one like the inhabitants of a small village) and $g_{ij} = g_{ji} = 1$ denotes that i and j are linked and can directly influence each other's opinions. For every $i \in N$, let

$$N_i = \{j \in N : g_{ij} = g_{ji} = 1\}$$

denote the set of neighbors of i and define her *degree* as the size of her neighborhood $d_i(G) = |N_i|$. After that the agents have observed their personal signal realization, they update their beliefs to obtain the profile of initial opinions $x^0 \in X$. Next, agents start a process of periodic revision of their own opinions by computing the averages of their neighbors' opinions. Formally,

$$x_i^t = \frac{1}{d_i(G)} \sum_{j\in N_i} x_j^{t-1} \qquad \forall t \in \mathbb{N}.$$

This process is akin to several interpretations. First, one may think that the agents understand that the opinions of their neighborhoods contain valuable information on $\omega \in \Omega$, but they are uncertain about both the full social structure and the signal distributions chosen by the sender

[12] Here, connectedness and aperiodicity of G are essential for reaching consensus and our interpretation. However, the hypothesis of undirected and unweighted network are purely made for expositional simplicity.

for all the other receivers.[13] With this, rather than updating their opinions in in a Bayesian fashion, the agents consider a heuristic method that linearly and uniformly aggregates the opinions observed.[14] Under the second viable interpretation, the aggregation process represents best-response dynamics among agents that are learning how to coordinate on a network. In this case, every agent i minimizes, in each period $t \in \mathbb{N}$, the quadratic loss between her current opinion and the opinions that her neighbors stated in period $t-1$.[15]

Given the assumptions made on G, it is well-known that the sequence of updates $\{x^t\}_{t\in\mathbb{N}} \subseteq X$ converges to a *constant* vector $x^\infty \in X$ whose unique entry is given by

$$T_G\left(x^0\right) = \frac{\sum_{i\in N} x_i^0 d_i\left(G\right)}{\sum_{i\in N} d_i\left(G\right)}.$$

It follows that T_G defines an affine aggregation mechanism with $s_i \propto d_i\left(G\right)$ for all $i \in N$. In this case, the weight s_i defines a measure of the centrality of agent i in the social network G, the so called *eigenvector centrality* of G.

Next, suppose that the sender has one-period *continuous* payoff function $v : X \times \Omega \to \mathbb{R}$ that depends on the entire profile of updates x^t in each period. Also, for all $(z,\omega) \in [0,1] \times \Omega$, define $u\left(z,\omega\right) = v\left(z,...,z,\omega\right)$. Given a sequence of updates $\{x^t\}_{t\in\mathbb{N}}$, the total payoff of the sender is

$$\left(1-\delta\right)\sum_{t=0}^{\infty}\delta^t v\left(x^t,\omega\right),$$

where $\delta \in (0,1)$ is a discount factor. Clearly, this setting is more general than the one described at the beginning of this section. However, it is next showed that when the sender is infinitely patient (i.e., $\delta \to 1$) then the previous criterion is consistent with the general model presented.

Remark 1 *For all priors $\pi \in \Delta\left(\Omega\right)$ and experiments $\left(Y,\sigma\right) \in \Sigma$,*

$$\lim_{\delta\to 1}\left\{\int_{\Omega\times Y}\left[\left(1-\delta\right)\sum_{t=0}^{\infty}\delta^t v\left(x^t\left(y\right),\omega\right)\right]d\sigma\left(y|\omega\right)d\pi\left(\omega\right)\right\} = \int_{\Omega\times Y} u\left(T_G\left(x\left(y\right)\right),\omega\right)d\sigma\left(y|\omega\right)d\pi\left(\omega\right).$$

Therefore, the general setting of the model is rich enough to describe the information design problem of a sender who faces a group of agents aggregating information on a network and who is interested on the long-run outcome of this aggregation. In Section 5, I will discuss how the infinite patience hypothesis might be relaxed. ▲

[13] This assumption is backed by the recent empirical findings in Breza et al (2018).

[14] The superiority of DeGroot updating over Bayesian updating in describing the dynamics of opinions in networks, especially in contexts with weak formal economic institutions, is showed by the recent empirical findings in Chandrasekhar et al (2020).

[15] These kind of dynamics have been considered, among many others, by Morris (2000) and Golub and Jackson (2012).

Example 2: Competitive markets with quadratic payoffs In this example, the receivers are consumers in a competitive economy. There is a unique tradeable good whose objective value is represented by the uncertain state $\omega \in \Omega$. The payoff of each consumer $i \in N$ is

$$U_i (q_i, \omega, r) = q_i (\theta_i \omega - r) - \frac{q_i^2}{2}$$

where $q_i \in \mathbb{R}$ is the quantity of the good acquired by consumer i, $\theta_i \in \mathbb{R}_+$ is the idiosyncratic value of i for the good, and $r \in \mathbb{R}$ is the competitive price. The interpretation is that, conditional on acquiring q_i units of the good, the agent has to suffer a quadratic cost $\frac{q_i^2}{2}$ that may represent transition costs, stocking costs or other forms of frictions in the market. Moreover, the agent is risk neutral obtaining a net unitary gain from the good equal to the difference of her subjective value for the good and the price paid.[16] For the sake of expositional simplicity, I normalize the value types of the agents in the economy such that $\bar{\theta} = \frac{\sum_{i \in N} \theta_i}{n} = 1$. Given risk neutrality, for each posterior belief $p_i \in \Delta (\Omega)$ with corresponding opinion x_i, the problem of consumer i is

$$\max_{q_i \in \mathbb{R}} \left\{ q_i (\theta_i x_i - r) - \frac{q_i^2}{2} \right\}.$$

The induced demand function of i is

$$q_i (x_i, \theta_i, r) = \theta_i x_i - r.$$

Let $\beta \in \mathbb{R}_+$ denote the fixed (perfectly inelastic) *unit* supply. It follows that the competitive price $r(x)$ solves

$$\sum_{i \in N} q_i (x_i, \theta_i, r(x)) = n\beta \implies r(x) = \frac{\sum_{i \in N} \theta_i x_i}{\sum_{i \in N} \theta_i} - \beta$$

With this, if the payoff of the sender depends only on the uncertain value ω and on the competitive price $r(x)$, then it will depend on x linearly as in our general setting with weights $s_i \propto \theta_i$. In this case, the relative importance of each agent i in the aggregation mechanism depends on her willingness to pay θ_i for the good. The general setting of the model is rich enough to describe the information design problem of a sender who tries to influence the final distribution of the market price of a good by disclosing information to the potential consumers. Finally, note that, in the example considered, the equilibrium price $r(x)$ acts as an information aggregation mechanism but does not convey additional information to the agents as in rational expectation equilibrium models.[17] I postpone to Section 5, a more detailed discussion on this topic. ▲

So far I have presented the primitive objects of the model, defined the information design problem of the sender, and showed how this framework can be used to model information disclosure in social networks and competitive markets. In the next sections, I will analyze the optimal value of the sender as well as some properties of optimal experiments.

[16] Even if restrictive, the quadratic payoff structure is a benchmark in the market microstructure literature. For a comprehensive treatment on these topics see Vives (2010).

[17] For similar disclosure models with informative prices see Vives (2010) and Goldstein and Yang (2017).

3 Unrestricted information design

In this section, I am going to focus on the case where the sender has no restriction in targeting the agents with personalized information. Formally, I assume that $k = n$ and let the set of feasible experiments be Σ. The main results of this section will provide general bounds on the optimal value for the sender that do not depend on the affine aggregation mechanism considered. This feature is even more relevant under the further assumption that the sender's payoff is state-independent. In this case, it is possible to characterize the optimal value of the sender as the optimal value of a sender facing a single receiver. The implications of this result are that the optimal value of the sender does not depend on the aggregation mechanism and can be attained by public experiments.

3.1 A revelation principle

I start by stating an ancillary lemma stating that a standard revelation principle holds in the current environment and that the problem of the sender can be represented as a linear program.

Lemma 1 *Fix a prior* $\pi \in \Delta$, *and an affine aggregator* $s \in \Delta(N)$. *The problem of the sender is equivalent to*

$$V_n(u, \pi, s) = \max_{\gamma \in \Delta(\Omega \times X)} \left\{ \int_{\Omega \times X} u(s \cdot x, \omega)\, d\gamma(\omega, x) \right\}$$

subject to

$$\mathbb{E}_\gamma \left[h_0(\omega) + \sum_{i \in N} h_i(x_i)(x_i - \omega) \right] = \mathbb{E}_\pi [h_0(\omega)] \tag{2}$$

for all $(h_i)_{i \in N_0} \in C([0,1])^{N_0}$.

The interpretation is that the sender can equivalently choose the joint distribution of states and opinions for the agents, under some additional constraints. On the one hand, the marginal distribution over states has to be consistent with the prior π, that is, the standard Bayes plausibility condition has to hold. On the other hand, the distribution over the opinion of each receiver i needs to be a mean-preserving contraction of the prior π (see for example Dworczak and Martini 2019, and Kolotilin 2018). Both these constraints are considered by the expression in eq. (2).

Before analyzing the general bounds on the optimal value of the sender, it will be useful to characterize the *restricted* value of the sender when he can only disclose public information, denoted by $V_1(u, \pi)$. Toward this characterization, define the indirect *posterior-dependent* payoff function $\hat{u} : \Delta(\Omega) \to \mathbb{R}$ as

$$\hat{u}(p) = \mathbb{E}_p [u(\mathbb{E}_p[\omega], \omega)] \qquad \forall p \in \Delta(\Omega).$$

The payoff function v describes the payoff of the sender whenever the *common* posterior belief induced by the public signal is $p \in \Delta(\Omega)$. Note that this also coincides with the indirect utility of the sender when he faces a single receiver. In this case, it is well known (see Kamenica and Gentzkow 2011) that the optimal value for the sender is given by the concavification of the indirect payoff function $\hat{u}$ evaluated at the prior π, denoted by $cav(\hat{u})(\pi)$. Given normalization, one can observe that this expression characterize the restricted value of the sender in the general model with n receivers and for any aggregation mechanism T (affine or not).

Lemma 2 *Fix a prior $\pi \in \Delta$, and an opinion aggregator T. The problem of the sender when restricted to public signals is*

$$V_1(u,\pi) = \max_{\nu \in \Delta(\Omega \times [0,1])} \left\{ \int_{\Omega \times [0,1]} u(z,\omega)\, d\nu(\omega,z) \right\}$$

subject to

$$\mathbb{E}_\nu \left[h_0(\omega) + \bar{h}(z)(z-\omega) \right] = \mathbb{E}_\pi \left[h_0(\omega) \right]$$

for all $\left(h_0,\bar{h}\right) \in C([0,1])^2$. Moreover, it holds $V_1(u,\pi) = cav(\hat{u})(\pi)$.

3.2 Irrelevance results

Next, I exploit the previous two ancillary lemmas to obtain aggregation-independent bounds on the optimal value for the sender. First, these bounds will be used to characterize the optimal value of the sender in the state-independent case. Second, given that the bounds do not depend on the aggregation mechanism, it follows that they provide a robust interval on the set of payoffs that can be reached by the sender. Indeed, by combining the last two observations, it is clear that the bounds will be tight if all the possible payoff functions of the sender (including state-independent ones) are considered.

