Belief Meddling in Social Networks: An Information-Design Approach

Simone Galperti
UC San Diego

Jacopo Perego
Columbia University

February 21, 2019

Preliminary and Incomplete

Abstract

Social media have become an increasingly important source of information about political, social and economic issues. While beneficial on many levels, the decentralized nature of these media may expose societies to novel risks of manipulation by third parties. To evaluate these risks, we study a model where a designer sends information to agents who interact in a game, so as to affect its outcome. The designer can communicate only with a limited number of agents, who then share information with each other on a network of social links before playing the game. We characterize the equilibrium outcomes that can be induced by seeding this social network with information. Our main result recasts this *constrained* information-design problem in terms of an equivalent linear program, which is particularly useful for applications. We show that a simple property of the network—the *depth* of communication—fully determines the scope for belief manipulation. Finally, we illustrate how a holistic use of linear-programming duality helps to characterize the solution to the optimal seeding problem. Our theory offers insights into the design of advertisement and political campaigns that are robust to (or leverage on) information spillovers.

We thank Laura Doval, Elliot Lipnowsky, Stephen Morris and seminar participants at Pittsburgh and Carnegie Mellon University, Yale University, CIDE, ITAM, University of Southern California, SWET 2019, and Pompeu Fabra for their comments and suggestions. We are grateful to the Cowles Foundation for hospitality.

1 Introduction

The information people get from their social networks shapes their beliefs as much as that directly provided by institutional sources, such as news outlets or government agencies. Information spillovers between agents therefore matter for any third party interested in influencing beliefs in societies. For instance, political candidates may not safely assume that private campaign events will not leak to the public. Foreign entities may try to meddle with domestic politics by exploiting social media, which may be easier to infiltrate than traditional media. The board of a company may communicate to its divisions and lower ranks both directly and indirectly via the organization hierarchy.

In this paper, we examine the problem of a third party who uses information to influence people's beliefs and behaviors when information can spill over social relations. We model such relations as an exogenous directed network and adopt an information-design approach. We are especially interested in settings where the third party (hereafter, the designer) can only target a subset of the population with her information campaign, relying on the network to spread her messages. We characterize the behavioral outcomes that the designer can achieve and how they depend on the structure of the social network. We provide insights into the qualitative properties of optimal information targeting.

The model has three phases. First, the designer chooses what information to provide to the agents she can target. For each of them, this information takes the form of a private signal about some underlying state of the world. In the second phase, the agents share information with their neighbors in the network. In the last phase, given what they learned directly from the designer and indirectly from their neighbors, the agents play a game whose payoffs depend on the underlying state. In contrast to the standard information-design framework, our model restricts the set of agents the designer can directly reach with signals and adds the second phase of information spillovers.

Modeling information spillovers on networks raises serious challenges, as they can occur in many ways (see the related literature and Section 2.1). For instance, communication may be myopic or strategic and may rely on rich, coarse, or noisy messages. As a first attempt at studying this kind of information-design problems, we ignore all these intricacies and assume the simplest form of information spillovers: If there is a path (i.e., a chain of links) from one agent to another, then the latter will learn the private signal received by the former—possibly through multiple rounds of communication, which we leave implicit. This assumption helps us focus on the trade-offs caused by information spillovers. Later in the paper, we consider richer and more realistic forms of spillovers.

Our first contribution is to characterize the set of feasible outcomes of the final game that the designer can induce. An outcome is a joint distribution between the agents' actions and the state.[1] We do this in two steps.

The first step is essential for the entire analysis. We show that every problem where the designer is constrained to targeting a subset of agents can be transformed into an auxiliary unconstrained problem which has the same feasible outcomes. In this problem, the designer can target all agents, but faces a richer network that is uniquely derived from the original network. Intuitively, the richer network captures the additional constraints of having to rely on targeted agents as intermediaries in order to reach all agents. To prove this result, we exploit an unexpected connection with the computer science literature on cryptography, from which we borrow a technique known as the *secret sharing* method (see Shamir (1979)).

The second step builds on the tradition of direct mechanisms (Myerson (1986)) and characterizes all feasible outcomes of any unconstrained problem by an obedience condition adapted to our settings. Indeed, we can view the unconstrained designer as directly recommending to each agent how to play in the final game. Obedience requires that each agent i be willing to follow his recommendation conditional on the information he has. In our setting, this is the information about the state and others' behavior revealed by the recommendation to i as well as to all agents connected to i through *some path*. We call these agents i's information sources. Importantly, obedience defines a system of linear inequalities. Together, these two steps dramatically simplifies finding the feasible outcomes for constrained problems, for which a direct-recommendation approach is invalid. In contrast to Bergemann and Morris (2016), in our obedience notion the agents can act on richer information sets and the designer may have to recommend *mixed* actions.

How do the feasible outcomes change depending on the social network and the set of target agents? To answer this, we introduce a new order on networks based on the idea of *followers*. Agent i is a follower of agent j if there is a path in the network *from j to i*. Given this, a network is deeper than another if, for each agent, his followers in the latter are also followers in the former. We show that this depth notion summarizes the degree to which the information spillovers alone constrain the designer's ability to influence the agents' behavior: Deeper networks shrink the set of feasible outcomes for every final game. A network's depth is also directly related to its capacity to aggregate information: At the end of the information-spillover phase, each agent is more informed (in a Blackwell sense) if and only if the network is deeper. We provide necessary and

[1] As another interpretation, the set of feasible outcomes describes everything that can happen in the final game, across all possible (and unknown to the analyst) forms of information which the agents may initially receive and then share over the social network.

sufficient conditions on the original network and target set for the network in the auxiliary unconstrained problem to be deeper.

Our second contribution is to study *optimal* outcomes—again relying on the unconstrained problem corresponding to our original problem. To this end, we extend to the present setting an approach based on duality of linear programming that we developed in a companion paper Galperti and Perego (2018) for standard information-design problems. The need to allow for mixed-action recommendations complicates the tasks of establishing existence of solutions and of formulating the dual. Once this is done, the interpretation, insights, and general optimality properties identified in Galperti and Perego (2018) also apply here. Some extend the properties of optimal Bayesian persuasion discovered by Kamenica and Gentzkow (2011) to settings with multiple receivers who interact strategically. Other are new, specific to this richer settings, and driven by the information spillovers. We show that a holistic use of the various parts of duality promises to be helpful in solving the kind of design problems studied here. We illustrate this approach with examples of investment games.

In ongoing work, we extend the analysis to settings where the agents may also receive information from exogenous sources, the designer is uncertain about the structure of the social network, and information may flow stochastically over links. In terms of applications, we use our theory to better understand optimal targeting with rich information policies and to study electoral campaigns that are robust to information leakage.

Related Literature. The closest to this paper is the literature on information design, based on the seminal work of Kamenica and Gentzkow (2011) and Bergemann and Morris (2016) (see Bergemann and Morris (2018) for a survey). This literature considers the problem of designing an information structure for one decision maker or a group of strategic agents. Bergemann and Morris (2016) highlight another important interpretation of this problem: The set of feasible outcomes in their setting—and in our setting as well—describes everything that can happen in equilibrium for a specific context if the analyst cannot or is not willing to make any assumption on the information available to the agents. This theory has found numerous applications to studying, for instance, political campaigns, rating systems, financial stress tests, and banking regulations. Other papers have applied linear-programming methods to analyze information design. For a comparison with this literature, see Galperti and Perego (2018) and references therein.

A key difference from this literature is that all papers assume that the designer can provide information privately and directly to each agent and that the agents never share any information coming from the designer. By contrast, we allow for information spillovers between agents and for constraints on the designer's ability to communicate directly with some agents. The first difference relaxes privacy and the designer's complete control over information; the second restricts the feasible information structures and introduces indirect information provision through the network. These seem important directions to extend the theory of information design for broader applicability to economically relevant settings. Also, common sense suggests that information policies are intrinsically fragile tools to shape incentives. One reason could be that agents share their information. Thus, our paper allows us to study the robustness of information policies to spillovers and how optimal policies change as the agents get better at sharing information.

The paper is also related to the vast literature on social or observational learning.[2] This literature is concerned with how people learn about some aspect of the world from how their peers act on their information about it. This is a prominent form of information spillovers. However, usually these papers assume a simple and exogenous information structure for the agents (for example, independence across them and some distributional assumption like normality). By contrast, in our paper not only information structures can take any form, but also they are chosen by a strategic third party. This literature also had to deal with the issue of how people learn from information that travels on social networks. Several models exist, which acknowledge and try to address the complexity of this issue (see, e.g., Golub and Jackson (2012) and Mueller-Frank (2013)). Our baseline assumption on how information spillovers work overlooks most of this complexity, but it allows us to focus on the information-design part of the problem. Relaxing this assumption seems a direction worth pursuing.

Last but not least, there is a rich literature about optimal seeding or targeting of the agents in some social network.[3] A broad theme of this literature is which network nodes (i.e., agents) should be targeted with an intervention that activates those nodes—viewed as an on-off switch—and starts a contagion in the network. Usually, contagion is modeled as a mechanical process and the third party's goal is to achieve maximal contagion (perhaps, within some deadline). A key difference of our paper is that the considered intervention takes the form of information provision. How the agents respond to information can be richer than simply being on or off and depends on their (often) heterogeneous preferences. Thus, in our paper the question is not only whom to target,

[2]See Banerjee (1992), Bikhchandani et al. (1992), Smith and Sørensen (2000), Acemoglu et al. (2011), Golub and Sadler (2017).

[3]See Granovetter (1978), Banerjee et al. (2013), Jackson and Storms (2018), Akbarpour et al. (2018), Morris (2000), Sadler (2017).

but also what information to provide to the targets taking into account what will leak through the links. In so doing, this paper bridges the so far separate literatures on information design and optimal seeding, thereby paving the way to a general theory of optimal information targeting.

2 Model

This section introduces the model and then discusses its main assumptions. In short, a third party called the "designer" (she) tries to influence the outcome of some game by providing information to a subset of its players (he), whom she can reach with private, direct, messages. These messages may then spread through social ties to the other players before the game is played.

Primitives. There is a finite set of players N, where N is also the number of players. There is a finite set of states Ω with typical element ω. Players have a common, full-support, prior belief $\mu \in \Delta(\Omega)$. An information structure, denoted by (S, π), consists of a finite and "sufficiently large" set of signals S_i for each player i and a function $\pi : \Omega \to \Delta(S)$, where $S = S_1 \times \cdots \times S_N$.[4] Without loss of generality, assume that for all $s \in S$ there exists $\omega \in \Omega$ such that $\pi(s|\omega) > 0$. We sometimes abuse notation by writing π for the information structure (S, π). Let Π be the class of all information structures.

Information Spillovers. Players are organized in a commonly-known *communication network* E, where $E \subseteq N^2$ is the set of directed links between players. Specifically, if $(i, j) \in E$, there is a link from i to j on which information can flow from i to j (but not vice versa).[5] Given E, player i's neighborhood, denoted by N_i, consists of all players j such that $(j, i) \in E$. A *directed path* from j to i is a sequence of players $i_1, \dots, i_m$ such that $i_1 = j$, $i_m = i$, and $(i_k, i_{k+1}) \in E$ for all $k = 1, \dots, m-1$.

Information spillovers occur through the communication network as follows. For every (S, π), as usual each player first observes privately his component of every realization $s \in S$. This private information, however, can then spill over between neighbors in E. More precisely, we adopt the following assumption.

Assumption 1 (Information Spillovers)**.** *If there exists a directed path from player j to player i, then i learns j's signal s_j.*

Under this assumption, spillovers are entirely governed by the network. As shown in Section 5.1, Assumption 1 is consistent with a more explicit communication model where,

[4]Later, we will formally identify what "sufficiently large" means in our setting (see Lemma 1).

[5]As a convention, we assume that $(i, i) \in E$ for every i.

after receiving his private signal from π, each player announces truthfully his belief about the vector (ω, s) to his neighbors over multiple rounds. For now, it is better to leave the communication details in the background and focus on the consequences of information spillovers. An immediate one is to transform any information structure π into a unique structure π'. Denote this transformation with the function $f_E : \Pi \to \Pi$, where we add the index E to emphasize the dependence on the network structure. To highlight the difference between the *initial* information structure, π, and the *final* one, $\pi' = f_E(\pi)$, we denote signals from the former by s and signals from the latter by t.

Game and Equilibrium. After the communication phase, players play a game. Each player $i \in N$ has a finite set of actions A_i and a utility function $u_i : A \times \Omega \to \mathbb{R}$, where $A = A_1 \times \cdots \times A_n$. We denote the *basic* game—namely, the structure comprising only the model primitives—with $G = (\Omega, (A_i, u_i)_{i=1}^N, \mu)$. This together with any information structure π induces an incomplete-information game, (G, π). We focus on Bayes-Nash equilibria (BNE). A (behavioral) strategy of player i in (G, π) is $\sigma_i : T_i \to \Delta(A_i)$. A profile σ is a BNE of (G, π) if for each i, $t_i \in T_i$, and $a_i \in A_i$ with $\sigma_i(a_i|t_i) > 0$,

$$\sum_{a_{-i} \in A_{-i}, t_{-i} \in T_{-i}, \omega \in \Omega} \left[u_i(a_i, a_{-i}, \omega) - u_i(a'_i, a_{-i}, \omega)\right] \sigma(a_i, a_{-i}|t_i, t_{-i}) \pi(t_i, t_{-i}|\omega) \mu(\omega) \geq 0$$

for all $a'_i \in A_i$, where $\sigma(a_i, a_{-i}|t_i, t_{-i}) := \prod_{j=1}^N \sigma_j(a_j|t_j)$. Let $\mathrm{BNE}(G, \pi)$ be the set of BNE of (G, π).

Targeting Problem. The designer chooses the initial information to provide to the subset of players she can reach, called the target set $M \subseteq N$.[6] The designer is constrained by not being able to communicate privately and directly to the players outside M. This is the only exogenous constraint on the designer. We can express this in terms of information structures. We assume that she can costlessly commit to any π in the set

$$\Pi_M = \{\pi \in \Pi : |S_i| = 1 \text{ for } i \notin M\}.$$

The designer's payoff function is $v : A \times \Omega \to \mathbb{R}$. Her prior belief is also μ as for the players. For every $\pi \in \Pi$, we define the value of π as

$$V(\pi) = \max_{\sigma \in \mathrm{BNE}(G, \pi)} \sum_{a \in A, t \in T, \omega \in \Omega} v(a, \omega) \sigma(a|t) \pi(t|\omega) \mu(\omega).$$

The equilibrium selection in this definition is common to most of the information-design literature (Bergemann and Morris (2018)). Given the information spillovers captured by f_E and the target set M, the value function of the designer's problem is

$$V_E^*(M) = \sup_{\pi \in \Pi_M} V(f_E(\pi)).$$

[6] We allow for the unconstrained problem, $M = N$, as a special case.

In some settings, the designer may have the possibility of choosing the target set. For instance, she may be constrained only in terms of the number m of players she can target, but not their identity or location. That is, she can choose any M with $m < N$ members. In general, let $\mathcal{M}$ be some collection of feasible target sets. The optimal-targeting problem consists in choosing $M \in \mathcal{M}$ to maximize $V_E^*(M)$. Since there can be only finitely many target sets, this problem reduces to comparing the finitely many values $V_E^*(M)$.

2.1 Discussion of the Model

A few comments on Assumption 1 are in order. Similarly to the literature on information diffusion,[7] in our model information flows between players mechanically and through an exogenous network. The players communicate myopically and non-strategically, without internalizing how communication may affect the play of the final game. This is clearly restrictive, yet we view it as a useful starting point for several reasons. First, assuming that information *will* flow whenever two players are linked lies at the polar opposite of the standard and well-studied model of information design where information *never* flows between players. This clean contrast helps better understand the qualitative implications of information spillovers for the designer's problem. A second reason is that, for the case of an *unconstrained* designer who can target all agents, Assumption 1 identifies a worst-case scenario: For a broad class of forms of communication between players—strategic and not—the value of the designer's problem is bounded below by $V_E^*(N)$ (see Section 5.1). Thus, our analysis takes a robustness connotation against the largest class of information structures (i.e., $\Pi_N = \Pi$) and uncertainty on how information may spill over due to strategic or technological considerations.

Another reason for Assumption 1 is that richer forms of communication are certainly interesting, but also challenging conceptually and methodologically. For instance, in a game of strategic communication on a network—followed by our final game—each player can be both a sender and a receiver of information over multiple rounds. Also, such games are likely to have multiple equilibria, including a babbling one. Such an equilibrium renders the problem equivalent to the case of an empty network and so is the best equilibrium from the designer's viewpoint. Therefore, adopting the standard selection of designer-preferred equilibrium leads to an uninteresting case. Identifying other selection criteria that are appropriate for the settings considered here is beyond the scope of this paper. Also, Bayesian updating on networks is known to be especially hard

[7] For example, see Akbarpour et al. (2018), Jackson and Storms (2018), and references therein.

when it requires each player to indirectly infer others' signals and avoid double-counting them (see, for example, Mueller-Frank (2013) and Eyster and Rabin (2010)). Given this, assuming myopic and non-strategic communication may be descriptively accurate in some cases. For instance, the network complexity may render strategic thinking hard for the players, or each may ultimately face a decision problem independent of others' decisions.

The model allows for the possibility that not all information structures are available to the designer. This adds a new kind of constraints to the information-design problem. Of course, different applications may call for different constraints on the available information structures. A natural one in the context of the present paper—especially for applications—is that the designer may be able to provide information only to a subset of the players. For large social networks, assuming that a third party can directly communicate with each player individually seems unrealistic. This may happen because reaching each player involves some cost, or some players pay no attention to the designer (consciously or not). The designer then has to choose which subset of players to target *and* what information to provide them. Information spillovers add an interesting spin, since targeted players become intermediaries that let the designer reach non-targeted players. Thus, this paper adds a novel perspective on the old problem of optimal seeding, studied extensively in sociology, economics, and computer science (see related literature).

