Information Systems

Simone Galperti
UC San Diego

Jacopo Perego
Columbia University

May 19, 2020

Abstract

An information system is a primitive structure that defines which agents can initially get information and how such information is then distributed to others. From political and organizational economics to privacy, information systems arise in various contexts and, unlike information itself, can be easily observed empirically. We introduce a methodology to characterize how information systems affect strategic behavior. This involves proving a revelation principle result for a novel class of constrained information design problems. We identify when such systems better distribute information and, as a result, impose more constraints on behavior. This leads to a novel notion of an agent's influence in the system. Finally, we apply our theory to examine how current patterns of news consumption from mass media may affect elections.

Keywords: Information, System, Design, Network, Seeding, Spillover, Privacy

JEL Codes: C72, D82, D83, D85, M3

A previous version of this paper circulated under the title "Belief Meddling in Social Networks." We thank Charles Angelucci, Odilon Câmara, Yeon-Koo Che, Laura Doval, Ben Golub, Emir Kamenica, Elliot Lipnowski, Stephen Morris, Xiaosheng Mu, Andrea Prat, Carlo Prato, Joel Sobel, Alex Wolitzky, and seminar participants at Pittsburgh, Yale, USC, Pompeu Fabra, UCSD, MIT, Harvard, Chicago, Stanford, UCLA, Queen Mary, Duke, Brown, PSE, Bonn, Georgetown, Maryland, Caltech, Carnegie Mellon, CIDE, ITAM, CETC 2019, SAET 2019, ESSET 2019, and SWET 2019, for comments and suggestions. We are grateful to the Cowles Foundation for hospitality. Aleksandr Levkun provided excellent research assistance. All remaining errors are ours. E-mail: sgalperti@ucsd.edu, jp3754@columbia.edu

1 Introduction

Modern societies feature increasingly complex informational environments. What agents know and how they act depend on information they get not only directly from original sources but also indirectly as spillovers from other agents. We think of such an environment as an *information system*, which defines (1) which agents can get information initially and (2) how such information is then distributed to others. These elements capture the bare bones of the informational environment in which the agents interact. As such, information systems are often stable and observable—unlike information itself, which may be hard to measure. This paper proposes a theory of information systems and their implications for the outcomes of strategic interactions. Such a theory is useful for obtaining sharper predictions that have empirical content, despite the lack of details about what information the agents might have.

Examples of information systems exist in various contexts. In the media industry, news outlets are sources of political information, which they distribute to their possibly overlapping audiences. This can be viewed as an information system. Its properties can affect the outcomes of elections or determine the market power of media conglomerates (Anderson et al. (2016)). In some companies, the organizational hierarchy and firewalls control how information that is sourced by specialized divisions is distributed to other divisions. This setup is another example of an information system and can affect the organization's productivity (Garicano and Van Zandt (2013)). In the digital economy, consumers often have to reveal personal data to various platforms, which may then distribute it to third parties. This environment can be captured with an information system to study the complex trade-offs that privacy, or lack thereof, creates for individual and societal welfare (Acquisti et al. (2016)).

Our goal is to characterize which equilibrium outcomes are compatible with a given information system. To provide an answer, we cast our question as a novel class of *constrained* information design problems. We imagine a metaphorical character, called the designer, who provides information to a group of Bayesian agents, but is constrained by the information system in two ways. First, the designer can provide information only to a subset of the agents, called *seeds*. Second, the agents share information with each other through a given communication *network*. These constraints weaken two standard postulates of the information design paradigm: They render it impossible to provide information either directly or privately to some agents. This precludes the use of standard solution methods. Our main methodological contribution is a revelation principle result that allows us to characterize all equilibrium outcomes compatible with an information system. With this tool, we then investigate how such outcomes depend on properties of the information system—namely, the seeds and the network—both in general and in specific applications. Our methodology can be useful in empirical work, as it allows the econometrician to better exploit observables, without making specific assumptions about what information the agents might have. Knowing the information system where firms operate, for example, can sharpen estimates of the effects of entry (e.g., Magnolfi and Roncoroni (2017)), the likelihood of collusion (e.g., Chassang et al. (2019)), or the market power of media outlets (Prat (2018)).

In our model, an information system consists of the set of seeds and the communication network, which is captured by a directed graph. There are three stages. First, the designer chooses what information to provide to the seeds. For each seed, information takes the form of a private signal about some underlying state of the world. Second, the agents share information with their neighbors in the network. Third, given what they learned, either directly from the designer or indirectly from others, the agents interact in an arbitrary finite game. The agents are fully Bayesian. Throughout, our key assumption concerns information spillovers. Modeling spillovers on networks raises serious conceptual challenges depending on which form of communication is assumed. Our baseline model makes a stark but simple assumption: If there is a path in the network from one agent to another, then the latter will learn the signal received by the former. Thus, the spillovers are mechanical and entirely governed by the network. Later, we relax this baseline assumption and consider richer and more realistic forms of spillovers.

Our revelation principle for constrained information design has two pillars. The first shows that, for every information system, there is an auxiliary "direct" information system that induces the same feasible outcomes for every game. In this auxiliary system the designer can seed *all* agents (hence the word "direct"), but information spills over a *richer* network. This network captures the additional constraints of having to rely on the seeds as information intermediaries in order to reach the other agents. In deriving this equivalence result, we exploit an unexpected connection with computer science, from which we borrow a cryptography technique known as "secret sharing" (Shamir (1979)). The second pillar builds on the tradition of direct mechanisms (Myerson (1986)) and characterizes the feasible outcomes of an information system via *robustly* obedient recommendations. Robust obedience requires that each agent be willing to follow his recommended behavior conditional on all the information he receives, either directly from the designer or indirectly from other agents.

This result allows us to study how information systems affect equilibrium play. What are the consequences of changing the network? When is a group of seeds more influential than another? To provide a holistic answer, we avoid focusing on the network or the set of seeds, but rather we introduce a notion of *more-connected* information systems.

This order is related to notions of connectedness from the network literature, but differs by explicitly considering the informational role played by seeds. We show that when an information system becomes more connected, the set of feasible outcomes shrinks for all games; we establish that the converse holds as well. We derive two implications of this result. First, increasing the information spillovers by simply adding links to the network has, in general, ambiguous effects on the set of feasible outcomes. This speaks to the fact that the network and the set of seeds should not be analyzed in isolation, but rather organically as a system. Second, for a given network, our result provides an operational notion to discern when one group of agents is more influential than another, and hence should be seeded. This connects our work to the literature on optimal seeding in networks.

We relax the assumption that information leaks mechanically on the network of the information system. What if information leaks only with some exogenous probability or the network is uncertain? What if the agents communicate strategically with their neighbors? Under some conditions, we can use our methodology to compute bounds for the set of feasible outcomes in these richer settings. Even if the designer does not know how the agents will communicate with each other, it is possible to find a policy that is "robust" to communication in the sense of securing a given probability for some specific behavior of the agents. These bounds are useful in practice because, when communication is endogenous, computing actual solutions can be infeasible.

Finally, we illustrate our theory with a political economy application. A set of news outlets provides information to ideologically polarized voters in a referendum. Voting is costly. An analyst wants to forecast the referendum outcome. Unfortunately, the analyst does not observe what information the voters receive. Instead, the analyst observes which news outlets are active in the market and how their audiences overlap. We study how the analyst's forecast depends on the ability of different outlets to convey news to broad or narrow audiences. Our findings suggest that outlets with narrow audiences may gain influence when voting costs are higher or polarization is lower.

Related Literature. Our work builds on and contributes to the literature on information design (see Bergemann and Morris (2019) and Kamenica (2019) for surveys), which has been applied to study political campaigns, rating systems, financial stress tests, banking regulations, and many other settings. One can view information systems as imposing specific informational constraints on the design problem. These constraints restrict, possibly in very complex ways, the set of available information structures.[1] Studying infor-

---

[1] In a different context, Doval and Skreta (2019) investigate how inter-temporal incentives under limited commitment lead to a constrained information design problem. Brooks et al. (2019) examine how information hierarchies—which differ from our information systems—constrain how much agents know about some underlying state and about other agents' beliefs regarding the state. Mathevet and

mation systems also allows us to assess the robustness of information policies to spillovers and how to modify them accordingly. This analysis highlights some of the limitations of using information, which is easy to replicate and share, as a tool to shape incentives.

Within the information design literature, the closest paper to ours is Bergemann and Morris (2016). However, the two papers ask fundamentally different questions. Bergemann and Morris (2016) examine how information that agents exogenously have constrains feasible outcomes; we ask how agents' ability to share any information they endogenously receive constrains feasible outcomes. In fact, to highlight the contrast, our analysis focuses on the case where the agents have no exogenous information.[2]

The class of problems introduced in this paper cannot be analyzed using the common solution methods in the literature. The method based on distributions over induced posterior beliefs of Kamenica and Gentzkow (2011)—extended to games by Mathevet et al. (2020)—cannot be productively used here because the new constraints severely complicate the structure of the possible posterior distributions. The approach based on recommendation mechanisms of Bergemann and Morris (2016) and Taneva (2019), which builds on Myerson (1986), is not applicable as we explain in Section 3.1. We develop an alternative method that preserves most of the convenient features of the standard revelation principle.

Our paper also provides a partial bridge between the literature on information design and the literature on optimal seeding in networks.[3] A broad theme in the latter is which network nodes should be targeted with an intervention to achieve maximal diffusion. This has led to several important notions of centrality. Our paper differs from papers in that literature in at least one of three dimensions: We import the Bayesian information model and specialize our interventions to information provision, we allow for arbitrary strategic interactions among agents, and we consider arbitrary objectives for the designer. This flexibility is common in information design, less so in the seeding literature. It allows us to revisit the notions of centrality and influence from a different perspective. We hope the bridge provided by this paper paves the way to a general theory of optimal information seeding.

Finally, our paper relates to the vast literature on observational social learning.[4] This

---

Taneva (2019) study strategic information spillovers, their implementable outcomes, and their optimality in binary-action environments with strategic complementarities.

[2] Building on our baseline model, Candogan (2020) studies optimal information structures in settings where all agents are seeded, non-strategic, and enter additively in the designer's objective.

[3] See, for example, Morris (2000), Ballester et al. (2006), Banerjee et al. (2013), Sadler (2017), Akbarpour et al. (2018), and Galeotti et al. (2019). Valente (2012) provides a review of the literature.

[4] See Banerjee (1992), Bikhchandani et al. (1992), Smith and Sørensen (2000), Acemoglu et al. (2011),

literature is concerned with how people learn from observing others' behavior, which is a prominent form of information spillovers. Usually, this literature assumes that the agents' initial information is given and simple, and considers rich, sometimes strategic, forms of spillovers. By contrast, we impose no restrictions on the initial information structures. In our baseline model, we consider a simple and tractable form of spillovers, which we then extend to richer—including strategic—spillovers in Section 5.

2 Model

We first outline the settings we are interested in modeling. There is a group of agents, hereafter called players, who are linked by social relations. Initially, some—possibly all—players privately observe a signal about some state of the world. The players then share their information through their social relations. Finally, using all the information obtained either initially or from their peers, the players interact in a game. As we discuss below, the initial signals can be interpreted in two ways: Either (1) they are part of the thought experiment of an analyst who is agnostic about the players' initial information, but who wants to understand what behaviors they may adopt (metaphorical interpretation), or (2) they are provided by a third party who wants to influence the players' behavior (literal interpretation).

The following elements of the model are standard. Let I be a finite set of players, where I also denotes the number of players. Let Ω be a finite set of states of the world. The players have a common, fully supported, prior belief $\mu \in \Delta(\Omega)$. Let A_i be a finite set of actions for player i whose utility function is $u_i : A \times \Omega \to \mathbb{R}$, where $A = A_1 \times \cdots \times A_I$. Fixing I, we denote the *base game* by $G = (\Omega, \mu, (A_i, u_i)_{i \in I})$. The information the players have when playing G is described by an information structure. This is denoted by (T, π_I) and consists of a finite signal space T_i for each player i and a function $\pi_I : \Omega \to \Delta(T)$, where $T = T_1 \times \cdots \times T_I$. Abusing notation, we sometimes write π_I instead of (T, π_I). Together, G and any π_I define a Bayesian game, denoted by (G, π_I). We focus on its Bayes-Nash equilibria. We denote strategies of player i by $\sigma_i : T_i \to \Delta(A_i)$ and the set of equilibria by $\mathrm{BNE}(G, \pi_I)$.

The novelty of our model is that the players are organized in an *information system*, which specifies which players can initially get original information about the state and how such information is then distributed to others. In this way, the system determines the players' information π_I in the game. Formally, an information system, denoted by (N, S), consists of a nonempty set of *seeds* $S \subseteq I$ and a *network* $N \subseteq I^2$—a set of directed

---

Golub and Sadler (2017), Golub and Jackson (2012), and Mueller-Frank (2013).

links between the players. The seeds are the players who can get a nontrivial initial signal about the state. We can model this *initial* information as an information structure that satisfies $|T_i| = 1$ for $i \notin S$. To emphasize this property, we denote the initial information structure by π_S. We call the information system *direct* if $S = I$ and *indirect* otherwise.

After the private realization of the initial signals, information spills over the network N. Being directed, a link from i to j can carry information from i to j, but not vice versa. A path from i to j is a sequence of players $(i_1, \dots, i_m)$ such that $i_1 = i$, $i_m = j$, and $(i_k, i_{k+1}) \in N$ for all $k = 1, \dots, m-1$. In the baseline version of the model, we start by assuming that information spills over between players in a mechanical and deterministic way that is governed by N.

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

We discuss and motivate this assumption shortly in Section 2.1. We relax it in Section 5.

For most of the analysis, we will focus on information systems that are *connected* in an informational sense. To define this, we call player j a *source* of player i if there is a path from j to i. Given N, let N_i be the set that contains i and all i's sources. We call j a *seed source* of i if $j \in N_i \cap S$. An information system is connected if every player has at least one seed source and therefore can receive some information. We relax this assumption in Appendix D.2.

Assumption 2. *The information system (N, S) is* connected*, that is, $N_i \cap S \neq \varnothing$ for all i.*

These assumptions imply that N transforms every *initial* π_S into a unique *final* information structure, which we denote by $\pi_I = f_N(\pi_S)$. Thus, every π_S leads to the Bayesian game $(G, f_N(\pi_S))$.

The main goal of this paper is to characterize what joint distributions of players' actions and states can arise across initial information structures given the game G and the information system (N, S). For the sake of presentation, it is convenient to imagine that an information designer chooses π_S. In some applications, we also specify a payoff function $v : A \times \Omega \to \mathbb{R}$ for the designer. For every π_S, define

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

where $f_N(\pi_S)(t|\omega)$ is the probability that the final information structure $f_N(\pi_S)$ assigns to the profile of signals t in state ω. The designer-preferred equilibrium selection is standard in the literature (Bergemann and Morris (2019)). The designer maximizes $V(G, N, \pi_S)$. This problem can be expressed formally, but doing so involves some technicalities tackled in Appendix A. We denote the value of this problem by $V^*(G, N, S)$. For convenience, we will refer to the designer using feminine pronouns and to a player using masculine pronouns.

2.1 Discussion of the Model

Information Spillovers. A key theme of this paper is to consider spillovers that are beyond the designer's control and occur following specific observable patterns. We find this of practical importance, especially in applications where it is difficult to enforce privacy because communication among agents is the norm. To analyze these issues, the information design framework needs to be extended to settings where the players can share the information they receive. In doing so, we highlight some limitations of using information—which (unlike money) is a good that is easy to replicate and can be non rivalrous—as a tool to shape incentives. To make progress, our baseline model assumes mechanical and deterministic spillovers on an exogenous network. Although descriptively relevant in several applications (e.g., see Section 4.2), this is a stark assumption. It is, however, a natural starting point for several reasons.

First, Assumption 1 lies at the polar opposite of the standard information design framework, which rules out communication among agents that is beyond the designer's control. By doing so, our baseline model transparently highlights the trade-offs and qualitative implications of information spillovers.

Second, mechanical spillovers are also a typical assumption in the literature on information diffusion and seeding.[5] The classic DeGroot model of information diffusion shares some of the mechanical and non-strategic features of our spillovers. A key difference is that in our model updating is Bayesian. From this perspective, our paper bridges that literature with the literature on information design.

Last but not least, our baseline model proves to be a useful steppingstone to obtaining robust predictions for more general forms of spillovers. Examples include strategic communication on the network, which can depend on the initial information structure itself, and mechanical but random spillovers. Although clearly important, such spillovers raise significant challenges that render them intractable, except in special cases. We discuss these extensions and issues in Section 5. In a nutshell, some challenges are computational: Calculating for every π_S how the players communicate over the network in equilibrium may be hard, and their strategies may be complex or unrealistic (for a discussion, see Eyster and Rabin (2010) and Mueller-Frank (2013)). Other challenges are con-

---

[5] For example, see DeMarzo et al. (2003), Golub and Jackson (2012), Akbarpour et al. (2018), and Jackson and Storms (2018).

ceptual. Many communication models (such as cheap talk) or solution concepts (such as communication equilibrium) would require a reasonable criterion to select among multiple equilibria in the communication stage of our model. For instance, in the case of cheap talk the designer-preferred equilibrium would be a babbling equilibrium, which effectively kills any spillover and thus defeats the purpose of this paper. Section 5 shows that Assumption 1 corresponds to the designer's worst-case scenario across a broad class of spillovers. Thus, the methods we develop under Assumption 1 become a tool for obtaining predictions that are robust against how linked players may share their information due to strategic or technological considerations.

Interpretations. An information system imposes constraints on the information design problem. It may be impossible to (1) directly provide information to some players (seeding), and (2) prevent information from leaking between players (spillovers). Thus, the system restricts what type of information the players may have before playing G. While in some settings different kinds of restrictions may be relevant (e.g., caps on informativeness), in many applications constraints (1) and (2) are reasonable and important. Moreover, they can be interpreted in two ways.

In a metaphorical interpretation, the designer is an analyst who wants to obtain robust predictions about the possible behavior of the players in G. She is agnostic about which information the players initially get, which is often hard to observe empirically. However, she may know the information system, which is a more primitive object (see Section 4.2 for a concrete example). That is, she may know which players can bring information into the system and how it usually spills over from these players to others. The system then becomes a tool to capture the analyst's knowledge of the informational environment in which the players interact. It represents a tractable way to model basic assumptions on what the players know about each others' information. Finally, note that the analyst can pick the function v instrumentally to learn specific aspects of the players' possible behavior. For example, if v equals 1 for specific action profiles and 0 otherwise, she can learn the maximal feasible probability of those profiles.

This metaphorical interpretation can be especially useful in empirical work. Our methodology allows the econometrician to better exploit observables in the data, without making specific assumptions about the information that agents might have. For example, in entry games (e.g., Magnolfi and Roncoroni (2017)), it can be reasonable that stores in the same chain share information, whereas stores in different chains do not. A similar observation can have useful implications when developing "safe tests" á la Chassang et al. (2019). Finally, explicitly allowing for information spillovers between voters can improve our estimates of media power (Prat (2018)).