Proposition 1 *Fix a prior $\pi \in \Delta$. For all sizes $n \in \mathbb{N}$ and affine aggregators $s \in \Delta(N)$,*

$$cav(\hat{u})(\pi) \leq V_n(u,\pi,s) \leq \max\left\{ \mathbb{E}_\nu[u] : \nu \in \Delta(\Omega \times [0,1]), marg_\Omega \nu = \pi, marg_{[0,1]} \nu \succsim_{cv} \pi \right\}$$

On the one hand, the lower bound identified in Proposition 1 is immediate given Lemma 2. Indeed, the sender can always commit himself to disclose a public signal. On the other hand, the upper bound is obtained by considering a relaxed version of the single-receiver persuasion problem. In particular, the constraints on the joint distribution $\nu \in \Delta(\Omega \times [0,1])$ require Bayes plausibility (i.e. consistency with the common prior π) and the marginal distribution over the single receiver's opinion dominates the prior in the concave order. The proof of Proposition 1 shows that these are *necessary* conditions that each feasible distribution $\nu \in \Delta(\Omega \times [0,1])$ over states and aggregated opinion induced by an experiment must satisfy. However, these conditions are not sufficient in general as they are ignoring the martingale constraint

$$\mathbb{E}_\nu[\omega | x_i] = x_i$$

that must hold for every $i \in N$. In particular, these martingale constraint will imply further restrictions on the induced joint distribution of states and aggregated opinion that are ignored in the upper bound

$$\max \left\{ \mathbb{E}_{\nu}[u] : \nu \in \Delta(\Omega \times [0,1]), \text{marg}_{\Omega}\nu = \pi, \text{marg}_{[0,1]}\nu \succsim_{cv} \pi \right\}.$$

However, if the payoff of the sender is state-independent, then the further constraint on the joint distributions between the state and the aggregate opinion have no bite. Indeed, in this case, only the induced marginal distribution on the aggregate opinion is payoff-relevant for the sender. This is formalized by the next irrelevance result.

Corollary 1 (Irrelevance result) *Fix a prior* $\pi \in \Delta(\Omega)$ *and a state-independent payoff function* $u : [0,1] \to \mathbb{R}$. *For all* $n \in \mathbb{N}$, *and affine opinion aggregator* $s \in \Delta(\{1, ..., n\})$,

$$V_n(u, \pi, s) = V_1(u, \pi) = cav(\hat{u})(\pi)$$

Moreover, there exists a public experiment $(Y^*, \sigma^*) \in \Sigma^P$ *that solves the sender's problem for all* n *and* s.

In words, under the state-independence hypothesis, the two bounds previously identified collapse to the same value $cav(\hat{u})(\pi)$ characterizing the optimal value for the sender. Moreover, this value is equivalent to the one obtained by a sender facing a single-receiver and it does not depend on the aggregation mechanism s. It follows that, without loss of optimality, the sender can disclose the same information to every agent avoiding to design personalized private signals for each receiver.

Finally, the optimal value of the sender can be expressed as

$$cav(\hat{u})(\pi) = \max \{\mathbb{E}_{\mu}[u] : \mu \in \Delta([0,1]), \mu \succsim_{cv} \pi\}. \tag{3}$$

The problem defined in the right-hand side of eq. (3) has been extensively analyzed in the recent literature on simple persuasion (see for example Gentzkow and Kamenica 2016, Kolotilin 2018, Dworczak and Martini 2019). Therefore, the irrelevance result permits to solve an apparently more involved multiple-receivers information design problem with known tools in the literature This observation will be also exploited in Section 4.3 to derive a comparative statics result in the restricted case.

Finally, note that the irrelevance result would still hold even if the sender is uncertain about the affine aggregator $s \in \Delta(N)$ characterizing the aggregation mechanism. Indeed, suppose the sender has a belief $\xi \in \Delta(\Delta(N))$ about the aggregation mechanisms. It follows that

$$\sup_{(Y,\sigma)\in\Sigma_k} \left\{ \int_{\Omega \times Y} \left(\int_{\Delta(N)} u(s \cdot x(y), \omega) \, d\xi(s) \right) d\sigma(y|\omega) \, d\pi(\omega) \right\} \leq \int_{\Delta(N)} V_n(u, \pi, s) \, d\xi(s) = cav(\hat{u})(\pi)$$

showing that $cav(\hat{u})(\pi)$ is an upper bound for the value of the sender under uncertainty. Given that an optimal public experiment $(Y^*, \sigma^*) \in \Sigma^P$ attains $cav(\hat{u})(\pi)$ also under uncertainty, it follows again that $cav(\hat{u})(\pi)$ characterizes the optimal value of the sender under ξ. In the next section, I show that a similar results hold even if the sender considers more general aggregation mechanism and he is uncertain averse.

3.3 The robustness of the irrelevance result

In this section, I am going to check the robustness of the irrelevance result with respect to more general aggregation mechanisms. In particular, I will relax the assumption of affinity by considering *robust opinion aggregators* introduced in Cerreia-Vioglio et al (2020).

Robust opinion aggregators and uncertainty aversion An aggregation mechanism T is called *robust* if it also satisfies *translation invariance*, that is,

$$T\left(x + (z, ..., z)'\right) = T(x) + z$$

for all $x \in X$ and $z \in [0,1]$ such that $x + (z, ..., z) \in X$. In other words, robust opinion aggregators relaxes affinity to translation invariance. The intuition is that now the relative importance of the agents in the aggregation mechanism is not fixed but it may depend on the particular profile of opinions x aggregated. Standard examples of robust aggregators are the median opinion, the minimum opinion, the maximum opinion, or in general all the order statistics of the profile of opinions.

Under this more general form of aggregators, the problem of the sender becomes

$$V_n(u, \pi, T) = \sup_{(Y,\sigma)\in\Sigma} \left\{ \int_{\Omega\times Y} u\left(T\left(x\left(y\right)\right), \omega\right) d\sigma\left(y|\omega\right) d\pi\left(\omega\right) \right\}$$

for all priors π and robust aggregators T. In what follows, I will focus on the state-independent case for the payoff function of the sender as to obtain sharper results that characterize the optimal value of the sender under robust aggregation. However, more general version of the bounds considered in proposition 1 can be obtained even if the payoff of the sender depends on the state.

The irrelevance result does not hold in general for robust aggregators. Indeed, it is enough to consider $T(x) = \max_{i\in N} x_i$ and a *strictly increasing* and state-independent payoff function $u : [0,1] \to \mathbb{R}$ for the sender. Suppose by contradiction that a public experiment σ^* inducing the common-opinion distribution $\mu^* \in \Delta([0,1])$ solves the sender's problem. Next, consider consider designing n iid copies of the candidate public experiment σ^* to be privately disclosed to each individual. Also, note that the obtained experiments are private and conditionally independent given the common state. With this, the induced marginal distribution over each agent's opinion will be μ^*. Moreover, conditional on the state, the aggregate opinion $z = T(x)$

will be distributed as the maximum of n independent draws from the same distribution. In general, the induced marginal distribution for the aggregate opinion z is going to strictly first-order stochastically dominate μ^* hence achieving an higher expected payoff for the sender and contradicting the optimality of σ^*.

Despite the previous counterexample, Lemma 2 suggests that $cav\,(\hat{u})\,(\pi)$ is still a valid lower bound for the unrestricted optimal value of the sender under any robust aggregator T, that is,

$$V_n\,(u, \pi, T) \geq cav\,(\hat{u})\,(\pi)\,.$$

Moreover, next proposition shows how under some further assumption on both u and T, the irrelevance result still holds. Recall that a robust aggregator T is concave (convex) if

$$T\,(\lambda x + (1-\lambda)\,x') \geq (\leq)\,\lambda T\,(x) + (1-\lambda)\,T\,(x')\,,$$

for all $x, x' \in X$ and $\lambda \in [0,1]$.

Proposition 2 *Fix a prior $\pi \in \Delta\,(\Omega)$ and a state-independent payoff function $u : [0,1] \to \mathbb{R}$. For all $n \in \mathbb{N}$,*

i) *If u is nondecreasing and T is concave, then $V_n\,(u, \pi, T) = cav\,(\hat{u})\,(\pi)$;*

ii) *If u is nonincreasing and T is convex, then $V_n\,(u, \pi, T) = cav\,(\hat{u})\,(\pi)$.*

Moreover, in both cases there exists a public experiment $(Y^*, \sigma^*) \in \Sigma^P$ that solves the Sender's problem for all n and T satisfying the stated assumptions.

It is useful to compare the previous proposition with the counterexample described above. In the latter, the sender had a strictly increasing utility and faced a strictly convex robust aggregator (i.e., the maximum opinion). Therefore, the sender benefits form higher values for the aggregated opinion and the receivers' aggregation process features an implicit attraction for higher opinions as well. It follows that the sender can exploit this concordance by avoiding the disclosure of public information. Conversely, as suggested by Proposition 2, whenever these two effects go in opposite directions (e.g., increasing u and concave T), then the optimal value of the sender cannot be higher than the one he would reach if restricted to public experiment. Also, note how this is capturing a robustness feature of public experiments both from the perspective of the sender and of an external analyst. Indeed, even if the sender suspects that the true aggregation mechanism characterizing the interaction among receivers in non-affine, it is still without loss of optimality for him to disclose public information as long as he considers adversarial mechanisms (depending on the monotonicity of his payoff function).

The idea of robustness of public signals can be even generalized to the entire class of robust opinion aggregators, under the assumption that the sender is averse to uncertainty over aggregation mechanisms. In particular, suppose that the sender considers a set $\mathcal{T}$ of potential opinion aggregators that may describe the interaction among agents. Moreover, suppose that the sender is uncertain averse and chooses the experiment that maximizes his worst-case expected payoff across all the conceivable aggregators in $\mathcal{T}$. Formally, the problem of the uncertain averse sender is

$$V_n(u, \pi, \mathcal{T}) = \sup_{(Y,\sigma)\in\Sigma} \inf_{T\in\mathcal{T}} \left\{ \int_{\Omega\times Y} u\left(T\left(x\left(y\right)\right)\right) d\sigma\left(y|\omega\right) d\pi\left(\omega\right) \right\}.$$

Next proposition states that, as long as the uncertainty set $\mathcal{T}$ includes at least one affine aggregator, it is without loss of optimality for the sender to restrict to public experiments and his optimal value does not depend on the set $\mathcal{T}$.

Proposition 3 *Fix a prior $\pi \in \Delta(\Omega)$, a state-independent payoff function $u : [0, 1] \to \mathbb{R}$, and a set $\mathcal{T}$ of robust opinion aggregators including at least an affine aggregator. It holds that*

$$V_n(u, \pi, \mathcal{T}) = cav(\hat{u})(\pi)$$

Moreover, there exists a public experiment $(Y^*, \sigma^*) \in \Sigma^P$ that solves the Sender's problem for all n and $\mathcal{T}$ satisfying the stated assumption.

The intuition behind the previous result is that, thanks to Lemma 2, the sender cannot obtain an expected payoff lower than $cav(\hat{u})(\pi)$. Moreover, the adversarial nature can always achieve such value by picking an affine aggregator s in $\mathcal{T}$. With this, the saddle value of the problem is characterized by $cav(\hat{u})(\pi)$ and this can be always attained through public experiments.

4 Restricted information design

In this section, I consider the general case with $k \in \{1, ..., n\}$. As anticipated, in this case the problem of the sender can be decomposed in two steps: first he has to choose a target set I with maximum size k, and then he commits to an experiment for all the agents in I. Therefore, it is necessary to add a further step 0 to the timeline illustrated in Section 2.1 where the sender targets a subset of receivers. Solving this two-step problem by backward induction, it is shown that the sender optimally target the set of agents maximizing the total sum of weights among the ones with size k. Then, I rely on this result to derive a comparative statics result that links the aggregation mechanism to the informativeness of the optimal experiment under some additional assumption on the payoff function of the sender. Finally, I extend the targeting result for a subclass of robust aggregators.