3 The Scope for Manipulation

A necessary step to solving our targeting problem is to characterize the outcomes of the game that the designer can induce with her initial information structure. These outcomes are intended as joint distributions between the players' actions and the state.

To proceed, we need some notation and terminology. Player j is called an information *source* of player i if there is a directed path from j to i. In this case, i is called a *follower* of j. Given E, we denote the set of i's sources by E_i and the set of i's followers by ${}_iE$.[8] Note that by convention i is a neighbor, source, and follower of himself. Given this, Assumption 1 implies that, before having to choose an action, each player i learns the signals of all his sources E_i. We will denote the vector of such signals by s_{E_i}. Thus, given any initial information structure (S, π), the final information structure $\pi' = f_E(\pi)$ is defined by $T_i = \times_{j \in E_i} S_j$ and, for every $t_i = s_{E_i}$, $\pi'(s_{E_i}|\omega) = \sum_{s_{-E_i}} \pi(s_{E_i}, s_{-E_i}|\omega)$, for

[8] To remember this, think of E_i as the set of players who "come before" i in the network E—hence his sources—and of ${}_iE$ as the set of players who "come after" i in the network E—hence his followers.

every $\omega \in \Omega$, and every player i, where $-E_i = N \setminus E_i$.

We can now formalize the outcomes of the game induced by the designer. Recall that each player acts on her signals by picking an action, possibly at random. Thus, we can view each information structure as inducing, for every ω, a distribution over mixed-action profiles. These profiles will then determine the actual final actions in a straightforward way. Since by assumption each π has finite support, we can only consider finite-support distributions. Denote typical elements of $R_i = \Delta(A_i)$ by α_i and let $R = \times_{i \in N} R_i$.

Definition 1 (Outcome Function)**.** An *outcome function* is a mapping $x : \Omega \to \Delta(R)$, where $x(\cdot|\omega)$ has finite support for every $\omega \in \Omega$.

Definition 2 (M-Feasible Outcome)**.** An outcome function x is *M-feasible* for (G, E) if there exists $\pi \in \Pi_M$ and $\sigma \in \text{BNE}(G, f_E(\pi))$ such that, for every $\omega \in \Omega$ and $\alpha \in R$,

$$x(\alpha_1, \ldots, \alpha_N|\omega) = \sum_{s \in S} \pi(s|\omega) \prod_{i \in N} \mathbb{I}\{\sigma_i(s_{E_i})) = \alpha_i\}, \tag{1}$$

where $\mathbb{I}\{\cdot\}$ is the indicator function. Let $X_M(G, E)$ be the set of M-feasible outcome functions for (G, E).

We denote the overall support of x by $\mathbf{x} = \{\alpha \in R : x(\alpha|\omega) > 0, \text{ for some } \omega \in \Omega\}$. Also, for every subset of players $N' \subset N$, let $\mathbf{x}_{N'}$ represent the projection of $\mathbf{x}$ on the components in N', that is, $\mathbf{x}_{N'} = \{\alpha \in R_{N'} : (\alpha, \alpha') \in \mathbf{x}, \text{ for some } \alpha' \in R_{-N'}\}$.

3.1 From Constrained to Unconstrained Seeding

The communication constraints imposed by the target set give rise to specific challenges for characterizing the feasible outcomes. By removing the ability to communicate directly and privately to each player, they render invalid the logic behind standard revelation-principle arguments which shows that feasible outcomes can be characterized in terms of action recommendations to the players. To overcome these challenges, we will show that our constrained-design problem is equivalent to another problem for which that characterization is valid. In this auxiliary problem, the designer can target all players (i.e., $M = N$), but faces a specific, richer, network directly derived from the original E.

To see the issue, consider the following case. There are three players. Player 1 and 2 are sources of player 3: $E_1 = \{1\}$, $E_2 = \{2\}$, and $E_3 = \{1, 2, 3\}$. Suppose $M = \{1, 2\}$. Given this, player 3's behavior is influenced only indirectly by his *joint* observation of player 1's and 2's signals through the network. Thus, it is possible that a_3 depends on (s_1, s_2) in such a way that neither s_1 nor s_2 alone is sufficient to pin a_3 down. This implies that neither player 1 nor player 2 can perfectly predict player 3's behavior. This uncertainty on the part of player 1 and 2 may be essential to sustain specific outcomes. If so, such outcomes cannot be sustained if the designer is restricted to communicating in a language that explicitly recommends how the players should play the game. This requires to reveal to either player 1 or 2 the behavior of player 3.

The challenge arises because targeted players serve as *information intermediaries* for the designer to reach non-targeted players. Thus, we can view the signal to a targeted player i as having two roles: It conveys to i the behavior the designer recommends him to play; it also conveys to i's followers *part* of the behavior the designer recommends them to play. She may distribute these parts among multiple sources of a non-targeted player, so that none of their signals can be reduced to equal the recommendation for that player.

This ability to communicate with a non-targeted player through multiple sources is also the key to solving the problem. Intuitively, suppose a non-targeted player has at least two sources who are not each other's source—we shall call them independent sources. Then, it may be possible to design signals for such sources so that together they allow their follower to learn the information sent by the designer, yet each signal alone reveals nothing about that information. Such "encrypted" information essentially allows the designer to restore a direct and private communication with non-targeted players. Therefore, it is as if these players are also in the target set.

This allows us to expand the target set in an auxiliary problem. The remaining hurdle is to identify which players can be reached with encrypted information. The rest of the players will essentially share all the information of some targeted player. We can therefore also add them to the target set, provided that we add to the network a path connecting each of them to the targeted player with whom they share all the information. The next definition formalizes these ideas.

Definition 3 (M-Expansion)**.** Given E, its M-expansion E^M is defined as the network such that, for all $i, j \in N$, $(i, j) \in E^M$ if $E_i \cap M \subseteq E_j$.[9]

Thus, E^M is always richer than the original network E and can be easily derived from it. The second condition in the definition covers several cases. First, suppose that $i \in M$, $j \notin M$, and $E_j = \{i, j\}$. In this case, j knows i's information because i is j's only source. At the same time, i knows j's information because j never receives information other than through i. Therefore, they commonly know each other's information, which we can capture in the auxiliary problem by adding a link backwards from j to i. Second, suppose that $i, j \notin M$, $i \notin E_j$, $j \notin E_i$, and $E_i \cap M = E_j \cap M = \{k, k'\}$. In this case, both i and j

[9] Note that if $(i, j) \in E$, then $E_i \cap M \subseteq E_j$. Therefore, $E \subseteq E^M$.

receive information only through the same targeted players k and k'. Therefore, again i and j know each other's information; we can capture this by adding a link between them in both directions. However, both i and j may learn information that neither k nor k' learn, because these sources are independent (i.e., $k' \notin E_k$ and $k \notin E_{k'}$) and that information is appropriately encrypted. Intuitively, this can be done by sending to k a signal containing an encoded message and to k' the key to decode it.

Before stating the main result of this section, we need to introduce a minor assumption. It is clear that if a non-targeted player has no source in the target set, it is impossible for the designer to reach this player. This case seems uninteresting and trivial to analyze. For this reason, we assume the following.

Assumption 2 (Reach)**.** *Given E, each possible M satisfies $E_i \cap M \neq \varnothing$ for all i.*

Under this assumption, we obtain the following result.

Theorem 1 (Unconstrained Equivalence)**.** $X_M(G, E) = X_N(G, E^M)$ *for all* (G, E).

Thus, by appropriately enriching the information spillovers, we can turn constrained problems into unconstrained problems where the designer can target all players. This result is crucial for reasons that will become clear in the next section. In proving this result (Lemma 6), we exploit a connection with the cryptography literature, from which we borrow a technique known as the *secret sharing* method. Secret sharing refers to the problem of distributing a "secret" among a group of m players, each of whom is allocated a "share" of the secret. The distribution is so that players learn the secret only if all m players pool their shares. If one or more shares are missing, nothing is learned about the secret. Thanks to this technique, the designer can exploit information intermediaries to *privately* communicate with agents even if they are not in M.

The intuition for the Theorem 1 is as follows. First, one can show that the M-expansion of a network does not change the targeted sources of every player. It follows that the M-expansion does not change the constraints imposed by information spillovers and so the set of M-feasible outcomes. The next step is the key. Suppose a non-targeted player i has independent sources in the expanded network. Then—as described before—it is possible to send encrypted information to i using only the flexibility allowed by the information structures constrained to targeting M. The consequence is that it is as if i can also be targeted by the designer. If a non-targeted player does not have independent sources in the expanded network, he must have a bi-directional path to either a player who has independent sources or to a targeted player. Either way, adding also this player to the target set does not allow the designer to send him information he would not already get otherwise. Repeating this argument, one can expand the set of players who can be targeted from M to the entire N.

3.2 Unconstrained Seeding and Robust Obedience

Motivated by Theorem 1, we now consider the problem of characterizing the feasible outcomes for an unconstrained designer who can target all players. That is, we want to characterize $X_N(G, E)$ for every basic game G and network E. Once the focus becomes the unconstrained problem, we can follow the tradition of the literature on direct mechanisms (in particular, Bergemann and Morris (2016)) and recast the problem *as if* the designer directly recommends to each player *how to play* in G. In general, this can be a pure or a mixed action. To state the following key definition, we extend each player i's utility function to mixed actions in the usual way, abusing notation to write $u_i(\alpha_i, \alpha_{-i}, \omega)$.[10]

Definition 4 (Robust Obedience)**.** The outcome function x is *spillover-robust obedient* for (G, E) if, for each $i = 1, \ldots, N$ and $\alpha_{E_i} \in \mathbf{x}_{E_i}$,

$$\sum_{\substack{\omega\in\Omega \\ \alpha_{-E_i}\in\mathbf{x}_{-E_i}}} [u_i(\alpha_i, \alpha_{-i}; \omega) - u_i(a_i, \alpha_{-i}; \omega)]x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i \in A_i.^{11}$$

To interpret this definition, imagine to divide both sides of the inequality by the total probability that the profile of recommendations α_{E_i} arises under x and μ. The left-hand side becomes the difference between player i's expected utilities from action α_i and from a_i, both *conditional* on learning the realization of α_j for himself *and* his sources under x. This realization may convey information about the α of the remaining players as well as the state. Thus, the definition says that, conditional on all this information, player i prefers to follow his recommendation to any other action.

We can now state our second main result. If the designer's recommendations are spillover-robust obedient (in short, robust obedient), no player will have an incentive to deviate from his, *irrespective* of the information spillovers occurring after all recommendations are released. Thus, robust obedience characterizes all outcomes the designer can implement—by directly recommending his behavior to each player—taking into account how information spillovers influence incentives.

[10]That is, $u_i(\alpha_i, \alpha_{-i}, \omega)$ is the expected payoff for player i in state ω, where the expectation is taken with respect to the distribution over action profiles given by (α_i, α_{-i}).

[11]Throughout the paper, we abuse the notation $\sum$ and write, $\sum_{\omega,\alpha_{N'}} u_i(\alpha_i, \alpha_{-i}; \omega)x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega)$ in place of $\sum_{\omega\in\Omega,\alpha_{N'}\in\mathbf{x}_{N'}} u_i(\alpha_i, \alpha_{-i}; \omega)x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega)$ for any $N' \subset N$.

Theorem 2 (Unconstrained Feasibility)**.** *The outcome function x is N-feasible for (G, E) (i.e., $x \in X_N(G, E)$) if and only if it is robust obedient for (G, E).*

The intuition is simple. Suppose π and a BNE σ induce x. Note that by learning the signals of his sources through the network, player i also learns the signals of *his sources'* sources. Knowing σ (by the equilibrium assumption), player i can then predict the mixed behavior of all his sources. By definition, in equilibrium he must best respond to this behavior, as well as to his belief about all other players' behavior and the state. But this is robust obedience. Conversely, suppose x is robust obedient. First, we can interpret x itself as an information structure. By robust obedience, it is a BNE of $(G, f_E(x))$ for each player to follow his recommendation, given what he learns through information spillovers and that the others follow their recommendations.

The novelty of our result is that, to account for information spillovers, recommendations have to be *robust* to communication between players as described by obedience. Thus, robust obedience captures the basic economic trade-off caused by information spillovers: As usual, the designer tries to directly influence the beliefs of each player, but now she also has to worry that her message for one player may alter his followers' beliefs. To see this, suppose there are two players, who are unconnected (i.e., $N = 2$ and $E = \{(1, 1), (2, 2)\}$). Then, the designer can span the entire spectrum from conveying directly to each player what the other will do, to keeping each player in the total dark about his opponent's behavior. Keeping player i uncertain about j's behavior can help relax i's incentives to choose specific actions. Now add to E a link from player 1 to player 2 (but not vice versa), so that the follower 2 will always learn the recommendation to his source 1. The entire spectrum mentioned before continues to be feasible with regard to player 1. With regard to player 2, instead, while before the designer could keep 2 uncertain about 1's action as much as allowed by its dependence on her recommendation to 2 himself and the state, now she can do so only for 1's behavior whose randomness is independent of both 2's recommendation and the state.

In short, robust obedience highlights that information spillovers remove the designer's ability to keep the followers of a player uncertain about his behavior when and *only* when she desires to link this behavior to the state or the behavior of others. While it remains possible to implement uncorrelated outcomes that are independent of the state, it becomes harder to implement outcomes that require some dependence on the state as well as mutual uncertainty among players.

Different target sets and communication networks entail different restrictions on reaching the players and patterns of information spillovers. However complex and rich these may be, by Theorem 1 and 2 we can always express the actual constraints imposed on the designer in terms of *linear* inequalities. This is operationally useful for two reasons, which will be developed later. First, it opens the door to tackling the designer's problem using powerful linear-programming methods—this is known for standard information design (Bergemann and Morris (2016)), but is a priori not obvious with target restrictions and information spillovers. Second, Theorem 2 uncovers the underlying structure of the designer's problem, so that we can study how the set of feasible outcomes varies across networks (see below).

Considering two extreme, though important, cases helps us further illustrate Theorem 2. First, suppose the communication network is empty (more precisely, $E = \{(i,i)\}_{i\in N}$). As one may expect, this reduces our model to that in Bergemann and Morris (2016). Since information spillovers are shut down, private signals never flow from one player to another. Thus, the designer does not have to worry about the effects that information conveyed to some player may *directly* have on others. The condition in Definition 4 simplifies to, for all i and $\alpha_i \in \mathbf{x}_i$,

$$\sum_{\substack{\omega\in\Omega \\ \alpha_{-i}\in\mathbf{x}_{-i}}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)] x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i \in A_i.$$

This condition is equivalent to the notion of obedience introduced by Bergemann and Morris (2016), which characterizes Bayes Correlated Equilibria (BCE). Indeed, the inequality can be written as

$$\sum_{\substack{\omega\in\Omega \\ \alpha_{-i}\in\mathbf{x}_{-i}}} \Big\{ \sum_{a\in A} [u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega)] \prod_{j\in N} \alpha_j(a_j) \Big\} x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i' \in A_i.$$

If for every $a \in A$ we define $x'(a|\omega) = \sum_{\alpha_{-i}\in\mathbf{x}_{-i}} \prod_{j\in N} \alpha_j(a_j) x(\alpha_i, \alpha_{-i}|\omega)$, we can express the latter condition as, for all i and $a_i \in A_i$,

$$\sum_{\omega\in\Omega, a\in A} [u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega)] x'(a_i, a_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i' \in A_i.$$ [12]

The other extreme case is the *complete* network (i.e., $E = N^2$). Now each player is a source for every other player (i.e., $E_i = N$ for all i). Thus, it is as if, after receiving his private signal, each player announces it publicly to everybody else. In terms of (obedient) recommendations, this means that each player i learns the profile α_{-i} everybody else is playing. This reduces the obedience condition to

$$\sum_{\omega\in\Omega} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)] x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i \in A_i,$$

[12] Some reader may recall that in Bergemann and Morris (2016) the designer recommends only pure actions, not mixed actions. This restriction is without loss of generality in their setting, but not in ours. When information spillovers are possible, in fact, the revelation principle needs to be slightly adapted to become useful. We will return to this point in Section 5.

for all i and $(\alpha_i, \alpha_{-i}) \in \mathbf{x}$. From this perspective, obedient outcome functions for complete networks can be viewed as public information structures, namely, as structures that allow each player to perfectly predict all the others' signals just from his private signal.

The constraints defining obedience can also be viewed as a combination of ex-ante and ex-post requirements, based on what each player knows upon taking his action. To see this, we can express all constraints in ex-ante form, but taking into account the information on which players can act. That is, obedience is equivalent to requiring that, for every i and $\delta_i : R_{E_i} \to A_i$,

$$\sum_{\substack{\omega\in\Omega \\ \alpha\in\mathbf{x}}}[u_i(\alpha_i, \alpha_{-i}; \omega) - u_i(\delta_i(\alpha_{E_i}), \alpha_{-i}; \omega)]x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \tag{2}$$

This formulation highlights that, due to information spillovers, players can base their deviations on richer information sets. In the extreme cases discussed above, for empty networks δ_i can depend only on player i's recommendation α_i; for complete networks δ_i can depend on the recommendations to everybody. These observations hint at a later result showing that when the communication network becomes richer (in a sense to be defined below), it shrinks the set of feasible outcomes, as it enlarges the set of possible deviations for each player.