In a literal interpretation, the designer is a third party interested in influencing the players' behavior. For instance, think of this party as the head of an organization. In many applications, this party provides all the information that the players have—she chooses the initial information structure. However, she may face explicit *constraints* that can be conveniently captured by an information system. For instance, the head of an organization may be able or allowed to directly provide information only to some of its members, or she may not be able to communicate to some members without others also listening. More generally, N can represent the hierarchical structure of the organization, describing the flow of information through the levels of its divisions.

3 Feasible Outcomes Under Information Systems

This section provides a general characterization of what behaviors of the players can be induced by the initial information structure given an arbitrary base game G and an information system (N, S). We will refer to such behaviors as the feasible outcomes. By Assumption 1, player i learns the vector t_{N_i}, the initial signal realizations of all his sources. By Assumption 2, this vector can have multiple possible realizations for every i under the final information structure. Therefore, given any initial (T, π_S), the final $\pi'_I = f_N(\pi_S)$ is defined by $T'_i = \times_{j \in N_i} T_j$ for all $i \in I$ and, for all $\omega \in \Omega$, $\pi'_I(t'_1, \ldots, t'_I|\omega) = \pi_S(t_1, \ldots, t_I|\omega)$ when $t'_i = t_{N_i}$ for all i.

It is convenient to formalize the feasible outcomes as distributions over *how the players play* in the base game. Since each player acts on his final information by picking an action, possibly at random, we can view information structures as inducing distributions over mixed-action profiles for every state. It is immediate to deduce how these profiles determine the final actions. Since by assumption each π_S has finite support, we can only consider finite-support distributions. Let $R_i = \Delta(A_i)$ and $R = \times_{i \in I} R_i$. We will refer to $\alpha = (\alpha_1, \ldots, \alpha_I) \in R$ as an outcome *realization*.[6]

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

Even if any initial information structure π_S is possible, only some outcomes are feasible. In words, x is feasible if an initial π_S and an equilibrium σ of the final game exist that induce the same distribution over mixed actions as x.

Definition 2 (Feasible Outcome). An outcome x is *feasible for* (G, N, S) if π_S and

---

[6] The reader may wonder why we do not define outcomes as mappings to *pure* actions only. Appendix A shows that, in contrast to unconstrained information design, such a definition involves a loss of generality when there are spillovers (independently of S).

$\sigma \in \mathrm{BNE}(G, f_N(\pi_S))$ exist such that, for every $\omega \in \Omega$ and $\alpha \in R$,

$$x(\alpha_1, \dots, \alpha_I|\omega) = \sum_{t \in T} \pi_S(t|\omega) \prod_{i \in I} \mathbb{I}\{\sigma_i(t_{N_i}) = \alpha_i\}, \tag{1}$$

where $\mathbb{I}\{\cdot\}$ is the indicator function. Let $X(G, N, S)$ be the set of such outcomes.

Information systems add new challenges to the characterization of feasible outcomes. On the one hand, the approach based on feasible distributions over induced posterior beliefs of Kamenica and Gentzkow (2011), which was extended to games by Mathevet et al. (2020), cannot be tractably used here. Our new constraints severely complicate calculating those posterior distributions. On the other hand, the approach based on recommendation mechanisms of Bergemann and Morris (2016) and Taneva (2019), which builds on Myerson (1986), is not applicable here for subtle reasons that we explain in Example 1 below.

Given these challenges, we develop an alternative approach, which has two parts. In Section 3.1, we transform the original information system into a direct one that leads to the same set of feasible outcomes. In Section 3.2, we characterize all feasible outcomes for this auxiliary problem in terms of recommendation mechanisms that are *robust* to information spillovers.

3.1 From Indirect to Direct Information Systems

We now show how to turn the original problem into an outcome-equivalent auxiliary problem where *all* players are seeds. This involves specifying a new network that is "richer" than N and depends on (N, S). We will introduce the logic of this construction using three examples and then present the general treatment.

The first example illustrates that the inability to seed all players is only part of the reason the standard recommendation-mechanism approach does not apply here. In some instances of our problem, we can describe a feasible outcome using the language of recommendations, but in others we simply cannot.

Example 1. Let $I = \{1, 2, 3\}$ and $S = \{1, 2\}$. As shown in Figure 1, player 1 and player 2 are the seeds (in gray) and are sources of player 3, who is not a seed. Suppose each player wants to match the state and let $\bar{a}$ be each players' optimal action under no information. Consider first an outcome where 3's behavior depends only on t_1. Concretely, suppose only 1 and 3 match the state, while 2 plays $\bar{a}$. In this case, we can replace t_1 with $t_1' = (a_1, a_3)$, where both actions equal the realized state, and interpret t_1' as a recommendation to both 1 and 3. By contrast, consider an outcome where 3's behavior depends on (t_1, t_2) in such a way that neither t_1 nor t_2 alone pins down a_3. Concretely, suppose only 3 matches the state, while 1 and 2 receive no information and play $\bar{a}$ (this can be done, as we show in the next example). If keeping 1 and 2 uncertain about a_3 is necessary to sustain such an outcome, we cannot restrict the designer to communicating in a language such that t'_1, or t'_2 for that matter, explicitly recommends an action for 3. △

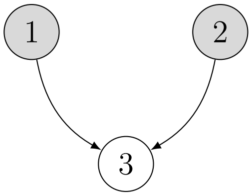

Figure 1: A Non-Seed Player (in White) with Unconnected Seed Sources (in Gray)

Example 1 illustrates the heart of the issue with recommendation mechanisms. A non-seed player j may receive from each of his sources only *part* of the information determining his behavior. These parts may be distributed among his seed sources so that none of their signals can be reduced to a recommendation to j without revealing too much information to j's sources and, thereby, changing their behavior.

The ability to divide a message intended for a non-seed player across multiple sources is also a key to relaxing the constraints imposed by seeding. Note that in Example 1, players 1 and 2 are not sources of each other; we therefore call them independent sources of player 3. It may then be possible to design 1's and 2's signals so that together they allow 3 to learn the message the designer intends to send him, yet each signal alone reveals nothing about that message (see Example 2). Such signals restore a direct and private communication from the designer to player 3. The crucial implication for us is that it is *as if* we can add 3 himself to the seeds.

Example 2. Consider again Figure 1. Suppose the state is binary: $\omega \in \{b, g\}$. We construct an information structure that allows player 3 to learn the state, while revealing nothing to players 1 and 2. Let $\kappa : \{b, g\} \to \{0, 1\}$ be any injective function. Let $T_1 = \{0, 1\}$. Player 1 receives either signal with equal probability, independently of ω and the signal of player 2. Let $T_2 = \{(\kappa, 0), (\kappa, 1)\}$, where $t_2 = (\kappa, 1)$ if and only if $\kappa(\omega) = t_1$, and $t_2 = (\kappa, 0)$ if and only if $\kappa(\omega) \neq t_1$. Clearly, t_1 is uninformative. Signal t_2 is also uninformative, because for every ω signal t_1 is equally likely to coincide with $\kappa(\omega)$. However, the pair (t_1, t_2) fully reveals ω. In words, κ is an encryption code, the second part of t_2 is an encrypted version of ω, and t_1 is the key to using the code to decipher t_2. △

This logic allows us to expand the seed set, but does not yet imply that we can expand it from S to I, thereby fully relaxing the seeding constraint. To see how to do this, consider the following example.

Figure 2: A Non-Seed Player with Connected Seed Sources

Example 3. Figure 2(a) modifies Figure 1 by adding link $(2, 1)$. Player 3's seed sources are no longer independent, so it is not possible to send encrypted messages to 3 that 1 will not decipher too. Even in such cases, we can add 3 to the seeds by first adding links to the original network. Here, 3's information either comes from 2, who is a source of 1, or directly from 1 himself, so 1 always knows what 3 knows. We can capture this by adding a link from 3 to 1, represented by the dashed arrow in Figure 2(b). We can then add player 3 to the seeds. In the expanded network where all players are seeds, independently of what signals they get initially, 1 and 3 will have the same final information, which is what ultimately happens in the original system. △

We now formalize and generalize these ideas.

Definition 3 (S-expansion)**.** The S-expansion of N is the network N^S that contains N and is obtained as follows: If $i \notin N_j$, we add link (i, j) to N if and only if $N_i \cap S \subseteq N_j$.

The logic behind this definition is that if all seed sources of i (i.e., $N_i \cap S$) are also sources of j, then j must infer all the information i could ever get. Adding a link from i to j simply makes this property explicit. Our notion of S-expansion adds exactly the links that are needed and no more. Figure 3 gives two other examples of expansions.

Our first main result follows. By appropriately enriching the network, we can turn any information system into a direct one that leads to the same feasible outcomes. This equivalence is crucial for the rest of our analysis.

Theorem 1 (Equivalence)**.** *$X(G, N, S) = X(G, N^S, I)$ for all (G, N, S).*

The proof builds on the insights from the examples above. First, we show that the S-expansion does not change the seed sources of any player and consequently it does not change the information on which each can ultimately act either. Second, we show that we can treat all players as seeds in the S-expansion. If a non-seed player i has independent sources—in the sense used before—then it is possible to provide information privately to i using only the seeds in S, as in Example 2. This may require dividing the message

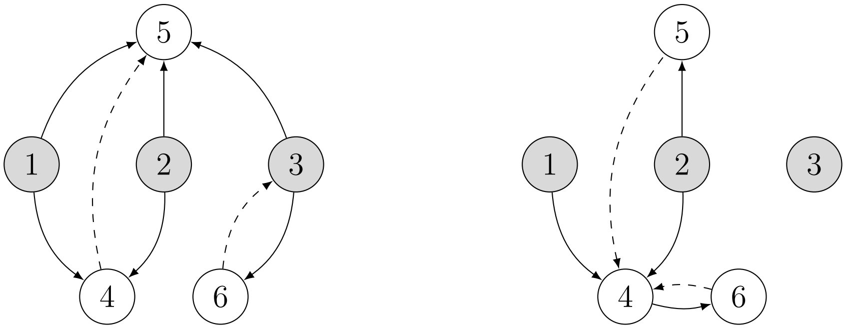

Figure 3: Examples of S-Expansions

intended for i across *all* his independent sources (not just two). If a non-seed player j does not have independent sources, he must have a bi-directional path to some player who is a seed or who has independent sources. Therefore, we can add j to the seeds as in Example 3.

In our argument, the designer can exploit the seeds as intermediaries to communicate *secretly* with non-seed players. This uses an encryption method called *secret sharing* (Shamir (1979)): A secret can be divided into "shares" so that pooling all shares reveals it, yet missing even one share renders the rest fully uninformative (see Example 2).[7] Of course, this method relies on the richness of information structures in our model. We do not interpret it in a descriptive way, as a prediction of how information is actually provided. More simply, we view this method as a proof technique. Nonetheless, once we abstract from the details, one substantive point emerges from the proof. When information goes through multiple independent channels, it is possible that some listeners will understand its intended meaning, while others who lack access to some of those channels will not. This is reminiscent of a kind of communication that is known in politics as *dog whistling*. This term refers to employing a coded language that only an intended subgroup of people will be able to understand, perhaps because they learned its "key" from sources that are not available to everyone.

3.2 Spillover-Robust Obedience

We can now proceed to the second part of our characterization approach. By Theorem 1, we can focus on direct systems where all players can receive information initially. Abusing

---

[7] Outside of the information design literature, Renou and Tomala (2012), Renault et al. (2014), and Rivera (2018) also use cryptography techniques. They seek to identify conditions of networks such that privacy of communication is guaranteed and the standard revelation principle holds. By contrast, we are interested in settings where privacy can fail and, thus, consider arbitrary information systems. Consequently, we have to develop new methods to characterize feasible outcomes.

notation, let

$$u_i(\alpha_i, \alpha_{-i}; \omega) = \sum_{a \in A} u_i(a; \omega) \prod_{j \in I} \alpha_j(a_j).$$

Given outcome x, let its support be $\boldsymbol{x} = \{\alpha \in R : x(\alpha|\omega) > 0, \omega \in \Omega\}$. Let the projection of $\boldsymbol{x}$ on any subset of players $I' \subseteq I$ be $\boldsymbol{x}_{I'} = \{\alpha \in R_{I'} : (\alpha, \hat{\alpha}) \in \boldsymbol{x}, \hat{\alpha} \in R_{-I'}\}$. Finally, let $-N_i = I \setminus N_i$.

Definition 4 (Spillover-Robust Obedience)**.** An outcome x is *spillover-robust obedient* for (G, N) if, for all $i = 1, \dots, I$ and $\alpha_{N_i} \in \boldsymbol{x}_{N_i}$,

$$\sum_{\substack{\omega \in \Omega \\ \alpha_{-N_i} \in \boldsymbol{x}_{-N_i}}} \Big(u_i(\alpha_i, \alpha_{-i}; \omega) - u_i(a_i, \alpha_{-i}; \omega)\Big) x(\alpha_i, \alpha_{-i}|\omega)\mu(\omega) \geq 0, \quad a_i \in A_i. \tag{2}$$

To interpret this definition, imagine dividing the left-hand side of condition (2) by the total probability that α_{N_i} arises under x and μ. We obtain an equivalent condition which requires that player i prefers to play α_i than any other action conditional on the information conveyed not only by knowing α_i, but also by knowing α_j for all his sources.

This leads to our second main result.

Theorem 2 (Feasibility)**.** *An outcome x is feasible for (G, N, I)—i.e., $x \in X(G, N, I)$—if and only if x is spillover-robust obedient for (G, N).*

Robust obedience captures the basic economic trade-off caused by information spillovers. As usual, the signal for each player directly influences his beliefs. In our setting, it can also influence the beliefs of a player's followers in the network. This curbs the scope for keeping them uncertain about that player's behavior, which may also reveal information about other players and the state. Put differently, spillovers render it harder—in the sense of incentive compatibility—to implement correlated behaviors that require some dependence on ω and mutual uncertainty among players. This is illustrated in our applications in Section 4.

Methodologically, the *combination* of Theorem 1 and 2 provides a revelation principle for constrained information design problems that lack both private and direct information provision. It allows us to still view feasible outcomes as if the designer directly recommends to each player how to play in G and each abides by her recommendations, irrespective of the spillovers after all recommendations are released. Due to the linearity of the obedience constraints, we can use powerful linear programming methods to solve our constrained information design problems and study comparative statics, as we illustrate in Section 4.[8]

---

[8] In Galperti and Perego (2018), we use linear programming duality to characterize optimal outcomes in standard information design problems. Due to the linearity of the robust obedience constraints, we can readily apply those techniques to the class of problems considered here.

The intuition for Theorem 2 is simple. Suppose π_I and a BNE σ induce x. Note that by learning his sources' signals through N, player i also learns the signals of *his sources'* sources and so on. Since in equilibrium i knows σ, he can predict the mixed action of all his sources. 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 property is robust obedience.[9] Conversely, suppose x is robust obedient. We can view x itself as an information structure. It is then a BNE of $(G, f_N(x))$ for each player to follow his recommendation, given what he learns through the spillovers and that the others follow their recommendations.

A by-product of Theorem 2 is that it offers a simple way of characterizing feasible outcomes when the designer can use only public information—a case often studied in the literature. It is easy to see that we can model public information using a complete N. In this case, condition (2) simplifies to requiring that, for all $i \in I$ and $(\alpha_i, \alpha_{-i}) \in \boldsymbol{x}$,

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

Three natural questions may arise at this point: Why are mixed recommendations necessary in Theorem 2? Is there no loss of generality in focusing on information structures and therefore on outcomes with finite support? Does our constrained information design problem always have a solution? Since these are rather technical questions, we postpone the answers to Appendix A.

4 The Effects of Information Systems

In this section, we put our theorems to work and analyze a number of applications that offer economic insights into information systems and showcase the flexibility of our framework. Section 4.1 introduces a notion of more-connected information systems and studies the effects on feasible outcomes. We then exploit this notion to describe the influence of different seeds and offer insights into how one might optimally choose them. Section 4.2 studies a political economy application where news outlets provide information to ideologically polarized voters in an upcoming referendum. The market structure of the news outlets and their audiences determine the information system. Our goal is to study its effects on the outcome of the referendum.

---

[9] In this light, spillover-robust obedience may seem related to the notion of peer-confirming equilibrium (Lipnowski and Sadler (2019)). However, its key idea is that a player has correct conjectures about the *strategies* (not actions) of the players to which he is linked, but faces *strategic* (not information) uncertainty about the remaining players in the spirit of rationalizability.

$(\hat{N}, \hat{S})$ (N, S)

Figure 4: Information Systems that are Not Ranked

4.1 More-Connected Information Systems

What does it mean for an information system to become more connected and how does this affect the feasible outcomes? Before presenting our answer, it is worth noting that seeding and spillovers give rise to nontrivial trade-offs. Given S, on the one hand, richer spillovers curb the scope for privately influencing the players' behavior, which should shrink the set of feasible outcomes. On the other hand, richer spillovers can open new channels to reach the players, which may expand the feasible set. To formalize these intuitions, we introduce the following order on information systems.

Definition 5. Fix I. (N, S) is *more connected* than $(\hat{N}, \hat{S})$—denoted by $(N, S) \trianglerighteq (\hat{N}, \hat{S})$—if, for all $i \in I$, i's sources in $\hat{N}^{\hat{S}}$ are also sources in N^S (i.e., $\hat{N}_i^{\hat{S}} \subseteq N_i^S$).

This order captures the idea that, in more-connected systems, players have information that is more correlated. It is related to notions of connectedness from the network literature, but differs by explicitly taking into account the informational role played by seeds. For example, consider the systems in Figure 4. Intuitively, $\hat{N} = \{(1, 3)\}$ is less connected than $N = \{(1, 3), (2, 3)\}$. However, $(N, S) \ntrianglerighteq (\hat{N}, \hat{S})$. To see this, note that in $(\hat{N}, \hat{S})$ the information players 1 and 3 have must be perfectly correlated, while in (N, S) it can be only partially correlated because s_1 and s_2 can be uncorrelated. Definition 5 is useful because it exactly characterizes when changes in the information system shrink the set of feasible outcomes.

Lemma 1. *$X(G, N, S) \subseteq X(G, \hat{N}, \hat{S})$ for all G if and only if $(N, S) \trianglerighteq (\hat{N}, \hat{S})$.*

When an information system becomes more connected, "local" information received by the seeds can more easily spread "globally." This shrinks the set of equilibria that can be achieved, irrespective of the game being played.[10] An interesting asymmetry emerges, which is particularly visible when $S = I$: While the designer can always replicate more spillovers by telling each player what he may learn from his new sources, she cannot undo

---

[10] Related to this, it can be shown that $(N, S) \trianglerighteq (\hat{N}, \hat{S})$ if and only if (N, S) "better aggregates" the information received by the players. We formalize this point in Appendix D.3.

them. This suggests an important difference between settings where players influence each other by sharing information and settings where they use side monetary transfers. Assuming transferable utility for simplicity, a central authority with no budget constraint can always undo side transfers through offsetting transfers of its own to the players. By contrast, it cannot undo information spillovers. In short, while money can be taken back, information—once leaked—cannot.