4.1 Targeting and persuasion

As for most part of the analysis, I focus on affine aggregators $T(x) = s \cdot x$ where $s \in \Delta(N)$ and state-independent payoff function for the sender $u : [0, 1] \to \mathbb{R}$. The main difference with respect to the analysis in Section 3 is that now only a subset I of up to $k \in \{1, ..., n\}$ receivers update their beliefs given the information disclosed by the sender. The remaining agents in I^c receive no information and their common opinion is equal to the prior mean:

$$x_i = \bar{\pi} \qquad \forall i \in I^c.$$

With this, given a restricted profile $(x_i)_{i \in I} \in [0,1]^I$ of opinions for the agents in I and aggregation weights s, the outcome of the aggregation process is

$$z = (1 - s_I)\,\bar{\pi} + \sum_{i \in I} s_i x_i$$

where $s_I = \sum_{i \in I} s_i$. Therefore, conditional on targeting the group of receivers $I \subseteq N$, the information design problem of the sender is

$$V_I(u, \pi, s) = \sup_{(Y,\sigma) \in \Sigma_k^I} \left\{ \int_{\Omega \times Y} u\left((1 - s_I)\,\bar{\pi} + \sum_{i \in I} s_i x_i(y_i) \right) d\sigma(y|\omega)\, d\pi(\omega) \right\} \tag{4}$$

where $\Sigma_k^I \subseteq \Sigma_k$ is the set of feasible experiments for the agents in I.

I am going to solve the sender's problem in two steps:

1. For each target set $I \subseteq N$, solve for the optimal experiment and compute the optimal value $V_I(u, \pi, s)$;

2. Solve for the optimal target set I^* that maximizes $V_I(u, \pi, s)$.

For step (1) I show that the irrelevance result still holds when the sender is restricted to a subset of agents I. Hence, the sender can without loss of optimality design a public experiment for agents in I where this experiment does not depend on the aggregation structure restricted on I. Next, for step (2) I show that the optimal value $V_I(u, \pi, s)$ is nondecreasing in the total influence of the agents in I. The following proposition formalizes these results. Before stating it, I will need another piece of definition: given a payoff function $u : [0,1] \to \mathbb{R}$, for all $I \subseteq N$ and $s \in \Delta(N)$, define the *distorted payoff function* $u_I(\cdot; s) : [0,1] \to \mathbb{R}$ as

$$u_I(z; s) = u((1 - s_I)\,\bar{\pi} + s_I z) \qquad \forall z \in [0,1].$$

Proposition 4 *Fix a prior $\pi \in \Delta(\Omega)$. For all subsets $I \subseteq N$ and affine aggregators $s \in \Delta(N)$,*

$$V_I(u, \pi, s) = \max_{\mu \succsim_{cv} \pi} \left\{ \int_{[0,1]} u_I(z; s)\, d\mu(z) \right\} = cav(\hat{u}_I(\cdot; s))(\pi),$$

and there exists a public experiment $(Y^*, \sigma^*) \in \Sigma_k^P$ that solves the sender's problem conditional on I. Moreover, for all $I, I' \subseteq N$ and $s, s' \in \Delta(N)$

$$s_I \geq s'_{I'} \implies V_I(u, \pi, s) \geq V_{I'}(u, \pi, s').$$

The first part of Proposition 4 states that the information design problem given an arbitrary target set I is equivalent to a single-agent persuasion problem with a degree of *stickiness* $(1 - s_I) \in [0, 1]$ to the prior mean $\bar{\pi}$. Moreover, the higher the total weight on agents in I and the higher the optimal value of the sender. In particular, for the second part of the statement, it is sufficient to show that all the distributions of aggregated opinion that are feasible given $s'_{I'}$ remain feasible when the total weight of the target set is higher $s_I \geq s'_{I'}$.

Next, relying on Proposition 4, I characterize the solution of the targeting problem of the sender. Toward this characterization, for all $k \in \{1, ..., n\}$, denote with $I(k, s)$ (one of) the subset of k agents with the highest total weight according to $s \in \Delta(N)$, that is,

$$I(k, s) \in \arg \max_{I \subseteq N : |I| \leq k} s_I.$$

Corollary 2 *For all $s \in \Delta(N)$, $I(k, s)$ solves the targeting problem of the sender.*

The previous results suggest an algorithmic procedure for solving the Sender's problem:

1. Order the receivers with respect to their weights $(s_i)_{i \in N}$;

2. Target the k most prominent agents in the society according to s;

3. Solve a single-agent simple persuasion problems with prior π and distorted utility $u_{I(k,s)}(\cdot; s)$.

Note that in step (2) of the suggested algorithm, the sender only needs to know the identity of the k most influential agents to solve for the optimal target set. Moreover, in step (3), the sender only needs to know the total weight of the k most influential receivers in order to solve for the optimal information policy. Therefore, under our social learning interpretation, the result is robust with respect to partial uncertainty of the sender about the network G. In general, this simple result provides a *foundation* for targeting information policies based on centrality measures in networks: a rational sender would choose to disclose the same information to the most prominent agents in the social network/market. Next, I leverage on the main targeting results of this section to derive some structural properties of the optimal information policy.

4.2 Information design in large networks and markets

The characterization of Proposition 4 can be used to easily solve for optimal experiments under large populations of agents such that the maximum weight among the agents in vanishing at infinity. Only for this section, I will consider a *sequence* of populations of receivers $\{N(n)\}_{n \in \mathbb{N}}$ increasing in size $n \in \mathbb{N}$. Following the two examples presented in Section 2.2, we interpret the results of this section as holding for large networks in which every agent is vanishingly influential and for large markets in which the willingness to pay of the consumers is bounded.

Formally, let $\{s(n)\}_{n\in\mathbb{N}}$ denote the sequence of aggregation mechanisms where, for every $n \in \mathbb{N}$, $s(n) \in \Delta(N(n))$. In particular, suppose that the sequence $\{s(n)\}_{n\in\mathbb{N}}$ is such that

$$\lim_n \left(\max_{i\leq n} s_i(n) \right) = 0. \tag{5}$$

The interpretation is that the relative importance of every agent in the aggregation mechanism is 0 in the limit. Under the social learning interpretation of the model, this condition coincides with the notion of *wisdom* introduced in Golub and Jackson (2010). In their framework, this condition is necessary and sufficient to have that the long-run consensus converges in probability to the true uncertain state. In my model, the original information of the receivers in endogenous and, in the restricted version of the problem, it depends on the aggregator $s(n)$. Yet, as formalized next, the wisdom condition is sufficient to characterize the optimal information structure chosen by the sender as $n \to \infty$.

As before, suppose that, for each $n \in \mathbb{N}$, the sender is constrained to target only k receivers. Given Corollary 2, for each $n \in \mathbb{N}$, the sender picks the k most prominent receivers $I(k, s(n))$. Combining this observation with a standard result in information design we have the following result.

Proposition 5 *Fix a prior $\pi \in \Delta([0,1])$ and a continuously differentiable payoff function $u : [0,1] \to \mathbb{R}$. It holds:*

i) *If u is strictly convex on a neighborhood of $\bar{\pi}$, then there exists $\bar{n} \in \mathbb{N}$ such that for all $n > \bar{n}$ full disclosure is optimal;*

ii) *If u is strictly concave on a neighborhood of $\bar{\pi}$, then there exists $\bar{n} \in \mathbb{N}$ such that for all $n > \bar{n}$ no disclosure is optimal;*

iii) *If u is affine on a neighborhood of $\bar{\pi}$, then there exists $\bar{n} \in \mathbb{N}$ such that for all $n > \bar{n}$ the sender is indifferent over all the information policies.*

Recall that, for each $n \in \mathbb{N}$, the distorted utility of the sender at the optimal targeting set $I = I(k, s(n))$ is $u((1 - s_I(n))\bar{\pi} + s_I(n) z)$. Therefore, as $s_I(n)$ converges to 0, the distorted utility of the Sender has the same curvature of the restriction of u to a neighborhood of $\bar{\pi}$. Therefore, by a standard result in information design (see for example Kamenica and Getnzkow 2011 or Dworczak and Martini 2019), either one between full disclosure and no disclosure is optimal, or the sender is indifferent over all the information structures. The interpretation of the result is that, eventually, either few "elected" receivers are perfectly informed or nobody is, depending on the prior belief of the society.

4.3 Aggregation structure and informativeness

In the previous section, I have showed how to characterize the optimal information structure for large populations. However, the representation obtained in Proposition 4 is flexible enough to explore the structure of the optimal experiment even for "small" populations. In particular, in this section, I link the informativeness of the optimal information policy with the underlying social structure for a relevant class of sender's payoffs and prior beliefs.

Fix an *absolutely continuous* prior $\pi \in \Delta(\Omega)$ (e.g., uniform in $[0,1]$), let the payoff function u be state-independent, nondecreasing, and let $\bar{z} \in (0,1)$ be such that u is convex on $[0,\bar{z}]$ and affine on $[\bar{z},1]$. For example, for some continuos, increasing and convex $\varphi : [0,1] \to \mathbb{R}$ and $\bar{z} \in (0,1)$, let

$$u(z) = \begin{cases} \varphi(z) & if \quad z \in [0,\bar{z}] \\ \varphi(\bar{z}) & if \quad z \in (\bar{z},1] \end{cases} \qquad \forall z \in [0,1]$$

The interpretation is that the sender can only gain from increasing the aggregated opinion of the agents up to $\bar{z} \in (0,1)$ (e.g., $\bar{z}$ is the threshold above which a consumer buys a good with certainty).

Under this restriction on the payoff function of the sender, in the single-receiver information design problem it is without loss of optimality to restrict on *upper censorship* information policies (see for example, Dworczak and Martini 2019 and Kolotilin et al 2019). These class of experiments are characterized by a unique threshold state $\omega^* \in \Omega$ such that the sender reveals the state every time that ω is below ω^* and pool all the states above ω^* with a unique signal realization. Dworczak and Martini (2019) in particular show that the optimal threshold satisfies

$$\mathbb{E}_\pi [\omega | \omega \geq \omega^*] = \bar{z}. \tag{6}$$

In this case, we can compare the informativeness of the optimal information structures as s varies just by comparing the optimal censorship thresholds ω^*: the higher the threshold ω^* and the more Blackwell-informative is the corresponding experiment.

Conditional on a target set I, the Sender solves a pure persuasion problem with utility $u((1-s_I)x_\pi + s_I z)$. This gives a characterization of the optimal threshold $\omega^*(s,I)$ for each target set I.

Proposition 6 *For all $I \subseteq N$ and $s \in \Delta(N)$, it is without loss of optimality for the sender to restrict to upper censorship. Moreover, an optimal censorship threshold is*

$$\omega^*(s,I) = \begin{cases} 0 & if \quad \bar{\pi} \geq \bar{z} \\ \min\{\omega' \in [0,1] : (1-s_I)\bar{\pi} + s_I \mathbb{E}_\pi[\omega|\omega \geq \omega'] = \bar{z}\} & if \quad \bar{\pi} < \bar{z} \leq (1-s_I)\bar{\pi} + s_I \\ 1 & if \quad \bar{z} \geq (1-s_I)\bar{\pi} + s_I \end{cases} \tag{7}$$

The result follows from the observation that the distorted utility $u((1-s_I)x_\pi + s_I z)$ is convex below

$$\max\left\{\min\left\{\frac{\bar{z}-(1-s_I)\bar{\pi}}{s_I},1\right\},0\right\}$$

and affine otherwise. Hence, one can apply the condition in eq. (6) to the distorted utility and obtain the threshold policy $\omega^*(s,I)$ defined in eq. (7). Finally, as an immediate corollary of Proposition 6, I obtain the desired comparative statics result.