Importantly, in our setting the players' richer information sets are statistically *and* physically interdependent from the designer's viewpoint. The private signal of one player eventually enters the information set of other players. Therefore, the designer has to take into account the direct effect of the same signal on multiple players, whose preferences can be very different. This is on top of the usual information that a private signal may convey about other players' signals and hence behavior, through its statistical dependence on those signals and the state. While the designer fully controls the latter—as in Bergemann and Morris (2016)—here she has no control on whether one player's private signal enters his followers' information sets. A consequence of this is that we have to allow the designer to recommend mixed actions. While Section 5 will clarify why, this grants the designer the possibility to keep i's followers uncertain about i's behavior at least by incentivizing i to randomize. However, such randomizations are by definition independent of others' randomizations and the state.

3.3 Network Depth and Information Aggregation

This section uses Theorem 2 to investigate how the set of feasible outcomes changes with the communication network. To this end, we introduce the following order on networks.

Definition 5. The communication network E is *deeper* than the communication network E' if, for all $i \in N$, i's followers in E' are also followers in E (i.e., ${}_iE' \subseteq {}_iE$).

This notion is more subtle than simply saying that E has "more links" than E'. In fact, this weaker order cannot rank communication networks in terms of how they affect the set of feasible outcomes, as the next proposition shows.

Proposition 1. *$X_N(G, E) \subseteq X_N(G, E')$ for all G if and only if E deeper than E'.*

Intuitively, a deeper communication network constrains more the designer's ability to prevent "local" information from spreading "globally" in the network. While the designer can always replicate information spillovers by essentially telling each player what he may learn from his sources, she cannot undo spillovers that start from a more localized information provision. This asymmetry determines the set inclusion in Proposition 1. More broadly, this result demonstrates that the two extreme cases of empty and complete networks are also *extreme* with respect to the set of feasible outcomes. The complete network is clearly deeper than any other network, which is in turn deeper than the empty network. Thus, the constraints imposed by information spillovers are maximal in one case and minimal in the other (i.e., $X_N(G, N^2) \subseteq X_N(G, E) \subseteq X_N(G, \varnothing)$). The following is an immediate implication of Proposition 1.

Corollary 1. *If E is deeper than E', then the designer is weakly worse off under E than under E' (i.e., $V^*_E(N) \leq V^*_{E'}(N)$).*

Proposition 1 also sheds light on how to measure the informational influence of a player. It shows that counting the number of his direct neighbors would not provide a satisfactory measure of influence. This measure should instead take into account the *global* informational impact that a player has in the communication network: This includes not only his neighbors, but also all players who indirectly follow him. While this influence measure may seem demanding and—perhaps—difficult to compute in large network, the "only if" part of our result shows that it is the right measure to use in order to understand how information spillovers per se constrain the designer.

To appreciate the subtle aspects of our depth order, note that if E is *more connected* than E' in the sense that $E' \subseteq E$—namely, E contains all the same links as does E' and possibly more—then E is also deeper than E'. The converse is not true, however. For example, suppose that $N = \{1, 2, 3, 4\}$, $E = \{(1, 2), (1, 3), (2, 3)\}$, and $E' = \{(1, 2), (2, 3), (4, 2)\}$. Then, neither network is contained in the other, yet E' is deeper than E. Despite this, the difference is immaterial for our purposes.

Remark 1. *Suppose E is deeper than E'. Let $\hat{E} = E \cup E'$. Then, ${}_iE = {}_i\hat{E}$ for all i.*[13]

Therefore, if E is deeper than E', we can without loss of generality assume that $E' \subseteq E$.

By determining how local spillovers become global phenomena, the depth of a network also controls its capacity to aggregate information at the social level. This connection between depth and information aggregation can be formalized as follows. Recall that spillovers through the network E transform every initial information structure π into a final information structure $f_E(\pi)$. Different networks may transform the same initial π into different information structures, which one may try to order by how informative they are for the players. To this end, we will say that $\pi \in \Pi$ is more informative than $\pi' \in \Pi$ for player i if the set of joint distributions between his action a_i and the state ω that i can achieve under π' is a subset of that under π. Formally, define $\Delta^\pi(\Omega \times A_i)$ as the set of joint distributions, denoted by y, such that

$$y(\omega, a_i) = \sum_{t_i, t_{-i}} \gamma(a_i|t_i)\pi(t_i, t_{-i}|\omega)\mu(\omega), \quad (\omega, a_i) \in \Omega \times A_i,$$

for some function $\gamma : T_i \to \Delta(A_i)$. Given this, π is more informative than π' for player i if $\Delta^{\pi'}(\Omega \times A_i) \subseteq \Delta^\pi(\Omega \times A_i)$.

Definition 6 (Information Aggregation)**.** E aggregates more information than does E' if, for all $\pi \in \Pi$, $f_E(\pi)$ is more informative than $f_{E'}(\pi)$ for every players.

Proposition 2. *Fix N. E is deeper than E' if and only if E aggregates more information than E'.*

Propositions 1 and 2 can be directly applied to unconstrained design problems, but they are also useful for constrained problems by Theorem 1. Consider two constrained problems with networks $\overline{E}$ and $\hat{E}$ and the same target set M. If $\overline{E}^M$ is deeper than $\hat{E}^M$, then $X_M(G, \overline{E}) \subseteq X_M(G, \hat{E})$ as well as $V^*_{\overline{E}}(M) \leq V^*_{\hat{E}}(M)$. Given this, it would be useful to know what properties of $\overline{E}$ and $\hat{E}$ determine that $\overline{E}^M$ is deeper than $\hat{E}^M$. It is not enough that $\overline{E}$ is deeper than $\hat{E}$. To see this, suppose there are three players, $N = \{1, 2, 3\}$, and $M = \{1, 2\}$. Let $\overline{E} = \{(1, 3)\}$ and $\hat{E} = \{(1, 3), (2, 3)\}$. Clearly, $\hat{E}$ is deeper than $\overline{E}$. In this case, $\hat{E}^M = \{(1, 3), (3, 1)\}$ and $\overline{E}^M = \overline{E}$, so neither is deeper than the other. The next result identifies the precise condition under which ranking of the original networks coincides with the same ranking of their M-expansions.

[13] The argument is simple. Clearly, ${}_iE \subseteq {}_i\hat{E}$. Now, consider arbitrary i and j such that $(i, j) \in E'$, but $(i, j) \notin E$. By construction, $j \in {}_i\hat{E}$. Because E is deeper than E', $j \in {}_iE' \subseteq {}_iE$. Since j was arbitrary, ${}_iE \supseteq {}_i\hat{E}$.

Proposition 3. *Suppose E is deeper than $\hat{E}$. Then, E^M is deeper than $\hat{E}^M$ if and only if for all $i \notin M$ and $i \notin E_j$ the following holds:*

$$\hat{E}_i \cap M \subseteq \hat{E}_j \Rightarrow E_i \cap M \subseteq E_j.$$

The intuition is as follows. Suppose i is not a targeted player. If in $\hat{E}$ all targeted sources of player i are also targeted sources of player j, then j knows i's information. This should be true also in E^M for it to be deeper. So, in E, either i is already a source of j, or all targeted sources of i must again be targeted sources of j.[14]

Another interesting comparative statics is in terms of the target set. Using Theorem 1 and Proposition 1, we immediately obtain the following.

Corollary 2. *Fix E and consider arbitrary target sets M and M' that satisfy Assumption 2. Then, $X_M(G, E) \subseteq X_{M'}(G, E)$ if and only if E^M is deeper than $E^{M'}$.*

That is, M constrains the designer more than does M' if and only if M leads to a deeper network expansion than M'. For instance, if $M \subseteq M'$, we certainly have $X_M(G, E) \subseteq X_{M'}(G, E)$. Proposition 1 helps us understand how having a smaller target set constrains the designer more. The smaller set forces the designer to rely more on information intermediation—and hence information spillovers—to influence the players she cannot target. This, in turn, tightens the obedience constraints and hence reduces the feasible outcomes. More interestingly, even if M and M' are neither a subset of the other, Proposition 1 still implies that M constrains the designer more than does M' if and only if their network expansions can be ranked in terms of depth. The next result allows us to predict this ranking from the primitives of the problem, similarly to Proposition 3.

Proposition 4. *Fix E and consider arbitrary target sets M and M' that satisfy Assumption 2. Then, E^M is deeper than $E^{M'}$ if and only if for all $i \notin E_j$*

$$E_i \cap M' \subseteq E_j \Rightarrow E_i \cap M \subseteq E_j.$$

The intuition is similar to that for Proposition 3.

4 Optimal Manipulation

4.1 A Dual Approach

We start by formulating the designer's problem given G and $E^M = E$ as a linear program. We follow here the dual approach in Galperti and Perego (2018) for information-design

[14] Based on Remark 1, we can consider the following analog of Proposition 3: If $\hat{E} \subseteq E$, then $\hat{E}^M \subseteq E^M$ if and only if $\hat{E}_i \cap M \subseteq \hat{E}_j$ implies $E_i \cap M \subseteq E_j$ for all $(i, j) \notin E$.

problems without spillovers and target restrictions, adapting it accordingly. Given the prior μ and any outcome function x, define the joint distribution over mixed-action profiles and states that they induce:

$$\chi(\alpha,\omega) = x(\alpha|\omega)\mu(\omega), \quad \alpha \in R, \omega \in \Omega.$$

Since by assumption $x(\cdot|\omega)$ has finite support for every ω, so does χ overall. The problem is then to choose a finite-support χ that solves

$$V_E^*(N) = \sup_{\chi} \sum_{\omega,\alpha} v(\alpha,\omega)\chi(\alpha,\omega)$$

subject to three sets of constraints:

- *Obedience*: for every $i \in N$ and $\alpha_{E_i} \in \operatorname{supp} \chi_{E_i}$,

$$\sum_{\omega,\alpha_{-E_i}} [u_i(\alpha_i,\alpha_{-i},\omega) - u_i(a_i,\alpha_{-i},\omega)]\chi(\alpha_i,\alpha_{-i},\omega) \geq 0, \quad a_i \in A_i; \tag{3}$$

- *Prior consistency*: for every $\omega \in \Omega$,

$$\sum_{\alpha} \chi(\alpha,\omega) = \mu(\omega); \tag{4}$$

- *Positivity*: for every α and ω,

$$\chi(\alpha,\omega) \geq 0.$$

Denote this primal problem by $\mathcal{P}$. Note that there exist χ's that satisfies the constraints of $\mathcal{P}$. This is because there exists at least one BNE following a completely uninformative information structure, and its equivalent representation in terms of recommendations has to satisfy all the constraints of $\mathcal{P}$.

Compared to standard information design, $\mathcal{P}$ raises some complications. In the standard problem it is without loss of generality to focus on pure-action recommendations. Consequently, the dimension of χ and the resulting set of constraints is exogenously determined: $\chi \in \mathbb{R}^{A\times\Omega}$. By contrast, in our case we have to allow for mixed-action recommendations. Thus, even if we know that the support of χ has finitely many slots, there are infinitely many ways to fill each with mixed recommendations. For these reasons, it helps to express obedience in the following equivalent form: For every $i \in N$, $\alpha_{E_i} \in \operatorname{supp} \chi_{E_i}$, and $a_i, a_i' \in A_i$,

$$\sum_{\omega,\alpha_{-E_i}} [u_i(a_i,\alpha_{-i},\omega) - u_i(a_i',\alpha_{-i},\omega)]\alpha_i(a_i)\chi(\alpha_i,\alpha_{-i},\omega) \geq 0.^{15} \tag{5}$$

[15] The term $\alpha_i(a_i)$ inside the brackets of (5) is superfluous when strictly positive, but its is a simple way to handle the case of $\alpha_i(a_i) = 0$ so that we can write a condition that always involves all a_i's.

This form highlights that for each player obedience ultimately involves the primitive pure actions. For a player to be willing to implement a specific mixed action, he must first deem all pure actions in its support optimal given his information. If this holds, he will be willing to implement any mixed action with that support (or a subset thereof). Expression (5) also highlights that α_{E_i} is really only an information set. Therefore, player i has to be willing to play a_i whenever called upon doing it—as in Bergemann and Morris (2016)—except that now this has to be the case conditional on also knowing α_j for $j \in E_i$. Thus, there is not a single obedience constraint for each a_i, but potentially multiple depending on this conditioning event.

Leveraging expression (5), we can show that it is without loss of generality to impose a finite, *exogenous*, upper bound on the dimension of χ.

Lemma 1. *Suppose $x \in X_N(G, E)$. There exists $x' \in X_N(G, E)$ such that $|\mathbf{x}'_i| \leq |2^{A_i}|$ for every i and x' induces the same joint distribution over $A \times \Omega$ as does x:*

$$\sum_{\alpha' \in \mathbf{x}'} \alpha'(a) x'(\alpha'|\omega)\mu(\omega) = \sum_{\alpha \in \mathbf{x}} \alpha(a) x(\alpha|\omega)\mu(\omega), \quad a \in A, \omega \in \Omega.$$

This lemma has two implications. On the one hand, it is without loss of generality to restrict attention to the subset of Π which contains only information structures with the property that $|S_i| \leq |2^{A_i}|$ for all i.[16] On the other hand, for analytical convenience we can allow for outcome functions with unbounded supports, knowing that this does not fictitiously inflate the designer's payoff (which ultimately depends on the induced distribution over $A \times \Omega$).

Despite the possibility of bounding the support of outcome functions, the elements of the support are profiles of mixed actions, which form an uncountable space. Therefore, extistance of a solution remains a non-trivial matter.

Proposition 5 (Existence)**.** *There exists χ such that*

$$\sum_{\omega, \alpha} v(\alpha, \omega)\chi(\alpha, \omega) = V^*_E(N).$$

Given this, we can proceed to characterize the solutions of our problem.

As shown in Galperti and Perego (2018), it is instructive to consider the dual of information-design problems. This is especially true for the richer problems considered here. For reasons that will become clear shortly, consider any finite grid of action profiles

[16] It is worth noting the difference from standard information-design problem a lá Bergemann and Morris (2016) where it is without loss of generality to restrict attention to information structure with the property that $|S_i| \leq |A_i|$.

$R_k \subset R$ (where $k = |R_k|$). Let $\mathcal{P}(R_k)$ be the primal problem restricted to outcome functions whose support must be contained in R_k. To state its dual, denoted by $\mathcal{P}^*(R_k)$, define two sets of variables. The first contains the dual variables for the obedience constraints and are denoted by λ. That is, $\lambda(a_i, a_i'|\alpha_{E_i})$ corresponds to the obedience constraint for player i when he is recommended to play α_i in the profile α_{E_i} and he contemplates deviating from a_i to a_i'. The second set contains the dual variables for the prior-consistency condition and are denoted by p. That is, $p(\omega)$ corresponds to the condition $\sum_\alpha \chi(\alpha, \omega) = \mu(\omega)$. The proof of the next lemma follows similar arguments as in Galperti and Perego (2018) and is therefore omitted.

Lemma 2. *Fix a grid R_k. The dual information-design problem $\mathcal{P}^*(R_k)$ consists of choosing $p \in \mathbb{R}^\Omega$ and $\lambda_i(a_i, a_i'|\alpha_{E_i})$ for every $a_i, a_i' \in A_i$, $\alpha \in R_k$, and $i \in N$ so as to minimize*

$$\sum_{\omega \in \Omega} p(\omega)\mu(\omega),$$

subject to, for all $i \in N$, $\alpha \in R_k$, and $a_i, a_i' \in A_i$

$$\lambda_i(a_i, a_i'|\alpha_{E_i}) \geq 0,$$

and for all $(\alpha, \omega) \in R_k \times \Omega$

$$p(\omega) \geq \sum_{a \in A} \left\{ v(a, \omega) + \sum_{i=1}^{N} \sum_{a_i' \in A_i} [u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega)] \lambda_i(a_i, a_i'|\alpha_{E_i}) \right\} \alpha(a). \quad (6)$$

As well known, the solutions to $\mathcal{P}(R_k)$ and $\mathcal{P}^*(R_k)$ are related to each other by complementary slackness conditions, which follow:

CS1. For all $i \in N$, $\alpha \in R_k$, and $a_i, a_i' \in A_i$,

$$\lambda_i(a_i, a_i'|\alpha_{E_i}) \left\{ \sum_{\omega, \alpha_{-E_i}} [u_i(a_i, \alpha_{-i}, \omega) - u_i(a_i', \alpha_{-i}, \omega)] \alpha_i(a_i) \chi(\alpha, \omega) \right\} = 0;$$

CS2. For all $\omega \in \Omega$,

$$p(\omega) \left\{ \sum_\alpha \chi(\alpha, \omega) - \mu(\omega) \right\} = 0;$$

CS3. For all $(\alpha, \omega) \in R_k \times \Omega$,

$$\chi(\alpha, \omega) \left\{ p(\omega) - \sum_{a \in A} \left\{ v(a, \omega) + \sum_{i=1}^{N} \sum_{a_i' \in A_i} [u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega)] \lambda_i(a_i, a_i'|\alpha_{E_i}) \right\} \alpha(a) \right\} = 0.$$

Of course, the challenge is that a priori we do not know which grid of action profiles allows us to solve either the primal or the dual. However, we can at least partially alleviate this issue using the dual and the complementary slackness conditions. Recall that if $\mathcal{P}$ has a solution, then so does its dual. In particular, as shown in the proof of Proposition 5, there exists an optimal R_k that leads to an overall optimal solution χ_{R_k} of $\mathcal{P}$ and an overall optimal solution (p_{R_k}, λ_{R_k}) of its dual $\mathcal{P}^*$. By Strong Duality, the value of the two problems must be the same. More importantly, by Weak Duality for every χ_{R_k} that satify the constraints of $\mathcal{P}$ and (p_{R_k}, λ_{R_k}) that satisfies the constraints of $\mathcal{P}^*$, we have

$$\sum_{\omega\in\Omega} p_{R_k}(\omega)\mu(\omega) \geq \sum_{\omega\in\Omega,\alpha} v(\alpha,\omega)\chi_{R_k}(\alpha,\omega).$$

By Lemma 1 and Proposition 5, there exists a finite $\overline{k}$ such that the largest value of the right-hand side of this inequality equals $V_E^*(N)$ for all $k \geq \overline{k}$ and so

$$\sum_{\omega\in\Omega} p_{R_k}(\omega)\mu(\omega) \geq V_E^*(N),$$

with equality for some feasible (p_{R_k}, λ_{R_k}). Importantly, this does not change if we let k grow without bounds. Note that, in terms of the dual this does not change its objective, as the state space remains fixed.