Before putting our order to work, we clarify its relationship with the primitives of our model, without resorting to network expansions. To this end, consider first direct systems: In this case, $(N, I) \trianglerighteq (\hat{N}, I)$ if and only if $\hat{N}_i \subseteq N_i$ for all $i \in I$.[11] Thus, when all players can be seeded, our notion of more-connected information systems coincides with an intuitive notion of more-connected networks. By contrast, when the system is indirect, our order takes into account that information originates from seeds only; hence, not all links in the network are equally important. The next result formalizes this intuition.

Proposition 1. *$(N, S) \trianglerighteq (\hat{N}, \hat{S})$ if and only if $\hat{N}_i \cap \hat{S} \subseteq \hat{N}_j$ implies $N_i \cap S \subseteq N_j$ for all $i, j \in I$.*

Intuitively, if in $(\hat{N}, \hat{S})$ all seed sources of player i are also sources of player j, then j knows i's information. This should also be true in (N, S) for it to be more connected. Therefore, in (N, S) either i is already a source of j, or all seed sources of i must again be sources of j.

4.1.1 The Effects of More-Connected Information Systems

To further understand the effects of more-connected information systems, we start again by considering direct systems. In this case, as confirmed by Lemma 1, having more channels to reach players should be irrelevant, so more connections should always shrink the feasible set. We illustrate this with an example.

Example 4 (Investment Game with Spillovers)**.** Two investment banks ($I = \{1, 2\}$) have to choose whether to invest ($a = y$) or not ($a = n$) in a pharmaceutical startup. Its business idea may be successful ($\omega = g$) or not ($\omega = b$) with equal probability. Table 1 shows the banks' payoff functions. To make the example more interesting, let the return from a successful investment differ between banks: $0 < \gamma_2 < \gamma_1$. Assume $\gamma_1 < 1$ so that under the prior neither bank invests. Joint investment causes externalities: If both banks sit on the startup's board, they may help or obstruct each other in running it. We discuss the case where the decisions to invest are strategic substitutes ($\varepsilon < 0$) here and the case of strategic complements ($\varepsilon > 0$) in Appendix C. Assume that ε is small (i.e., $|\varepsilon| \approx 0$).

---

[11] Note that $\hat{N} \subseteq N$ implies $(N, I) \trianglerighteq (\hat{N}, I)$, but the converse is not true. For example, neither $\hat{N} = \{(1, 2), (1, 3), (2, 3)\}$ nor $N = \{(1, 2), (2, 3), (4, 2)\}$ is contained in the other, yet $(N, I) \trianglerighteq (\hat{N}, I)$. This is because the primitive network may not be transitive.

	y_2	n_2
y_1	$\gamma_1 + \varepsilon\ ,\ \gamma_2 + \varepsilon$	$\gamma_1\ ,\ 0$
n_1	$0\ ,\ \gamma_2$	$0\ ,\ 0$

$\omega = g$

	y_2	n_2
y_1	$-1 + \varepsilon\ ,\ -1 + \varepsilon$	$-\,1\ ,\ 0$
n_1	$0\ ,\ -1$	$0\ ,\ 0$

$\omega = b$

Table 1: Investment-Game Payoffs

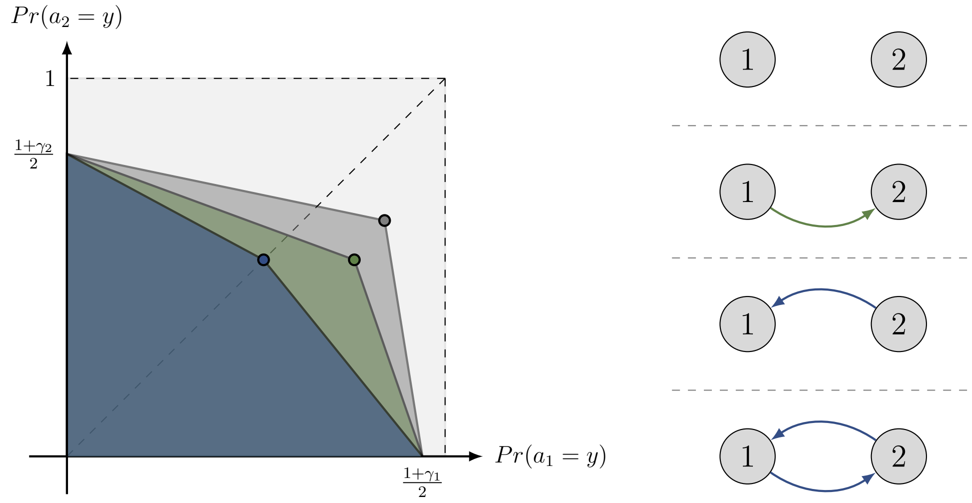

Figure 5: Investment-Game Feasible Outcomes

Suppose the startup (the designer) wants to persuade the banks to invest: for all ω,

$$v(a_1, a_2; \omega) = \mathbb{I}\{a_1 = y_1\} + \mathbb{I}\{a_2 = y_2\}.$$

To this end, it can develop in-house tests of product prototypes and submit the results to each bank. We interpret this as choosing π_S. Since the startup chooses which test to send to each bank, we have $S = I$.

We are interested in the feasible outcomes for the networks in the right panel of Figure 5. One way to interpret the link (i, j) is that bank j has an informant in bank i who leaks the information i gets about the startup to j. We will describe outcomes in terms of the total probability that each bank invests. The formal steps are in Appendix C.

Figure 5 can be understood as follows. The intercepts of the sets with each axis represent the solutions to the Bayesian-persuasion problem of maximizing the probability that a single bank invests. With regard to the boundaries, as the probability that one bank invests rises, the maximal probability the other can be induced to invest falls due to strategic substitutability. Spillovers from i to j shrink the feasible set for two reasons. First, j is better informed. Second, knowing part of j's information allows i to better predict when j invests, which reduces i's willingness to invest due to strategic substitutability. Interestingly, the link $(2,1)$ shrinks the feasible set more than the link $(1,2)$ does. This is because, due to $\gamma_2 < \gamma_1$, persuading bank 2 to invest requires a more informative signal, whose leakage to bank 1 is therefore more damaging.

The solution to the startup's problem corresponds to the dotted corner of each set in Figure 5. Appendix C reports the actual optimal outcomes, whose main features are as follows. For $N = \varnothing$ and $N = \{(1,2)\}$, it is optimal to use private information: When it is recommended that a bank invest, that bank is kept uncertain about whether the other invests. This exploits strategic substitutability to relax the obedience constraints. Instead, for $N = \{(2,1)\}$ it is optimal to use public information (i.e., perfectly correlated recommendations). Thus, the more-constraining spillover from 2 to 1 removes all the value from providing information privately to 1, even though the network per se does not rule out this option (as when $N = \{(1,2),(2,1)\}$). △

This asymmetry between links illustrates a general feature of our class of problems. Because it may require different information to persuade two players to take an action, the direction of spillovers between them can have significantly different consequences for the designer. In contrast to standard seeding problems, here the question is not only who to provide with information, but also what information to provide.

Next, we illustrate the effects of richer spillovers for indirect information systems. When $S \subsetneq I$, there is a trade-off between losing privacy and gaining channels to reach the players. Building on Example 4, we show that more connections can change the set of feasible outcomes non-monotonically.

Example 5 (Investment Game with Indirect Information Systems)**.** Suppose the startup is unable to develop in-house tests due to the nature of its product. Instead, it needs to hire independent labs to test its prototypes. Labs have existing relationships with the banks, to which they promptly release the test results. Thus, the startup can no longer directly communicate with the banks, but can selectively choose which labs to hire. We can model this as a link from the labs to the banks. There are three possible labs: $S = \{3,4,5\}$. Theorem 1 allows us to immediately see the effects of which lab is linked to which bank by looking at the S-expansions. Suppose lab 3 is the only source of both banks: $N = \{(3,1),(3,2)\}$ as in Figure 6(a). Since $N^S = N \cup \{(1,2),(2,1)\}$, the feasible set is the smallest in Figure 5. Now suppose lab 5 is also a source of bank 2: $\tilde{N} = \{(3,1),(3,2),(5,2)\}$ as in Figure 6(b). Since $\tilde{N}^S = \tilde{N} \cup \{(1,2)\}$, the feasible set becomes

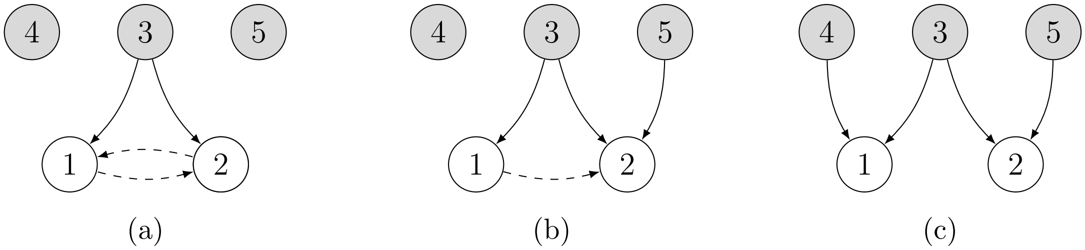

Figure 6: Investment Game with Indirect Information Systems

the intermediate one in Figure 5.[12] Finally, suppose lab 4 is also a source of bank 1: $\bar{N} = \{(3,1),(3,2),(4,2),(5,1)\}$ as in Figure 6(c). Now the startup can provide information directly and privately to each bank, which results in the largest set in Figure 5. It is easy to see that adding links to $\bar{N}$ will shrink the feasible set back to the smallest one in Figure 5. Thus, the startup prefers an intermediate level of connectedness. △

One may wonder whether more-connected information systems benefit the players. When information provision is endogenous, the effect of more connections on the players is ambiguous. The resulting lack of privacy and new ways to reach the same players may lead the designer to provide more or less information in the first place, which in turn may help or hurt the players.[13]

Understanding which effect prevails may be of particular interest in applications of information design to elections. Heese and Lauermann (2019) ask how much a designer can undermine the information-aggregation result of Feddersen and Pesendorfer (1997) by providing information to voters in addition to their exogenous signals. Perhaps surprisingly, the designer can entirely break the result. Crucially for us, she does so by providing information privately to each voter. In Chan et al. (2019), a designer can improve the chances that voters support her desired policy by again relying on private information. The key is to keep pivotal voters uncertain about *which* other voters support the policy. However, among other things, spillovers on social networks exactly limit private information provision. For instance, for complete networks, which essentially force the designer to use public information, the result of Feddersen and Pesendorfer (1997) survives (see Heese and Lauermann (2019)). These remarks suggest that information sharing on social networks need not be only detrimental to elections, in contrast to the common wisdom. Once we realize that private information (i.e., tailored messages) can be

---

[12] If $\hat{N} = \{(3,1),(3,2),(4,1)\}$, $\hat{N}^S = \hat{N} \cup \{(2,1)\}$ and therefore the feasible set is still the smallest in Figure 5.

[13] In Example 4 the players are better off (sometimes strictly) as (N, I) becomes more connected. It is possible to build examples where, even when players face independent decision problems, more-connected information systems lead to less information provision and Pareto-worse outcomes.

a key part of strategies to manipulate voters, we can see that social networks may limit the scope for manipulation.

4.1.2 Seeds and Their Influence

Our framework allows us to shed some light on the following classic question: Fixing N, which seeds between S and S' would a designer choose to achieve some specific objective? Relatedly, which group of players is more *influential*? The problem of "optimal seeding" has received considerable attention in the network literature, in economics as well as computer science and sociology.[14] Building on various measures of network centrality (like Bonacich's), this literature has proposed several notions of nodes' influence that can be operationally computed based on N only. However, these notions are often derived under specific assumptions on the players' interactions and the designer's objective. For example, a typical framework assumes that players do not interact strategically and the designer maximizes diffusion. In the spirit of this literature, we introduce a notion of influence in our framework that depends only on N, but not G or v.

Definition 6 (Influence)**.** Fix N. S is *more influential* than S' if $X(G,N,S) \supseteq X(G,N,S')$ for all G.

Accordingly, when S is more influential than S', the designer is better off seeding players in S, irrespective of the game. Our influence order is demanding, as its condition has to hold across all games. However, its characterization is simple, and hence, operational. Indeed, Lemma 1 implies that there is a tight relation between influence and the notion of more-connected information systems.

Corollary 1. *Fix N. S is more influential than S' if and only if $(N,S') \trianglerighteq (N,S)$.*

The question of which seed set is more influential then boils down to which leads to a less-connected system. Clearly, when $S' \subseteq S$, S is more influential than S'. Yet, S can be more influential than S' even when $S' \not\subseteq S$. Consider Figure 7. On the left panel $S' = \{1,2\}$ and $N^{S'} = I^2$; on the right panel $S = \{3,4\}$ and $N^S \subsetneq I^2$. Therefore, $(N,S') \trianglerighteq (N,S)$. Corollary 1 implies that, irrespective of G and of the designer's objective v, she prefers S to S'.

Our notion of influence can disagree with classic notions of influence that rely on network centrality. This is because we allow for arbitrary strategic interactions among

---

[14] See, for example, Morris (2000), Ballester et al. (2006), Banerjee et al. (2013), Sadler (2017), Akbarpour et al. (2018), and Galeotti et al. (2019). Valente (2012) provides a broad review of the literature outside economics.

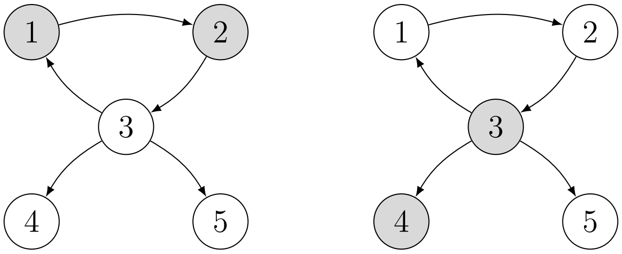

Figure 7: Influence Order: $S = \{3, 4\}$ is More Influential than $S' = \{1, 2\}$.

players in N. For example, in Figure 8, player 2 is strictly more central than player 1 according to Bonacich centrality, but 1 and 2 are equally influential according to our notion. Another difference is that our notion can be applied to groups of players and not just to single players.

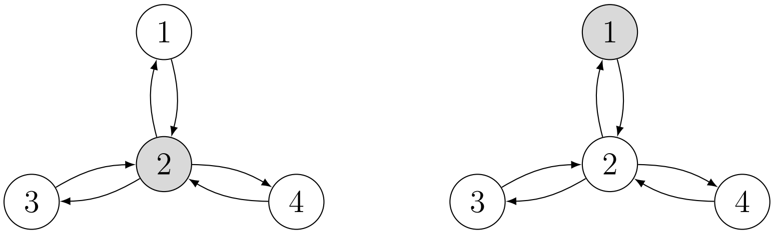

Figure 8: Strategic Influence and Bonacich Centrality

4.2 Application: News Consumption and Elections

News outlets often serve as information gatekeepers for voters and can therefore affect elections.[15] This depends not only on what information they convey, but also who their audiences are, namely, on voters' news consumption *patterns* (Kennedy and Prat (2019)). Arguably, the audience of an outlet depends on what news it conveys. However, to isolate the effects of the channels whereby voters get their news, one would want to remove the effects of the specific news conveyed by each outlet. Our methodology does exactly this. We can view news consumption patterns as an information system, where the outlets are the seeds and the network describes their audiences. We can then span all the news they may convey by varying the initial information structure. This allows us to identify how the breadth and overlaps of the outlets' audiences determine how the news from each outlet—whatever that is—can affect voting behavior. Thus, by exploiting such observables, our approach can offer insights into what determines the power of news outlets that complements the role of what news they convey.

---

[15] For a review of the literature on mass media and elections, see Prat and Strömberg (2013), Gentzkow et al. (2015), and Anderson et al. (2016).

Viewed as an information system, news consumption patterns have specific features. Let S be the set of outlets and let N describe their audiences: Voter i is in the audience of outlet j if there is a link from j to i in N. For instance, building on data from the Pew Research Center (2014; 2020), Figure 9(a) depicts a selection of US news outlets as the seeds and their audiences according to their political ideology (Democrats versus Republicans). Each link reports the share of voters in the respective ideological group that receives information from each outlet. These data document a stark asymmetry in the news consumption patterns between ideological groups.[16] This asymmetry is an observable, persistent feature that is important to understand the role of the channels of news consumption. Indeed, while all voters in the audience of one outlet *must* receive the same news (perfect correlation), this news *can* differ—and possibly be uncorrelated—from the news received by other audiences. By Theorem 1, we can study this complex problem where information flows from outlets to audiences as the simpler problem where we can ignore the outlets and *all* voters directly receive information, but share it according to the S-expansion N^S. In this way, N^S captures the specific and unavoidable correlations in the voters' information caused by this media information system.

We illustrate the effects of such systems on voters' behavior through a stylized model, represented in Figure 9(b). Consider an electorate split into two ideological groups. We view this underlying heterogeneity as resulting from socioeconomic status, not information. Suppose there are M voters in one ideological group and m in the other, where $M > m > 0$; we will refer to the former as the majority voters and to the latter as the minority voters. There are three news outlets. The first is a *broad* outlet whose audience is the entire electorate, which we denote by i_B. The second is a *one-sided* outlet whose audience is only the majority voters, which we therefore call the majority outlet and denote by i_M. The third is a *one-sided* outlet whose audience is only the minority voters, which we call the minority outlet and denote by i_m. Note that this outlet classification is based on audiences, not the kind of information provided—although the two may be related.

The voters face an upcoming referendum on a bill. Its content is known, but its consequences are uncertain and represented by an equally likely $\omega \in \{-1, 1\}$. Each voter i can vote yes, no, or abstain: $A_i = \{y, n, h\}$. Casting a vote costs $c > 0$. The bill passes if, of the *cast* votes, the ayes exceed the nays. The voters vote expressively. That is, they derive utility from supporting the bill if they think it is favorable to them and from opposing it if they think it is not. They turn out to vote if and only if the

---

[16] The asymmetry remains even if people share news with others, which we omit in this section. In fact, the Pew Research Center's (2014; 2020) reports show that ideologically extreme voters mostly interact with like-minded people. Figure 9(a) omits links for audiences smaller than 1 percent.