Corollary 3 *For all $I, I' \subseteq N$ and $s, s' \in \Delta(N)$, $s_I \geq s'_{I'}$ implies that $\omega^*(s,I) \leq \omega^*(s',I')$, that is, there exist two public experiments (Y,σ) and (Y',σ') that are optimal respectively for (s,I) and (s',I) and such that (Y,σ) is Blackwell-dominated by (Y',σ').*

In words, the sender will react with a less informative policy to an increase in the total influence of the k most prominent receivers that are targeted. The intuition of the result goes as follows: if $\bar{\pi} \geq \bar{z}$, then s_I has no effect on the policy described in (6), if $\bar{\pi} < \bar{z} \leq (1-s_I)\bar{\pi}+s_I$, then an increase in s_I still leaves $\bar{z}$ in the intermediate region and implies that the threshold $\omega^*(s,I)$ has to decrease, finally if $\bar{z} \geq (1-s_I)\bar{\pi}+s_I$, then an increase in s_I can only reduce the informativeness of the policy as it was starting from full-disclosure. In particular, note that, in the social learning example, every time that more individuals become neighbors with the k most connected agents, the total influence of the latter will increase inducing a reduction in the informativeness of the signal that they are going to observe.

4.4 Targeting and robust aggregation

I close this part of the paper by generalizing the targeting result for a subclass of robust opinion aggregators. For simplicity, I will assume that the payoff function of the sender u is state-independent and nondecreasing. The additional properties that I am going to require on the robust opinion aggregators considered are *concavity* and *constant affinity*, that is,

$$T\left(\lambda x + (1-\lambda)(z,...,z)'\right) = \lambda T(x) + (1-\lambda) z \qquad \forall x \in X, \forall \lambda, z \in [0,1]$$

In words, the aggregation process of the receivers' opinions capture an attraction for lower opinions (i.e., opposite to the preference of the sender) and it is invariant with respect to constant shifts and rescaling of the opinions. In this case, the aggregator admits the following representation

$$T(x) = \min_{s \in S}(s \cdot x) \qquad \forall x \in X$$

for some convex and closed $S \subseteq \Delta(N)$ (see Ghirardato et al, 2004). In words, the aggregation mechanism returns the lower outcome consistent with a set of different affine opinion aggregators S. Clearly, whenever $S = \{s\}$ is a singleton, the analysis reduces to the one in the previous sections.

Given a target set $I \subseteq N$, the problem of the sender reads

$$V_I(u,\pi,S) = \sup_{(Y,\sigma)\in\Sigma_k^I} \left\{ \int_{\Omega\times Y} u\left(\min_{s\in S}\left[(1-s_I)\,\bar{\pi} + \sum_{i\in I} s_i x_i(y_i) \right] \right) d\sigma(y|\omega)\, d\pi(\omega) \right\}. \qquad (8)$$

Note that considering a convex and constant affine aggregator T is equivalent to consider an uncertain averse sender. Indeed, the problem stated in eq. (8) captures the idea that the sender plays a game against a malevolent nature that chooses the worst (from the point of view of the sender) affine aggregator *after* that the receivers have observed the realizations of the chosen experiments. Formally, it is easy to show that the sender's problem is equivalent to

$$V_I(u,\pi,S) = \sup_{(Y,\sigma)\in\Sigma_k^I} \inf_{(s(y))_{y\in Y}\in\Delta(N)^Y} \left\{ \int_{\Omega\times Y} u\left((1-s_I)\,\bar{\pi} + \sum_{i\in I} s_i x_i(y_i) \right) d\sigma(y|\omega)\, d\pi(\omega) \right\}. \qquad (9)$$

With this, the particular case considered in this section is flexible enough to capture robustness concerns of the sender.

Next, define the set function $\underline{S} : 2^N \to [0,1]$ as

$$\underline{S}(I) = \min_{s\in S}\left(\sum_{j\in I} s_j \right) \qquad \forall I \subseteq N.$$

In words, $\underline{S}(I)$ is a "pessimistic" measure of the total importance of the subgroup I (centrality, under the network interpretation). As formalized in the next proposition, the targeting result of Section 4.1 extends to the aggregators considered in this section when we replace the linear measure induced by a single s with the (non-linear) measure $\underline{S}$. Define,$u_I(\cdot;s) : [0,1] \to \mathbb{R}$ as

$$u_I(z;S) = u\left(\left(1-\underline{S}(I)\right)\bar{\pi} + \underline{S}(I)\, z \right) \qquad \forall z \in [0,1].$$

Proposition 7 *Fix a prior $\pi \in \Delta(\Omega)$. For all subsets $I \subseteq N$ and set of affine aggregators $S \subseteq \Delta(N)$,*

$$V_I(u,\pi,S) = \max_{\mu\succsim_{cv}\pi} \left\{ \int_{[0,1]} u_I(z;S)\, d\mu(z) \right\} = cav\left(\hat{u}_I(\cdot;S)\right)(\pi)$$

and the optimal target set solves

$$I^*(k,S) \in \arg\max_{I\subseteq N:|I|\leq k} \underline{S}(I).$$

In this case, the sender picks the group of k agents that maximize the nonlinear measure of influence $\underline{S}$. A completely symmetric result holds whenever u is nonincreasing and T is convex.

5 Conclusion and discussion

In this paper, I have proposed a model that combines information design and aggregation to study how optimal information targeting is influenced by an aggregation process such as social networks or a competitive markets. I have first focused on unrestricted information design and reduced-form linear opinion aggregation models to consider the simplest setting possible. Under these assumptions, I have derived an irrelevance result showing that the Sender can restrict without loss of optimality to public information policies that do not depend on the structure of the aggregation. The irrelevance result is robust to more general aggregators and to the uncertainty of the sender about the aggregation mechanism. Next, I have restricted the sender to target only a subgroup of the receivers with his disclosure policy. In this case, the Sender optimally targets the k most prominent agents still disclosing a public experiment to them.

The findings of this paper represent an initial step toward a theory describing information design and aggregation but several important questions remain unanswered. Next, I discuss the most relevant extensions of the baseline model.

Finitely patient sender In the social learning interpretation of the model, I have assumed that the sender is infinitely patient. However, under a separability assumption on the one-period state-independent payoff function $v : X \to \mathbb{R}$, the irrelevance result still holds for all discount factor $\delta \in [0, 1]$. Formally, define

$$v(x) = \sum_{i \in N} \alpha_i u(x_i) \qquad \forall x \in X$$

for some $\alpha = (\alpha_i)_{i \in N} \in \Delta(N)$ representing the relative weight of i in the payoff of the sender. It follows that the for every initial opinion $x \in X$, the total payoff of the sender is

$$(1 - \delta) \sum_{t=0}^{\infty} \delta^t \left(\sum_{i \in N} \alpha_i u \left(g_i^{(t)} \cdot x \right) \right)$$

where, for every $(i, t) \in N \times \mathbb{N}_0$, $g_i^{(t)} \in \Delta(N)$ is the set of weights representing the updating rule for i at period t induced by the DeGroot process described in Section 2.2. Finally, note that the induced information design problem is equivalent to a static problem where the sender is uncertain about the affine aggregation mechanism and places weight $\xi(i, t) = \delta^t \alpha_i$ to each $g_i^{(t)} \in \Delta(N)$. As argued at the end of Section 3.2, in this case the irrelevance result still holds. The analysis for more general one-period payoff functions v is left for future research.

Bayesian updating and rational expectation equilibrium In the social learning interpretation of the model, I have assumed that the receivers naively update their beliefs from what they observe from their neighbors following the DeGroot model. Conversely, suppose that the receivers update their beliefs in a full Bayesian fashion conditional on the (truthful) stated opinions of their neighbors. In this case, as showed by Muller-Frank (2013) the agents will eventually agree and state a common opinion obtained as the conditional expectation with respect to a common (but not necessarily efficient) information structure.[18] Similarly to my

[18] In the sense that the common information structure does not need to synthesize the information contained in all the private signals of the receivers.

model, if the sender anticipates this aggregation process, he can always disclose to the receivers a public signal that is equivalent to the common information structure induced by the periodic communication and (Bayesian) updating. With this, given that all the opinions are common knowledge, there will not be any updating process and the induced distribution over the consensus will be the same. It follows that, without loss of optimality the sender can restrict to public signals, confirming the validity of the irrelevance result. Also, note that in this case no linearity restriction are necessary as the functional form of the aggregator is endogenous and depends on the experiment designed by the sender.

Similarly, in the market interpretation of the model, one can consider rational expectation equilibria (REE) rather than competitive equilibria. In this case, the opinion of each receiver i would satisfy

$$\mathbb{E}_{\pi,\sigma}\left[\omega|x_i, r\left(x\right)\right] = x_i$$

that is, the conditional expectation is computed also with respect to the endogenous variable $r\left(x\right)$. In general, the agents will not share the same opinion, even conditional on $r\left(x\right)$. However, as showed by McKelvey and Page (1986) and Nielsen et al (1990), if the map $x \mapsto r\left(x\right)$ is *stochastically monotone*, then the properties of the REE imply that the agents share the same opinion and that the latter is obtained as the conditional expectation with respect to a common information structure. With this, as argued above, the sender can replicate any such information structure by disclosing a public information in the first place, hence the irrelevance result would still hold.

The commitment assumption In the entire analysis, I have assumed that the sender has the ability to commit himself to generate a given experiment *before* observing the state. Relaxing this hypothesis implies that an additional *incentive compatibility* constraint for the sender. Formally, given a state-independent payoff function u, the distribution $\gamma \in \Delta\left(\Omega \times X\right)$ over states and opinions chosen by the sender has to satisfy

$$\mathbb{E}_{\gamma}\left[u\left(T\left(x\right)\right)|\omega\right] = \mathbb{E}_{\gamma}\left[u\left(T\left(x\right)\right)|\omega'\right] \qquad \forall\omega,\omega' \in \Omega,$$

in addition to the usual constraints defined in Lemma 1. One can note that the proofs of Proposition 1 and Corollary 1 may be adapted in a straightforward way as to take the new constraint into account. In particular, in the state-independent case, the irrelevance result holds and the optimal value of the sender coincides with the *quasi-concave envelope* of $\hat{u}$ evaluated at the prior π. The quasi-concave envelope has been recently introduced by Lipnowski and Ravid (2020) in disclosure games and characterizes the optimal value of a sender in single-receiver cheap talk, exactly as the concave envelope characterizes the optimal value of a sender in single-receiver Bayesian persuasion.