Thus, the complexity of identifying the grid is removed, as we can now study the following relaxed dual to find $V_E^*(N)$, which allows for all action profiles. The relaxed dual information-design problem $\mathcal{P}^*$ consists of choosing $p \in \mathbb{R}^\Omega$ and $\lambda_i(a_i, a_i'|\alpha_{E_i})$ for every $a_i, a_i' \in A_i$, $\alpha \in R$, and $i \in N$ so as to minimize

$$\sum_{\omega\in\Omega} p(\omega)\mu(\omega),$$

subject to, for all $i \in N$, $\alpha \in R$, and $a_i, a_i' \in A_i$

$$\lambda_i(a_i, a_i'|\alpha_{E_i}) \geq 0,$$

and for all $(\alpha,\omega) \in R \times \Omega$

$$p(\omega) \geq \sum_{a\in A} \left\{ v(a,\omega) + \sum_{i=1}^{N} \sum_{a_i'\in A_i} [u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega)]\lambda_i(a_i, a_i'|\alpha_{E_i}) \right\} \alpha(a). \quad (7)$$

Finally, consider condition (CS3) for the relaxed problem: For all $(\alpha,\omega) \in R \times \Omega$,

$$\chi(\alpha,\omega) \left\{ p(\omega) - \sum_{a\in A} \left\{ v(a,\omega) + \sum_{i=1}^{N} \sum_{a_i'\in A_i} [u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega)]\lambda_i(a_i, a_i'|\alpha_{E_i}) \right\} \alpha(a) \right\} = 0.$$

It is easy to see that, for every α, $\chi(\alpha,\omega)$ can be positive only if it is possible to find λ's that render the term multiplying $\chi(\alpha,\omega)$ equal to zero. Otherwise, $\chi(\alpha,\omega) = 0$. Also, note that equation (7) can impose a lower bound on $p(\omega)$. For instance, one can easily see that if the set of pure-strategy Nash equilibria of the game identified by ω (denoted by $NE(G_\omega)$) is not empty, then

$$p(\omega) \geq \max_{a \in NE(G_\omega)} v(a, \omega).$$

These steps can help us identify the support of an optimal χ.[17]

In the companion paper Galperti and Perego (2018), we offer an economic interpretation of the dual that can be easily extended to the present contexts. This interpretation offers insights into the trade-offs of designing information structures for games. One important difference introduced by information spillovers involves CS1. It states that the designer can set $\lambda_i(a_i, a_i'|\alpha_{E_i}) > 0$ only if, given the information conveyed by α_{E_i}, player i is *indifferent* between a_i and a_i'. Setting $\lambda_i(a_i, a_i'|\alpha_{E_i}) > 0$ may help the designer relax (7) so as to achieve a lower dual objective. Thus, CS1 conveys a general principle on how to optimally design information, as already noted in Galperti and Perego (2018). Clearly, the stronger the information spillovers are, the harder it is to satisfy this indifference condition. So, CS1 emerges as *the* condition where we can see how information spillovers complicate the problem of designing an optimal information structure.

In Galperti and Perego (2018), we also derive some general properties of optimal solutions that can be easily translated to the problem with target restrictions and informaiton spillovers through its dual. We also provide a necessary condition for the solution to take the form of full information about the state as well as the action of everybody. Since such information structure is public, information spillovers become irrelevant and the condition applies directly to the present settings.

In general, we view duality not as a substitute of the primal approach, rather as a complement. The analyst should combine primal, dual, complementary slackness conditions, and Strong Duality to enhance the understanding of information-design problems. We now illustrate these methods with an example.

4.2 An Illustrative Example: Investment Games

This example is borrowed from Bergemann and Morris (2018). Its goal is to illustrate the dual approach to the information-design problems considered in the present paper.

[17] Another possible solution strategy is to conjecture that the optimal χ never recommends mixed actions to any player. One can solve the dual under this conjecture, use the CS conditions to derive χ, and then check whether the value of the primal and the dual coincide. If so, Strong Duality implies that we have a solution. Otherwise, the solution must rely on mixed-action recommendations. The example below illustrates this.

The basic game is as follows. Two firms have to choose whether to invest (y) in a project or not (n): $A_i = \{y, n\}$ for $i \in N$. The project can be either good or bad: $\Omega = \{g, b\}$. Firm i's payoffs $u_i(a_1, a_2, \omega)$ are described in Table 1. The designer (the "government") wants to foster investment, irrespective of the quality of the project:

$$v(a_1, a_2) = \mathbb{I}\{a_1 = y\} + \mathbb{I}\{a_2 = y\},$$

where $\mathbb{I}\{\cdot\}$ is the indicator function.

$\omega = b$ (Firm 1: rows; Firm 2: columns)

	y	n
y	$\varepsilon - 1, \varepsilon - 1$	$-1, 0$
n	$0, -1$	$0, 0$

$\omega = g$ (Firm 1: rows; Firm 2: columns)

	y	n
y	$\varepsilon + q, \varepsilon + q$	$q, 0$
n	$0, q$	$0, 0$

Table 1: Firms' payoffs in G

The prior satisfies $\mu(g) = \mu(b) = \frac{1}{2}$ and $0 < q < 1$. Thus, without additional information neither firm wants to invest. If $\varepsilon > 0$, the game features strategic complementarities; if $\varepsilon < 0$, the game features strategic substitutabilities. Finally, ε is small so that y is dominant in state g and n is dominant in state b: $|\varepsilon| \leq q - \frac{1}{2}$.

Bergemann and Morris (2018) derived the optimal information structure when there are no information spillovers (see also Galperti and Perego (2018) for a derivation using duality). Table 2 presents the optimal x for the case of $\varepsilon > 0$, which features public information. Intuitively, this is because with strategic complementarities each firm is more willing to invest when recommended y knowing that, if the state is b, also the other firm will invest, which reduces the loss of the bad decision. This higher willingness to invest allows the designer to pool more the unfavorable state b with the favorable g in the recommendation y, thereby increasing the overall chances of investment. Table 3 presents the optimal x for the case of $\varepsilon < 0$, which features private information. Intuitively, this is because with strategic substitutabilities instead each firm is more willing to invest when recommended y knowing that, if the state is b, the other firm will not invest.

Now, suppose that there are information spillovers and the designer is unrestricted in terms of the firms she can target. To consider an interesting case, suppose that E has a link from firm 1 to firm 2 (but not vice versa): $E_1 = \{1\}$ and $E_2 = \{1, 2\}$. It is immediate to see that the x in Table 2 remains optimal in the case of strategic complementarities. This is because public information is not affected by information spillovers. Therefore,

		Firm 2 y	Firm 2 n
Firm 1	y	$\frac{q+\varepsilon}{1-\varepsilon}$	0
	n	0	$1-\frac{q+\varepsilon}{1-\varepsilon}$

$x(a_1, a_2|\omega = b)$

		Firm 2 y	Firm 2 n
Firm 1	y	1	0
	n	0	0

$x(a_1, a_2|\omega = g)$

Table 2: Optimal x for strategic complements (i.e., $\varepsilon > 0$)

		Firm 2 y	Firm 2 n
Firm 1	y	0	$q+\varepsilon$
	n	$q+\varepsilon$	$1-2q-2\varepsilon$

$x(a_1, a_2|\omega = b)$

		Firm 2 y	Firm 2 n
Firm 1	y	1	0
	n	0	0

$x(a_1, a_2|\omega = g)$

Table 3: Optimal x for strategic substitutes (i.e., $\varepsilon < 0$)

it remains feasible and optimal. By contrast, x in Table 3 is no longer feasible. When x produces the recommendations (n, y) in state b, firm 2 learns that firm 1 is recommended not to invest and so that the state must be b. Consequently, it is no longer obedient for firm 2 to follow the recommendation to invest.

We can solve for the new optimal x for the latter case using our duality approach. For simplicity, we conjecture that it is enough to consider only pure-action recommendations: (y, y), (y, n), (n, y), and (n, n). We will derive a candidate solution under this conjecture and then use Strong Duality to prove its optimality.

We start by calculating the right-hand side of the dual constraint (7). Let λ be the vector of all dual variables of the form $\lambda_1(a_1, a_1'|a_1)$ and $\lambda_2(a_2, a_2|a)$ for all $a \in A$. Also, for every a, ω, and λ, define

$$\begin{aligned} w(a, \omega, \lambda) \quad = \quad & v(a) + [u_1(a_1, a_2, \omega) - u_1(a_1', a_2, \omega)]\lambda_1(a_1, a_1'|a_1) \\ & +[u_2(a_1, a_2, \omega) - u_2(a_1, a_2', \omega)]\lambda_2(a_2, a_2'|a). \end{aligned}$$

We have

$$w(a,g,\lambda) = \begin{cases} 2 + [q+\varepsilon]\lambda_1(y,n|y) + [q+\varepsilon]\lambda_2(y,n|y,y) & \text{if } a = (y,y) \\ 1 + q\lambda_1(y,n|y) - [q+\varepsilon]\lambda_2(n,y|y,n) & \text{if } a = (y,n) \\ 1 - [q+\varepsilon]\lambda_1(n,y|n) + q\lambda_2(y,n|n,y) & \text{if } a = (n,y) \\ -q\lambda_1(n,y|n) - q\lambda_2(n,y|n,n) & \text{if } a = (n,n), \end{cases}$$

$$w(a,b,\lambda) = \begin{cases} 2 - [1-\varepsilon]\lambda_1(y,n|y) - [1-\varepsilon]\lambda_2(y,n|y,y) & \text{if } a = (y,y) \\ 1 - \lambda_1(y,n|y) + [1-\varepsilon]\lambda_2(n,y|y,n) & \text{if } a = (y,n) \\ 1 + [1-\varepsilon]\lambda_1(n,y|n) - \lambda_2(y,n|n,y) & \text{if } a = (n,y) \\ \lambda_1(n,y|n) + \lambda_2(n,y|n,n) & \text{if } a = (n,n). \end{cases}$$

We first refine the set of candidate solutions for the primal and the dual with a sequence of claims.

Claim 1. *$p(g) \geq 2$ and $p(b) \geq 0$.*

Proof. Note that (y,y) is the unique Nash equilibrium of G_g and (n,n) is the unique Nash equilibrium of G_b. Therefore, $p(g) \geq v(y,y) = 2$ and $p(b) \geq v(n,n) = 0$. □

Claim 2. *$\chi(n,n,g) = 0$.*

Proof. Since $p(g) \geq 2$ and $w((n,n),g,\lambda) \leq 0$ for all λ, the claim follows from CS3. □

Claim 3. *No information and full information are not a solution.*

Proof. The lower bounds for $p(g)$ and $p(b)$ imply that the value of the problem, V_E^*, cannot be zero. This means that providing the firms with no information is not optimal, because that would induce the outcome (n,n) for sure, contradicting $\chi(n,n,g) = 0$. It is also not optimal to provide full information, that is, $\chi(y,y,g) = \chi(n,n,b) = \frac{1}{2}$. In this case, no firm is ever indifferent between the two actions and so $\lambda \equiv 0$. Given this, CS3 requires that $p(b) = 0$ and $p(g) = 2$. However, note that $w((y,y),b,0) = 2 > p(b)$, which violates the dual constraint. □

Claim 4. *Either $\lambda_1(y,n|y) > 0$ or $\lambda_2(y,n|y,y) > 0$.*

Proof. If $\lambda_1(y,n|y) = \lambda_2(y,n|y,y) = 0$, we have $p(b) \geq w((y,y),b,\lambda) = 2$. This means the value of the dual and hence V_E^* must be at least 2. This is impossible, as the designer cannot get both firms to invest for sure. □

Claim 5. $[q+\varepsilon]\chi(y,y,g) + q\chi(y,n,g) = [1-\varepsilon]\chi(y,y,b) + \chi(y,n,b)$.

Proof. Since firm 2 always has as much information as firm 1, if the latter is never indifferent when recommended y, so is firm 2. This would imply $\lambda_1(y,n|y) = \lambda_2(y,n|y,y) = 0$ in contradiction with Claim 4. Therefore, firm 1 must be rendered indifferent after recommendation y. This requires the following:

$$[q+\varepsilon]\chi(y,y,g) + q\chi(y,n,g) = [1-\varepsilon]\chi(y,y,b) + \chi(y,n,b).$$

□

Claim 6. *$\chi(y,y,g) > 0$ and $p(g) = 2 + [q+\varepsilon]\lambda_1(y,n|y) + [q+\varepsilon]\lambda_2(y,n|y,y) > 2$.*

Proof. The fact that either $\lambda_1(y,n|y) > 0$ or $\lambda_2(y,n|y,y) > 0$ implies that $p(g) > 2$—recall that $|\varepsilon| < q$—and hence $V_E^* > 1$. Therefore, $\chi(y,y,\omega) > 0$ for some ω. This cannot be only $\omega = b$. Otherwise, firm 2 learns that the state is b from the recommendation (y,y) and so it would not follow its recommendation y, violating feasibility. Therefore, we must have $\chi(y,y,g) > 0$. The equation for $p(g)$ follows from CS3. □

Claim 7. *$\lambda_1(n,y|n) = 0$.*

Proof. Suppose that firm 1 weakly prefers y when recommended y, which is necessary for obedience, and is indifferent between the two actions when recommended n, which is necessary for $\lambda_1(n,y|n) > 0$ by CS1. This requires

$$[q+\varepsilon]\chi(y,y,g) + q\chi(y,n,g) \geq [1-\varepsilon]\chi(y,y,b) + \chi(y,n,b)$$

and

$$q\chi(n,y,g) = [1-\varepsilon]\chi(n,y,b) + \chi(n,n,b).$$

Summing the two conditions we get (using $\chi(n,n,g) = 0$)

$$\varepsilon\chi(y,y,g) + q \geq 1 - \varepsilon\chi(y,y,b) - \varepsilon\chi(n,y,b) \geq 1.$$

This is impossible since $q < 1$ and $\varepsilon < 0$. Therefore, we conclude that firm 1 must strictly prefer n when recommended n and so $\lambda_1(n,y|n) = 0$. Note that firm 1 must be recommended n with positive probability: Otherwise, it always receives recommendation y, which is uninformative and so can never be obediently followed by firm 1, contradicting $\chi(y,y,g) > 0$. □

Claim 8. *$\chi(n,n,b) > 0$, so $p(b) = 0$ and $\lambda_2(n,y|n,n) = 0$.*

Proof. We now argue by contradiction. Suppose that $\chi(n,n,b) = 0$. Then, by the previous argument we must have $\chi(n,y,b) > 0$. Otherwise, $\chi(n,y,\omega)$ is positive only for $\omega = g$, which means that firm 1 is recommended n only in state g and so this recommendation cannot be obediently followed. Since $\chi(n,y,b) > 0$, by CS3 we must have

$$0 = p(b) = 1 - \lambda_2(y,n|n,y) \Rightarrow \lambda_2(y,n|n,y) = 1.$$

This implies that

$$w((n,y),g,\lambda) = 1 + q\lambda_2(y,n|n,y) < 2 < p(g).$$

Therefore, we must have $\chi(n,y,g) = 0$ by CS3. This implies that the profile (n,y) is recommended only in state b. Since firm 2 learns this from learning (n,y), it cannot obediently follow the recommendation y. This establishes the desired contradiction. □

Claim 9. $\chi(n,y,b) = \chi(n,y,g) = 0$, *so* $\lambda_2(y,n|n,y) \geq 1$ *as needed by the dual constraint.*

Proof. That $\chi(n,y,b) = 0$ follows from the argument for the previous claim. Using this, we can also show that $\chi(n,y,g) = 0$. If not, firm 2 learns from the recommendation (n,y) that the state must be g. Given this, it strictly prefers y to n and hence we should have $\lambda_2(y,n|n,y) = 0$, which is not compatible with $p(b) = 0 \geq 1 - \lambda_2(y,n|n,y) + [1 - \varepsilon]\lambda_1(n,y|n)$. □

We have narrowed the set of candidate solutions enough to now guess one. We know that firm 2 must be always recommended y in state g. Recommending firm 1 with some probability to not invest in g seems counterproductive, as it does not help to convince firm 2 to invest and is missing the opportunity to induce firm 1 to invest in the favorable state. Therefore, suppose that $\chi(n,y,g) = 0$. Then, we must have $\chi(n,y,b) = 0$ because otherwise the recommendation (n,y) reveals that the sate is b and firm 2 will not follow the recommendation y obediently. In this case, we must have $\chi(y,y,g) = \frac{1}{2}$ and, since we know that $\chi(n,n,b) < \frac{1}{2}$, we must also have $\chi(y,y,b) > 0$. Given this, the required indifference of firm 1 and firm 2 after recommendation (y,y) implies that $\chi(y,y,b) = \frac{1}{2}\left[\frac{q+\varepsilon}{1-\varepsilon}\right]$ (see Claim 5). Thus, the value of the primal is

$$2\chi(y,y,g) + 2\chi(y,y,b) = 1 + \frac{q+\varepsilon}{1-\varepsilon}.$$

Consider now the dual. Its candidate value is

$$\frac{1}{2}p(b) + \frac{1}{2}p(g) = 1 + \frac{1}{2}[q+\varepsilon]\lambda_1(y,n|y) + \frac{1}{2}[q+\varepsilon]\lambda_2(y,n|y,y) > 1.$$

Given this, a candidate solution involves $\lambda_1(y,n|y) = 1$, $\lambda_2(n,y|y,n) = 0$, and

$$0 = p(b) = 2 - [1-\varepsilon]\lambda_1(y,n|y) - [1-\varepsilon]\lambda_2(y,n|y,y) \Rightarrow \lambda_2(y,n|y,y) = \frac{1+\varepsilon}{1-\varepsilon}.$$

This yields the value

$$\frac{1}{2}p(b) + \frac{1}{2}p(g) = 1 + \frac{1}{2}[q+\varepsilon]\left[\frac{1+\varepsilon}{1-\varepsilon} + 1\right] = 1 + \frac{q+\varepsilon}{1-\varepsilon}.$$

By Strong Duality, we have a solution.