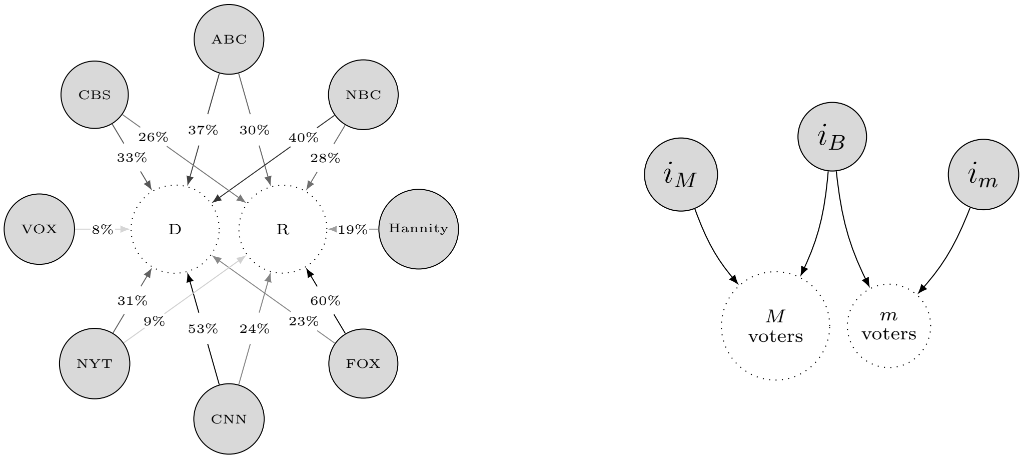

(a) (b)

Figure 9: News Outlets and their Audiences as Information Systems

expected utility from voting for the preferred alternative exceeds that from abstaining.[17] Since robust obedience depends only on payoff differences between actions, it suffices to specify for voter i and state ω

$$\begin{aligned} u_i(y;\theta_i,\omega) - u_i(h;\theta_i,\omega) &= 1-(\omega-\theta_i)^2 - c, \\ u_i(n;\theta_i,\omega) - u_i(h;\theta_i,\omega) &= -\big(1-(\omega-\theta_i)^2\big) - c, \end{aligned}$$

which can be used to calculate $u_i(y;\theta_i,\omega) - u_i(n;\theta_i,\omega)$. The constant 1 is a payoff normalization and θ_i is i's known ideological bliss point. The minority voters have $\theta_i = \theta \in (0,1]$, and the majority voters have $\theta_i = -\theta$. Suppose c is sufficiently low that all voters turn out if they know the state: $c < 1-(1-\theta)^2$. To focus on the asymmetric news consumption of majority and minority voters, we assume that each ideological group is homogeneous. Given this, one can interpret θ as a measure of ideological polarization. Indeed, a higher θ increases voter i's incentive to vote yes when ω is favorable to him and no when ω is unfavorable to him. The outlets do not take actions. These elements define the base game G.

An analyst wants to predict the probability that the bill passes (hereafter, passage

---

[17] We use expressive voting to keep this illustrative application short and simple. Our focus is on news consumption patterns and their interaction with voting costs and polarization. Expressive voting based on ethics, identity, or ideology has been used to explain why people vote despite its cost (Brennan and Lomasky (1997); Kan and Yang (2001); Glaeser et al. (2005); Hamlin and Jennings (2011); Pons and Tricaud (2018); Spenkuch (2018)). Similarly to our setup, in Feddersen and Sandroni (2006), ethical agents vote for their preferred candidate because they feel morally obligated to do their part, even though they understand each vote individually has no effect on the final outcome. Feddersen (2004) surveys various challenges encountered in modeling strategic *and* costly voting. This being said, our framework can handle more general voting games, which we leave for future research.

probability) and how it depends on where the voters get their news. In addition, she is interested in how this dependence is affected by the voting cost and polarization. Not knowing what exact information the outlets provide, she looks for the range of passage probabilities. Its lower bound is zero, which occurs if no information is provided. Its upper bound, denoted by P, captures the extent to which the outlets' news can affect the referendum outcome. To find P, let $v(a,\omega) = \mathbb{I}\{a \in A_{\text{pass}}\}$ where $A_{\text{pass}} \subseteq A$ contains all action profiles such that the ayes exceed the nays. We then have $P = V^*(G, N, S)$ where $S = \{i_M, i_B, i_m\}$. It is easy to show that we can focus on outcomes x such that all majority voters behave the same and all minority voters behave the same. We then refer collectively to the former as "the majority" and to the latter as "the minority" hereafter.

Several voting scenarios can lead the bill to pass. An obvious one is that the majority votes yes. However, costly voting creates another scenario where the minority votes yes while the majority abstains. It is not immediately clear how the probability of these scenarios depends on what information comes from which outlet and on c and θ. Figure 10 represents P as the circled black line, where panel (a) illustrates its dependence on c fixing $\theta = 3/10$ and panel (b) illustrates its dependence on θ fixing $c = 1/5$ (see Online Appendix D.6 for alternative parametrizations). To understand this complex dependence and the role of each outlet, a useful thought experiment is to examine what would happen if we shut down some outlets. The corresponding passage probabilities are represented in Figure 10 for $\hat{S} \in \{\{i_B\}, \{i_B, i_M\}, \{i_B, i_m\}\}$. Comparing these cases reveals some interesting findings.[18]

Finding 1: *News consumption patterns do not matter if the voting cost is zero.*

As shown in Figure 10(a), if the voting cost is zero, the analyst's predictions about the passage probability do not depend on the channels from which the voters get their news. The probability can only depend on the content of the news. The reason is that, if $c = 0$, the majority never abstains, so it will always carry the bill. Thus, whether the majority receives news from the broad outlet or the majority outlet is irrelevant. Also, since the minority can never affect whether the bill passes, where it gets its news is also irrelevant.

Finding 2: *The minority outlet has more informational power than the majority outlet if and only if the voting cost is sufficiently high.*

That is, news coming from one-sided outlets can raise the passage probability above the level it could achieve if voters got news only from the broad outlet. Moreover, for the

---

[18] To be concise, we only report numerical solutions and some commentary, which rely on Theorem 1 and 2. More details about the corresponding outcomes x appear in Appendix D.6. The analytical solutions can be computed using Galperti and Perego's (2018) method, which we already used for the examples in Section 4.1. They are available upon request.

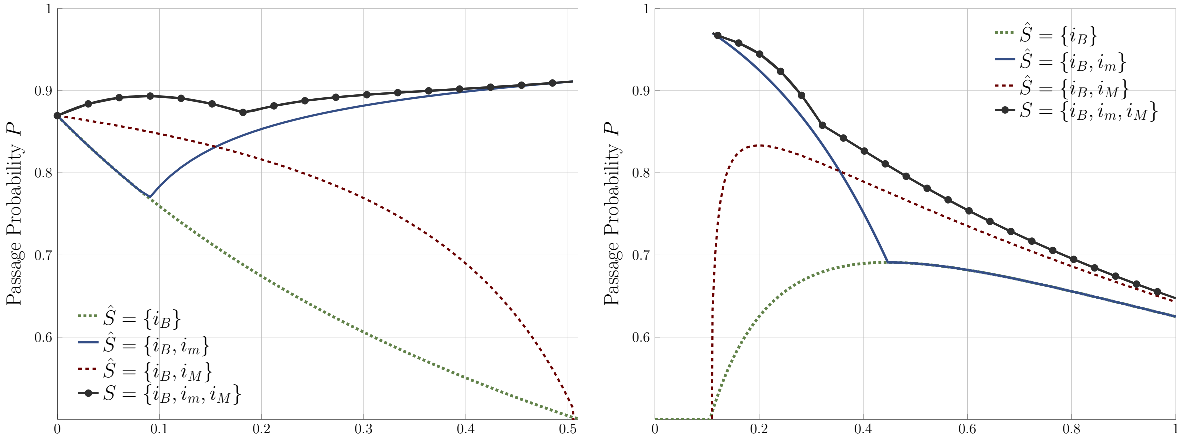

(a) Voting Cost c (b) Polarization θ

Figure 10: Robust Passage Probabilities

minority outlet, this effect is positive and increasing in c if and only if c is large enough; for the majority outlet, it is hump-shaped in c and eventually vanishes. These properties determine the shape of the P line in Figure 10(a). To explain the intuition, we proceed in steps.

First, imagine that we shut down both one-sided outlets: $\hat{S} = \{i_B\}$. In such a world, all voters would get the same news; also, simple steps show that if that news leads one ideological group to vote yes, the other must vote no. The bill can then pass only in the scenario in which the majority supports it, rendering the minority irrelevant. Thus, we essentially have a Bayesian persuasion problem with one receiver: the majority. Concretely, the broad outlet's news that maximizes the passage probability either reveals that the state is $\omega = 1$—leading the majority to vote no—or it pools the states and induces the majority to barely prefer yes to no.

Second, suppose we shut down only the minority outlet: $\hat{S} = \{i_B, i_M\}$. As Figure 10(a) shows, the maximal passage probability is always higher than with the broad outlet alone. This can only happen if sometimes the bill passes in the scenario in which the minority supports it, a vote which hinges on news from the broad outlet. However, how can such a piece of news induce the minority to vote yes without also inducing the majority to vote no? The key is that that news can be only partially informative. Thus, although by itself that news would lead the majority to block the bill, it leaves room for additional news from the majority outlet to induce the majority to sometimes abstain. This news combination expands the possible scenarios in which the bill passes and highlights the role of the majority outlet's ability to convey information privately to its audience. A higher c requires the broad outlet's news to be more informative for the minority to vote yes. In turn, this always curbs the scope for the majority outlet's news to induce its audience to abstain, thereby lowering the passage probability.

Third, suppose we instead shut down the majority outlet: $\hat{S} = \{i_B, i_m\}$. As Figure 10(a) shows, the maximal passage probability is higher than with the broad outlet alone only for sufficiently large c and is then strictly increasing. This again requires that sometimes the bill pass with the minority's support while the majority abstains. For this to happen, suppose that sometimes the news from the broad outlet is uninformative. The majority then votes no if c is small, preventing any news of the minority outlet from affecting the bill's passage. But if c is large, the majority abstains after receiving uninformative news. In this case, there is room for news from the minority outlet to sometimes induce its narrow audience to carry the bill. A higher c renders it more likely that the broad outlet' news fails to mobilize the majority. In turn, this widens the scope for the minority outlet's news to induce its audience to vote yes, thereby raising the passage probability.[19]

Finding 3: *The minority outlet has more informational power than the majority outlet if and only if polarization is sufficiently low.*

That is, higher polarization curbs the ability of news from the minority outlet to raise the passage probability. Moreover, for high polarization, this ability falls below that of news from the majority outlet, which is hump-shaped in θ. Overall, higher polarization lowers P, as shown in Figure 10(b).

To gain intuition, we can again examine what would happen if we shut down some outlets. Recall that a higher θ strengthens how much the voters like and dislike the bill in the two states. Under no news all voters abstain if θ is sufficiently small and vote no if θ is sufficiently large. Thus, higher polarization renders voters easier to mobilize for low θ, but harder to dissuade from voting no for large θ. If $\hat{S} = \{i_m, i_B\}$, both effects render it less likely that the broad outlet's news leads the majority to abstain, which is essential for the minority outlet to have informational power as we saw before. By contrast, if $\hat{S} = \{i_M, i_B\}$, higher polarization implies that the broad outlet's news has to be more informative to induce the minority to vote yes for high θ, but it can be less informative for low θ. Thus, it widens the scope left for the majority outlet to have informational power for low θ, but curbs it for high θ.

The complex dependence of P on the voting cost and polarization illustrates one last point of this application. When formal analyses reach unequivocal conclusions about how voting costs and preference polarization drive elections, such conclusions are likely

---

[19] These remarks suggest that the informational power of one-sided outlets weakens if their ability to convey information privately to their audiences is curbed by spillovers between voters in different ideological groups. We refer the reader to our discussion in Section 4.1.

to depend on specific assumptions about what information the voters get, not only where they get it. Any such conclusion may therefore warrant careful scrutiny.

5 General Spillovers and Random Networks

This section considers information systems in which spillovers are more complex and less mechanical than under Assumption 1. We analyze general spillovers on a fixed network, including strategic communication and random networks. It is well-known that even simple communication games can have rich equilibrium structures. In some cases, this richness makes it hard or impossible to directly characterize the set of feasible outcomes. In other cases, the richness of predictions defeats the purpose of allowing for spillovers in the first place, as we explain shortly. We show how our methodology enables us to compute bounds for the designer's payoff. By specifying her payoff function v instrumentally, we can also use these bounds to obtain robust predictions on the feasible outcomes, such as the maximal or minimal probability the players will play specific actions. Thus, our methodology can bypass the complexity of general spillovers while retaining much of their economic relevance.

5.1 General Spillovers in Direct Information Systems

This section focuses on direct information systems ($S = I$). Later, we will mention some issues with general spillovers in indirect information systems.

Consider the following general model of communication on a network N, which subsumes the one of Section 2. For all $i \in I$, let $\bar{N}_i = \{j \in I : (j, i) \in N\}$ be the set of i's *direct* sources and ${}_i\bar{N} = \{j \in I : (i, j) \in N\}$ be the set of i's *direct* followers. Let K be a finite number of communication rounds. In every round, player i sends a message to player j if and only if $j \in {}_i\bar{N}$. For every initial (T, π_I), let $M_{ij}(T, \pi_I)$ be the finite set of messages i can send to j at every round. We assume that the message spaces per se do not constrain the players' ability to convey all the information they may get initially. One way to ensure this is to assume that $T \subseteq M_{ij}(T, \pi_I)$ for all i, j. Let i's initial histories be of the form $h_i^0 = (\pi_I, t_i)$, where t_i is privately observed by i. Thus, i's set of initial histories is $H_i^0 = \{(\pi_I, t_i) : t_i \in \mathbf{supp}\,\pi_I\}$.[20] For every $k \geq 1$, denote i's histories at round k recursively by $H_i^k = \{(h_i^{k-1}, ({}_im, m_i)) : h_i^{k-1} \in H_i^{k-1}, {}_im \in {}_iM, m_i \in M_i\}$, where ${}_iM = \times_{j \in {}_i\bar{N}} M_{ij}(T, \pi_I)$ and $M_i = \times_{j \in \bar{N}_i} M_{ji}(T, \pi_I)$. Let the set of all i's histories be $\mathcal{H}_i = \cup_{k=0}^{K} H_i^k$. Player i's communication is described by a function $\xi_i : \mathcal{H}_i \to \Delta({}_iM)$.

---

[20] We include π_I in the histories to allow for the possibility that communication behavior depends on π_I and to easily keep track of this dependence.

The profile $\xi = (\xi_i)_{i=1}^I$ defines a *spillover process.* We assume that ξ is common knowledge among the players—for example, it is pinned down by some equilibrium concept.

Every spillover process transforms initial information structures into final information structures. Indeed, for every ω, π_I determines a distribution over finitely many profiles of initial histories $h^0 = (h_i^0)_{i=1}^I$. For every h^0, ξ induces a distribution over finitely many profiles of histories $h^K = (h_i^K)_{i=1}^I$. Interpreting h_i^K as i's final signal realization from these compounded distributions, we get that every ω induces a distribution over such signal profiles. Therefore, ξ maps every initial π_I into a final $\pi'_I = f_{\xi,N}(\pi_I)$.[21]

The next result shows that our baseline approach from Section 3 identifies the worst-case scenario for the designer across arbitrary forms of spillover on networks. This result gives our baseline analysis a connotation of robustness to players' sharing information in more general ways. For every G, N, and ξ, define

$$V_\xi^*(G, N, I) = \sup_{\pi_I} V(G, f_{\xi,N}(\pi_I)),$$

where $V(G, f_{\xi,N}(\pi_I))$ is the expected payoff under the designer-preferred equilibrium in $\text{BNE}(G, f_{\xi,N}(\pi_I))$. Let $V^*(G, N, I)$ be the value of the designer's problem in the baseline model with network N.

Proposition 2 (Payoff Bounds for General Spillovers)**.** *Fix G and (N, I). For every spillover process ξ, $V^*(G, N, I) \leq V_\xi^*(G, N, I) \leq V^*(G, \varnothing, I)$.*

Intuitively, the second inequality arises because, in a direct information system with $N = \varnothing$, the designer could always replicate the spillovers induced by ξ by choosing $\pi'_I = f_{\xi,N}(\pi_I)$. The first inequality arises because spillovers are always weaker under ξ: Given π_I, some information may not leak between some linked players under ξ, while it will leak under our baseline assumption.

The key aspect of Proposition 2 is not the inequalities per se, which are intuitive, but its practical usefulness. For example, it can be extremely complicated, if not impossible, to characterize ξ and compute $V_\xi^*(G, N, I)$ when players communicate strategically on the network N. By contrast, computing both payoff bounds involves a relatively standard linear program.

The lower bound of Proposition 2 is invalid in indirect information systems. To see this, consider $N = \{(1, 3), (2, 3)\}$ and $S = \{1, 2\}$. For the sake of conciseness, assume neither seed will ever volutarily convey any information to player 3 in any ξ. Suppose the designer cares only about 3's action and wants it to match the state, which is also 3's goal.

---

[21] We conjecture that relaxing the assumption that ξ is commonly known does not change our results. Intuitively, it should be possible to build the uncertainty player i has about how others communicate into ξ itself.

It follows that $V_\xi^*(G, N, I) < V^*(G, N, I)$ in this case. More generally, one complication with general spillovers in indirect information systems is that whether Assumption 2 holds may depend on π_I and ξ.

Several canonical spillover processes fit in our general model of network communication (see Appendix D.1 for further details). A vast literature in economics, surveyed by Golub and Sadler (2017), studies how agents learn about some underlying state from their peers' actions in social contexts. We can interpret these actions as the messages the agents send to their neighbors in our model. Unlike most of the social learning literature, Proposition 2 covers settings where no assumption is made about the agents' initial information structure. As another example, suppose that after the designer chooses π_I, every player truthfully reports his belief about the state *and* everybody's signal to his neighbors over multiple rounds. This process provides a foundation for our baseline Assumption 1.

Last but not least, the model presented above covers strategic communication over networks. This may follow the rules of cheap talk (Crawford and Sobel (1982); Aumann and Hart (2003)), verifiable information (Milgrom and Roberts (1986)), or Bayesian persuasion (Kamenica and Gentzkow (2011); Gentzkow and Kamenica (2016)). Although important, strategic communication raises specific issues, which we leave for future research. One is equilibrium multiplicity and selection. Especially striking is the case of cheap talk. Every π_I can be followed by a babbling equilibrium and therefore no spillovers between players, which removes the heart of the matter for this paper. In fact, this implies that the set of feasible outcomes for every (G, N) *coincides* with Bergemann and Morris' (2016) BCE set for G.[22] A more interesting approach is to select the "worst" equilibrium in the communication phase and the ensuing final game. Although promising, this maxmin approach raises its own challenges even without spillovers (Bergemann and Morris (2019); Mathevet et al. (2020)). Communication via verifiable information or Bayesian persuasion may render nontrivial spillovers unavoidable and therefore some BCE outcomes infeasible. For these forms of communication, however, characterizing the equilibrium that lead to a spillover process ξ remains challenging—even focusing on designer-preferred equilibria. Proposition 2 allows us to bypass these issues, at least as a first pass.

5.2 Information Systems with Random Networks

Another extension of our baseline model is to allow for uncertainty about the network. The analyst may not know the network and each player may only know the part that

---

[22] Some readers may wonder whether it is possible to develop a revelation principle type of argument to characterize all feasible outcomes under cheap-talk spillovers. Based on some preliminary analysis, this seems to us a promising approach, although it does not change the point made here.

directly involves him. To model this, fix the players I and seeds S. Let φ be a commonly known probability distribution over networks with nodes I, whose support is $\Phi = \mathbf{supp}\,\varphi$. We assume that every player i learns the realization of his information sources N_i and direct followers ${}_i\bar{N} = \{j \in I : (i,j) \in N\}$. For convenience, we write $\nu_i = (N_i, {}_i\bar{N})$ and interpret ν_i as a signal that player i receives about the realized N, whose distribution follows from φ. We assume that φ is independent of ω to avoid that ν_i conveys exogenous information about ω, which—although possible—only distracts from the focus of the paper.