Heterogenous priors The common prior assumption played a key role in the results of the analysis, but it seems natural to ask how these findings would change in a setting with heterogeneous prior beliefs. Suppose for a moment that the uncertainty state is binary $\Omega = \{0,1\}$ and the agents have heterogeneous prior beliefs $(\pi_i)_{i\in N_0} \in (0,1)^{N_0}$.[19] In this case, as showed by Alonso and Camara (2016), for every $i \in N$, the posterior beliefs $p_i \in [0,1]$ and $p_0^i \in [0,1]$ of receivers i and the sender conditional on the private information of i are in a one-to-one relationship:

$$p_i = Q_{\ell_i}\left(p_0^i\right) = \frac{p_0^i}{p_0^i + \left(1-p_0^i\right)\ell_i}$$

where $\ell_i = \left(\frac{\pi_0}{1-\pi_0}\right)\left(\frac{1-\pi_i}{\pi_i}\right)$. With this, the problem of the sender can be restated by pretending that all the receivers share the same common prior π_0 and that their posterior beliefs are distorted by the maps Q_{ℓ_i} before aggregation, that is,

$$\sup_{(Y,\sigma)\in\Sigma}\left\{\int_{\Omega\times Y} u\left(\sum_{i\in N} s_i Q_{\ell_i}\left(p_0^i\left(y_i\right)\right), \omega\right) d\sigma\left(y|\omega\right) d\pi\left(\omega\right)\right\}$$

where, for every $i \in N$, $p_0^i(y_i) \in [0,1]$ is the posterior belief over Ω induced by prior π_0 and signal realization y_i. Given that each Q_{ℓ_i} is nonlinear, it is not possible to apply in general the results of this paper to the problem just defined. However, this shows as the framework analyzed is flexible enough to study the heterogenous prior case. In particular, it is reasonable to conjecture that the sender would now target the agents that, at the same time, have higher weights s_i and more favorable priors π_i. The detailed analysis of this more challenging problem is left for future research.

Targeting and uninformed agents In a recent contribution, Banerjee et al (2019) proposed a variation of the DeGroot model where only a (exogenously determined) subset of agents $I \subseteq N$ receive initial private information about the relevant state. Differently from the setting of this paper, the agents $i \in I^c$ that do not receive any information have a "null" opinion $\emptyset$ and do not participate to the average updating as long as none of their neighbors $j \in N_i$ has a "proper" opinion x_i^t. As soon as one of their neighbors becomes "active" stating a proper opinion, they start updating according to the standard DeGroot model. With this, the authors are able to combine both information diffusion and aggregation in a single parsimonious model. However, they do not study targeting problems or endogenous information design. Hence, it seems natural to ask how the findings of Section 4 would change under this different hypothesis. The answer turns out to be quite simple if, as in the most part of this paper, we consider a sender that is only interested in the long-run consensus. In this case, there is no incentive for the sender to try to persuade more that 1 receiver. Indeed, regardless of the number of receivers targeted, the sender can without loss of optimality design a public signal and, given the absence of stickiness to the prior of the remaining agents, the common opinion induced coincides with the long-run

[19] Note that a binary uncertain state can always be modeled in the setting of this paper by considering $\Omega = [0,1]$ and restricting to prior beliefs supported on $\{0,1\}$.

outcome. Moreover, this is independent on the number of agents originally targeted. Therefore, in this case, any targeting policy would be optimal.[20] Conversely, if the sender is finitely patient, then the problem of the sender would become very similar to a standard seeding problem under network diffusion augmented with endogenous design of information. The formal analysis of this interesting case is left for future research but one can easily conjecture that the decay centrality (see Jackson 2008) of the receivers with respect to the discount factor δ of the sender would be important both in the targeting step and the information design step.

6 Appendix

Before stating the formal proofs of the results in the main text let me add some pieces of notation. For all priors $\pi \in \Delta(\Omega)$ and experiments $(Y, \sigma) \in \Sigma_k$, with $k \in \{1, ..., n\}$, let $\pi \otimes \sigma \in \Delta(\Omega \times Y)$ denote the induced joint distribution on $\Omega \times Y$.

Proof of Remark 1 Fix $\pi \in \Delta(\Omega)$ and $(Y, \sigma) \in \Sigma$. Note that, for all $x^0 \in X$, the induced sequence $\{x^t\}_{t \in \mathbb{N}}$ converges to $(T_G(x^0), ..., T_G(x^0))'$. Therefore, for all $(x^0, \omega) \in X \times \Omega$, continuity of v implies that the sequence $\{v(x^t, \omega)\}_{t \in \mathbb{N}}$ converges to $u(T_G(x^0), \omega)$. With this, it holds

$$\lim_{\delta \to 1} \left[(1-\delta) \sum_{t=0}^{\infty} \delta^t v\left(x^t, \omega\right) \right] = \lim_{\tau \to \infty} \left[\frac{1}{\tau} \sum_{t=0}^{\tau} v\left(x^t, \omega\right) \right] = \lim_{t \to \infty} v\left(x^t, \omega\right) = u\left(T_G\left(x^0\right), \omega\right).$$

Finally, by the Dominated Convergence Theorem we have

$$\begin{aligned} & \lim_{\delta \to 1} \left\{ \int_{\Omega \times Y} \left[(1-\delta) \sum_{t=0}^{\infty} \delta^t v\left(x^t(y), \omega\right) \right] d\sigma(y|\omega)\, d\pi(\omega) \right\} \\ = & \int_{\Omega \times Y} \lim_{\delta \to 1} \left[(1-\delta) \sum_{t=0}^{\infty} \delta^t v\left(x^t(y), \omega\right) \right] d\sigma(y|\omega)\, d\pi(\omega) \\ = & \int_{\Omega \times Y} u\left(T_G\left(x^0\right), \omega\right) d\sigma(y|\omega)\, d\pi(\omega) \end{aligned}$$

as desired. ■

Proof of Lemma 1 Let $\Delta_M(\Omega \times X)$ denote the set of joint distributions satisfying the constraints of eq. (2) in the statement of the lemma. Fix a prior $\pi \in \Delta$, and an affine aggregator $s \in \Delta(N)$. It is well known that, without loss of generality, the sender can restrict to experiments $(Y, \sigma) \in \Sigma_n$ such that $Y = X$ and σ is such that, for all $i \in N$,

$$\mathbb{E}_{\pi,\sigma}[\omega | x_i] = x_i \qquad \pi \otimes \sigma \text{ a.e.} \tag{10}$$

[20] Differently, in the case of costly targeting, the sender strictly prefers to target only one receiver. However, the sender is indifferent on the identity of the receiver.

To see this, fix an arbitrary experiment $(Y,\sigma) \in \Sigma_n$ and consider the measurable map $f_\sigma : Y \to X$ defined as

$$f_\sigma(y)_i = \mathbb{E}_{\pi,\sigma}[\omega | y_i].$$

Next, consider $Y' = X$ and $\sigma' : \Omega \to \Delta(Y')$ defined by

$$\sigma'(E|\omega) = \sigma(\{y \in Y : f_\sigma(y) \in E\} | \omega) \qquad \forall E \in \mathcal{B}_X, \forall \omega \in \Omega.$$

Then Bayesian updating implies that $\mathbb{E}_{\pi,\sigma'}[\omega | x_i] = x_i$. Given that (Y,σ) and (Y',σ') induce the same distribution on the payoff relevant space $\Omega \times X$, then the sender obtains the same expected payoff with both the experiments. Let Σ_n^* denote the space of experiments satisfying the condition in eq. (10). Next, I show that the set of joint distributions over $\Omega \times X$ induced by all $\sigma \in \Sigma_n^*$ is equal to $\Delta_M(\Omega \times X)$. Let $\gamma = \pi \otimes \sigma$ for some $\sigma \in \Sigma_n^*$ and fix an arbitrary $(h_i)_{i \in N_0} \in C([0,1])^{N_0}$. It follows that

$$\begin{aligned}
\mathbb{E}_\gamma\left[h_0(\omega) + \sum_{i \in N} h_i(x_i)(x_i - \omega)\right] &= \mathbb{E}_\gamma[h_0(\omega)] + \sum_{i \in N} \mathbb{E}_\gamma[h_i(x_i)(x_i - \omega)] \\
&= \mathbb{E}_\pi[h_0(\omega)] + \sum_{i \in N} \mathbb{E}_\gamma[h_i(x_i)\mathbb{E}_\gamma[(x_i - \omega)|x_i]] \\
&= \mathbb{E}_\pi[h_0(\omega)] + \sum_{i \in N} \mathbb{E}_\gamma[h_i(x_i)0] = \mathbb{E}_\pi[h_0(\omega)].
\end{aligned}$$

Conversely, let $\gamma \in \Delta_M(\Omega \times X)$ and note that, for all $i \in N$ and $h_i \in C([0,1])$,

$$\mathbb{E}_\gamma[h_i(x_i)\mathbb{E}_\gamma[x_i - \omega | x_i]] = 0$$

that is $\mathbb{E}_\gamma[x_i - \omega | x_i] = 0$ γ a.e. It follows that γ is a Bayes correlated equilibrium of an ancillary game where players feasible actions are the reports of their conditional expectations and their payoff is given by $u_i(x_i,\omega) = -(x_i - \omega)^2$. Therefore, there exists an information structure $\sigma \in \Sigma_n^*$ that induces γ (see Bergemann and Morris 2013, 2016). Finally, I show that the problem in the statement admits solutions. First, note that the map

$$\gamma \mapsto \int_{\Omega \times X} u(s \cdot x, \omega)\, d\gamma(\omega, x)$$

is continuous. Second let $\{\gamma_n\}_{n \in \mathbb{N}} \subseteq \Delta_M(\Omega \times X)$ be such that $\gamma_n \to \gamma$. Fix an arbitrary $(h_i)_{i \in N_0} \in C([0,1])^{N_0}$ and note that, for all $n \in \mathbb{N}$,

$$\mathbb{E}_{\gamma_n}\left[h_0(\omega) + \sum_{i \in N} h_i(x_i)(x_i - \omega)\right] = \mathbb{E}_\pi[h_0(\omega)].$$

It then follows by definition of weak convergence that

$$\mathbb{E}_\gamma\left[h_0(\omega) + \sum_{i \in N} h_i(x_i)(x_i - \omega)\right] = \mathbb{E}_\pi[h_0(\omega)],$$

proving that $\Delta_M(\Omega \times X)$ is closed, hence compact. With this, the maximum is attained in $\Delta_M(\Omega \times X)$. ■

Proof of Lemma 2 The proof of the first part follows the same steps of the proof of Lemma **1**. The second part follows from Corollary 2 in Kamenica and Gentzkow (2011). ■

Proof of Corollary 1 Fix $\pi \in \Delta(\Omega)$, $n \in \mathbb{N}$, $s \in \Delta(\{1, ..., n\})$, and a continuous $u : [0,1] \to \mathbb{R}$. It is well known (see for example Kolotilin, 2018) that

$$V_1(u, \pi) = cav(\hat{u})(\pi) = \max \{\mathbb{E}_\mu [u] : \mu \in \Delta([0,1]), \mu \succsim_{cv} \pi\}.$$

For all $\gamma \in \Delta_M(\Omega \times X)$ define $\nu_{\gamma,s}$ and $\hat{\nu}_{\gamma,s}$as in the proof of Proposition 1. Finally, observe that

$$\begin{aligned}
V_1(u, \pi) &\leq V_n(u, \pi, s) = \max_{\gamma \in \Delta_M(\Omega \times X)} \int_{\Omega \times X} u(s \cdot x)\, d\gamma(\omega, x) \\
&= \max_{\nu_{\gamma,s} : \gamma \in \Delta_M(\Omega \times X)} \int_{\Omega \times [0,1]} u(z, \omega)\, d\nu_{\gamma,s}(z, x) \\
&= \max_{\hat{\nu}_{\gamma,s} : \gamma \in \Delta_M(\Omega \times X)} \int_{[0,1]} u(z)\, d\hat{\nu}_{\gamma,s}(z) \\
&\leq \max \{\mathbb{E}_\mu [u] : \mu \in \Delta([0,1]), \mu \succsim_{cv} \pi\} = V_1(u, \pi)
\end{aligned}$$

proves the first part of the statement. Finally, given that there exists a public experiment $(Y^*, \sigma^*) \in \Sigma^P$ that attains the optimal value $V_1(u, \pi)$ in the restricted problem, it follows that the same public experiment solves the unrestricted problem of the sender. ■