Remarkably, the unidirectional information spillovers render public information optimal also for the case of strategic substitutabilities. Note that this is a priori not obvious: In order to preserve some private information, the designer could recommend some firm to play mixed actions conditional on some states. In this example, such recommendations are not useful.

Finally, suppose the designer can only target one firm. Under our reach assumption, there must exist at least a link from the targeted firm to the other firm. This implies that the M-expansion is a complete network, so the designer is essentially constrained to use public information structures. Thus, the optimal structure remains the one in Table 2.

5 Discussion and Extensions

5.1 Unconstrained Designer and Richer Forms of Spillovers

In this section, we consider richer forms of information spillovers than those assumed in the baseline model. We show that for all these forms our previous results provide bounds for the payoff of the information designer when she can target all players (i.e., for the case of $M = N$). This is not true for a constrained designer (i.e., $M \subsetneq N$). However, this bounds result may be useful, at least if one is interested in the metaphorical interpretation of information design as a description of all feasible outcomes across all possible information structures, which by definition impose no restriction on the information each agent can receive.

Information spillovers in a social network can be described in the following general way, which includes our baseline model as a special case. There is a finite number K of communication rounds, where K is at least as large as the shortest path between the two most distant players in the network (i.e., its diameter). At every round, each player i can send (possibly at random) a message to player j if and only if j is a neighbor of i. For every (π, S) let $M_{ij}(\pi, S)$ be the *finite* set of messages that player i can send to his neighbor j at every round. We assume that $S \subseteq M_{ij}(\pi, S)$, so that each player can convey at least as much information as what he can receive from the designer under (π, S). More formally, let the initial history of player i be of the form $h_i^0 = (\pi, s_i)$, where s_i is a realization of π observed by i privately. Thus, the set of initial histories for player i is $H_i^0 = \{(\pi, s_i) : \pi \in \Pi, s_i \in \mathbf{supp}\,\pi\}$.[18] For every round $k \geq 1$, player i also observes

[18]We include π in the history for each player to allow for the possibility that communication behavior

the profiles of messages from his neighbors. Given this, denote i's histories of length k recursively by $H_i^k = \{(h_i^k, m^i) : h_i^{k-1} \in H_i^{k-1}, m^i \in M^i(h_i^0)\}$, where $M^i = \times_{j \in N_i} M_{ji}(h_i^0)$. Let the collection of all i's histories by $\mathcal{H}_i = \cup_{k=0}^K H_i^k$. Player i's communication behavior can then be described as a general function ξ_i that assigns to every history $h_i \in \mathcal{H}_i$ a distributions over messages, i.e., an element of $\Delta(\times_{j \in N_i} M_{ij}(\pi, S))$ for every (π, S) chosen by the designer.

It can be easily seen that every information-spillover process of this form transforms every initial information structure chosen by the designer into a final information structure. Indeed, fix the profile $\xi = (\xi_i)_{i=1}^N$. Then, for every ω, every $\pi \in \Pi$ determines a distribution over finitely many profiles of initial histories $h^0 = (h_i^0)_{i=1}^N$ across the players. For every such profile, ξ induces a distribution over finitely many profiles of finite histories $h^K = (h_i^K)_{i=1}^N$. Interpreting every h_i^K as the signal realization of player i from these compounded distributions, we obtain that every ω induces a finite distribution over such profiles of final signals. This is again an element of Π. Therefore, ξ induces a mapping $f_{\xi,E} : \Pi \to \Pi$ for the network E. Note that, for this to be true, it is not necessary to know where ξ comes from; also, if ξ is commonly known among the players, so is the final information structure $f_{\xi,E}(\pi)$ for every π.[19]

Using this formalism, we can express the designer's problem under general information spillovers. Fixing $f_{\xi,E}$, this problem can be written as

$$V_{\xi,E}^* = \sup_{\pi \in \Pi} V(f_{\xi,E}(\pi)),$$

where recall that

$$V(\pi') = \max_{\sigma \in \mathrm{BNE}(G,\pi')} \sum_{a \in A, t \in T, \omega \in \Omega} v(a,\omega)\sigma(a|t)\pi'(t|\omega)\mu(\omega).$$

Given this, let V_E^* be the value of the designer's problem for the network E under Assumption 1, holding fixed the basic game G.

Theorem 3. *Fix the basic game G and the network E. Let ξ be any profile of communication functions as described above. Then,*

$$V_E^* \leq V_{\xi,E}^* \leq V_\varnothing^*.$$

The result implies that the analysis under our baseline assumption provides a lower bound of the designer's payoff for a large class of communication protocols in social

depends on π and to easily keep track of this dependence.

[19]We conjecture that the assumption that ξ is commonly known among the players can be relaxed without changing our results. Intuitively, it seems that the uncertainty that player i has about the communication behavior of others can be built into ξ itself.

networks. Importantly, while such protocols can be very intricate to describe and analyze, the lower bound of Theorem 2 is amenable to a relatively standard linear-programming characterization. Of course, the implicit assumption here is that ξ describes the players' communication behavior according to some well-defined equilibrium notion. But this seems a minimal requirement for any analysis of communication in social networks.

The rest of this section presents in more detail some of the communication protocols allowed by our formalism. It is also possible that some players in the network communicate according to one rule and others according to a different rule (e.g., some players may be strategic, while others are not).

Truthful Belief Announcement

Suppose that at each round of communication after the designer chooses (π, S), every player truthfully reveals to all his neighbors his current, Bayesian, posterior belief over the space $\Omega \times S$. Note that S is included in the belief announcements, because learning about the information of others is important for predicting their actions in the final game following the communication phase.

It is easy to see that this communication protocol fits into the general model outlined above. Since S is finite, at every round each player can have at most finitely many different beliefs that he can announce. Moreover, this communication process converges after finitely many rounds, whose number cannot exceed the diameter of the network. To see this intuitively, note that in the first round, each player i announces a degenerate belief with regard to his own private signal s_i to his neighbors. In the following round, all i's neighbors announce a degenerate belief about s_i to *their* neighbors. Continuing this way, all followers of player i will hear a degenerate announcement of s_i within a number of rounds that cannot exceed the diameter of the network. Therefore, all players' beliefs will be degenerate about the signals of their sources and so they will stop evolving in at most as many rounds as the network diameter.

The following lemma verifies these intuitive observations. It shows that this communication model can be a foundation for Assumption 1, where we directly assumed that each player i learns the signals received by all players for which there exists a path connecting them to player i.

Lemma 3. *Fix a (S, π) and a signal realization s. For every i and $h_i^K \in \mathcal{H}_i$ consistent with (s, π),*

$$Pr_\pi(\omega, \hat{s}|h_i^K) = Pr_\pi(\omega, \hat{s}|s_{E_i}), \quad (\omega, \hat{s}) \in \Omega \times S.$$

Observational Social Learning

A vast literature in economics studies how people in social environments learn about some underlying state of the world by observing the actions of their peers (see Golub and Sadler (2017) for a literature review). For example, people choosing whether to attend a political rally may learn about its appeal and importance by observing whether their friends are planning to attend it or not. This social-learning process can be modeled through a network framework where initially each of its members receives some signal about the underlying state (e.g., by reading a newspaper article about the rally's lineup of speakers) and then chooses an action (e.g., to attend or not the rally), which is observed by his neighbors in the network. To the extent that the actions of i's neighbors reflect their information—from external sources as well as their peers—they can be interpreted as messages sent to i in our general communication model. Therefore, social learning can be embedded in our framework in the communication phase that precedes the final game. In the example, this game may be a local or national election. The social-learning phase shapes the voters' information about the candidates, which they may then use in the voting booth. Note that their decision to attend political rallies may or may not take into account its information spillovers on peers and any strategic consideration related to the actual election. Either way, the resulting communication falls in the general family covered by Theorem 3.

The approach of this paper and most of the social-learning literature differ in one important respect. We envision a third party who may purposefully design the initial signals that the agents receive before engaging in social learning.[20] In the example, this third party may choose what information to reveal about the rally speakers or agenda and to whom to reveal it given a specific structure of the social connections in a country. By contrast, most of the social-learning literature takes the agents' initial information structure as given and usually assumes for it a very specific form (e.g., independently and identically distributed signals across all agents). We instead impose no restrictions on the information structures the designer can choose. This flexibility, combined with the richness of social-learning processes, may render it hard to provide any prediction on what the designer can achieve. The lower bound in Theorem 3 can prove to be a useful tool to bypass these intricacies.

[20]Of course, an interesting and important extension is to consider multiple competing designers of initial signals.

Strategic Communication

Another possibility is to assume that the players are fully strategic during the communication phase of our model, anticipating how they will use their information in the final game. This seems a reasonable scenario in settings with a limited number of players, a simple network structure (e.g., a star), and experienced players. Though reasonable, strategic communication involves specific modeling challenges in order to define the equilibrium notion that determines the ξ played by the players. To see why, suppose ξ^* is a candidate equilibrium strategy in the communication phase. As long as players stick to their strategy ξ^*, every possible path of play induces a final information structure according to the corresponding function $f_{\xi^*,E}$ as defined above. Given this, for every π we can specify the continuation behavior through the BNE of $(G, f_{\xi^*,E}(\pi))$ preferred by the designer as in the baseline model. However, suppose player i deviates from ξ_i^*. This deviation may or may not be detectable by the other players once the communication phase ends. If not, a question arises of how to define the continuation of the game once the players reach the basic game G: What beliefs should the players have? What actions should they play?

These are important questions. Theorem 3 allows us to ignore them, however, at least on a first pass. Suppose there exists a way of specifying the players' off-path behavior and beliefs so that ξ^* describes their on-path strategic behavior in the communication phase of the model, followed by the BNE of the resulting final game selected according to the usual criterion. Depending on the specific setting, such equilibrium communication may be very complex and may admit of multiple solutions. Nonetheless, Theorem 3 provides a way to bound what the designer can achieve when facing strategic information spillovers, which holds for all those solutions and does not require calculating equilibria.

Because of its generality, Theorem 3 also applies across a variety of scenarios with strategic communication. For instance, communication between any two connected players may follow the rules of cheap talk à la Crawford and Sobel (1982) and Aumann and Hart (2003) in each round. Note that in this case, each player can be simultaneously a "sender" and a "receiver" and should take into account the entire network structure and the equilibrium strategy of all players in the network. Alternatively, communication may follow the rules à la Milgrom and Roberts (1986), where players use verifiable information. Finally, communication may follow the model of Gentzkow and Kamenica (2016), where each player strategically commits to a commonly observable communication strategy at the beginning of the communication phase, taking into account that he competes with all other sources of his followers to determine their beliefs for the final game G they will all play. It is reasonable to expect that such models are hard to analyze explicitly, perhaps except in specific cases with very few players and simple networks. This is why Theorem 3 seems a useful result.

5.2 A Direct Characterization of Feasibility for the Constrained-Designer Problem

In this section, we provide a direct characterization of the feasible outcomes for an information designer constrained to target only the agents in $M \subsetneq N$. Besides for completeness, this clarifies why the direct-mechanism approach is not valid in this case.

Definition 7 (M-Mediated Obedience)**.** The outcome function $x : \Omega \to R$ is M-mediated obedient for (G, E) if there exists $\kappa : \Omega \times R \to \Delta(Z)$ with the following properties:
(1) $Z = \times_{i \in N} Z_i$, where Z_i is finite for every i and $|Z_j| = 1$ for $j \notin M$;
(2) for all i and z_{E_i}, there exists at most one $\alpha_i \in \mathbf{x}_i$ such that $\kappa(z_{E_i}, z_{-E_i}|\alpha_i, \alpha_{-i}, \omega) > 0$ for some $(z_{-E_i}, \alpha_{-i}, \omega)$;
(3) for every i, α_i, and z_{E_i},

$$\sum_{\omega, \alpha_{-i}, z_{-E_i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a'_i, \alpha_{-i}, \omega)]\kappa(z_{E_i}, z_{-E_i}|\alpha, \omega)x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a'_i \in A_i.$$

Theorem 4. *Fix a basic game G, a network E, and a target set M. The outcome function x is feasible if and only if it is M-mediated obedient.*

This result shows that, in general, a direct characterization of feasibility for constrained problems requires a supplementary tool, κ, besides a direct-recommendation mechanism, x. This tool keeps track of information intermediation, namely, what non-targeted players learn from targeted players. Condition (1) captures the fact that only targeted players can provide messages to non-targeted players. Condition (2) requires that every player i should unambiguously figure out which mixed action the designer wants him to implement from the messages i hears from his sources. Finally, condition (3) requires that, given these messages, player i should not have a profitable deviation from the recommended behavior.

One may wonder why κ is necessary and has to take this form. First, why is it not enough to rely on the recommendations to targeted players in order to convey to non-targeted players how they should play in the game? To see this, consider a situation with only two players and a network that contains only a link from player 1 to player 2. Player 1 has a strictly dominant action independently of the state, and the designer can only target player 1. In this case, any feasible outcome must involve player

1's taking that action. However, the designer can "use" player 1 to convey information to player 2. Clearly, this requires a richer communication language than the constant recommendation of the dominant action to player 1. A second question may be why it is not enough, or correct, to let the designer recommend to the targeted players also the suggested behavior for their non-targeted followers and simply let κ pass along these recommendations. As noted earlier, the reason is that implementing this behavior may require some non-targeted player to *combine* the pieces of information coming from multiple targeted players, without them being able to exactly predict his behavior based only on the information each receives from the designer. Our approach based on the equivalence in Theorem 1 allows us to entirely avoid these issues.

5.3 Relation to Bayes Correlated Equilibrium and Individual Sufficiency

In this section, we briefly compare our work with Bergemann and Morris (2016) and, in particular, their comparative statics results that relate the set of feasible outcomes with the players *exogenous* information. Bergemann and Morris (2016) propose a ranking of such information structures, called individual sufficiency, and show that a higher degree of individual sufficiency corresponds to a smaller set of feasible outcomes. At a conceptual level, the present paper and Bergemann and Morris (2016) ask fundamentally different questions. They examine how information that players already have constrains what the designer can do; we ask how the players' ability to share any information they may receive—especially from the designer—constrains what the designer can do. In fact, our analysis focuses on the case where the players have no exogenous information. Different networks affect what information each player receives from his neighbors, in addition to the information he receives from the designer. But the important difference is that the designer is also in control of the information those neighbors receive. Thus, the information that player i receives from his neighbors is not entirely beyond the designer's control, in contrast to the exogenous information in Bergemann and Morris (2016).

Despite these differences, it is instructive to uncover the relation between our notion of more information aggregation through the network and individual sufficiency. To this end, consider any initial information structure (S, π) chosen by the designer. We say that (S, π) is communication invariant under the network E if, for all $i \in N$, s_i is a sufficient statistics for s_{E_i}. That is, letting $(\pi', T) = f_E(\pi)$, we have

$$Pr_{\pi'}(\omega, t|t_i) = Pr_{\pi}(\omega, s|s_{E_i}) = Pr_{\pi}(\omega, s|s_i), \tag{8}$$

where we leave the dependence of these probabilities on the prior μ implicit. Note that the first equality holds by the definition of how information spillovers work in our model (Assumption 1), while the last equality is the substantive condition on the informativeness of the initial and final information structures. Communication invariance captures the idea that the designer already gives each player as much information as he can learn from the signals of his sources about both ω and the signals of all the other players. So, in particular, given a communication-invariant π, its *physical* sctructure is formally changed by information spillovers, but π is essentially unchanged in terms of its information content. Given this, since the designer cannot prevent the players from communicating, we can conclude that essentially the designer is restricted to choosing information structures that are communication invariant under E.

Now take any initial (S, π). Consider two networks E' and E where E is deeper than E' and denote the respective functions f and f'. Can we say that the information structures $f(\pi)$ and $f'(\pi)$ are ranked according to individual sufficiency as defined in Bergemann and Morris (2016)? Note that, for every i, his signal space under $f(\pi)$ is $T_i = \times_{j \in E_i} S_j$ and his signal space under $f'(\pi)$ is $T'_i = \times_{j \in E'_i} S_j$. Given this, let's consider

$$Pr(t'_i | t_i, t_{-i}, \omega) = Pr(s_{E'_i} | s_{E_i}, (s_{E_j})_{j \neq i}, \omega).$$

Since $E'_i \subseteq E_i$ for all i when E is deeper than E', it follows that the right-hand side is independent of $(s_{E_j})_{j \neq i}$ and ω given s_{E_i}, which means that the left-hand side is independent of t_{-i} and ω given t_i. Therefore, $f(\pi)$ is individually sufficient for $f'(\pi)$ in a very trivial and stark way: t'_i is independent of t_{-i} and ω given t_i because t'_i is deterministically pinned down by t_i (as opposed to being t_i plus some noise). Clearly, the notion of individual sufficiency of Bergemann and Morris (2016) is much more general.