As in Assumption 1, spillovers occur mechanically: If N has a path from player i to player j, then j learns all the information i initially gets, namely, (t_i, ν_i).[23] Since spillovers are mechanical, albeit random, an equivalent interpretation for this model of information systems with random networks is that player i always attempts to share information with every other player j, but succeeds only if link (i,j) exists in the realized N. We extend Assumption 2 by requiring that, for each network realization, the resulting information system is connected.

Assumption 3. *For all $N \in \Phi$, (N, S) is connected.*

This of course restricts the designer's uncertainty about N: She knows there is always a channel to reach every player with information, but not which channel is available. This restriction is weaker when the number of players and therefore channels is larger.

As in Section 5.1, we focus on characterizing bounds on the value of the designer's problem. Because she does not know the realization of N, she cannot condition her choice of π_S on it. We can then view φ and the mechanical spillovers as inducing a specific spillover process ξ^φ, as defined in Section 5.1, over the fixed network $N^\varphi = \cup_{N\in\Phi} N$ (Appendix D.1 explains the construction of ξ^φ in detail). Given such ξ^φ, the value of the designer's problem is

$$V^*(G,\varphi,S) \equiv V^*_{\xi^\varphi}(G,N^\varphi,S) = \sup_{\pi_S} V(G, f_{\xi^\varphi,N^\varphi}(\pi_S)).$$

To construct our bounds, we need to define two sets of outcomes. An outcome is now described by $x^\varphi : \Omega \times \Phi \to \Delta(R)$, where the probability assigned to α can depend on N. Let $\bar{X}(G,\varphi,I)$ be the set of outcomes that satisfy the following property: For every $N \in \Phi$, $i \in I$, $\alpha_{N_i} \in \boldsymbol{x}^\varphi_{N_i}$, and $a_i \in A_i$

$$\sum_{\omega\in\Omega,\alpha_{-N_i}\in\boldsymbol{x}^\varphi_{-N_i},N\in\Phi} \Big(u_i(\alpha_i,\alpha_{-i},\omega) - u_i(a_i,\alpha_{-i},\omega)\Big) x^\varphi(\alpha_i,\alpha_{-i}|N,\omega)\mu(\omega)\hat{\varphi}(N|\nu_{N_i}) \geq 0, \quad (3)$$

---

[23] Note that the more each player knows about the network, the better he can predict other players' behavior, which tightens incentive constraints. Thus, letting the spillovers include the information about N provides a conservative analysis.

where $\hat{\varphi}(\cdot|\nu_{N_i})$ is i's belief about the realized N conditional on $\nu_{N_i} = (\nu_j)_{j\in N_i}$. Also, let

$$\underline{X}(G, \varphi, S) = \cap_{N\in\Phi} X(G, N, S),$$

where $X(G, N, S)$ is the set of feasible outcomes of the baseline problem (G, N, S). Since each $X(G, N, S)$ contains any common BNE-outcome of G under the prior μ and Φ is finite, $\underline{X}(G, \varphi, S)$ is nonempty. The set $\bar{X}(G, \varphi, I)$ contains outcomes for which it is optimal for every player to follows his recommendation, knowing the recommendations for his sources in the realized N and his exogenous information about N. The set $\underline{X}(G, \varphi, S)$ contains outcomes that are feasible for each baseline problem defined by each realized N.

Proposition 3 (Payoff Bounds for Random Networks)**.** *Fix G and network distribution φ. We have $\underline{V}(G, \varphi, S) \le V^*(G, \varphi, S) \le \bar{V}(G, \varphi, I)$, where*

$$\underline{V}(G, \varphi, S) = \begin{cases} V^*(G, N^\varphi, I) & \text{if } S = I, \\ \sup\limits_{x^\varphi \in \underline{X}(G,\varphi,S)} \sum\limits_{\alpha\in \boldsymbol{x}^\varphi, \omega\in\Omega, N\in\Phi} v(\alpha, \omega) x^\varphi(\alpha|N, \omega)\mu(\omega)\varphi(N) & \text{if } S \subsetneq I, \end{cases}$$

and

$$\bar{V}(G, \varphi, I) = \sup_{x^\varphi \in \bar{X}(G,\varphi,I)} \sum_{\alpha\in \boldsymbol{x}^\varphi, \omega\in\Omega, N\in\Phi} v(\alpha, \omega) x^\varphi(\alpha|N, \omega)\mu(\omega)\varphi(N).$$

To see the intuition, consider first $S = I$. The lower bound follows immediately from Proposition 2. The upper bound follows because condition (3) corresponds to robust obedience in the fictitious setting where the designer issues her recommendations after learning N. By contrast, in the original setting, the designer's π_I and therefore the signals cannot depend on N. Thus, her ability to correlate the players' behavior with N is limited to how the network realizations *exogenously* shape the spillovers and consequently the players' final information. Now consider $S \subsetneq I$. Since seeding constrains the initial information structures, $\bar{V}(G, \varphi, I)$ remains a valid upper bound. Regarding the lower bound, outcomes that are feasible for each realized N must be implemented by some "common" π_S that cannot rely on differences between Ns. Such π_S then also works when N is uncertain.[24]

To conclude, it is worth emphasizing that Proposition 3 exploits the method and insights of our baseline analysis to provide useful bounds for an otherwise challenging problem. All bounds are again expressed in terms of recommendation mechanisms (as opposed to information structures) and involve solving some linear programs.

6 Concluding Remarks

In this paper, we studied information systems and their effects on strategic interactions. We defined such a system as a set of seed players, who can get original information

---

[24] It is easy to construct examples such that $V^*(G, \varphi, S) < \bar{V}^*(G, \varphi, I)$ when $S \subsetneq I$.

initially, and a network, which determines how such information is then distributed across players and thus limits information privacy. This flexible framework can be used to describe the bare bones of the players' informational environment in various contexts, from the flow of information in organizations, to voters' news consumption about political issues, to the consequences of privacy protection or breaches, to the use of influencers and viral messages in marketing.

The main pillars of our methodology involve an equivalence between indirect and direct systems and a characterization of feasible outcomes via spillover-robust obedience. We used this to introduce a notion of more-connected information systems, which identifies when spillovers are more constraining and seeds are more influential. Although we develop our methodology for the case of mechanical and deterministic spillovers between players, it also proves to be useful for cases of random and strategic spillovers, which are otherwise intractable. Our results can serve both theorists interested in studying how information flows between players shape their behavior, and empiricists interested in exploiting observables about the information system to sharpen their estimates.

References

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

Acquisti, A., C. Taylor, and L. Wagman (2016): "The economics of privacy," *Journal of Economic Literature*, 54, 442–92.

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

Anderson, S. P., J. Waldfogel, and D. Stromberg (2016): *Handbook of Media Economics, vol 1A*, Elsevier.

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

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006): "Who's Who in Networks. Wanted: The Key Player," *Econometrica*, 74, 1403–1417.

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," *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.

——— (2019): "Information Design: A Unified Perspective," *Journal of Economic Lit-*

erature, 57(1), pp. 44-95).

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.

Blackwell, D. (1951): "Comparison of experiments," *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability.*

Brennan, G. and L. Lomasky (1997): *Democracy and decision: The pure theory of electoral preference*, Cambridge University Press.

Brooks, B., A. P. Frankel, and E. Kamenica (2019): "Information Hierarchies," *Manuscript, University of Chicago.*

Candogan, O. (2020): "On Information Design with Spillovers," *Working Paper.*

Chan, J., S. Gupta, F. Li, and Y. Wang (2019): "Pivotal persuasion," *Journal of Economic Theory*, 180, 178–202.

Chassang, S., N. J. Kawai, Kei, and J. Ortner (2019): "Data Driven Antitrust: Theory and Application to Missing Bids," *working paper.*

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

DeMarzo, P. M., D. Vayanos, and J. Zwiebel (2003): "Persuasion bias, social influence, and unidimensional opinions," *The Quarterly journal of economics*, 118, 909–968.

Doval, L. and V. Skreta (2019): "Mechanism Design with Limited Commitment," May, 2019.

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

Feddersen, T. and W. Pesendorfer (1997): "Voting behavior and information aggregation in elections with private information," *Econometrica*, 1029–1058.

Feddersen, T. and A. Sandroni (2006): "A theory of participation in elections," *American Economic Review*, 96, 1271–1282.

Feddersen, T. J. (2004): "Rational choice theory and the paradox of not voting," *Journal of Economic perspectives*, 18, 99–112.

Galeotti, A., B. Golub, and S. Goyal (2019): "Targeting Interventions in Networks," *Mimeo.*

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

Garicano, L. and T. Van Zandt (2013): "Hierarchies and the Division of Labor," in *Handbook of Organizational Economics*, ed. by R. Gibbons and J. Roberts, Princeton University Press.

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

Gentzkow, M., J. M. Shapiro, and D. F. Stone (2015): “Media bias in the marketplace: Theory,” in *Handbook of media economics*, Elsevier, vol. 1, 623–645.

Glaeser, E. L., G. A. Ponzetto, and J. M. Shapiro (2005): “Strategic extremism: Why Republicans and Democrats divide on religious values,” *Quarterly journal of economics*, 120, 1283–1330.

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

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

Hamlin, A. and C. Jennings (2011): “Expressive political behaviour: Foundations, scope and implications,” *British Journal of Political Science*, 41, 645–670.

Heese, C. and S. Lauermann (2019): “Persuasion and Information Aggregation in Elections,” Tech. rep., Bonn University.

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

Kamenica, E. (2019): “Bayesian persuasion and information design,” *Annual Review of Economics*, 11, 249–272.

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

Kan, K. and C.-C. Yang (2001): “On expressive voting: Evidence from the 1988 US presidential election,” *Public Choice*, 108, 295–312.

Kennedy, P. J. and A. Prat (2019): “Where do people get their news?” *Economic Policy*, 34, 5–47.

Lipnowski, E. and E. Sadler (2019): “Peer-Confirming Equilibrium,” *Econometrica*, 87, 567–591.

Magnolfi, L. and C. Roncoroni (2017): “Estimation of discrete games with weak assumptions on information,” *Manuscript, University of Wisconsin-Madison.[1663]*.

Mathevet, L., J. Perego, and I. Taneva (2020): “On Information Design in Games,” *Journal of Political Economy*, 128(4), 1370–1404.

Mathevet, L. and I. Taneva (2019): “Organized Information Transmission,” *Working Paper*.

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

Morris, S. (2000): “Contagion,” *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*, 323–358.

Pew Research Center (2014): "Political Polarization and Media Habits," October, 2014.

——— (2020): "U.S. Media Polarization and the 2020 Election: A Nation Divided," January, 2020.

Pons, V. and C. Tricaud (2018): "Expressive voting and its cost: Evidence from runoffs with two or three candidates," *Econometrica*, 86, 1621–1649.

Prat, A. (2018): "Media power," *Journal of Political Economy*, 126, 1747–1783.

Prat, A. and D. Strömberg (2013): "The political economy of mass media," *Advances in economics and econometrics*, 2, 135.

Renault, J., L. Renou, and T. Tomala (2014): "Secure message transmission on directed networks," *Games and Economic Behavior*, 85, 1–18.

Renou, L. and T. Tomala (2012): "Mechanism design and communication networks," *Theoretical Economics*, 7, 489–533.

Rivera, T. J. (2018): "Incentives and the Structure of Communication," *Journal of Economic Theory*, 175, 201–247.

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.

Spenkuch, J. L. (2018): "Expressive vs. strategic voters: An empirical assessment," *Journal of Public Economics*, 165, 73–81.

Taneva, I. (2019): "Information Design," *American Economic Journal: Microeconomics*, 11(4), 151–85.

Valente, T. (2012): "Network Interventions," *Science*, 337, 49–52.

Appendix

A Three Technical Remarks

This appendix answers the three questions raised at the end of Section 3. Why are mixed recommendations necessary in Theorem 2? Is there no loss of generality in focusing on information structures and so outcomes with finite support? Does our constrained information design problem admit a solution?

Mixed Recommendations in Theorem 2

To understand the need for mixed recommendations, it is instructive to consider a simple example. Let G be a matching pennies game with complete information (i.e., $|\Omega| = 1$). In the unique Nash equilibrium, each player mixes uniformly over the two actions. First, suppose the network is empty. In this case, we can describe the players' behavior as follows: A designer flips a fair coin for each player and *privately* recommends to each player the *pure* action corresponding to the realization of the coin toss. Such recommendations are obedient, as they replicate the unique Nash equilibrium. Conversely, suppose the network is complete. The previous recommendation is not spillover-robust obedient. In fact, with positive probability player 1 is recommended to play Tails after learning that his opponent will play Heads. Because of information spillovers, the designer's recommendation are no longer private. Yet, the fact that initial information may spill over should not compromise the existence of an equilibrium in this game or, equivalently, should not prevent players to privately randomize just before interacting in the game. This apparent impasse has a simple solution: Allow the designer to recommend each player a *mixed* action. In other words, the designer hands over one coin to each player for him to use when choosing his action. Such recommendations are robustly obedient. Similar issues can arise in more general settings, whenever mixing by some player is essential to sustain an equilibrium outcome.

Finite-support Outcomes x

Our model assumes that information structures are finite and so the induced outcomes have finite support (Definition 1). However, it is natural to ask whether this assumption is without loss of generality, since we have to allow for mixed-action recommendations. That is, given a game G, do the primitives of G imply a finite upper bound on the support of outcomes such that allowing for a larger support does not allow the designer to achieve higher payoffs. This is not an issue in standard information design (i.e., if

$N = \varnothing$), because it is without loss of generality to focus on outcomes x that involve only pure-action recommendations (Bergemann and Morris (2016)). Given this, the finite base game sets exogenously the maximal size of x's support: $x \in \mathbb{R}^{\Omega \times A}$ where $A = \times_{i \in I} A_i$. By contrast, in our case (i.e., if $N \neq \varnothing$) there is a priori no exogenous upper bound on the support size, as $x \in \mathbb{R}^{\Omega \times R}$ where $R = \times_{i \in I} \Delta(A_i)$. In short, we can fill each of the finitely many slots of x's support in infinitely many ways using mixed recommendations.

In the following, we use Theorem 2 to show that, without loss of generality, we can identify a finite, *exogenous*, upper bound on the outcomes' support. To show this, it helps to rewrite the robust-obedience condition (2) as follows: For every $i \in I$, $\alpha_{N_i} \in \boldsymbol{x}_{N_i}$, and $a_i, a_i' \in A_i$,

$$\sum_{\omega \in \Omega, \alpha_{-N_i} \in \boldsymbol{x}_{-N_i}} \Big(u_i(a_i, \alpha_{-i}, \omega) - u_i(a_i', \alpha_{-i}, \omega) \Big) \alpha_i(a_i) x(\alpha_i, \alpha_{-i} | \omega) \mu(\omega) \geq 0. \tag{4}$$

This highlights that for each player i robust obedience ultimately involves his primitive pure actions. To willingly implement any mixed action, i must deem all pure actions in its support optimal given his information. This information is provided by the realization of α_{N_i}. Each α_{N_i} then pins down a subset of optimal actions for i. Since A_i is finite, there can only be finitely many such subsets.

Lemma 2. *Suppose $x \in X(G, N, I)$. There exists $x' \in X(G, N, I)$ such that $|\boldsymbol{x}_i'| \leq |2^{A_i}|$ for every $i \in I$ and x and x' induce the same joint distribution over A for every $\omega \in \Omega$:*

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

Proof. See Online Appendix D.5. □

This also implies that it is without loss of generality to focus on information structures such that $|T_i| \leq |2^{A_i}|$ for all $i \in I$.

Existence of a Solution

Building on Lemma 2, we can express formally the problem of choosing an optimal information structure π_S and show that it always has a solution. For every $i \in S$, fix a signal space $\bar{T}_i$ that satisfy $|\bar{T}_i| = |2^{A_i}|$. For $i \notin S$, let $\bar{T}_i$ be singleton. Given this, let $\Pi_S = \{\pi_S : \Omega \to \Delta(\bar{T})\}$, where $\bar{T} = \times_{i \in I} \bar{T}_i$. The designer's problem can be written as

$$V^*(G, N, I) = \sup_{\pi_S \in \Pi_S} V(G, N, \pi_S).$$

By Theorem 2, we have that

$$V^*(G, N, I) = \sup_{x \in X(G,N,I)} \sum_{\omega \in \Omega, \alpha \in \boldsymbol{x}} v(\alpha, \omega) x(\alpha | \omega) \mu(\omega),$$

where $v(\alpha, \omega)$ is the designer's expected payoff in ω from the profile α.

We now turn to the existence of solutions. In standard information design, $X(G, \varnothing, I)$ is compact subset of $\mathbb{R}^{\Omega \times A}$, so existence follows from standard arguments. By contrast, $X(G, N, I)$ remains quite complex for $N \neq \varnothing$. Despite the exogenous bound on the outcomes' support, its elements are profiles of mixed actions, which form an uncountable space. The next result establishes the existence of feasible outcomes.

Lemma 3 (Existence)**.** *For every G and N, there exists $x \in X(G, N, I)$ such that*

$$\sum_{\omega \in \Omega, \alpha \in \boldsymbol{x}} v(\alpha, \omega) x(\alpha|\omega) \mu(\omega) = V^*(G, N, I).$$

Proof. See Online Appendix D.5. □

In some settings with information spillovers (i.e., $N \neq \varnothing$), the solution to the designer's problem may involve only pure-actions recommendations. One simple example of this is when the base game is actually a collection of single-agent decision problems: $u_i(a_i, a_{-i}; \omega)$ does not depend on a_{-i} for all $i \in I$. Intuitively, in this case we cannot relax condition (4) by keeping player i uncertain about other players' behavior whose randomness is independent of the state. Thus, mixed-action recommendations are useless. More generally, we can always search for a candidate solution within the space of outcomes that only recommend pure actions. The companion paper Galperti and Perego (2018) shows how to verify that this candidate solves the overall problem using linear-programming duality.

B Main Proofs

To prove Theorem 1, we first introduce and prove Lemmas 4, 6 and 7, and the intermediate equivalence result of Lemma 5. Lemma 4 characterizes N^S. It shows that in N^S, while each player may have new sources relative to N (formally, $N_i \subseteq N_i^S$), none of these new sources is a seed, i.e. $N_i^S \cap S = N_i \cap S$.