Proof of Proposition 2 Fix $\pi \in \Delta(\Omega)$, $n \in \mathbb{N}$, a continuous $u : [0,1] \to \mathbb{R}$ and a robust opinion aggregator T. First, let u be nondecreasing and T concave. Following the same step of Lemma **1** it is possible to prove that

$$V_n(u, \pi, T) = \max_{\gamma \in \Delta(\Omega \times X)} \int_{\Omega \times X} u(T(x))\, d\gamma(\omega, x).$$

where the affine aggregator has to be replaced by the more general T. Moreover, as showed in Ghirardato et al (2004), given that T is monotone, translation invariant, and concave there exist $s_T \in \Delta(N)$ such that

$$T(x) \leq s_T \cdot x \qquad \forall x \in X.$$

Next, fix an arbitrary $\gamma \in \Delta(\Omega \times X)$ and, given that u is nondecreasing, it follows that

$$\int_{\Omega \times X} u(T(x))\, d\gamma(\omega, x) \leq \int_{\Omega \times X} u(s_T \cdot x)\, d\gamma(\omega, x).$$

Therefore,

$$V_n(u, \pi, T) = \max_{\gamma \in \Delta(\Omega \times X)} \int_{\Omega \times X} u(T(x))\, d\gamma(\omega, x) \leq \max_{\gamma \in \Delta(\Omega \times X)} \int_{\Omega \times X} u(s_T \cdot x)\, d\gamma(\omega, x) = cav(\hat{u})(\pi)$$

where the last step follows from Corollary **1**. The case for u nonincreasing and T convex is completely symmetric. Finally, there exists a public signal $(Y^*, \sigma^*) \in \Sigma^P$ such that

$$\begin{aligned}\int_{\Omega\times Y^*} u\left(T\left(\mathbb{E}_{\pi,\sigma^*}[\omega|y],...,\mathbb{E}_{\pi,\sigma^*}[\omega|y]\right)\right) d\sigma^*(y|\omega)\,d\pi(\omega) &= \int_{\Omega\times Y^*} u\left(\mathbb{E}_{\pi,\sigma^*}[\omega|y]\right) d\sigma^*(y|\omega)\,d\pi(\omega) \\ &= cav(\hat{u})(\pi)\end{aligned}$$

where the first equality follows from normalization of T and the second equality from Lemma **2**. ■

Proof of Proposition 1 Fix $\pi \in \Delta$, $n \in \mathbb{N}$, and an affine aggregator $s \in \Delta(N)$. First, given Lemma **2**, it follows that

$$cav(\hat{u})(\pi) = V_1(u,\pi) \leq V_n(u,\pi,s).$$

Indeed, the sender can always commit to design a public experiment in the unrestricted case. Next, fix an arbitrary $\gamma \in \Delta_M(\Omega \times X)$ and define $\nu_\gamma \in \Delta(\Omega \times [0,1])$ as

$$\nu_{\gamma,s}(E) = \gamma\left(\{(\omega,x) \in \Omega \times X : (\omega, s\cdot x) \in E\}\right) \qquad \forall E \in \mathcal{B}_{\Omega\times X}.$$

First, note that

$$\text{marg}_\Omega \nu_{\gamma,s} = \text{marg}_\Omega \gamma = \pi.$$

Second, fix an arbitrary continuous and concave function $\phi : [0,1] \to \mathbb{R}$, let $\hat{\nu}_{\gamma,s} = \text{marg}_{[0,1]}\nu_{\gamma,s}$, and observe that

$$\begin{aligned}\int_{[0,1]} \phi(z)\,d\hat{\nu}_{\gamma,s}(z) &= \int_{\Omega\times[0,1]} \phi(z)\,d\nu_{\gamma,s}(\omega,z) = \int_{\Omega\times X} \phi(s\cdot x)\,d\gamma(\omega,x) \\ &\geq \int_{\Omega\times X}\left(\sum_{i\in N} s_i\phi(x_i)\right) d\gamma(\omega,x) = \sum_{i\in N} s_i\left(\int_{\Omega\times X}\phi(x_i)\,d\gamma(\omega,x)\right) \\ &= \sum_{i\in N} s_i\left(\int_{\Omega\times X}\phi\left(\mathbb{E}_\gamma[\omega|x_i]\right)d\gamma(\omega,x)\right) \geq \sum_{i\in N} s_i\left(\int_{\Omega\times X}\mathbb{E}_\gamma[\phi(\omega)|x_i]\,d\gamma(\omega,x)\right) \\ &= \sum_{i\in N} s_i\left(\int_{\Omega\times X}\phi(\omega)\,d\gamma(\omega,x)\right) = \int_\Omega \phi(\omega)\,d\pi(\omega).\end{aligned}$$

Given that ϕ was arbitrarily chosen, it follows that $\text{marg}_{[0,1]}\nu_{\gamma,s} = \hat{\nu}_{\gamma,s} \succsim_{cv} \pi$. Given that γ was arbitrarily chosen, it follows that for all $\gamma \in \Delta_M(\Omega \times X)$,

$$\nu_{\gamma,s} \in \left\{\nu \in \Delta(\Omega\times[0,1]) : \text{marg}_\Omega\nu = \pi, \text{marg}_{[0,1]}\nu \succsim_{cv} \pi\right\}.$$

Finally, given that, for all $\gamma \in \Delta_M(\Omega \times X)$,

$$\int_{\Omega\times X} u(s\cdot x,\omega)\,d\gamma(\omega,x) = \int_{\Omega\times[0,1]} u(z,\omega)\,d\nu_{\gamma,s}(z,x),$$

it follows from Lemma **1** that

$$\begin{aligned} V_n(u,\pi,s) &= \max_{\gamma\in\Delta_M(\Omega\times X)}\int_{\Omega\times X} u(s\cdot x,\omega)\,d\gamma(\omega,x) \\ &= \max_{\nu_{\gamma,s}:\gamma\in\Delta_M(\Omega\times X)}\int_{\Omega\times[0,1]} u(z,\omega)\,d\nu_{\gamma,s}(\omega,z) \\ &\leq \max\left\{\mathbb{E}_\nu[u] : \nu\in\Delta(\Omega\times[0,1]),\mathrm{marg}_\Omega\nu=\pi,\mathrm{marg}_{[0,1]}\nu\succsim_{cv}\pi\right\} \end{aligned}$$

proving the validity of the upper-bound. ■

Proof of Proposition 3 Fix $\pi\in\Delta(\Omega)$, $n\in\mathbb{N}$, a continuous $u:[0,1]\to\mathbb{R}$ and a set of robust opinion aggregator $\mathcal{T}$ containing at least an affine aggregator $s\in\Delta(N)$. First note that for every public experiment $(Y^*,\sigma^*)\in\Sigma^P$ and for each $T\in\mathcal{T}$, the value

$$\int_{\Omega\times Y^*} u(T(x(y)))\,d\sigma^*(y|\omega)\,d\pi(\omega)=\int_{\Omega\times Y^*} u(\mathbb{E}_{\pi,\sigma^*}[\omega|y])\,d\sigma^*(y|\omega)\,d\pi(\omega)$$

does not depend on T. Given that the sender can always commit to a public experiment it must be the case that

$$cav(\hat{u})(\pi)\leq V_n(u,\pi,\mathcal{T}).$$

Next, suppose by contradiction that

$$cav(\hat{u})(\pi)< V_n(u,\pi,\mathcal{T}).$$

But then it must hold that

$$\begin{aligned} V_n(u,\pi,\mathcal{T}) &> cav(\hat{u})(\pi)=\sup_{(Y,\sigma)\in\Sigma_k}\left\{\int_{\Omega\times Y} u(s\cdot x(y))\,d\sigma(y|\omega)\,d\pi(\omega)\right\} \\ &\geq \inf_{T\in\mathcal{T}}\sup_{(Y,\sigma)\in\Sigma_k}\left\{\int_{\Omega\times Y} u(T(x(y)))\,d\sigma(y|\omega)\,d\pi(\omega)\right\} \\ &\geq \sup_{(Y,\sigma)\in\Sigma_k}\inf_{T\in\mathcal{T}}\left\{\int_{\Omega\times Y} u(T(x(y)))\,d\sigma(y|\omega)\,d\pi(\omega)\right\}=V_n(u,\pi,\mathcal{T}) \end{aligned}$$

offering a contradiction. Therefore, it must be that $cav(\hat{u})(\pi)=V_n(u,\pi,\mathcal{T})$. The last part of the statement follows since normalization implies that there exists a public experiment $(Y^*,\sigma^*)\in\Sigma^P$ that does not depend on n and $\mathcal{T}$ and attains $cav(\hat{u})(\pi)$. ■

Proof of Proposition 4 Fix $\pi\in\Delta(\Omega)$, $I\subseteq N$, and $s\in\Delta(N)$. Note that the distorted payoff function $u_I(\cdot;s)$ is continuous, hence, by Corollary 1, it holds

$$V_I(u,\pi,s)=\max_{\mu\succsim_{cv}\pi}\left\{\int_{[0,1]} u_I(z;s)\,d\mu(z)\right\}=cav(\hat{u}_I(\cdot;s))(\pi)$$

and there exists a public experiment $(Y^*,\sigma^*)\in\Sigma_k^P$ that attains $cav(\hat{u}_I(\cdot;s))(\pi)$. For the second part of the statement, it will be enough to show that for all $\lambda,\lambda'\in[0,1]$ with $\lambda'\leq\lambda$,

$$\max_{\mu\succsim_{cv}\pi}\left\{\int_{[0,1]} u((1-\lambda')\bar{\pi}+\lambda' z)\,d\mu(z)\right\}\leq\max_{\mu\succsim_{cv}\pi}\left\{\int_{[0,1]} u((1-\lambda)\bar{\pi}+\lambda z)\,d\mu(z)\right\}.$$

Fix

$$\mu' \in \arg\max_{\mu \succsim_{cv} \pi} \left\{ \int_{[0,1]} u\left(\left(1-\lambda'\right)\bar{\pi} + \lambda' z\right) d\mu\left(z\right) \right\}$$

and define $\alpha\left(\lambda', \lambda\right) = \frac{\lambda'}{\lambda} \in [0,1]$. It follows that

$$\begin{aligned}
\int_{[0,1]} u_I\left(\left(1-\lambda'\right)\bar{\pi} + \lambda' z\right) d\mu'\left(z\right) &= \int_{[0,1]} u\left(\left(1-\lambda'\right)\bar{\pi} - \lambda\bar{\pi} + \lambda\bar{\pi} + \lambda' z\right) d\mu'\left(z\right) \\
&= \int_{[0,1]} u\left(\left(1-\lambda\right)\bar{\pi} - \lambda\left(\bar{\pi} + \frac{\lambda'}{\lambda}\left(z-\bar{\pi}\right)\right)\right) d\mu'\left(z\right) \\
&= \int_{[0,1]} u\left(\left(1-\lambda\right)\bar{\pi} - \lambda\left(\left(1-\alpha\left(\lambda',\lambda\right)\right)\bar{\pi} + \alpha\left(\lambda',\lambda\right) z\right)\right) d\mu'\left(z\right) \\
&= \int_{[0,1]} u\left(\left(1-\lambda\right)\bar{\pi} - \lambda z\right) d\mu'_{\alpha(\lambda',\lambda)}\left(z\right)
\end{aligned}$$

where

$$\mu'_{\alpha(\lambda',\lambda)}\left(E\right) = \mu'\left(\left\{z \in [0,1] : \left(1-\alpha\left(\lambda',\lambda\right)\right)\bar{\pi} + \alpha\left(\lambda',\lambda\right) z \in E\right\}\right) \qquad \forall E \in \mathcal{B}_{[0,1]}.$$