References

Acemoglu, D., M. A. Dahleh, I. Lobel, and A. Ozdaglar (2011): "Bayesian learning in social networks," *The Review of Economic Studies*, 78, 1201–1236.

Akbarpour, M., S. Malladi, and A. Saberi (2018): "Just a Few Seeds More: Value of Network Information for Diffusion," *Working Paper*.

Aumann, R. J. and S. Hart (2003): "Long cheap talk," *Econometrica*, 71, 1619–1660.

Banerjee, A., A. G. Chandrasekhar, E. Duflo, and M. O. Jackson (2013): "The diffusion of microfinance," *Science*, 341, 1236498.

Banerjee, A. V. (1992): "A simple model of herd behavior," *The quarterly journal of economics*, 107, 797–817.

Bergemann, D. and S. Morris (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

——— (2018): "Information Design: A Unified Perspective," *Journal of Economic Literature*, Forthcoming.

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

Crawford, V. P. and J. Sobel (1982): "Strategic information transmission," *Econometrica: Journal of the Econometric Society*, 1431–1451.

Eyster, E. and M. Rabin (2010): "Naive herding in rich-information settings," *American economic journal: microeconomics*, 2, 221–43.

Galperti, S. and J. Perego (2018): "A Dual Perspective on Information Design," .

Gentzkow, M. and E. Kamenica (2016): "Competition in persuasion," *Review of Economic Studies*, 84, 300–322.

Golub, B. and M. O. Jackson (2012): "How homophily affects the speed of learning and best-response dynamics," *The Quarterly Journal of Economics*, 127, 1287–1338.

Golub, B. and E. Sadler (2017): "Learning in social networks," *Working Paper*.

Granovetter, M. (1978): "Threshold models of collective behavior," *American journal of sociology*, 83, 1420–1443.

Jackson, M. O. and E. Storms (2018): “Behavioral communities and the atomic structure of networks,” *Available at SSRN 3049748*.

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

Milgrom, P. and J. Roberts (1986): “Relying on the information of interested parties,” *The RAND Journal of Economics*, 18–32.

Morris, S. (2000): “Contagion,” *The Review of Economic Studies*, 67, 57–78.

Mueller-Frank, M. (2013): “A general framework for rational learning in social networks,” *Theoretical Economics*, 8, 1–40.

Myerson, R. B. (1986): “Multistage games with communication,” *Econometrica: Journal of the Econometric Society*, 323–358.

Sadler, E. (2017): “Diffusion games,” *Working Paper*.

Shamir, A. (1979): “How to share a secret,” *Communications of the ACM*, 22(11), 612–613.

Smith, L. and P. Sørensen (2000): “Pathological outcomes of observational learning,” *Econometrica*, 68, 371–398.

A Proofs

A.1 Proof of Theorem 1

To prove Theorem 1, we first introduce and prove Lemmas 4, 5 and 6, and the intermediate equivalence result of Proposition 10. Lemma 4 characterizes the M-expansion of E by showing that, in the process of expansion, while each player may gain new sources, i.e. $E_i \subseteq E_i^M$, none of these sources are original targets in M, i.e. $E_i^M \cap M = E_i \cap M$.

Lemma 4. *Let E^M be the M-expansion of E. Then, $E_i^M \cap M = E_i \cap M$, for all $i \in N$.*

Proof of Lemma 4. Fix i and let E^M be the M-expansion of E. We first show that $E_i \cap M \subseteq E_i^M \cap M$. To see this, note that $E \subseteq E^M$, by definition of M-expansion. This implies that $E_i \subseteq E_i^M$. Hence, $E_i \cap M \subseteq E_i^M \cap M$. Conversely, in order to show that $E_i^M \cap M \subseteq E_i \cap M$, it is enough to show that $E_i^M \cap M \subseteq E_i$. Suppose not, that is, suppose that there is $j \in M$ such that $j \in E_i^M$ but $j \notin E_i$. Since $j \in E_i^M$, there exists a E^M-path $P = (k_1, \ldots k_m)$ from j to i, that is a sequence such that $k_1 = j$, $k_m = i$, and $(k_l, k_{l+1}) \in E^M$, for all $l \leq m-1$. Since $j \notin E_i$, it must be that $(k_l, k_{l+1}) \notin E$, for *some* $l \leq m-1$. Call these l's the *gaps* of P. Choose path P so that the number gaps is the smallest. This object is well-defined and we denote it $\underline{P} = (\underline{k}_1, \ldots \underline{k}_m)$. Along path $\underline{P}$, denote $\underline{l}$ the first gap from j to i, that is the gap l with the smallest index. By construction, we have that (1) $j \in E_{\underline{k}_{\underline{l}}}$, (2) $(\underline{k}_{\underline{l}}, \underline{k}_{\underline{l}+1}) \notin E$, and (3) $(\underline{k}_{\underline{l}}, \underline{k}_{\underline{l}+1}) \notin E^M$. Moreover, $j \notin E_{\underline{k}_{\underline{l}+1}}$. If this was not the case, $\underline{P}$ would not be the E^M-path from j to i with the smallest number of gaps. Definition 3 and points (2) and (3) above imply that $E_{\underline{k}_{\underline{l}}} \cap M \subseteq E_{\underline{k}_{\underline{l}+1}}$. Yet, $j \in E_{k_{\underline{l}}} \cap M$ and $j \notin E_{k_{\underline{l}+1}}$. We conclude that E^M is not the expansion of E, a contradiction. $\square$

Thanks to Lemma 4, we know that the sets of original sources don't change when expanding the network from E to E^M. Since $\pi \in \Pi_M$ requires signals to non-targeted players to be singletons, this means that from an informational point of view the information structures $f_E(\pi)$ and $f_{E^M}(\pi)$ are equivalent. The next result formalizes this idea.

Proposition 10. *Fix (E, M) and let E^M be the M-expansion of E. Then for all G, $X_M(G, E) = X_M(G, E^M)$.*

Proof of Proposition 10. Fix $\pi \in \Pi_M$, $i \in N$ and let $s := (s_1, \ldots, s_n)$ be a possible realization of π. We want to show that:

$$Pr_\pi(s_1, \ldots, s_n | s_{E_i}) = Pr_\pi(s_1, \ldots, s_n | s_{E_i^M})$$

Note that $(E_i^M \setminus E_i) \cap M = (E_i^M \cap M) \setminus (E_i \cap M) = \emptyset$. The first equality derives from the distributive property of set intersection over set difference. The second equality derives from Lemma 4. This implies that, since $\pi \in \Pi_M$, the set $S_{E_i^M \setminus E_i}$ is a singleton and, equivalently, its only element $s_{E_i^M \setminus E_i}$ realizes with probability 1. Therefore, s_{E_i} and $s_{E_i^M}$ are identical up to elements $s_{E_i^M \setminus E_i}$ that realize with probability 1. Hence $Pr_\pi(s_1, \ldots, s_n | s_{E_i}) = Pr_\pi(s_1, \ldots, s_n | s_{E_i^M})$. Since $i \in N$ and s were arbitrary, we have that $BNE(G, f_E(\pi)) = BNE(G, f_{E^M}(\pi))$. Since $\pi \in \Pi_M$ was arbitrary, we conclude that $X_M(G, E) = X_M(G, E^M)$. $\square$

Lemma 4 is important in establishing an another important property of the M-expansion of network E. E^M is the M-expansion of itself. The next result formalizes this idea.

Lemma 5. *$(i, j) \in E^M$ if and only if $E_i^M \cap M \subseteq E_j^M$.*

Proof of Lemma 5 *Only if.* Let $(i, j) \in E^M$. Then, $E_i^M \subseteq E_j^M$, hence $E_i^M \cap M \subseteq E_j^M$. *If.* Suppose $E_i^M \cap M \subseteq E_j^M$. By Lemma 4, $E_i^M \cap M = E_i \cap M$ and $E_j^M \cap M = E_j \cap M$. Thus, $E_i \cap M \subseteq E_j \cap M$. By Definition 3, this implies $(i, j) \in E^M$. $\square$

The next result constitutes the building block of the proof of our main theorem. To prove this result, we borrow a technique from the computer science literature on cryptography, known as the *secret sharing* method. Secret sharing refers to the problem of distributing a "secret" among a group of n players, each of whom is allocated a "share" of the secret. The distribution is so that players learn the secret if all players pool their shares. If one or more shares are missing, nothing is learned about the secret.

Lemma 6. *Fix (G, E, M) and some $M' \supseteq M$. Let $i \in M'$ and $(i, j) \in E^M$. Then $X_{M'}(G, E^M) = X_{M' \cup \{j\}}(G, E^M)$.*

Proof of Lemma 6.
($\subseteq$). This direction is trivial since, by definition, $\Pi_{M'} \subseteq \Pi_{M' \cup \{j\}}$.
($\supseteq$) If $j \in M'$ there is nothing to show as, in such case, $M' \cup \{j\} = M'$. Therefore, let $j \notin M'$. Fix $\pi \in \Pi_{M' \cup \{j\}}$ and denote π's signal space $S_1 \times \ldots \times S_n$. We want to find a $\hat{\pi} \in \Pi_{M'}$ such that $BNE(G, f_{E^M}(\pi)) = BNE(G, f_{E^M}(\hat{\pi}))$. Denote the signal space of $\hat{\pi}$, $\hat{S}_1 \times \ldots \times \hat{S}_n$. We will construct $\hat{\pi}$ using the technique behind Shamir's *secret sharing* scheme (Shamir (1979)). Let $B(\kappa) := \{0, 1\}^\kappa$ be the set of all binary number of length k. Define $\underline{\kappa} := \min\{\kappa \in \mathbb{N} : |S_j| \leq |B(\kappa)|\}$. For notational convenience, let's denote $B := B(\underline{\kappa})$ and choose an injective $pub : S_j \to B$. This map represents the "public key"

that allows players to translates binary numbers into signals for player j. We denote $\oplus$ the bitwise XOR operation.[21] Let $Q := E_j \cap M'$ and let $q \geq i$, for all $i \in Q$. Let $\hat{S}_i := S_i \times B \times \{pub\}$ for all $i \in Q$; $\hat{S}_j := \{\bar{s}_j\}$; and $\hat{S}_i := S_i$ for all $i \notin N \setminus (Q \cup \{j\})$. Note that $\hat{\pi} \in \Pi_{M'}$. For all $s \in S$, define $\hat{s}(s)$ in the following way: $\hat{s}_j(s) := \bar{s}_j$; for all $i \notin N \setminus (Q \cup \{j\})$, $\hat{s}_i(s) = s_i$; for all $i \in Q$, $\hat{s}_i(s) = (s_i, b_i, pub)$. $b_i \in B$ is random and independent of ω and s and it is determined as follows: If $i \neq q$, elements of b_i are independent realizations of fair coin toss; if $i = q$, $b_q := pub(s_j) \oplus \big(\oplus_{i \in Q: i \neq q} b_i \big)$. The probability $\hat{\pi}(\hat{s}(s)|\omega)$ is determined from $\pi(s|\omega)$. If $|Q| > 1$, b_q is random and independent of s and ω. Moreover, the construction of this secret sharing scheme implies that $(b_i)_{i \in L} \perp \omega$ and $(b_i)_{i \in L} \perp s$ for all $L \in 2^Q \setminus Q$. That is, receiving all but one share of the secret carries no information about s_j. Instead, the vector $(b_i)_{i \in Q}$ fully reveals s_j because

$$pub^{-1}\Big(\oplus_{i \in Q} b_i \Big) = s_j.$$

This implies that, player $l \in N$ learns s_j from the observation of $\hat{s}_{E_l^M}$ if and only if $j \in E_l^M$. To see this, note that, by construction, player $l \in N$ learns s_j if and only if $Q \subseteq E_l^M$. We are left to show that $Q \subseteq E_l^M$ if and only if $j \in E_l^M$. Let's first assume that $j \in E_l^M$. Then, $E_j^M \subset E_l^M$ and $Q = E_j^M \cap M \subset E_l^M$. Conversely, assume $Q = E_j^M \cap M \subseteq E_l^M$. By Lemma 5, $(j, l) \in E^M$ and thus $j \in E_l^M$. We now discuss the special case $|Q| = 1$ and show that it is taken care of. To see this, let $Q = E_j^M \cap M' = \{i\}$. By construction, b_i reveals s_j. Therefore, we need to show that $j \in E_i^M$. To this purpose, note that $E_j^M \cap M' = \{i\}$ implies that $E_j^M \cap M \subset E_j^M \cap M' = \{i\} \subseteq E_i^M$. By Lemma 5, $(j, i) \in E^M$, hence $j \in E_i^M$.

In summary, our construction is so that, for all $s \in S$, either $\hat{s}_{E_l^M}(s)$ perfectly reveals s_j or it is completely uninformative about s_j. Therefore, we have shown that for all $s \in S$ and $l \in N$, $Pr_\pi(s|s_{E_l^M}) = Pr_{\hat{\pi}}(\hat{s}(s)|\hat{s}_{E_l^M}(s))$. This implies that any outcome x induced by π can be also induced by $\hat{\pi}$. Since π was arbitrary, this shows that $X_{M' \cup \{j\}}(G, E^M) \subseteq X_{M'}(G, E^M)$.

Proof of Theorem 1. Fix G, E and M. By Proposition 10, $X_M(G, E) = X_M(G, E^M)$. We are left to show that $X_M(G, E^M) = X_N(G, E^M)$. The following induction argument proves the claim.

Basis Step. Let $M_1 = M$. If $M_1 = N$ there is nothing to prove and the procedure stops. Hence, let $M_1 \subsetneq N$. Because of Assumption 2, there exists $(i, j) \in E$ such that $i \in M_1$ and $j \notin M_1$. Since $E \subseteq E^M$, $(i, j) \in E^M$. Let $M_2 := M_1 \cup \{j\}$. Since $M_2 \supseteq M_1$, $i \in M_2$ and $(i, j) \in E^M$, we can invoke Lemma 6 to conclude that $X_{M_1}(G, E^M) =$

[21]Given two vectors $b, b' \in B$, $b \oplus b \in B$. Component-by-component, elements of b and b' are added using the following rules: $0 \oplus 0 = 1 \oplus 1 = 0$ and $0 \oplus 1 = 1 \oplus 0 = 1$.

$X_{M_2}(G, E^M)$. Finally, note that E and M_2 satisfy Assumption 2.

Inductive Step. Suppose that $X_{M_1}(G, E^M) = X_{M_t}(G, E^M)$. If $M_t = N$ there is nothing to prove and the procedure stops. Hence, let $M_t \subsetneq N$. Because of Assumption 2, there exists $(i, j) \in E$ such that $i \in M_t$ and $j \notin M_t$. Since $E \subseteq E^M$, $(i, j) \in E^M$. Let $M_{t+1} := M_t \cup \{j\}$. Since $M_{t+1} \supseteq M_t$, $i \in M_{t+1}$ and $(i, j) \in E^M$, we can invoke Lemma 6 to conclude that $X_{M_{t+1}}(G, E^M) = X_{M_t}(G, E^M)$.

Because N is finite, this procedure stops after $\bar{t} = |N \setminus M|$ steps, concluding that $X_{M_{\bar{t}}}(G, E^M) = X_N(G, E^M)$. □

A.2 Remaining Proofs

Proof of Theorem 2. Part 1 (⇒): Suppose (S, π) and $\sigma \in \text{BNE}(G, f(\pi))$ induce x. Then, for every i and $s \in S$,

$$\sum_{\omega, s'} [u_i(\sigma_i(s'_{E_i}), \sigma_{-i}(s'_{E_{-i}})), \omega) - u_i(a_i, \sigma_{-i}(s'_{E_{-i}}), \omega)] Pr_\pi(\omega, s'|(s'_{E_i})) \geq 0, \quad a_i \in A_i.$$

where $\sigma_{-i}(s'_{E_{-i}}) = (\sigma_j(s'_{E_j}))_{j \neq i}$. Using π, we can write this condition as

$$\frac{\sum_{\omega, s'_{-E_i}} [u_i(\sigma_i(s'_{E_i}), \sigma_{-i}(s'_{E_{-i}})), \omega) - u_i(a_i, \sigma_{-i}(s'_{E_{-i}}), \omega)] \pi(s'|\omega)\mu(\omega)}{\sum_{\omega', s''_{-E_i}} \pi(s''|\omega')\mu(\omega')} \geq 0, \quad a_i \in A_i,$$

or equivalently,

$$\sum_{\omega, s'_{-E_i}} [u_i(\sigma_i(s'_{E_i}), \sigma_{-i}(s'_{E_{-i}})), \omega) - u_i(a_i, \sigma_{-i}(s'_{E_{-i}}), \omega)] \pi(s'|\omega)\mu(\omega) \geq 0, \quad a_i \in A_i.$$

Note that, for every $i \in N$ and s, by knowing s_{E_i} player i knows the mixed action $\sigma_j(s_{E_j})$ for all $j \in E_i$.

Given this and using the definition of x in (1), the last family of inequalities can be written as follows: For all i and α_{E_i},

$$\sum_{\omega, \alpha_{-E_i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)] x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i \in A_i.$$

Thus, we conclude that if x is feasible, then it is obedient*.