Lemma 4. *Fix (N, S). For all i, $N_i^S \cap S = N_i \cap S$.*

Proof of Lemma 4. Fix N and i. First, we show that $N_i \cap S \subseteq N_i^S \cap S$. To see this, note that $N \subseteq N^S$, by definition of S-expansion. This implies that $N_i \subseteq N_i^S$. Hence, $N_i \cap S \subseteq N_i^S \cap S$. Second, we show that $N_i^S \cap S \subseteq N_i \cap S$. Note that it is enough to show that $N_i^S \cap S \subseteq N_i$. Suppose not, $N_i^S \cap S \nsubseteq N_i$, there is $j \in N_i^S \cap S$ such that $j \notin N_i$. Since $j \in N_i^S$, there exists a path in N^S from j to i. That is, a sequence $P = (k_1, \ldots k_m)$ of distinct k_l for $1 \leq l \leq m$, such that $k_1 = j$, $k_m = i$, and $(k_l, k_{l+1}) \in N^S$, for all $l \leq m - 1$. Since $j \notin N_i$, it must be that $(k_l, k_{l+1}) \notin N$, for at least one $l \leq m - 1$. We refer to these l's as the *gaps* of P. Let $\underline{P} = (\underline{k}_1, \dots \underline{k_m})$ be a path from j to i in N^S with the property that its number of gaps is smaller or equal than the number of gaps in any other path P from j to i in N^S. Note that $\underline{P}$ is well-defined since I is finite. Denote $\underline{l}$ the gap in $\underline{P}$ with the smallest index. By construction, we have that (1) $j \in S$, (2) $j \in N_{\underline{k}_{\underline{l}}}$, (3) $(\underline{k}_{\underline{l}}, \underline{k}_{\underline{l}+1}) \in N^S$, (4) $j \notin N_{\underline{k}_{\underline{l}+1}}$, and (5) $(\underline{k}_{\underline{l}}, \underline{k}_{\underline{l}+1}) \notin N$. Points (1) and (2) imply that $j \in N_{k_{\underline{l}}} \cap S$. By Definition 3, Point (3) implies that $N_{\underline{k}_{\underline{l}}} \cap S \subseteq N_{\underline{k}_{\underline{l}+1}}$. However, $j \notin N_{k_{\underline{l}+1}}$, by point (4). Thus, $N_{\underline{k}_{\underline{l}}} \cap S \not\subseteq N_{\underline{k}_{\underline{l}+1}}$. Finally, by (5), $(\underline{k}_{\underline{l}}, \underline{k}_{\underline{l}+1}) \notin N$. We conclude that N^S is not the expansion of N, a contradiction. $\square$

Lemma 4 shows that the set of i's seeded sources is the same in N and N^S. Under any π_S, non-seeded players receive singleton signals. This suggests that, from an informational point of view, the information structures $f_N(\pi_S)$ and $f_{N^S}(\pi_S)$ are equivalent. The next result formalizes this idea.

Lemma 5. *Fix (N, S). Then for all G, $X(G, N, S) = X(G, N^S, S)$.*

Proof of Lemma 5. Fix G, i and the information structure (T, π_S). Note that $(N_i^S \setminus N_i) \cap S = (N_i^S \cap S) \setminus (N_i \cap S) = \varnothing$. The first equality derives from the distributive property of set intersection over set difference. The second equality derives from Lemma 4. This implies that $T_{N_i^S \setminus N_i}$ is a singleton. Fix $t := (t_1, \dots, t_I)$, such that $\sum_\omega \mu(\omega)\pi_S(t|\omega) > 0$. We want to show that $\Pr_{\pi_S}(t|t_{N_i}) = \Pr_{\pi_S}(t|t_{N_i^S})$. Namely, conditioning on $t_{N_i^S}$ rather than t_{N_i} does not change the probability assessment over t. Thus, vectors t_{N_i} and $t_{N_i^S}$ are identical up to $t_{N_i^S \setminus N_i}$, which realizes with probability 1 under π_S, since $T_{N_i^S \setminus N_i}$ is a singleton. Hence $\Pr_{\pi_S}(t|t_{N_i}) = \Pr_{\pi_S}(t|t_{N_i^S})$. Since $i \in I$ and t were arbitrary, we have that $BNE(G, f_N(\pi_S)) = BNE(G, f_{N^S}(\pi_S))$. Since π_S was arbitrary, we conclude that $X(G, N, S) = X(G, N^S, S)$. $\square$

The next result shows that N^S is the S-expansion of itself, thus proving the uniqueness of the expansion of a network.

Lemma 6. *$(i, j) \in N^S$ if and only if $N_i^S \cap S \subseteq N_j^S$.*

Proof of Lemma 6 *Only if.* Let $(i, j) \in N^S$. Then, $N_i^S \subseteq N_j^S$, hence $N_i^S \cap S \subseteq N_j^S$. *If.* Suppose $N_i^S \cap S \subseteq N_j^S$. Then, $N_i^S \cap S \subseteq N_j^S \cap S$. By Lemma 4, $N_i \cap S \subseteq N_j \cap S$. Thus, $N_i \cap S \subseteq N_j$. By Definition 3, this implies $(i, j) \in N^S$. $\square$

The next result will play an important role in the proof of our main theorem. To prove this result, we use a technique known as *secret sharing*.

Lemma 7. *Fix* (G, N, S) *and* $S \subseteq S' \subseteq I$. *Let* $i \in S'$ *and* $(i, j) \in N^S$. *Then* $X(G, N^S, S') = X(G, N^S, S' \cup \{j\})$.

Proof of Lemma 7.
($\subseteq$). This direction is trivial since, by definition, every π_S can be written as a π_I.
($\supseteq$). If $j \in S'$ there is nothing to show since, in such case, $S' \cup \{j\} = S'$. Therefore, let $j \notin S'$. Fix any $(T, \pi_{S' \cup \{j\}})$, where $T = T_1 \times \ldots \times T_I$. Using a *secret sharing*'s technique (Shamir (1979)), we will construct a $(\hat{\pi}_{S'}, \hat{T})$ such that $BNE(G, f_{N^S}(\pi_{S' \cup \{j\}})) = BNE(G, f_{N^S}(\hat{\pi}_{S'}))$. Let $B(\kappa) := \{0, 1\}^\kappa$ and $\underline{\kappa} := \min\{\kappa \in \mathbb{N} : |T_j| \leq |B(\kappa)|\}$. For notational convenience, denote $B := B(\underline{\kappa})$. Let $\mathcal{P} : T_j \to B$ be an arbitrary injective function. It represents a "public key," that univocally transforms j's signals into binary numbers. Denote $\oplus$ the following logical operator (exclusive or). Bits are added according to the following rule: $0 \oplus 0 = 1 \oplus 1 = 0$ and $0 \oplus 1 = 1 \oplus 0 = 1$. For all $b, b' \in B$, $b \oplus b' = (b_1 \oplus b'_1, \ldots, b_{\underline{\kappa}} \oplus b'_{\underline{\kappa}}) \in B$. For notational convenience, denote $Q := N^S_j \cap S'$, the seeded sources of player j. Recall that, by assumption, $j \notin Q$. We now construct the type space in $(\hat{T}, \hat{\pi}_{S'})$. Let $\hat{T}_i := T_i$ for all $i \notin Q$ and $\hat{T}_i := T_i \times B \times \{\mathcal{P}\}$ for all $i \in Q$. Note that, by construction, $\hat{T}$ is such that $\hat{T}_i = \{\hat{t}_i\}$ for all $i \notin S'$. That is $(\hat{T}, \pi_{S'})$ seeds only players in S'. Next, we construct $\hat{\pi}_{S'}$ from $\pi_{S' \cup \{j\}}$. Let $q := \max\{i : i \in Q\}$. For each realization $t \in T$ under $\pi_{S' \cup \{j\}}$, the realization $\hat{t} \in \hat{T}$ under $\hat{\pi}_{S'}$ is such that $\hat{t}_i = t_i$ for $i \notin Q$, with the same conditional distribution of $\pi_{S' \cup \{j\}}$. Instead, if $i \in Q$, $\hat{t}_i = (t_i, b_i, \mathcal{P})$. More specifically, if $i \in Q \setminus \{q\}$, $b_i \in B$ is drawn at uniform random from B, independently of (ω, t); instead, if $i = q$, $b_q := \mathcal{P}(t_j) \oplus \big(\oplus_{i \in Q \setminus \{q\}} b_i\big)$. There are two cases to consider, $|Q| = 1$ and $|Q| > 1$.

- If $|Q| = 1$, $b_q = \mathcal{P}(t_j)$ and observing $\hat{t}_q$ reveals t_j. Thus, player i learns t_j if and only if $\{q\} = Q \subseteq N^S_i$.
- If $|Q| > 1$, instead, observing all but one element in $(b_i)_{i \in Q}$ carries no information about t_j. Instead, observing the whole sequence $(b_i)_{i \in Q}$ fully reveals t_j. This is because:

$$\begin{aligned} \mathcal{P}^{-1}\Big(\oplus_{i \in Q} b_i\Big) &= \mathcal{P}^{-1}\Big(\big(\oplus_{i \in Q \setminus \{q\}} b_i\big) \oplus b_q\Big) \\ &= \mathcal{P}^{-1}\Big(\big(\oplus_{i \in Q \setminus \{q\}} b_i\big) \oplus \big(\mathcal{P}(t_j) \oplus \big(\oplus_{i \in Q \setminus \{q\}} b_i\big)\big)\Big) \\ &= \mathcal{P}^{-1}\Big(\mathcal{P}(t_j)\Big) \\ &= t_j. \end{aligned}$$

  The third equality comes from the fact that $b_i \oplus b_i = \mathbf{0}$ and $\mathcal{P}(t_j) \oplus \mathbf{0} = \mathcal{P}(t_j)$. Thus, player i learns t_j if and only if $Q \subseteq N^S_i$.

Therefore, irrespective of whether or not Q is a singleton, player i learns t_j if and only if $Q \subseteq N^S_i$. However, note that $Q \subseteq N^S_i$ if and only if $j \in N^S_i$. In fact, if $Q \subseteq N^S_i$, $N^S_j \cap S \subseteq$

$N_j^S \cap S' = Q \subseteq N_i^S$ and, by Lemma 6, $(j,i) \in N^S$ and $j \in N_i^S$. Conversely, if $j \in N_i^S$, then $N_j^S \subseteq N_i^S$, and therefore $Q \subseteq N_i^S$. We conclude that under the constructed $(\hat{T}, \hat{\pi}_{S'})$ player i learns t_j if and only if $j \in N_i^S$, just like under the given $(T, \pi_{S' \cup \{j\}})$. Therefore, any outcome x induced by $(T, \pi_{S' \cup \{j\}})$ can be also induced by the constructed $(\hat{T}, \hat{\pi}_{S'})$. Since $(T, \pi_{S' \cup \{j\}})$ was arbitrary, this shows that $X(G, N^S, S' \cup \{j\}) \subseteq X(G, N^S, S')$. ☐

Proof of Theorem 1. Fix (G, N, S). By Lemma 5, $X(G,N,S) = X(G,N^S,S)$. We are left to show that $X(G,N^S,S) = X(G,N^S,I)$. If $S = I$ there is nothing to prove, so let $S \subsetneq I$. The following induction argument proves the claim.

Basis Step. Let $S_1 = S$. Since (N, S) is connected (by Assumption 2), there are $i, j \in I$ such that $i \in S_1$, $j \notin S_1$, and $(i,j) \in N$. Since $N \subseteq N^S$, $(i,j) \in N^S$. Let $S_2 := S_1 \cup \{j\}$. Since $S_1 \subseteq S_2 \subseteq I$, $i \in S_2$ and $(i,j) \in N^S$, we can invoke Lemma 7 to show that $X(G,N^S,S_1) = X(G,N^S,S_2)$. Finally, it is straightforward to see that (N, S_2) is connected.

Inductive Step. Suppose that $X(G,N^S,S_1) = X(G,N^S,S_k)$ for $S_k := S_1 \cup \{j_1, \ldots, j_k\}$. If $S_k = I$ there is nothing to prove. Hence, let $S_k \subsetneq I$. (N, S_k) is connected. Hence, there are $i, j \in I$ such that $i \in S_k$, $j \notin S_k$, and $(i,j) \in N$. Since $N \subseteq N^S$, $(i,j) \in N^S$. Denote $S_{k+1} := S_k \cup \{j\}$. Since $S_{k+1} \supseteq S_k$, $i \in S_{k+1}$, and $(i,j) \in N^S$, we can invoke Lemma 7 to show that $X(G,N^S,S_{k+1}) = X(G,N^S,S_k) = X(G,N^S,S_1)$.

Since I is finite, this procedure stops after $\bar{k} = |I \setminus S|$ steps. We conclude that $X(G,N^S,S) = X(G,N^S,I)$. ☐

In what follows, a (behavioral) strategy of player i in $(G,(\pi,T))$ is $\sigma_i : T_i \to \Delta(A_i)$. We write $\sigma_i(a_i|t_i)$ instead of $\sigma(t_i)[a_i]$. A profile $\sigma = (\sigma_i)_{i \in I}$ belongs to $BNE(G,\pi)$ if for each i, $t_i \in T_i$, and $a_i \in A_i$ with $\sigma_i(a_i|t_i) > 0$,

$$\sum_{a_{-i}, t_{-i}, \omega} \Big(u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega) \Big) \sigma(a_i, a_{-i}|t_i, t_{-i}) \pi(t_i, t_{-i}|\omega) \mu(\omega) \geq 0$$

for all $a_i' \in A_i$, where $\sigma(a_i, a_{-i}|t_i, t_{-i}) := \prod_{j=1}^{I} \sigma_j(a_j|t_j)$.

Proof of Theorem 2. **Part 1 (⇒)**: Suppose (T, π_I) and $\sigma \in \text{BNE}(G, f(\pi_I))$ induce x. Then, for every i and $t_{N_i} \in T_{N_i}$,

$$\sum_{\omega, t'} \Big(u_i(\sigma_i(t_{N_i}), \sigma_{-i}(t'_{N_{-i}}), \omega) - u_i(a_i, \sigma_{-i}(t'_{N_{-i}}), \omega) \Big) \Pr_{\pi_I}(\omega, t'|t_{N_i}) \geq 0, \quad a_i \in A_i.$$

where $\sigma_{-i}(t'_{N_{-i}}) = (\sigma_j(t'_{N_j}))_{j \neq i}$. Using π_I, we can write this condition as, for every i and t_{N_i},

$$\sum_{\omega, (t'_{N_j})_{j \neq i}} \Big(u_i(\sigma_i(t_{N_i}), \sigma_{-i}(t'_{N_{-i}}), \omega) - u_i(a_i, \sigma_{-i}(t'_{N_{-i}}), \omega) \Big) \frac{\pi_I(t_{N_i}, (t'_{N_j})_{j \neq i}|\omega)\mu(\omega)}{\sum_{\omega', (t''_{N_j})_{j \neq i}} \pi_I(t_{N_i}, (t''_{N_j})_{j \neq i}|\omega')\mu(\omega')} \geq 0,$$

for all $a_i \in A_i$, or equivalently,

$$\sum_{\omega,(t'_{N_j})_{j\neq i}} \Big(u_i(\sigma_i(t_{N_i}),\sigma_{-i}(t'_{N_{-i}}),\omega) - u_i(a_i,\sigma_{-i}(t'_{N_{-i}}),\omega)\Big)\pi_I(t_{N_i},(t'_{N_j})_{j\neq i}|\omega)\mu(\omega) \geq 0,$$

for all $a_i \in A_i$. Note that, for every i and t, by knowing t_{N_i} player i knows the mixed action $\sigma_j(t_{N_j})$ for all $j \in N_i$.

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_{N_i},

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

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

Part 2 (⇐): Suppose x is robustly obedient. Recall that $\boldsymbol{x} = \{\alpha \in R : x(\alpha|\omega) > 0, \omega \in \Omega\}$ is finite. Note that x can be thought of as an information structure, by viewing each $\alpha_i \in \boldsymbol{x}_i$. Given this, for every i, consider the strategy $\sigma_i : \boldsymbol{x}_{N_i} \to \Delta(A_i)$ defined by $\sigma_i(\alpha_{N_i}) = \alpha_i$, for all $\alpha_{N_i} \in \boldsymbol{x}_{N_i}$. Optimality for each i requires that, for every α_{N_i},

$$\sum_{\omega,\alpha'_{-N_i}} \Big(u_i(\sigma_i(\alpha_{N_i}),\sigma_{-i}(\alpha'_{N_{-i}}),\omega) - u_i(a_i,\sigma_{-i}(\alpha'_{N_{-i}})),\omega)\Big)\mathrm{Pr}_x(\omega,\alpha'|\alpha_{N_i}) \geq 0, \quad a_i \in A_i,$$

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

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

which holds because x is robustly obedient. □

Proof of Lemma 1. We begin with two preliminary observations. First, it is easy to see that by Theorem 1 we only need to show that $X(G,N,I) \subseteq X(G,N',I)$ for all G if and only if $(N,I) \trianglerighteq (N',I)$. Second, spillover-robust obedience is equivalent to requiring that, for every i and $\delta_i : R_{N_i} \to A_i$,

$$\sum_{\omega\in\Omega,\alpha\in\boldsymbol{x}} \Big(u_i(\alpha_i,\alpha_{-i};\omega) - u_i(\delta_i(\alpha_{N_i}),\alpha_{-i};\omega)\Big)x(\alpha_i,\alpha_{-i}|\omega)\mu(\omega) \geq 0. \tag{5}$$

Part 1 (⇐): Suppose $(N,I) \trianglerighteq (N',I)$ and $x \in X(G,N,I)$ for some G. Then, by Theorem 2, x satisfies (5) for all i and $\delta_i \in D_i = \{\hat{\delta}_i : R_{N_i} \to A_i\}$. Let $D'_i = \{\delta_i : R_{N'_i} \to A_i\}$. To prove that $x \in X(G,N',I)$, 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 $\hat{\delta}_i : R_{N_i} \to A_i$ as $\hat{\delta}_i(\alpha_{N'_i},\alpha_{N_i\setminus N'_i}) = \delta_i(\alpha_{N'_i})$, for all $\alpha_{N_i} \in R_{N_i}$. Since $N_i \supseteq N'_i$ for all i, $\hat{\delta}_i$ is a well-defined function and $\hat{\delta}_i \in D_i$.

Part 2 (⇒): We prove this with a contrapositive argument. The only relevant case to consider is that $(N,I) \not\trianglerighteq (N',I)$ and $(N,I) \not\trianglelefteq (N',I)$. This implies that for some i, there exists a k such that $k \in N'_i$ and $k \notin N_i$, and for some j (possibly $i = j$), there exists m such that $m \in N_j$ and $m \notin N'_j$. It follows that there exists a player i_k such that $i_k \neq k$ and there is a direct link from k to i_k in N' but not in N, and there exists a player i_m such that $i_m \neq m$ there is a direct link from m to i_m in N but not in N'. 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 $j \notin \{k, m, i_k, i_m\}$, let the payoff function u_j be such that 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 $a = \frac{1}{2}$. Thus, $x \in X(G, N, I)$, but clearly does not belong to $X(G, N', I)$. This is because in N' player i_k has to choose $a_{i_k} = \frac{1}{2}$ after learning $a_k = \omega$, which renders $a_{i_k} = \frac{1}{2}$ strictly suboptimal. Thus, x violates robust obedience for (G, N', I). Now consider x' such that player m always matches the state, while all the other players choose $a = \frac{1}{2}$. This x' belongs to $X(G, N', I)$, but clearly does not belong to $X(G, N, I)$. This is because in N 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 arguments work 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 $X(G, N, I)$, but clearly does not belong to $X(G, N', I)$. This is because in N' player i_k has to choose $a_{i_k} = \frac{1}{2}$ after learning $a_k = \omega$, which renders $a_{i_k} = \frac{1}{2}$ strictly suboptimal. Thus, x violates obedience for (G, N', I). Alternatively, 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' belongs to $X(G, N', I)$, but clearly does not belong to $X(G, N, I)$. This is because in N player i_m has to choose $a_{i_m} = \frac{1}{2}$ after learning that $a_m = \omega$, which renders $a_{i_m} = \frac{1}{2}$ strictly suboptimal. Thus, x' violates obedience for (G, N, I). □

Proof of Proposition 1. **Part 1 (⇒)**: Suppose $\hat{N}_i^{\hat{S}} \subseteq N_i^S$ for all i. Consider $i, j \in I$ that satisfy $\hat{N}_i \cap \hat{S} \subseteq \hat{N}_j$ and so $i \in \hat{N}_j^{\hat{S}}$ by Definition 3. It follows that $i \in N_j^S$. If $i \in N_j$, then $N_i \cap S \subseteq N_j$ holds automatically. If $i \notin N_j$, we must have added links to N according to Definition 3 that result in $i \in N_j^S$. 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 $N_{j_k} \cap S \subseteq N_{j_{k+1}}$. This implies $N_i \cap S \subseteq N_j$.