Next, I show that $\mu'_{\alpha(\lambda',\lambda)} \succsim_{cv} \pi$. Indeed, for all continuos and concave functions $\phi : [0,1] \to \mathbb{R}$ it holds

$$\begin{aligned}
\int_{[0,1]} \phi\left(z\right) d\mu'_{\alpha(\lambda',\lambda)}\left(z\right) &= \int_{[0,1]} \phi\left(\left(1-\alpha\left(\lambda',\lambda\right)\right)\bar{\pi} + \alpha\left(\lambda',\lambda\right) z\right) d\mu'\left(z\right) \\
&\geq \left(1-\alpha\left(\lambda',\lambda\right)\right)\int_{[0,1]} \phi\left(\bar{\pi}\right) d\mu'\left(z\right) + \alpha\left(\lambda',\lambda\right)\int_{[0,1]} \phi\left(z\right) d\mu'\left(z\right) \\
&\geq \left(1-\alpha\left(\lambda',\lambda\right)\right)\phi\left(\bar{\pi}\right) + \alpha\left(\lambda',\lambda\right)\int_{[0,1]} \phi\left(z\right) d\pi\left(z\right) \\
&\geq \left(1-\alpha\left(\lambda',\lambda\right)\right)\int_{[0,1]} \phi\left(z\right) d\pi\left(z\right) + \alpha\left(\lambda',\lambda\right)\int_{[0,1]} \phi\left(z\right) d\pi\left(z\right) = \int_{[0,1]} \phi\left(z\right) d\pi\left(z\right).
\end{aligned}$$

Finally, by combining the previous two steps it follows that

$$\begin{aligned}
\max_{\mu \succsim_{cv} \pi} \left\{ \int_{[0,1]} u\left(\left(1-\lambda'\right)\bar{\pi} + \lambda' z\right) d\mu\left(z\right) \right\} &= \int_{[0,1]} u_I\left(\left(1-\lambda'\right)\bar{\pi} + \lambda' z\right) d\mu'\left(z\right) \\
&= \int_{[0,1]} u\left(\left(1-\lambda\right)\bar{\pi} - \lambda z\right) d\mu'_{\alpha(\lambda',\lambda)}\left(z\right) \\
&\leq \max_{\mu \succsim_{cv} \pi} \left\{ \int_{[0,1]} u\left(\left(1-\lambda\right)\bar{\pi} + \lambda z\right) d\mu\left(z\right) \right\}
\end{aligned}$$

proving the claim. ■

Proof of Corollary 2 Suppose by contradiction that there exists $I^* \neq I\left(k,s\right)$ such that $|I^*| \leq k$ and $V_{I^*}\left(u,\pi,s\right) > V_{I(k,s)}\left(u,\pi,s\right)$. However, by construction it holds that $s_{I(k,s)} \geq s_{I^*}$, hence, by Proposition 4 $V_{I(k,s)}\left(u,\pi,s\right) \geq V_{I^*}\left(u,\pi,s\right)$ obtaining a contradiction. ■

Proof of Proposition 5 Fix $\pi \in \Delta([0,1])$, a continuously differentiable $u : [0,1] \to \mathbb{R}$ and recall that the capacity $k \in \mathbb{N}$ is fixed and does not depend on the population size n. Suppose that there exists $\varepsilon > 0$ such that u is strictly convex when restricted to $(\bar{\pi} - \varepsilon, \bar{\pi} + \varepsilon)$. Next, note that the assumption in eq. (5) implies that for all sequence of subsets $\{I(n)\}_{n\in\mathbb{N}} \subseteq \{N(n)\}_{n\in\mathbb{N}}$ with $|I(n)| \leq k$, it holds $s(n)_{I(n)} \to 0$. Therefore, there exists $\bar{n} \in \mathbb{N}$ such that for all $n > \bar{n}$ and $z \in [0,1]$

$$\left(1 - s(n)_{I(n)}\right)\bar{\pi} + s(n)_{I(n)} z \in (\bar{\pi} - \varepsilon, \bar{\pi} + \varepsilon).$$

With this, for all $n > \bar{n}$, the distortion payoff function $u_{I(n)}\left(\cdot; s(n)_{I(n)}\right) : [0,1] \to \mathbb{R}$ is strictly convex. By Part (ii) of Corollary 2 in Kolotilin (2018), full disclosure is a solution of the problem

$$\max_{\mu \succsim_{cv} \pi} \int_{[0,1]} u_{I(n)}\left(z; s(n)_{I(n)}\right) d\mu(z)$$

hence, by Proposition 4, full-disclosure solves the problem of the sender conditional on target set $I(n)$ for all $n > \bar{n}$. Finally, given Corollary 2, for all $n \in \mathbb{N}$, $I(k, s(n))$ solves the targeting problem of the sender and the sequence of subsets $\{I(k, s(n))\}_{n\in\mathbb{N}} \subseteq \{N(n)\}_{n\in\mathbb{N}}$ satisfies $|I(n)| \leq k$, implying Part (i) of the statement in the main text. The proofs of Parts (ii) and (iii) of the statement in the main text follows the same steps and exploits Parts (i) and (iii) of Corollary 2 in Kolotilin (2018). ■

Before proving the results of Section 4.3, I will need two ancillary lemmas.

Lemma 3 *Fix a continuous payoff function $u : [0,1] \to \mathbb{R}$ satisfying the assumptions of Section 4.3. For all $I \subseteq N$ and $s \in \Delta(N)$,the distorted payoff function $u_I(\cdot; s)$ is convex on $[0, \bar{z}(s_I)]$ and affine on $[\bar{z}(s_I), 1]$, where*

$$\bar{z}(s_I) = \max\left\{\min\left\{\frac{\bar{z} - (1 - s_I)\bar{\pi}}{s_I}, 1\right\}, 0\right\}.$$

Proof of Lemma 3 Fix $I \subseteq N$ and $s \in \Delta(N)$. Recall that

$$u_I(z; s) = u((1 - s_I)\bar{\pi} + s_I z) \qquad \forall z \in [0,1]$$

and that u is convex in $[0, \bar{z}]$ and affine in $[\bar{z}, 1]$. Next, fix $z, z' \in [0, \bar{z}(s_I)]$ and $\lambda \in [0,1]$. If $\bar{z}(s_I) = 0$, then trivially $z = z' = 0$ and

$$u_I(\lambda z + (1 - \lambda) z'; s) = 0 = \lambda u_I(z; s) + (1 - \lambda) u_I(z'; s).$$

If $\bar{z}(s_I) > 0$ note that

$$((1 - s_I)\bar{\pi} + s_I z), ((1 - s_I)\bar{\pi} + s_I z') \in [0, \bar{z}]$$

hence

$$(1-s_I)\,\bar{\pi}+s_I\,(\lambda z+(1-\lambda)\,z')=\lambda\,((1-s_I)\,\bar{\pi}+s_I z)+(1-\lambda)\,((1-s_I)\,\bar{\pi}+s_I z')\in[0,\bar{z}]\,,$$

Therefore,

$$\begin{array}{rcl} u_I\,(\lambda z+(1-\lambda)\,z';s) & = & u\,((1-s_I)\,\bar{\pi}+s_I\,(\lambda z+(1-\lambda)\,z')) \\ & = & u\,(\lambda\,((1-s_I)\,\bar{\pi}+s_I z)+(1-\lambda)\,((1-s_I)\,\bar{\pi}+s_I z')) \\ & \leq & \lambda u\,(1-s_I)\,\bar{\pi}+s_I z+(1-\lambda)\,u\,((1-s_I)\,\bar{\pi}+s_I z') \end{array}$$

showing that $u_I\,(\cdot;s)$ is convex on $[0,\bar{z}\,(s_I)]$. Next, fix $z,z'\in[\bar{z}\,(s_I)\,,1]$ and $\lambda\in[0,1]$. If $\bar{z}\,(s_I)=1$, the trivially $z=z'=1$ and

$$u_I\,(\lambda z+(1-\lambda)\,z';s)=u_I\,(1;s)=\lambda u_I\,(z;s)+(1-\lambda)\,u_I\,(z';s)\,.$$

If $\bar{z}\,(s_I)<1$ note that

$$((1-s_I)\,\bar{\pi}+s_I z)\,,((1-s_I)\,\bar{\pi}+s_I z')\in[\bar{z},1]$$

hence

$$(1-s_I)\,\bar{\pi}+s_I\,(\lambda z+(1-\lambda)\,z')=\lambda\,((1-s_I)\,\bar{\pi}+s_I z)+(1-\lambda)\,((1-s_I)\,\bar{\pi}+s_I z')\in[\bar{z},1]\,,$$

Therefore,

$$\begin{array}{rcl} u_I\,(\lambda z+(1-\lambda)\,z';s) & = & u\,((1-s_I)\,\bar{\pi}+s_I\,(\lambda z+(1-\lambda)\,z')) \\ & = & u\,(\lambda\,((1-s_I)\,\bar{\pi}+s_I z)+(1-\lambda)\,((1-s_I)\,\bar{\pi}+s_I z')) \\ & = & \lambda u\,(1-s_I)\,\bar{\pi}+s_I z+(1-\lambda)\,u\,((1-s_I)\,\bar{\pi}+s_I z') \end{array}$$

showing that $u_I\,(\cdot;s)$ is affine on $[\bar{z}\,(s_I)\,,1]$. ■

Lemma 4 *Let $u:[0,1]\to\mathbb{R}$ be continuous and let $\bar{z}\in[0,1]$ be such that the restriction $u_{|[0,\bar{z}]}$ (resp. $u_{|[\bar{z},1]}$) of u over $[0,\bar{z}]$ (resp. $[\bar{z},1)$ is convex (resp. affine). Moreover, let $\omega^*\in[0,1]$ be a solution of*

$$\mathbb{E}_\pi\,[\omega|\omega\geq\omega^*]=\bar{z}$$

The following facts are true:

(i) *If $\bar{\pi}\geq\bar{z}$, then it is optimal to reveal no information;*

(ii) *If $\bar{\pi}<\bar{z}$, then it is optimal to disclose ω whenever $\omega<\omega^*$ and to reveal only that $\omega\geq\omega^*$ whenever $\omega\geq\omega^*$.*

Proof of Lemma 4 Suppose that $\bar{\pi} \geq \bar{z}$. Define $\varphi : [0,1] \to \mathbb{R}$ as the affine extension of $u_{|[\bar{z},1]}$ over $[0,1]$ and $\mu = \delta_{\bar{\pi}}$, i.e., the Dirac measure that puts mass 1 over $\bar{\pi}$. It follows that $\varphi \geq u$, supp$\mu = \{\bar{\pi}\} \subseteq \{\varphi = u\}$, $\mu \succsim_{cv} \pi$, and

$$\int_{[0,1]} \varphi(z)\, d\mu(z) = \varphi(\bar{\pi}) = \int_{[0,1]} \varphi(z)\, d\pi(z)$$

since φ is affine. By Theorem 1 in Dworczak and Martini (2019), μ solves the sender's problem. Next, suppose that $\bar{\pi} < \bar{z}$ and define

$$\varphi(z) = \begin{cases} u(z) & if \quad z < \omega^* \\ u(\omega^*) + \left(\frac{u(\bar{z}) - u(\omega^*)}{\bar{z} - \omega^*}\right)(z - \omega^*) & else \end{cases} \qquad \forall z \in [0,1].$$