Part 2 (⇐): Suppose x is obedient*. Note that x can be thought of as an information structure in Π, by viewing each α_i as a signal first. Given this, for every i, consider the strategy $\sigma_i : \times_{j \in E_i} \Delta(A_j) \to \Delta(A_i)$ defined by

$$\sigma_i(\alpha_{E_i}) = \alpha_i, \qquad \alpha_{E_i} \in \times_{j \in E_i} \Delta(A_j).$$

By our assumption that player i learns the signals of his parents in E_i, optimality for each player i requires that, for every α_{E_i},

$$\sum_{\omega,\alpha'_{-E_i}} [u_i(\sigma_i(\alpha'_{E_i}), \sigma_{-i}(\alpha'_{E_{-i}}), \omega) - u_i(a_i, \sigma_{-i}(\alpha'_{E_{-i}})), \omega)] Pr_x(\omega, \alpha'|\alpha_{E_i}) \geq 0, \quad a_i \in A_i,$$

where $\sigma_{-i}(\alpha'_{E_{-i}}) = (\sigma_j(\alpha'_{E_j}))_{j\neq i}$. Given our construction of σ, this is equivalent to, for every α_{E_i} and $a_i \in A_i$,

$$\sum_{\omega,\alpha'_{-E_i}} [u_i(\alpha_{E_i}, \alpha'_{-E_i}, \omega) - u_i(a_i, \alpha_{E_i\setminus i}, \alpha'_{-E_i}, \omega)] \frac{x(\alpha_{E_i}, \alpha'_{-E_i}|\omega)\mu(\omega)}{\sum_{\omega',\alpha''_{-E_i}} x(\alpha_{E_i}, \alpha''_{-E_i}|\omega')\mu(\omega')} \geq 0,$$

which holds because x is obedient*. □

Proof of Proposition 1. Part 1 (⇐): Suppose E exhibits more influence than does E' and $x \in X(G, E)$ for some G. Then, by Theorem 2 and Remark 2, x satisfies for every $i = 1, \ldots, N$,

$$\sum_{\substack{\omega\in\Omega \\ \alpha\in\mathbf{x}}} u_i(\alpha_i, \alpha_{-i}; \omega) x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq \sum_{\substack{\omega\in\Omega \\ \alpha\in\mathbf{x}}} u_i(\delta_i(\alpha_{E_i}), \alpha_{-i}; \omega) x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega),$$

for all $\delta_i \in D_i = \{\hat{\delta}_i : \times_{j\in E_i}\Delta(A_j) \to A_i\}$. Note that ${}_iE \supseteq_i E'$ for every $i \in N$ if and only if $E_j \supseteq E'_j$ for every $j \in N$. Indeed, if $i \in E'_j$, then $j \in_i E' \subseteq_i E$ and so $i \in E_j$; if $j \in_i E'$, then $i \in E'_j \subseteq E_j$ and so $j \in_i E$. Therefore, if E exhibits more influence than E', then $E_i \supseteq E'_i$ for all $i \in N$. Let $D'_i = \{\delta_i : \times_{j\in E'_i}\Delta(A_j) \to A_i\}$. To prove that $x \in X(G, E')$ it suffices to show that the set of available deviations D'_i is smaller than D_i for all $i \in N$. To show this, consider any $\delta_i \in D'_i$ and define $\hat{\delta}_i : \times_{j\in E_i}\Delta(A_j) \to A_i$ as $\hat{\delta}_i(\alpha_{E'_i}, \alpha_{E_i\setminus E'_i}) = \delta_i(\alpha_{E'_i})$ for all $\alpha_{E_i} \in \times_{j\in E_i}\Delta(A_j)$. Since $E_i \supseteq E'_i$ for all i, $\hat{\delta}_i$ is a well-defined function and $\hat{\delta}_i \in D_i$.

Part 2 (⇒): We prove this by contrapositive. The only relevant case to consider is that E does not exhibit more influence than E' (and vice versa). This implies that for some i, there exists a k such that $k \in E'_i$ and $k \notin E_i$, and for some j (possibly $i = j$), there exists m such that $m \in E_j$ and $m \notin E'_j$. It follows that theres exists a player i_k such that $i_k \neq k$ and there is a direct link from k to i_k in E' but not in E, and there exists a player i_m such that $i_m \neq m$ there is a direct link from m to i_m in E but not in E'.

Now consider the following game G. Let $\Omega = \{0, 1\}$ and $\mu(0) = \mu(1) = \frac{1}{2}$. Let $A_i = \{0, \frac{1}{2}, 1\}$ for all $i \in N$. For all $j \notin \{k, m, i_k, i_m\}$, the payoff function u_j is such that the action $a_j = \frac{1}{2}$ is strictly dominant. For $j \in \{k, m, i_k, i_m\}$, the payoff function is $u_j(a, \omega) = -(a_j - \omega)^2$.

Consider the following two cases.

Case 1: Suppose that all players in $\{k, m, i_k, i_m\}$ are distinct. Consider x such that player k always matches the state, while all other players choose $a = \frac{1}{2}$. This x belong to $X(G, E)$, but clearly does not belong to $X(G, E')$ because in E' player i_k has to choose $a = \frac{1}{2}$ after learning $a_k = \omega$, which renders $a = \frac{1}{2}$ strictly suboptimal and so x violates obedience*. Consider x' such that player m always matches the state, while all the other players choose $a = \frac{1}{2}$. This x' belong to $X(G, E')$, but clearly does not belong to $X(G, E)$ because in E player i_m has to choose $a = \frac{1}{2}$ after learning $a_m = \omega$, which renders $a = \frac{1}{2}$ strictly suboptimal and so x' violates obedience*. The same argument works for the same x and x' for the following four alternative configurations of the network that satisfy the aforementioned properties: (1) $m = i_k$ and $k \neq i_m$; (2) $m \neq k$ and $i_k = i_m$; (3) $k = i_m$ and $m = i_k$; (4) $i_m = k$ and $m \neq i_k$.

Case 2: Suppose that $m = k$ and $i_k \neq i_m$. Consider x such that m and i_m always match the state, while all other players choose $a = \frac{1}{2}$. This x belongs to $X(G, E)$, but clearly does not belong to $X(G, E')$ because in E' player i_k has to choose $a = \frac{1}{2}$ after learning $a_k = \omega$, which renders $a = \frac{1}{2}$ strictly suboptimal and so x violates obedience*. Consider x' such that player m and i_k always match the state, while all the other players choose $a = \frac{1}{2}$. This x' belong to $X(G, E')$, but clearly does not belong to $X(G, E)$ because in E player i_m has to choose $a = \frac{1}{2}$ after learning $a_m = \omega$, which renders $a = \frac{1}{2}$ strictly suboptimal and so x' violates obedience*. $\square$

Proof of Proposition 2. Part 1 ($\Leftarrow$): Suppose that E exhibits more influence than E'. Recall that this implies that $E_j \supseteq E'_j$ for every $j \in N$. We want to show that, for every $\pi \in \Pi$, $\Delta^{f_E(\pi')}(\Omega \times A_i) \subseteq \Delta^{f_E(\pi)}(\Omega \times A_i)$ for all i. Fix any i. Recall that, for every π, i's set of types under $f_E(\pi)$ is $T_i^{f_E(\pi)} = \times_{j \in E_i} S_j^\pi$ and under $f_E(\pi')$ it is $T_i^{f_{E'}(\pi)} = \times_{j \in E'_i} S_j^\pi$, where S_j^π is the set of types of player j under the common initial π. Take $y \in \Delta^{f_{E'}(\pi)}(\Omega \times A_i)$. We will show that $y \in \Delta^{f_E(\pi)}(\Omega \times A_i)$. Since $y \in \Delta^{f_{E'}(\pi)}(\Omega \times A_i)$, there exists $\gamma' : T_i^{f_{E'}(\pi)} \to \Delta(A_i)$ such that

$$y(\omega, a_i) = \sum_s \gamma'(a_i | s_{E'_i}) \pi(s|\omega) \mu(\omega), \quad (\omega, a_i) \in \Omega \times A_i.$$

Now define $\gamma : T_i^{f_E(\pi)} \to \Delta(A_i)$ from γ' by letting

$$\gamma(a_i | s_{E_i}) = \gamma'(a_i | s_{E'_i})$$

whenever $(s_{E_i})_{j \in E'_i} = s_{E'_i}$—that is, γ depends only on the components in E_i that also belong to E'_i and in the same way that γ' depends on them. This is well defined because $E_i \supseteq E'_i$. Clearly, this γ gives rise to the distribution y induced by γ', which therefore belongs to $\Delta^{f_E(\pi)}(\Omega \times A_i)$.

Part 2 (⇒): We prove this by contrapositive. Again, the only relevant case to consider is that E does not exhibit more influence than E' (and vice versa). This implies that for some i, there exists a k such that $k \in E'_i$ and $k \notin E_i$, and for some j (possibly $i = j$), there exists m such that $m \in E_j$ and $m \notin E'_j$. It follows that theres exists a player i_k such that $i_k \neq k$ and there is a direct link from k to i_k in E' but not in E, and there exists a player i_m such that $i_m \neq m$ there is a direct link from m to i_m in E but not in E'.

First, take an information structure π_1 such that k gets full information and all other players always get fully uninformative signals. Then, under $f_{E'}(\pi_1)$ player i_k gets full information, while under $f_E(\pi_1)$ he still gets no information. Therefore, f_E does not aggregate more information than does $f_{E'}$. Now, take an information structure π_2 such that m gets full information and all other players always get fully uninformative signals. Then, under $f_E(\pi_2)$ player i_m gets full information, while under $f_{E'}(\pi_2)$ he still gets no information. Therefore, $f_{E'}$ does not aggregate more information than does f_E. This proves the contrapositive. □

Proof of Proposition 3. Suppose that E^M is deeper than $\hat{E}^M$, $i \notin M$, and $i \notin E_j$. If $\hat{E}_i \cap M \subseteq \hat{E}_j$, then $i \in \hat{E}^M_j$ by the definition of M-expansion. Since $\hat{E}^M_j \subseteq E^M_j$,[22] we have $i \in E^M_j$ as well. Since $i \notin E_j$, we must have added links to E according to the definition of M-expansion that render $i \in E^M_j$. For this to be the case, there must exist some sequence $\{j_k\}_{k=0}^m$ which satisfies $j_0 = i$, $j_m = j$, and $E_{j_k} \cap M \subseteq E_{j_{k+1}}$. In turn, this implies that $E_i \cap M \subseteq E_j$.

Now suppose that the condition in the proposition holds. Then, E^M is deeper than $\hat{E}^M$ if $i \in \hat{E}^M_j$ implies $i \in E^M_j$. Fix any i and j that satisfy $i \in \hat{E}^M_j$. If $i \in \hat{E}_j$, then since $\hat{E}_j \subseteq E_j \subseteq E^M_j$ we have $i \in E^M_j$. Next, suppose $i \in \hat{E}^M_j \setminus \hat{E}_j$ and therefore it must be that $\hat{E}_i \cap M \subseteq \hat{E}_j$. In this case, we must have $i \notin M$: Indeed, if $i \in M$, $\hat{E}_i \cap M \subseteq \hat{E}_j$ implies that $i \in \hat{E}_j$. If i belongs already to E_j, then again $i \in E^M_j$. If instead $i \notin E_j$, by assumption $E_i \cap M \subseteq E_j$ and therefore $(i, j) \in E^M$, which implies $i \in E^M_j$. □

Proof of Proposition 4. Suppose that $i \notin E_j$. If $E_i \cap M' \subseteq E_j$, then $(i, j) \in E^{M'}$ by the definition of M'-expansion and so $i \in E^{M'}_j$. By assumption, we also have $E_i \cap M \subseteq E_j$, so again $i \in E^M_j$. This implies that, for all j, $E^{M'}_j \subseteq E^M_j$ and therefore E^M is deeper than $E^{M'}$.

Suppose that $E^{M'}_j \subseteq E^M_j$ for all j. Consider i such that $i \notin E_j$. Then, $E_i \cap M' \subseteq E_j$ implies by the same argument as before that $i \in E^{M'}_j$. Since then $i \in E^M_j \setminus E_j$, we have

[22] Note that i is a source of j if and only if j is a follower of i.

added links to E according to the definition of M-expansion that render $i \in E_j^M$. For this to be the case, there must exist some sequence $\{j_k\}_{k=0}^m$ which satisfies $j_0 = i$, $j_m = j$, and $E_{j_k} \cap M \subseteq E_{j_{k+1}}$. In turn, this implies that $E_i \cap M \subseteq E_j$. $\square$

Proof of Lemma 1. **Step 1.** Consider any finite-support (S, π) and let σ be the designer-preferred equilibrium of $(G, f_E(\pi))$. For every i, every s_{E_i} determines a non-empty subset of optimal actions $A_i(s_{E_i})$:

$$A_i(s_{E_i}) = \arg\max_{a_i \in A_i} \mathbb{E}_{\pi,\sigma}[u_i(a_i, a_{-i}, \omega)|s_{E_i}].$$

Since A_i is finite, every (π, σ) can determine at most finitely many subsets $A_i(s_{E_i})$ for every player i. This requires no more than $|2^{A_i}|$ signals for player i. Therefore, every (π, σ) can determine at most finitely many profiles of optimal-action sets of the form $A(s) = \times_{i \in N} A_i(s_{E_i})$. We conclude that if we are interested in only such profiles, it is enough to consider information structures that satisfy $|S_i| = |2^{A_i}|$ for every i.

Step 2. We now need to transition from profiles of optimal-action sets to distributions over pure-action profiles, which is what ultimately matters for the designer. To this end, we use Theorem 2. Recall that each recommendation profile α can be interpreted, first of all, as a signal realization from the information structure x. Step 1 shows that, if we are interested only in spanning the profiles of optimal-action sets, it is enough to consider xs with finite support. But this may not be enough for the entire set of feasible outcomes intended as joint distributions between actions and states that satisfy obedience.

Suppose that x is a feasible outcomes/satisfies obedience. That is, for every i, $\alpha_{E_i} \in \mathbf{x}_{E_i}$, and $a_i, a_i' \in A_i$,

$$\sum_{\omega, \alpha_{-E_i}} \left\{ \sum_{a_{-i}} [u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega] \alpha_{E_i}(a_{E_i}) \alpha_{-E_i}(a_{-E_i}) \right\} x(\alpha_{E_i}, \alpha_{-E_i}|\omega)\mu(\omega) \geq 0,$$

where $\alpha_{E_i}(a_{E_i}) = (\alpha_j(a_j))_{j \in E_i}$ and $\alpha_{-E_i}(a_{-E_i}) = (\alpha_j(a_j))_{j \notin E_i}$. We want to construct an alternative x' that is also feasible/satisfies obedience and induces the same joint distribution between *pure*-action profiles and states as does x.

From step 1, we know that we can identify finitely many profiles of sets $A^x(\alpha) = \times_{i \in N} A_i^x(\alpha_{E_i})$, where we treat each α as a signal realization from x. Let $\mathcal{A}^x$ be the finite collection of such profiles determined by x. In particular, we know that $|\mathcal{A}^x| \leq \prod_{i \in N} |2^{A_i}|$ independently of x. For every ω, construct x' as follows. For every $A^x \in \mathcal{A}^x$, define

$$\alpha^{A^x, \omega}(a) = \sum_{\alpha \in A^x} \alpha(a) \frac{x(\alpha|\omega)}{\sum_{\alpha' \in A^x} x(\alpha'|\omega)}, \quad a \in A.$$

This is the average mixed-action profile in state ω, conditional on α belonging to A^x. Given this, for every $\alpha^{A^x,\omega}$ so identified, let

$$x'(\alpha^{A^x,\omega}|\omega) = \sum_{\alpha \in A^x} x(\alpha|\omega), \quad \omega \in \Omega.$$

It is immediate to see that x and x' induce the same joint distribution over pure-action profiles and states: For every a and ω,

$$\begin{aligned}\sum_{\alpha' \in \mathbf{x}'} \alpha'(a)x'(\alpha'|\omega)\mu(\omega) &= \sum_{A^x \in \mathcal{A}^x} \alpha^{A^x,\omega}(a)x'(\alpha^{A^x,\omega}|\omega)\mu(\omega) \\ &= \sum_{A^x \in \mathcal{A}^x} \left[\sum_{\alpha \in A^x} \alpha(a)\frac{x(\alpha|\omega)}{\sum_{\alpha' \in A^x} x(\alpha'|\omega)}\right] \sum_{\hat{\alpha} \in A^x} x(\hat{\alpha}|\omega)\mu(\omega) \\ &= \sum_{A^x \in \mathcal{A}^x} \left[\sum_{\alpha \in A^x} \alpha(a)x(\alpha|\omega)\mu(\omega)\right] = \sum_{\alpha \in \mathbf{x}} \alpha(a)x(\alpha|\omega)\mu(\omega).\end{aligned}$$