Part 2 (⇐): Now suppose that $\hat{N}_i \cap \hat{S} \subseteq \hat{N}_j$ implies $N_i \cap S \subseteq N_j$ for all $i, j \in I$. We need to show that $i \in \hat{N}_j^{\hat{S}}$ implies $i \in N_j^S$. Fix any i and j that satisfy $i \in \hat{N}_j^{\hat{S}}$. If $i \in \hat{N}_j$, then we automatically have $\hat{N}_i \cap \hat{S} \subseteq \hat{N}_j$ and so $N_i \cap S \subseteq N_j$ by assumption. Thus, if $i \notin N_j$ (which is the only relevant case), we must add (i, j) to N according to Definition 3, implying $i \in N_j^S$. Next, suppose that $i \in \hat{N}_j^{\hat{S}} \setminus \hat{N}_j$. By Definition 3, there must exist a sequence $\{j_k\}_{k=0}^m$ which satisfies $j_0 = i$, $j_m = j$, and $\hat{N}_{j_k} \cap \hat{S} \subseteq \hat{N}_{j_{k+1}}$. Therefore, it must be that $\hat{N}_i \cap \hat{S} \subseteq \hat{N}_j$ and, by assumption, $N_i \cap S \subseteq N_j$. If again $i \notin N_j$, it follows that $(i, j) \in N^S$, which implies $i \in N_j^S$. □

Proof of Proposition 2. We begin by proving the upper bound, i.e., $V_\xi^*(G, N, I) \leq V^*(G, \varnothing, I)$. Fix ξ. For every π_I, $f_{\xi,N}(\pi_I)$ is itself an information structure, possibly seeding all players in I. It follows that,

$$V_\xi^*(G, N, I) = \sup_{\pi_I} V(G, f_{\xi,N}(\pi_I)) \leq \sup_{\pi_I'} V(G, \pi_I') = V^*(G, \varnothing, I),$$

where the inequality follows from the fact that the final $\pi_I' = f_{\xi,N}(\pi_I)$ is just one among all possible information structures leading to a Bayesian game among the players.

Next, we prove the lower bound, i.e., $V^*(G, N, I) \leq V_\xi^*(G, N, I)$. Fix π_I and consider the information structure $(T, f_N(\pi_I))$. It has the property that, for every player i and signal $t \in T$, Bayesian posteriors satisfy,

$$\Pr_{f_N(\pi_I)}(\omega, t|t_{N_i}) = \Pr_{f_N(\pi_I)}(\omega, t|t_i);$$

that is, there is nothing that player i can learn from his sources that is not already contained in his private signal t_i. Therefore, for every spillover process ξ, we have

$$\Pr_{f_{\xi,N}(f_N(\pi_I))}(\omega, t|h_i) = \Pr_{f_{\xi,N}(f_N(\pi_I))}(\omega, t|h_i^0) = \Pr_{f_N(\pi_I)}(\omega, t|t_i).$$

for all players i, histories h_i and h_i^0. That is, additional communication according to ξ adds nothing to what players already know. Therefore, for all π_I,

$$\sigma \in \text{BNE}(G, f_N(\pi_I)) \qquad \Longrightarrow \qquad \sigma \in \text{BNE}(G, f_{\xi,N}(f_N(\pi_I))).$$

Therefore,

$$\bigcup_{\pi_I} \text{BNE}(G, f_N(\pi_I)) \subseteq \bigcup_{\pi_I} \text{BNE}(G, f_{\xi,N}(f_N(\pi_I)) \subseteq \bigcup_{\pi_I} \text{BNE}(G, f_{\xi,N}(\pi_I))$$

from which it follows that $V^*(G, N, I) \leq V_\xi^*(G, N, I)$. □

Proof of Proposition 3.

Case 1: Direct Provision ($S = I$). Theorem 2 immediately implies the inequality $V^*(G, N^\varphi, I) \leq V^*(G, \varphi, I)$. The other inequality $V^*(G, N^\varphi, I) \leq \bar{V}(G, \varphi, I)$ follows from showing that every π_I and $\sigma \in \text{BNE}(G, f_{\xi^\varphi,N^\varphi}(\pi_I))$ leads to an outcome x^φ that must satisfy condition (3) and therefore belongs to $\bar{X}(G, \varphi, I)$. First, given such π_I and σ, let

$$x^\varphi(\alpha_1, \ldots, \alpha_I|\omega, N) = \sum_{t\in T} \pi_I(t|\omega) \prod_{i\in I} \mathbb{I}\{\sigma_i(\nu_i, \hat{t}_{N_i}) = \alpha_i\}, \quad (\alpha, N, \omega) \in R \times \Phi \times \Omega,$$

where $\hat{t}_{N_i} = (t_j, \nu_j)_{j \in N_i}$. Let $\pi^{\varphi} = f_{\xi^{\varphi}, N^{\varphi}}(\pi_I)$. Since σ is a BNE, for all $N \in \Phi$, $i \in I$, $\hat{t}_{N_i}$, and $a_i \in A_i$,

$$\sum_{\omega \in \Omega, N \in \Phi, t \in T} u_i(\sigma_i(\nu_i, \hat{t}_{N_i}), (\sigma_j(\nu_j, \hat{t}_{N_j}))_{j \neq i}, \omega) \mathrm{Pr}_{\pi^{\varphi}, \mu}(\omega, t | t_{N_i}) \hat{\varphi}(N | \nu_{N_i})$$

$$\geq \sum_{\omega \in \Omega, N \in \Phi, t \in T} u_i(a_i (\sigma_j(\nu_j, \hat{t}_{N_j}))_{j \neq i}, \omega) \mathrm{Pr}_{\pi^{\varphi}, \mu}(\omega, t | t_{N_i}) \hat{\varphi}(N | \nu_{N_i}),$$

which is equivalent to

$$\sum_{\omega \in \Omega, N \in \Phi, t \in T, \alpha_{-i}} u_i(\alpha, \omega) \prod_{j \in I} \mathbb{I}\{\sigma_j(\nu_j, \hat{t}_{N_j}) = \alpha_j\} \mathrm{Pr}_{\pi^{\varphi}, \mu}(\omega, t | t_{N_i}) \hat{\varphi}(N | \nu_{N_i})$$

$$\geq \sum_{\omega \in \Omega, N \in \Phi, t \in T, \alpha_{-i}} u_i(a_i, \alpha_{-i}, \omega) \prod_{j \in I} \mathbb{I}\{\sigma_j(\nu_j, \hat{t}_{N_j}) = \alpha_j\} \mathrm{Pr}_{\pi^{\varphi}, \mu}(\omega, t | t_{N_i}) \hat{\varphi}(N | \nu_{N_i}),$$

which is equivalent (once we multiply both sides by the total probability of t_{N_i}) to

$$\sum_{\omega \in \Omega, N \in \Phi, t \in T, \alpha_{-i}} u_i(\alpha, \omega) \prod_{j \in I} \mathbb{I}\{\sigma_j(\nu_j, \hat{t}_{N_j}) = \alpha_j\} \mathrm{Pr}_{\pi^{\varphi}, \mu}(\omega, t_{N_i} (t_{N_j})_{j \neq i}) \hat{\varphi}(N | \nu_{N_i})$$

$$\geq \sum_{\omega \in \Omega, N \in \Phi, t \in T, \alpha_{-i}} u_i(a_i, \alpha_{-i}, \omega) \prod_{j \in I} \mathbb{I}\{\sigma_j(\nu_j, \hat{t}_{N_j}) = \alpha_j\} \mathrm{Pr}_{\pi^{\varphi}, \mu}(\omega, t_{N_i} (t_{N_j})_{j \neq i}) \hat{\varphi}(N | \nu_{N_i}),$$

which is equivalent (once we explicitly write $\mathrm{Pr}_{\pi^{\varphi}, \mu}(\cdot)$) to

$$\sum_{\omega \in \Omega, N \in \Phi, \alpha_{-i}} u_i(\alpha, \omega) \left\{ \sum_{t \in T} \prod_{j \in I} \mathbb{I}\{\sigma_j(\nu_j, \hat{t}_{N_j}) = \alpha_j\} \pi_I(t_{N_i}, t_{-N_i} | \omega) \mu(\omega) \right\} \hat{\varphi}(N | \nu_{N_i})$$

$$\geq \sum_{\omega \in \Omega, N \in \Phi, \alpha_{-i}} u_i(a_i, \alpha_{-i}, \omega) \left\{ \sum_{t \in T} \prod_{j \in I} \mathbb{I}\{\sigma_j(\nu_j, \hat{t}_{N_j}) = \alpha_j\} \pi_I(t_{N_i}, t_{-N_i} | \omega) \mu(\omega) \right\} \hat{\varphi}(N | \nu_{N_i}),$$

By knowing N_i and $\hat{t}_{N_i}$, player i also knows ν_j and $\hat{t}_{N_j}$ for all $j \in N_i$ and so also $\alpha_j = \sigma_j(\nu_j, \hat{t}_{N_j})$. That is, conditioning on $\hat{t}_{N_i}$ implies that α_{N_i} is fixed. Using the definition of x^{φ}, the last inequality requires that for all $N \in \Phi$, $i \in I$, $\alpha_{N_i} \in \mathbf{x}^{\varphi}_{N_i}$, and $a_i \in A_i$,

$$\sum_{\omega \in \Omega, \alpha_{-N_i} \in \mathbf{x}^{\varphi}_{-N_i}, N \in \Phi} \Big(u_i(\alpha_i, \alpha_{-i}, \omega) - u_i(a_i, \alpha_{-i}, \omega) \Big) x^{\varphi}(\alpha_i, \alpha_{-i} | N, \omega) \mu(\omega) \hat{\varphi}(N | \nu_{N_i}) \geq 0.$$

Case 2: Indirect Provision ($S \subsetneq I$). The inequality $V^*(G, \varphi, S) \leq \bar{V}(G, \varphi, I)$ follows from observing that every outcome x^{φ} the designer can induce with some π_S followed by $\sigma \in \mathrm{BNE}(G, f_{\xi^{\varphi}, N^{\varphi}}(\pi_S))$ is also feasible when she can choose a π_I, since $S \subsetneq I$. Therefore, $x^{\varphi} \in \bar{X}(G, \varphi, I)$ and, hence, the designer does weakly better when maximizing over $\bar{X}(G, \varphi, I)$ relative to the original problem.

Now consider the inequality $\underline{V}(G, \varphi, I) \leq V^*(G, \varphi, S)$. Let $\underline{x}^{\varphi} \in \underline{X}(G, \varphi, S)$ achieve $\underline{V}(G, \varphi, I)$.[25] It suffices to argue that there exists π_S and $\sigma \in \mathrm{BNE}(G, f_{\xi^{\varphi}, N^{\varphi}}(\pi_S))$ such that (π_S, σ) induces $\underline{x}^{\varphi}$, that is,

$$\underline{x}^{\varphi}(\alpha_1, \ldots, \alpha_I | \omega, N) = \sum_{t \in T} \pi_S(t, \omega) \prod_{i \in I} \mathbb{I}\{\sigma_i(\nu_i, \hat{t}_{N_i}) = \alpha_i\}, \quad (\alpha, N, \omega) \in R \times \Phi \times \Omega,$$

---

[25] One can show the existence of a solution $\underline{x}^{\varphi}$ using an argument similar to the proof of Lemma 3.

where $\hat{t}_{N_i} = (t_j, \nu_j)_{j \in N_i}$. Since $\underline{x}^\varphi \in X(G, N, S)$ for all $N \in \Phi$, $\underline{x}^\varphi$ can depend at most on the network aspects that are common to all $N \in \Phi$, namely, the set of connections in $\cap_{N \in \Phi} N$. Since $\cap_{N \in \Phi} N$ is non-random, $\underline{x}^\varphi$ statistically depends only on ω, i.e., $\underline{x}^\varphi : \Omega \to \Delta(R)$.

By Theorem 1 and Assumption 3, this implies that $\underline{x}^\varphi \in X(G, N^S, I)$ for all $N \in \Phi$.

By Theorem 2, there exists a pair $(\underline{\pi}_I, \sigma^{N^S})$ for every $N \in \Phi$ such that $\sigma^{N^S} \in \text{BNE}(G, f_{N^S}(\underline{\pi}_I))$ and the pair induces $\underline{x}^\varphi$. Namely, $\underline{\pi}_I = \underline{x}^\varphi$ and σ^{N^S} is obedient given $(G, f_{N^S}(\underline{\pi}_I))$ for all $N \in \Phi$. But then there also exists a common $\pi' \in \Pi_S$ that conveys as much information to every player as does $\underline{\pi}_I$ for every $N \in \Phi$ (by Theorem 1 and Assumption 3) and a profile $\sigma^N \in \text{BNE}(G, f_N(\pi'))$ such that, for every $N \in \Phi$,

$$\underline{x}^\varphi(\alpha_1, \ldots, \alpha_I | \omega) = \sum_{t \in T} \pi'(t|\omega) \prod_{i \in I} \mathbb{I}\{\sigma_i^N(t_{N_i}) = \alpha_i\}, \quad (\alpha, \omega) \in R \times \Omega.$$

Given this, we can let σ' be defined, for every $i \in I$ and $N \in \Phi$, by

$$\sigma_i'(\nu_i, \hat{t}_{N_i}) = \sigma_i^N(t_{N_i}), \quad t \in T.$$

By construction, σ' is measurable with respect to the information each player i can receive from π' for each realization of N. Also, for each N and t, the realization of (ν_i, t_{N_i}) provides player i weakly less information than (N, t_{N_i}). Therefore, if σ_i^N is a best reply for every N in $(G, f_N(\pi'))$, so is σ' in $(G, f_{\xi^\varphi, N^\varphi}(\pi'))$. That is, $\sigma \in \text{BNE}(G, f_{\xi^\varphi, N^\varphi}(\pi'))$.

□

C Analysis of Examples 4 and 5

We characterize the feasible sets in Figures 5 and calculate the solutions to the designer's problem for the networks $N \in \{\varnothing, \{(1,2)\}, \{(2,1)\}, \{(1,2),(2,1)\}\}$.

C.1 Empty network

It is without loss of generality to consider mechanisms that recommend pure actions. Since $\mu(g) = \mu(b) = \frac{1}{2}$, the obedience constraints can be written as follows: Given recommendation y_i,

$$(\gamma_i + \varepsilon)x(y_i, y_{-i}|g) + \gamma_i x(y_i, n_{-i}|g) + (\varepsilon - 1)x(y_i, y_{-i}|b) - x(y_i, n_{-i}|b) \geq 0,$$

and given recommendation n_i,

$$-(\gamma_i + \varepsilon)x(n_i, y_{-i}|g) - \gamma_i x(n_i, n_{-i}|g) - (\varepsilon - 1)x(n_i, y_{-i}|b) + x(n_i, n_{-i}|b) \geq 0.$$

In Figure 5, we represent the total probability that bank i chooses y_i, namely, $\chi_i = \sum_{\omega,a_{-i}} x(y_i, a_{-i}|\omega)\mu(\omega)$. It is notationally convenient to denote $x(y_i|\omega) = x(y_1, y_2|\omega) + x(y_i, n_{-i}|\omega)$. Non-negativity of probabilities requires that, for all ω,

$$\max\{0, x(y_1|\omega) + x(y_2|\omega) - 1\} \leq x(y_1, y_2|\omega) \leq \min\{x(y_1|\omega), x(y_2|\omega)\}. \tag{6}$$

Using this, we can rewrite the obedience conditions as

$$x(y_i|g) + x(y_i|b) \leq (\gamma_i + 1)x(y_i|g) + \varepsilon[x(y_1, y_2|g) + x(y_1, y_2|b)],$$

$$x(y_i|g) + x(y_i|b) + \varepsilon(x(y_{-i}|g) + x(y_{-i}|b)) \leq (\gamma_i + 1)x(y_i|g) + \varepsilon[x(y_1, y_2|g) + x(y_1, y_2|b)] + 1 - \gamma_i.$$

Since we assume that ε is small, $|\varepsilon| \approx 0$, the second constraint holds whenever the first holds. To obtain the intercepts on the horizontal axis, set $\chi_2 = 0$, which implies that $x(y_2|\omega) = 0$ and so $x(y_1, y_2|\omega) = 0$ for all ω. In this case, the remaining constraint for bank 1 is

$$x(y_1|g) + x(y_1|b) \leq (\gamma_1 + 1)x(y_1|g).$$

Thus, maximizing χ_1 subject to this constraint involves setting $x(y_1|g) = 1$ and $x(y_1|b) = \gamma_1$, which delivers $\chi_1 = \frac{\gamma_1+1}{2}$. A similar argument delivers the intercept on the vertical axis equal to $(0, \chi_2)$ where $\chi_2 = \frac{\gamma_2+1}{2}$.