Note that φ is convex and such that $\varphi \geq u$. Finally, define $\mu \in \Delta([0,1])$ with distribution function

$$F_\mu(z) = \begin{cases} F_\pi(z) & if \quad z < \omega^* \\ F_\pi(\omega^*) + \mathbf{1}[z \geq \bar{z}](1 - F_\pi(\omega^*)) & else \end{cases} \qquad \forall z \in [0,1].$$

where F_π is the distribution function of π. Clearly, it holds

$$\text{supp}\mu = [0, \omega^*] \cup \{\bar{z}\} \subseteq \{\varphi = u\}$$

and

$$\begin{aligned} \int_{[0,1]} \varphi(z)\, d\mu(z) &= \mathbb{E}_\pi[\varphi | \omega \leq \omega^*] F_\pi(\omega^*) + \varphi(\bar{z})(1 - F_\pi(\omega^*)) \\ &= \int_{[0,1]} \varphi(z)\, d\pi(z). \end{aligned}$$

Finally, μ is a mean-preserving spread of π as it only moves mass from $[\omega^*, 1]$ to $\{\bar{z}\}$ with respect to π. By Theorem 1 in Dworczak and Martini (2019), μ solves the sender's problem.■

Proof of Proposition 6 From Lemma 3, the distorted payoff $u_I(\cdot, s)$ is convex below $\bar{z}(s_I)$ and affine above this threshold. Moreover, Lemma 4 implies that if $\bar{\pi} \geq \bar{z}(s_I)$, then it is optimal to reveal no information, otherwise the censorship policy characterized by $\omega^*(s, I) \in [0,1]$ defined as

$$\mathbb{E}_\pi[\omega | \omega \geq \omega^*(s, I)] = \bar{z}(s_I)$$

solves the sender's problem with payoff function $u_I(\cdot, s)$. There are four cases:

1. If $\bar{\pi} \geq \frac{\bar{z}}{1 - s_I}$,that is $\bar{z}(s_I) = 0$, then always $\bar{\pi} \geq \bar{z}(s_I)$ and it is optimal to reveal no information (i.e., $\omega^*(s, I) = 0$);

2. If $\frac{\bar{z}}{1 - s_I} > \bar{\pi} \geq \bar{z}$, that is $\bar{z}(s_I) \in (0,1)$ and $\bar{\pi} \geq \bar{z}(s_I)$, then is optimal to reveal no information (i.e., $\omega^*(s, I) = 0$);

3. If $\bar{\pi} < \bar{z} \leq (1 - s_I)\bar{\pi} + s_I 1$, that is $\bar{z}(s_I) \in (0,1]$ and $\bar{\pi} < \bar{z}(s_I)$, then the optimal threshold $\omega^*(s,I)$ solves

$$\mathbb{E}_{\pi}[\omega|\omega \geq \omega^*(s,I)] = \frac{\bar{z} - (1 - s_I)\bar{\pi}}{s_I} \iff (1 - s_I)\bar{\pi} + s_I \mathbb{E}_{\pi}[\omega|\omega \geq \omega^*(s,I)] = \bar{z};$$

4. If $\bar{z} \geq (1 - s_I)\bar{\pi} + s_I 1$, that is $\bar{z}(s_I) = 1$, then the optimal threshold $\omega^*(s,I)$ solves

$$\mathbb{E}_{\pi}[\omega|\omega \geq \omega^*(s,I)] = 1 \iff \omega^*(s,I) = 1.$$

Together the four cases prove the statement. ■

Proof of Corollary 3 Fix $I, I' \subseteq N$ and $s, s' \in \Delta(N)$ such that $s_I \geq s'_{I'}$. There are four cases:

1. If $\bar{\pi} \geq \bar{z}$, then $\omega^*(s,I) = \omega^*(s',I') = 0$;

2. If $\bar{\pi} < \bar{z} \leq (1 - s'_{I'})\bar{\pi} + s'_{I'}1$, then

$$(1 - s'_{I'})\bar{\pi} + s'_{I'}\mathbb{E}_{\pi}[\omega|\omega \geq \omega^*(s',I')] = \bar{z} = (1 - s_I)\bar{\pi} + s_I \mathbb{E}_{\pi}[\omega|\omega \geq \omega^*(s,I)]$$

   which necessarily implies that $\omega^*(s,I) \leq \omega^*(s',I')$;

3. If $(1 - s'_{I'})\bar{\pi} + s'_{I'}1 < \bar{z} \leq (1 - s_I)\bar{\pi} + s_I 1$, then $\omega^*(s,I) \leq 1 = \omega^*(s',I')$;

4. If $(1 - s_I)\bar{\pi} + s1 < \bar{z}$, then $\omega^*(s,I) = \omega^*(s',I') = 1$.

Together the four cases prove the statement. ■

Proof of Proposition 7 Fix $\pi \in \Delta(\Omega)$, $I \subseteq N$, $S \subseteq \Delta(N)$. Given the equivalent representation of the sender's problem of eq. (9), with the same steps of Proposition 3, it is possible to show that it is without loss of optimality for the sender to design public signals and that his problem is

$$\begin{aligned} V_n(u,\pi,S) &= \max_{\mu \succsim_{cv} \pi} \left\{ \int_{[0,1]} u\left(\min_{s \in S} ((1 - s_I)\bar{\pi} + s_I z) \right) d\mu(z) \right\} \\ &= \max_{\mu \succsim_{cv} \pi} \left\{ \int_{[0,1]} u((1 - \underline{S}(I))\bar{\pi} + \underline{S}(I) z) \, d\mu(z) \right\} \\ &= \max_{\mu \succsim_{cv} \pi} \left\{ \int_{[0,1]} u_I(z;S) \, d\mu(z) \right\} = cav(\hat{u}_I(\cdot;S))(\pi) \end{aligned}$$

where the second equality follows from eq. 3 in Ghirardato et al. (2004). Finally, the proof of second part of the statement follows exactly the same steps of the proofs of the second part of Proposition 4 and Corollary 2. ■

References

[1] M. Akbarpour, S. Malladi, and A. Saberi, Just a few seeds more: value of network information for diffusion, mimeo, 2018.

[2] C. D. Aliprantis and K. C. Border, *Infinite Dimensional Analysis*, 3rd ed., Springer Verlag, Berlin, 2006.

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

[4] I. Arieli, Y. Babichenko, F. Sandomirskiy, and O. Tamuz, Feasible Joint Posterior Beliefs, mimeo, 2020.

[5] C. Ballester, A. Calvo-Armengol, and Y. Zenou, Who's who in networks.Wanted: The key player, *Econometrica*, 73, 1403–1417, 2006.

[6] A. Banerjee, E. Breza, A. G. Chandrasekhar, and M. Mobius, Naive learning with uninformed agents, *National Bureau of Economic Research*, 25497, 2019.

[7] A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson, The diffusion of microfinance, *Science*, 341, 2013.

[8] A. Banerjee, A. G. Chandrasekhar, E. Duflo, and M. O. Jackson, Using gossip to spread information: theory and evidence from two randomized controlled trials, *Review of Economic Studies*, 86, 2453-2490, 2019.

[9] D. Bergemann and S. Morris, Robust predictions in games of incomplete information, *Econometrica*, 81, 1251-1308, 2013.

[10] D. Bergemann and S. Morris, Bayes correlated equilibrium and the comparison of information structures in games, *Theoretical Economics*, 2, 487-522, 2016.

[11] D. Bergemann and S. Morris, Information Design: A Unified Perspective, *Journal of Economic Literature*, 57, 1–57, 2019.

[12] E. Breza, A. G. Chandrasekhar, B. Golub, and A. Parvathaneni, Networks in economic development, *Oxford Review of Economic Policy*, 35, 678–721, 2019.

[13] E. Breza, A. G. Chandrasekhar, and A. Tahbaz-Salehi, Seeing the forest for the trees? An investigation of network knowledge, *National Bureau of Economic Research*, 24359, 2018.

[14] S. Cerreia-Vioglio, R. Corrao, and G. Lanzani, Robust opinion aggregation and its dynamics, IGIER Working Paper, 662, 2020.

[15] A. G. Chandrasekhar, H. Larreguy, and J. P. Xandri. Testing models of social learning on networks: Evidence from two experiments, *Econometrica*, 88, 1–32, 2020.

[16] M. DeGroot, Reaching a Consensus, *Journal of the American Statistical Association*, 69: 118–21, 1972.

[17] P. M. DeMarzo, D. Vayanos, and J. Zwiebel, Persuasion Bias, Social Influence, and Unidimensional Opinions, *Quarterly Journal of Economics*, 118, 909–968, 2003.

[18] P. Dworczak, and G. Martini, The simple economics of optimal persuasion, *Journal of Political Economics*, 127, 1993-2048, 2019.

[19] A. Galeotti, B. Golub, and S. Goyal, Targeting interventions in networks, *Econometrica*, forthcoming, 2020.

[20] S. Galperti and J. Perego, Belief meddling in social networks: An information-design approach, mimeo, 2019.

[21] M. Gentzkow and E. Kamenica, A Rothschild-Stiglitz approach to Bayesian persuasion, *American Economic Review Papers & Proceedings*, 106, 597-601, 2016.

[22] P. Ghirardato, F. Maccheroni, and M. Marinacci, Differentiating ambiguity and ambiguity attitude, *Journal of Economic Theory*, 118, 133–173, 2004.

[23] I. Goldstein and L. Yang, Information disclosure in financial markets, *Annual Review of Financial Economics*, 9, 101–125, 2017.

[24] B. Golub and M. O. Jackson, Naïve Learning in Social Networks and the Wisdom of Crowds, *American Economic Journal: Microeconomics*, 2, 112–149, 2010.

[25] B. Golub and M. O. Jackson, How Homophily Affects the Speed of Learning and Best-Response Dynamics, *Quarterly Journal of Economics*, 127, 1287–1338, 2012.

[26] M. O. Jackson, *Social and Economic Networks*, Princeton University Press, Princeton, 2008.

[27] E. Kamenica, M. Gentzkow, Bayesian Persuasion, *American Economic Review*, 101, 2590-2615, 2011.

[28] D. Kempe, J. Kleinberg, and É. Tardos, Maximizing the spread of influence through a social network, *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining*, 2003.

[29] A. Kolotilin, Optimal information disclosure: A linear programming approach, *Theoretical Economics*, 13, 607-635, 2018.

[30] A. Kolotilin, T. Mylovanov, and A. Zapechelnyuk, Censorship as optimal persuasion, mimeo, 2019.

[31] E. Lipnowski and D. Ravid, Cheap Talk with Transparent Motives, *Econometrica*, forthcoming, 2020.

[32] L. Mathevet, J. Perego, and I. Taneva, On Information Design in Games, *Journal of Political Economy*, forthcoming, 2020.

[33] R. McKelvey and T. Page, Common knowledge, consensus, and aggregate Information, *Econometrica*, 54, 109-127, 1986.

[34] S. Morris, Contagion, *Review of Economic Studies*, 67, 57-78, 2000.

[35] M. Muller-Frank, A general framework for rational learning in social networks, *Theoretical Economics*, 8, 1-40, 2013.

[36] E. T. Nielsen, A. Brandenburger, J. Geanakoplos, R. McKelvey, and T. Page, Common knowledge of an aggregate of expectations, *Econometrica*, 58, 1235-1239, 1990.

[37] X. Vives, *Information and learning in markets: the impact of market microstructure*, Princeton University Press, 2010.