Let's now consider obedience. If we can show that x' also satisfies obedience, we are done. Fix any player i, any $\alpha'_{E_i} \in \mathbf{x}'_{E_i}$, and $a_i, a'_i \in A_i$. Note that α'_{E_i} must equal $\alpha_{E_i}^{A^x,\omega}$ for some some A^x and ω. Let $\mathcal{A}^x(\alpha'_{E_i})$ contain all the profiles A^x that are compatible with α'_{E_i}, i.e., that satisfy $\alpha_{E_i}^{A^x,\omega} = \alpha'_{E_i}$. Letting $\Delta u_i(a_i, a'_i; a_{-i}, \omega) = u_i(a_i, a_{-i}, \omega) - u_i(a'_i, a_{-i}, \omega)$, we have

$$\begin{aligned}&\sum_{\omega,\alpha'_{-E_i}} \left\{\sum_{a_{-i}} \Delta u_i(a_i, a'_i; a_{-i}, \omega)\alpha'_{E_i}(a_{E_i})\alpha'_{-E_i}(a_{-E_i})\right\} x'(\alpha'_{E_i}, \alpha'_{-E_i}|\omega)\mu(\omega) \\ &= \sum_{\omega, A^x \in \mathcal{A}^x(\alpha'_{E_i})} \left\{\sum_{a_{-i}} \Delta u_i(a_i, a'_i; a_{-i}, \omega)\alpha_{E_i}^{A^x,\omega}(a_{E_i})\alpha_{-E_i}^{A^x,\omega}(a_{-E_i})\right\} x'(\alpha_{E_i}^{A^x,\omega}, \alpha_{-E_i}^{A^x,\omega}|\omega)\mu(\omega) \\ &= \sum_{\omega, A^x \in \mathcal{A}^x(\alpha'_{E_i})} \left\{\sum_{a_{-i}} \Delta u_i(a_i, a'_i; a_{-i}, \omega) \sum_{\alpha \in A^x} \alpha_{E_i}(a_{E_i})\alpha_{-E_i}(a_{-E_i})\frac{x(\alpha_{E_i}, \alpha_{-E_i}|\omega)}{\sum_{\alpha' \in A^x} x(\alpha'|\omega)}\right\} \times \\ &\quad \times \sum_{\alpha \in A^x} x(\alpha_{E_i}, \alpha_{-E_i}|\omega)\mu(\omega) \\ &= \sum_{\omega, A^x \in \mathcal{A}^x(\alpha'_{E_i})} \sum_{\alpha \in A^x} \left\{\sum_{a_{-i}} \Delta u_i(a_i, a'_i; a_{-i}, \omega)\alpha_{E_i}(a_{E_i})\alpha_{-E_i}(a_{-E_i})x(\alpha_{E_i}, \alpha_{-E_i}|\omega)\right\} \mu(\omega) \\ &= \sum_{A^x \in \mathcal{A}^x(\alpha'_{E_i})} \sum_{\alpha \in A^x} \left\{\sum_{\omega, a_{-i}} \Delta u_i(a_i, a'_i; a_{-i}, \omega)\alpha_{E_i}(a_{E_i})\alpha_{-E_i}(a_{-E_i})x(\alpha_{E_i}, \alpha_{-E_i}|\omega)\mu(\omega)\right\}.\end{aligned}$$

Now, recall that for every $\alpha \in A^x$, we have that the set of optimal actions for player i conditional on α_{E_i} is the same. Since x satisfies obedience for player i, his α_i assigns positive probability only to actions that are optimal conditional on α_{E_i}. Therefore, the entire sum must be non-negative. This shows that x' satisfies obedience for player i and every $\alpha'_{E_i} \in \mathbf{x}'_{E_i}$. By the same argument, x' satisfies obedience for all players. $\square$

Proof of Proposition 5. Let Y be the set of outcome functions χ that satisfy all the constraints of the primal $\mathcal{P}$. Since $V_E^*(N) = \sup_\chi \sum_{\omega,\alpha} v(\alpha,\omega)\chi(\alpha,\omega)$ where χ satisfies $|\mathbf{supp}\,\chi(\cdot,\omega)| \leq \prod_{i\in N} |2^{A_i}|$ for all ω, for every $n \geq 1$ there exists χ^n that satisfies the same support property and

$$V_E^*(N) \geq \sum_{\omega,\alpha} v(\alpha,\omega)\chi^n(\alpha,\omega) \geq V_E^*(N) - \frac{1}{n}.$$

Let $K_i = |2^{A_i}|$ for every i. To every such χ^n, there correspond finite subsets $A_i^n \subset \Delta(A_i)$ such that $|A_i^n| = K_i$ for all i which define a grid in $\times_{i\in N}\Delta(A_i)$ over which we can restrict the construction of χ^n itself. (Note that the support of χ^n may not use the entire grid, but it is without loss to allow for these extra elements that receive zero probability). Thus, for every i and $k_i = 1,\ldots,K_i$, there is a sequence $\alpha_i^{k_i,n} \in \Delta(A_i)$ where $\alpha_i^{k_i,n} \in A_i^n$ is an element of the grid of player i with (fixed) K_i elements to construct χ^n. Also, for each ω and every $(k_1,\ldots,k_N)$ where $k_i = 1,\ldots,K_i$ for every i, we have a sequence of elements $\chi^n(\alpha_1^{k_1,n},\ldots,\alpha_N^{k_N,n},\omega) \in [0,1]$. Since all these sequences belong to a compact space, each has a converging subsequence. Moreover, since we have finitely many sequences because each K_i is fixed and finite, there exists an overall subsequence of indexes $\tilde{n}$ such that the following holds:

$$\lim_{\tilde{n}\to\infty} \alpha_i^{k_i,\tilde{n}} = \hat{\alpha}_i^{k_i} \in \Delta(A_i), \quad k_i = 1,\ldots,K_i, i \in N;$$
$$\lim_{\tilde{n}\to\infty} \chi^{\tilde{n}}(\alpha_1^{k_1,\tilde{n}},\ldots,\alpha_N^{k_N,\tilde{n}},\omega) = \hat{\chi}(\hat{\alpha}_1^{k_1},\ldots,\hat{\alpha}_N^{k_N},\omega), \quad k_i = 1,\ldots,K_i, \omega \in \Omega.$$

Since $\chi^{\tilde{n}} \in Y$ for all $\tilde{n}$ by assumption, it is easy to see that $\hat{\chi} \in Y$ by continuity of the functions in the constraints that define Y. Finally, we have that

$$\begin{aligned} V_E^*(N) \geq \sum_{\omega,\hat{\alpha}} v(\hat{\alpha},\omega)\hat{\chi}(\hat{\alpha},\omega) &= \lim_{\tilde{n}\to\infty} \sum_{\omega,\alpha^{\tilde{n}}} v(\alpha^{\tilde{n}},\omega)\chi^{\tilde{n}}(\alpha^{\tilde{n}},\omega) \\ &\geq \lim_{\tilde{n}\to\infty} \left(V_E^*(N) - \frac{1}{\tilde{n}}\right) = V_E^*(N). \end{aligned}$$

Therefore, $\sum_{\omega,\hat{\alpha}} v(\hat{\alpha},\omega)\hat{\chi}(\hat{\alpha},\omega) = V_E^*(N)$, which implies that $\hat{\chi}$ is a solution to $\mathcal{P}$. $\square$

Proof of Theorem 3. Part 1: $V_{\xi,E} \leq V_\varnothing$**.** Fix ξ. Since for every π, we have $f_{\xi,E} \in \Pi$, if follows that

$$\sup_{\pi\in\Pi} V(f_{\xi,E}(\pi)) = \sup_{\pi'\in f_{\xi,E}(\Pi)} V(\pi') \leq \sup_{\pi\in\Pi} V(\pi) = V_\varnothing.$$

Part 2: $V_E \leq V_{\xi,E}$**.** Let $f_E(\Pi)$ be the set of all information structures that can result from information spillovers under our baseline assumption given network E. For every

$\pi \in f_E(\Pi)$, the Bayesian posteriors of every player i satisfy the property that, for all signal realizations s,

$$Pr_\pi(\omega, s|s_{E_i}) = Pr_\pi(\omega, s|s_i);$$

that is, there is nothing that player i can learn from his sources that is not already contained in his private signal s_i. For such information structures, also any other form of communication captured by ξ cannot add anything to what each player i learns from s_i. That is, if $\pi \in f_E(\Pi)$, then

$$Pr_{f_{\xi,E}(\pi)}(\omega, s|h_i) = Pr_{f_{\xi,E}(\pi)}(\omega, s|h_i^0) = Pr_\pi(\omega, s|s_i).$$

Therefore, if $\sigma \in \text{BNE}(G, \pi')$ for some $\pi' \in f_E(\Pi)$, then we also have $\sigma \in \text{BNE}(G, f_{\xi,E}(\pi'))$. Since $f_E(\Pi) \subset \Pi$, we have

$$\bigcup_{\pi \in f_E(\Pi)} \text{BNE}(G, \pi) \subseteq \bigcup_{\pi \in f_{\xi,E}(\Pi)} \text{BNE}(G, \pi).$$

It follows that

$$\begin{aligned} V_E &= \sup_{\pi \in f_E(\Pi)} \{ \max_{\sigma \in \text{BNE}(G,\pi)} \sum_{a \in A, t \in T, \omega \in \Omega} v(a, \omega)\sigma(a|t)\pi(t|\omega)\mu(\omega) \} \\ &\leq \sup_{\pi \in f_{\xi,E}(\Pi)} \{ \max_{\sigma \in \text{BNE}(G,\pi)} \sum_{a \in A, t \in T, \omega \in \Omega} v(a, \omega)\sigma(a|t)\pi(t|\omega)\mu(\omega) = V_{\xi,E}. \end{aligned}$$

□

Proof of Lemma 3. Define $N_i^0 = \{i\}$ and $N_i^n = \cup_{j \in N_i^{n-1}} N_j$ for $n = 1, \ldots, N$. Note that $N_i^N = E_i$. Fix a signal realization $\overline{s}$ and the corresponding unique $h^K(\overline{s}) = h^K$. For every player i,

$$\xi_i(h_i^0)(\omega, s) = Pr_\pi(\omega, s|\overline{s}_i) = Pr_\pi(\omega, s|\overline{s}_{N_i^0}), \quad (\omega, s) \in \Omega \times S.$$

Note that

$$\sum_{\omega, s_{-i}} Pr_\pi(\omega, s|\overline{s}_i) = \begin{cases} 1 & \text{if } s_i = \overline{s}_i \\ 0 & \text{otherwise.} \end{cases}$$

Fix $n \geq 1$. Given h_i^n, suppose that for every player j,

$$\xi_j(h_j^{n-1})(\omega, s) = Pr_\pi(\omega, s|\overline{s}_{N_j^{n-1}}), \quad (\omega, s) \in \Omega \times S.$$

Note that

$$\sum_{\omega, s_{-j}} Pr_\pi(\omega, s|\overline{s}_{N_j^{n-1}}) = \begin{cases} 1 & \text{if } s_{N_j^{n-1}} = \overline{s}_{N_j^{n-1}} \\ 0 & \text{otherwise.} \end{cases}$$

Therefore,

$$\xi_i(h_i^n)(\omega, s) = Pr_\pi(\omega, s|\overline{s}_{N_i^n}), \quad (\omega, s) \in \Omega \times S.$$

Since this is true for every i, by induction we have that

$$\xi_i(h_i^K)(\omega, s) = Pr_\pi(\omega, s|\overline{s}_{N_i^N}) = Pr_\pi(\omega, s|\overline{s}_{E_i}), \quad (\omega, s) \in \Omega \times S,$$

thus concluding the proof. □

Proof of Theorem 4. Part 1 ($\Rightarrow$): Suppose $(S, \pi) \in \Pi_M$ and $\sigma \in \mathrm{BNE}(G, f(\pi))$ induce x. Then, for every i, $h_i^\infty(s)$, and $s \in S$,

$$\sum_{\omega, s'} [u_i(\sigma_i(h_i^\infty(s')), \sigma_{-i}(h_{-i}^\infty(s')), \omega) - u_i(a_i, \sigma_{-i}(h_{-i}^\infty(s')), \omega)] Pr_\pi(\omega, s' | h_i^\infty(s)) \geq 0, \quad a_i \in A_i.$$

By Lemma 3, each agent i learns the signals of all his parents (i.e., s_{E_i}) for all $s \in S$. Therefore, the previous inequalities become, for all i and s_{E_i},

$$\sum_{\omega, s'_{-E_i}} [u_i(\sigma_i(h_i^\infty(s')), \sigma_{-i}(h_{-i}^\infty(s')), \omega) - u_i(a_i, \sigma_{-i}(h_{-i}^\infty(s')), \omega)] Pr_\pi(\omega, s' | s_{E_i}) \geq 0, \quad a_i \in A_i,$$

where $-E_i = N \setminus E_i$. The last inequalities can be written as, for all i and s_{E_i},

$$\sum_{\omega, s'_{-E_i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)] \prod_{j \in N} \mathbb{I}\{\sigma_j(h_j^\infty(s')) = \alpha_j\} Pr_\pi(\omega, s' | s_{E_i}) \geq 0, \quad a_i \in A_i.$$

Now note that

$$\prod_{j \in N} \mathbb{I}\{\sigma_j(h_j^\infty(s')) = \alpha_j\} Pr_\pi(\omega, s' | s_{E_i}) = \frac{Pr_{\pi,\sigma}((\sigma_j(h_j^\infty(s'_{-E_i}, s_{E_i})) = \alpha_j)_{j \in N}, \omega, s'_{-E_i}, s_{E_i})}{Pr_\pi(s_{E_i})}.$$

Using this, we can write the last inequalities as, for all i, α_i, and s_{E_i},

$$\sum_{\omega, s'_{-E_i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)] Pr_{\pi,\sigma}((\sigma_j(h_j^\infty(s'_{-E_i}, s_{E_i})) = \alpha_j)_{j \in N}, \omega, s'_{-E_i}, s_{E_i}) \geq 0, \quad a_i \in A_i,$$

or equivalently,

$$\sum_{\omega, s'_{-E_i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)] Pr_{\pi,\sigma}(s'_{-E_i}, s_{E_i} | (\sigma_j(h_j^\infty(s'_{-E_i}, s_{E_i})) = \alpha_j)_{j \in N}, \omega) \times$$
$$Pr_{\pi,\sigma}((\sigma_j(h_j^\infty(s'_{-E_i}, s_{E_i})) = \alpha_j)_{j \in N}, \omega) \geq 0, \quad a_i \in A_i,$$

or equivalently,

$$\sum_{\omega, s'_{-E_i}, \alpha_{-i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)] Pr_{\pi,\sigma}(s'_{-E_i}, s_{E_i} | (\sigma_j(h_j^\infty(s'_{-E_i}, s_{E_i})) = \alpha_j)_{j \in N}, \omega) x(\alpha_i, \alpha_{-i} | \omega) \mu(\omega) \geq 0$$

Now define $X_i = S_i$ for all i and $\kappa : \Omega \times (\times_{j \in N} \Delta(A_j))) \to \Delta(X)$ as

$$\kappa(x | \alpha, \omega) = Pr_{\pi,\sigma}(x | (\sigma_j(h_j^\infty(x)) = \alpha_j)_{j \in N}, \omega),$$

for every ω and $\alpha \in \mathbf{supp}\, x(\cdot|\omega)$. If $\alpha \notin \mathbf{supp}\, x(\cdot|\omega)$ for any ω, κ can be defined in any arbitrary way. Note that, as for Theorem 2, by knowing $x_{E_i} = s_{E_i}$, player i knows the mixed action of all $j \in E_i$, in particular the unique α_i that can give rise to x_{E_i} under κ. In other words, for all $i \in N$ and x_{E_i} in the support of κ,

$$Pr_\kappa(\alpha_i | x_{E_i}) = \begin{cases} 1 & \text{if } \alpha_i = \sigma_i(h_i^\infty(x_{-E_i}, x_{E_i})), \\ 0 & \text{else.} \end{cases}$$

This condition implies property (2) of κ. Given this, we obtain property (3) that, for all i, α_i, and all x_{E_i},

$$\sum_{\omega, x'_{-E_i}, \alpha_{-i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega)]\kappa(x'_{-E_i}, x_{E_i}|\alpha, \omega)x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i \in A_i.$$

Part 2 (⇐): Suppose x and κ satisfy conditions (1) and (2). First, we construct $\pi_{x\kappa} \in \Pi_M$, by letting $S_i = X_i$ for all i and

$$\pi_{x\kappa}(s|\omega) = \sum_{\alpha} \kappa(s|\alpha, \omega)x(\alpha|\omega), \quad \omega \in \Omega \text{ and } s \in S.$$

Since $|X_i| = 1$ for all $i \notin M$, $\pi_{x\kappa} \in \Pi_M$. Now, we construct the candidate equilibrium strategies σ. Recall that by Lemma 3, for every realization s from $\pi_{x\kappa}$, every $h_i^\infty(s)$ fully reveals $s_{E_i} = x_{E_i}$ to player i, which by property (2) of κ fully reveals the realized α_i under x. Let $\alpha_i(s_{E_i})$ be such a recommendation given s_{E_i}. For every i, consider $\sigma_i : H_i^\infty \to \Delta(A_i)$ defined by

$$\sigma_i(h_i^\infty(s)) = \alpha_i(s_{E_i}),$$

for every realization s under $\pi_{x\kappa}$.

Given this, by property (3), for every i, α_i, and x_{E_i},

$$\sum_{\omega, \alpha_{-i}, x_{-E_i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a'_i, \alpha_{-i}, \omega)]\kappa(x_{E_i}, x_{-E_i}|\alpha, \omega)x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a'_i \in A_i,$$

which is equivalent to

$$\sum_{\omega, \alpha_{-i}, x_{-E_i}} [u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a'_i, \alpha_{-i}, \omega)] \frac{\kappa(x_{E_i}, x_{-E_i}|\alpha, \omega)x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega)}{\sum_{\omega', \alpha'_{-i}, x'_{-E_i}} \kappa(x_{E_i}, x'_{-E_i}|\alpha_i, \alpha'_{-i}, \omega')x(\alpha_i, \alpha'_{-i}|\omega')\mu(\omega')} \geq 0, \quad a'_i \in A_i$$

or equivalently (using property (2) again) for all i and s_{E_i},

$$\sum_{\omega, s'_{-E_i}} [u_i((\alpha_j(s_{E_j}))_{j \in N}, \omega) - u_i(a'_i, (\alpha_j(s_{E_j}))_{j \neq i}, \omega)] Pr_\pi(\omega, s'|s_{E_i}) \geq 0, \quad a'_i \in A_i.$$

Using the definition of σ, these last inequalities boil down to the equilibrium conditions that, for all i and s_{E_i},

$$\sum_{\omega, s'_{-E_i}} [u_i(\sigma_i(h_i^\infty(s')), \sigma_{-i}(h_{-i}^\infty(s')), \omega) - u_i(a_i, \sigma_{-i}(h_{-i}^\infty(s')), \omega)] Pr_x(\omega, s'|s_{E_i}) \geq 0, \quad a_i \in A_i.$$

Therefore, $\sigma \in \text{BNE}(G, f(\pi_{x\kappa}))$. □