Another extreme point of the feasible set can be found by maximizing the designer's objective, which corresponds to

$$\chi_1 + \chi_2 = \frac{1}{2}[x(y_1|g) + x(y_1|b) + x(y_2|g) + x(y_2|b)].$$

Since $|\varepsilon| \approx 0$, the solution involves $x(y_1|g) = x(y_2|g) = 1$, which implies $x(y_1, y_2|g) = 1$. Thus, the constraints simplify to

$$x(y_1|b) \leq \gamma_1 + \varepsilon + \varepsilon x(y_1, y_2|b) \quad \text{and} \quad x(y_2|b) \leq \gamma_2 + \varepsilon + \varepsilon x(y_1, y_2|b).$$

We have to consider two cases, depending on the sign of ε. If $\varepsilon > 0$, then $x(y_1, y_2|b) = \min\{x(y_1|b), x(y_2|b)\} = x(y_2|b)$, where the last equality follows from $\gamma_2 < \gamma_1$ and the last two constraints. Therefore, we conclude that

$$x(y_2|b) = \frac{\gamma_2 + \varepsilon}{1 - \varepsilon} \quad \text{and} \quad x(y_1|b) = \gamma_1 + \varepsilon + \frac{\varepsilon}{1 - \varepsilon}(\gamma_2 + \varepsilon).$$

In the (χ_1, χ_2) space, this leads to the extreme point

$$(\chi_1^*, \chi_2^*) = \left(\frac{1 + \gamma_1 + \varepsilon}{2} + \frac{\varepsilon(\gamma_2 + \varepsilon)}{2(1 - \varepsilon)}, \frac{\gamma_2 + 1}{2(1 - \varepsilon)}\right).$$

If instead $\varepsilon < 0$, it is optimal to set $x(y_1, y_2|b) = 0$, which implies that $x(y_i|b) = \gamma_i + \varepsilon$ for all i. In the (χ_1, χ_2) space, this leads to the extreme point

$$(\chi_1^*, \chi_2^*) = \left(\frac{1 + \gamma_1 + \varepsilon}{2}, \frac{1 + \gamma_2 + \varepsilon}{2}\right).$$

Now we need to argue that the feasible set is the convex combination of the four points consisting of the origin, the intercepts on the axes, and (χ_1^*, χ_2^*). This is the case if these four are the only extreme points of the feasible set. Suppose there exists another extreme point with $\chi_2 \in (0, \chi_2^*)$. Then, $\chi_2 = \beta\chi_2^* + (1-\beta)0 = \beta\chi_2^*$ for some $\beta \in (0,1)$. The corresponding largest possible χ_1 must satisfy

$$\chi_1 \leq \frac{1}{2}\Big((\gamma_1 + 1) + \varepsilon\beta[1 + x_{\chi_2}(y_1, y_2|b)]\Big),$$

where $x_{\chi_2}(y_1, y_2|b)$ is the value of $x(y_1, y_2|b)$ associated to χ_2. Since the right-hand side describes all convex combinations across $\beta \in (0,1)$ between (χ_1^*, χ_2^*) and the intercept on the horizontal axis—where obedience binds—it is not possible to have an extreme point to the right of this line. A similar argument shows that there is no other extreme point with $\chi_1 \in (0, \chi_1^*)$.

C.2 Network $N = \{(1,2)\}$

We start by restricting attention to mechanisms that recommend only pure actions. We relegate to Online Appendix D.4 the proof that shows that this is without loss of generality. The obedience constraints for bank 1 involve only

$$x(y_1|g) + x(y_1|b) \leq (\gamma_1 + 1)x(y_1|g) + \varepsilon[x(y_1, y_2|g) + x(y_1, y_2|b)],$$

just as before. For bank 2, instead, we have:

$$\begin{aligned}
(\gamma_2 + \varepsilon)x(y_2, y_1|g) + (\varepsilon - 1)x(y_2, y_1|b) &\geq 0,\\
\gamma_2 x(y_2, n_1|g) - x(y_2, n_1|b) &\geq 0,\\
-(\gamma_2 + \varepsilon)x(n_2, y_1|g) - (\varepsilon - 1)x(n_2, y_1|b) &\geq 0,\\
-\gamma_2 x(n_2, n_1|g) + x(n_2, n_1|b) &\geq 0.
\end{aligned}$$

Using the redefined variables $x(y_i|\omega)$ as before, these constraints become:

$$\begin{aligned}
(\gamma_2 + \varepsilon)x(y_2, y_1|g) + (\varepsilon - 1)x(y_2, y_1|b) &\geq 0,\\
\gamma_2[x(y_2|g) - x(y_2, y_1|g)] - [x(y_2|b) - x(y_2, y_1|b)] &\geq 0,\\
-(\gamma_2 + \varepsilon)[x(y_1|g) - x(y_1, y_2|g)] - (\varepsilon - 1)[x(y_1|b) - x(y_1, y_2|b)] &\geq 0,\\
-\gamma_2[1 + x(y_1, y_2|g) - x(y_1|g) - x(y_2|g)] + [1 + x(y_1, y_2|b) - x(y_1|b) - x(y_2|b)] &\geq 0.
\end{aligned}$$

First, let's argue that the intercepts on the axes are the same as before (clearly, for these extreme points mixed recommendations are irrelevant). Consider bank 2. Using $x(y_1|\omega) = x(y_1, y_2|\omega) = 0$ for all ω, we obtain the remaining constraints

$$\gamma_2 x(y_2|g) - x(y_2|b) \geq 0 \quad \text{and} \quad -\gamma_2[1 - x(y_2|g)] + [1 - x(y_2|b)] \geq 0.$$

Clearly, $x(y_2|g) = 1$ and $x(y_2|b) = \gamma_2$ satisfy both constraints. Now consider bank 1. Clearly, $x(y_1|g) = 1$, $x(y_1|b) = \gamma_1$, and $x(y_1, y_2|\omega) = 0$ for all ω satisfy the obedience constraints for bank 1 as before. The non-trivial obedience constraints for bank 2 evaluated at $x(y_2|\omega) = 0$ for all ω become

$$(1-\varepsilon)\gamma_1 \geq \gamma_2 + \varepsilon \quad \text{and} \quad 1 - \gamma_1 \geq 0.$$

The latter constraint always holds; the former always holds when $\varepsilon < 0$ and when $\varepsilon > 0$ provided that ε is small enough (which we assume).

To find the extreme point (χ_1^*, χ_2^*), let's ignore the obedience constraint of bank 2 following recommendation n_2 for now. The remaining constraints can be written as

$$\begin{aligned} x(y_2, y_1|b) &\leq \tfrac{\gamma_2+\varepsilon}{1-\varepsilon} x(y_2, y_1|g), \\ x(y_2|g) + x(y_2|b) &\leq (\gamma_2+1)x(y_2|g) - \gamma_2 x(y_2, y_1|g) + x(y_2, y_1|b). \end{aligned}$$

It is again optimal to set $x(y_1|g) = x(y_2|g) = x(y_1, y_2|g) = 1$. Therefore, the second constraint becomes $x(y_2|b) \leq x(y_2, y_1|b)$. Since the opposite inequality must also hold by (6), we conclude that $x(y_2|b) = x(y_2, y_1|b)$. This implies that

$$x(y_1|b) = \gamma_1 + \varepsilon + \varepsilon x(y_2|b) \quad \text{and} \quad x(y_2|b) \leq \frac{\gamma_2 + \varepsilon}{1-\varepsilon}.$$

Since $\varepsilon < 1$, we always want to maximize $x(y_2|b)$, which leads to

$$x(y_2|b) = \frac{\gamma_2+\varepsilon}{1-\varepsilon} \quad \text{and} \quad x(y_1|b) = \gamma_1 + \varepsilon + \varepsilon\frac{\gamma_2+\varepsilon}{1-\varepsilon}.$$

Thus, we have

$$(\chi_1^*, \chi_2^*) = \left(\frac{1+\gamma_1+\varepsilon}{2} + \frac{\varepsilon(\gamma_2+\varepsilon)}{2(1-\varepsilon)}, \frac{\gamma_2+1}{2(1-\varepsilon)}\right).$$

Let's now check the remaining constraints for bank 2:

$$\begin{aligned} 0 &\leq -(\gamma_2+\varepsilon)[x(y_1|g) - x(y_1, y_2|g)] - (\varepsilon - 1)[x(y_1|b) - x(y_1, y_2|b)] \\ &= (1-\varepsilon)[x(y_1|b) - x(y_1, y_2|b)], \end{aligned}$$

which holds by definition since $\varepsilon < 1$;

$$\begin{aligned} 0 &\leq -\gamma_2[1 + x(y_1, y_2|g) - x(y_1|g) - x(y_2|g)] + [1 + x(y_1, y_2|b) - x(y_1|b) - x(y_2|b)] \\ &= 1 - x(y_1|b), \end{aligned}$$

which again holds by definition. Denote the outcome associated to (χ_1^*, χ_2^*) by x^*.

Suppose x^* is also the solution of the designer's problem if we allow for mixed-action recommendations—which we show in Online Appendix D.4. Using the argument for the case of $N = \varnothing$, we can conclude that the four points consisting of the origin, the intercepts on the axes, and (χ_1^*, χ_2^*) are the only extreme points of the whole feasible set, which is therefore defined as all convex combinations of those points.

C.3 Network $N = \{(2, 1)\}$

A similar argument to the previous case applies with the following differences. First, we have to swap the player indexes. When calculating (χ_1^*, χ_2^*), the main change is when we conlcude that $x(y_1|b) \leq x(y_1, y_2|b)$. Since we also have to satisfy $x(y_1, y_2|b) \leq \min\{x(y_1|b), x(y_2|b)\} = x(y_2|b)$, it follows that now $x(y_1|b) = x(y_2|b)$. Therefore, we have

$$x(y_2|b) = \gamma_2 + \varepsilon + \varepsilon x(y_2|b),$$

which implies

$$x(y_1|b) = x(y_2|b) = \frac{\gamma_2 + \varepsilon}{1 - \varepsilon}.$$

Thus, we have

$$(\chi_1^*, \chi_2^*) = \left(\frac{\gamma_2 + 1}{2(1 - \varepsilon)}, \frac{\gamma_2 + 1}{2(1 - \varepsilon)} \right).$$

The rest of the analysis is the same.

C.4 Network $N = \{(1, 2), (2, 1)\}$

Again, a similar argument to the previous two cases applies with the following differences. Now the obedience constraints after recommendation y_i can be written as

$$\begin{aligned} x(y_2, y_1|b) &\leq \tfrac{\gamma_1+\varepsilon}{1-\varepsilon} x(y_2, y_1|g), \\ x(y_1|g) + x(y_1|b) &\leq (\gamma_1 + 1)x(y_1|g) - \gamma_1 x(y_2, y_1|g) + x(y_2, y_1|b), \\ x(y_2, y_1|b) &\leq \tfrac{\gamma_2+\varepsilon}{1-\varepsilon} x(y_2, y_1|g), \\ x(y_2|g) + x(y_2|b) &\leq (\gamma_2 + 1)x(y_2|g) - \gamma_2 x(y_2, y_1|g) + x(y_2, y_1|b). \end{aligned}$$

By the same argument as in the previous case, the second constraint and condition (6) again imply that $x(y_1|b) = x(y_2|b) = x(y_1, y_2|b)$. The first constraint is redundant given the third constraint. Moreover, if the last two constraints hold, then their 50-50 combination holds, which corresponds to the obedience constraint of bank 2 for the network $N = \{(2, 1)\}$. This implies that the cases of network $N = \{(1, 2), (2, 1)\}$ and $N = \{(2, 1)\}$ are the same. The rest of the analysis is as before.

C.5 Investment Game with Strategic Complements

For completeness, we illustrate how the set of feasible outcomes change when $\varepsilon > 0$, that is the investment game features strategic complements. Figure 11 is the analog of Figure 5 that is presented in the main text. Due to strategic complementarities, the maximal probability that j can be induced to invest rises as the probability that i invests rises. The spillovers from i to j shrink the feasible set because j is better informed (as before) and because knowing part of j's information allows i to better predict when j invests.

This is why the feasible set collapses to the main diagonal as we add links. As in the case of strategic substitutes, link $(1, 2)$ constrains the designer less than $(2, 1)$ in terms of feasible outcomes. However, any non-empty network leads to the same optimal outcome. This is represented by the red dot on the main diagonal, which can be interpreted as public information.

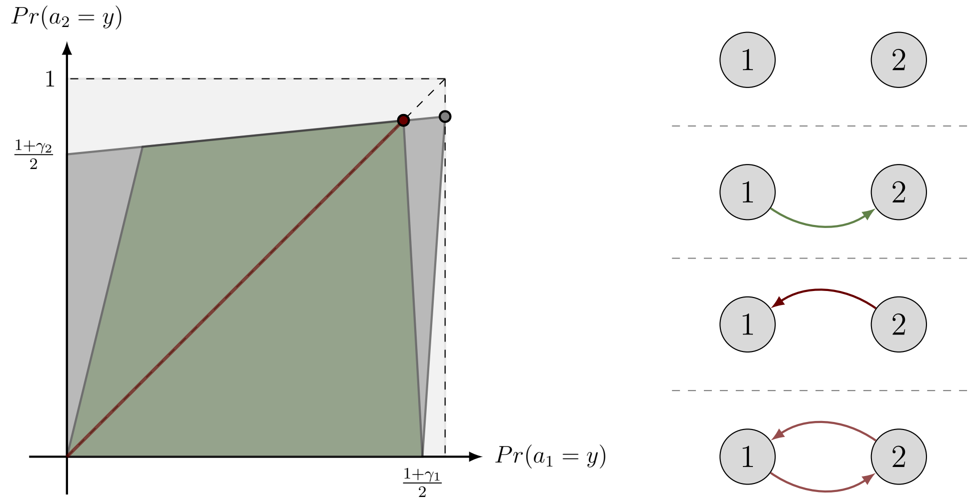

Figure 11: Investment-Game Feasible Outcomes

D Supplemental Appendix (For Online Publication Only)

D.1 Examples of General Spillovers

Truthful Belief Announcement. Suppose at each round of communication after the designer chooses (T, π_I) every player truthfully reports to all his neighbors his current, Bayesian, posterior over $\Omega \times T$, namely, about the state and all players' initial signal realizations. We include the signal realizations in the reports because learning about others' information matters for predicting their actions in the final game. Assume that the number of communication rounds K is at least as large as the shortest path between the two most distant players in N (i.e., its diameter).

This communication process fits into the general model outlined above. Since T is finite, at every round each player can have at most finitely many different posteriors. In the first round, each i reports a degenerate posterior about his private signal t_i to his direct followers in ${}_i\bar{N}$. In the next round, all i's direct followers report a degenerate posterior about t_i to *their* direct followers. Continuing this way, all players for which i is an information source in N will hear a degenerate report of t_i within a number of rounds that cannot exceed the diameter of N. After these rounds, all players' posteriors will be degenerate about their sources' signals and will stop evolving. Thus, this process of truthful belief announcement can be a foundation for our baseline Assumption 1, as confirmed by the next lemma.

Lemma 8. *Fix (T, π_I) and a signal realization t. For every i and $h_i^K \in \mathcal{H}_i$ consistent with (t, π_I), $Pr_{\pi_I}(\omega, \hat{t}|h_i^K) = Pr_{\pi_I}(\omega, \hat{t}|t_{N_i})$ for all $(\omega, \hat{t}) \in \Omega \times T$.*

Proof. Define $N_i^0 = \{i\}$ and $N_i^n = \cup_{j \in N_i^{n-1}} \bar{N}_j$ for $n = 1, \ldots, I$. Note that $N_i^I = N_i$. Fix a signal realization $\bar{t}$ and the corresponding unique $h^K(\bar{t}) = h^K$. For every player i,

$$\xi_i(h_i^0)(\omega, t) = \Pr_{\pi_I}(\omega, t|\bar{t}_i) = \Pr_{\pi_I}(\omega, t|\bar{t}_{N_i^0}), \quad (\omega, t) \in \Omega \times T.$$

Note that

$$\sum_{\omega, t_{-i}} \Pr_{\pi_I}(\omega, t|\bar{t}_i) = \begin{cases} 1 & \text{if } t_i = \bar{t}_i \\ 0 & \text{otherwise.} \end{cases}$$

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

$$\xi_j(h_j^{n-1})(\omega, t) = \Pr_{\pi_I}(\omega, t|\bar{t}_{N_j^{n-1}}), \quad (\omega, t) \in \Omega \times T.$$

Note that

$$\sum_{\omega, t_{-j}} \Pr_{\pi_I}(\omega, t|\bar{t}_{N_j^{n-1}}) = \begin{cases} 1 & \text{if } t_{N_j^{n-1}} = \bar{t}_{N_j^{n-1}} \\ 0 & \text{otherwise.} \end{cases}$$

Therefore,

$$\xi_i(h_i^n)(\omega, t) = \mathrm{Pr}_{\pi_I}(\omega, t|\bar{t}_{N_i^n}), \quad (\omega, t) \in \Omega \times T.$$

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

$$\xi_i(h_i^K)(\omega, t) = \mathrm{Pr}_{\pi_I}(\omega, t|\bar{t}_{N_i^I}) = \mathrm{Pr}_{\pi_I}(\omega, t|\bar{t}_{N_i}), \quad (\omega, t) \in \Omega \times T,$$

thus concluding the proof. □

Observational Social Learning. A vast literature in economics—recently surveyed by Golub and Sadler (2017)—studies how agents learn about some underlying state of the world by observing their peers' actions in social contexts. For example, they may learn about the importance of some policy issue from their friends' decision to join protests about it. This social-learning process can fit into our general communication model. The agents belong to a network N of social relations. Each agent can initially receive some signal about the state (e.g., a newspaper article about the policy issue) and then chooses an action (e.g., to protest or not), which is observed by his neighbors in N. We can interpret this action as the message the agent sends to his neighbors in our model—provided that each agent's action-message space is sufficiently rich. Social learning may occur over multiple rounds. Some agents may act simultaneously, others sequentially. Also, they can have any degree of strategic sophistication in this phase. In the example, they may or may not take into account how their actions convey information that influences their neighbors' vote in an upcoming referendum on the policy issue (the base game). Either way, under appropriate richness assumptions, observational social learning is covered by Proposition 2. To see its potential usefulness, recall that the social-learning literature often assumes that the agents' initial information structure takes a very specific and simple form (e.g., independently and identically distributed signals across agents). By contrast, we impose no restrictions on the information structures. The cost of this richness in information and social-learning processes is that it may be hard to characterize ξ and $V_\xi^*(G, N, I)$. Proposition 2 helps us bypass these intricacies and provide robust predictions.

Random Networks as a Form of General Spillovers. Section 5.2 claimed that we can think of random networks distributed according to φ as a spillover process ξ^φ over the fixed network $N^\varphi = \cup_{N \in \Phi} N$. Such ξ^φ is constructed as follows. Let the number of communication rounds K be at least as large as the diameter of N^φ. For each $i \in I$, the strategy ξ_i^φ depends on the realization of $\nu_i = (N_i, {}_i\bar{N})$ according to φ as a randomization device. Given n_i, at each round of the communication phase, player i truthfully reports his current Bayesian posterior about $\Omega \times T$ to every $j \in {}_i\bar{N}$—as in the truthful-belief-announcement process described above—and a constant, history-independent, message to every $j \in {}_i\bar{N}^\varphi \setminus {}_i\bar{N}$. Such a ξ^φ induces a process that converges in a finite number of rounds, bounded above by the diameter of N^φ.

D.2 Disconnected Information Systems

Our analysis extends to settings that violate Assumption 2, i.e., to information systems that are not connected. Formally, let $\hat{I} = \{i \in I : N_i \cap S \neq \varnothing\}$ and suppose $I \setminus \hat{I} \neq \varnothing$. For $i \notin \hat{I}$, $N_i \cap S = \varnothing$ and so $|T_i| = 1$ for every initial *and* final information structure. Under the assumption that the chosen π_S is common knowledge to all players, the question is how to characterize the feasible outcomes.

The equivalence result in Theorem 1 goes through almost unchanged. The definition of S-expansion is adapted by requiring that the link (i, j) be added to N if and only if $N_i \cap S \neq \varnothing$ and $N_i \cap S \subseteq N_j$. The same argument behind Theorem 1 then implies that $X(G, N, S) = X(G, N^S, \hat{I})$ for all G.

With regard to the characterization of Theorem 2, it is clear that robust obedience is necessary for an outcome to be feasible for (G, N, S). However, sufficiency has to take into account that with isolated players feasible outcomes have additional statistical properties: If $i \notin \hat{I}$, then α_i cannot depend on the state nor on the action of any other player. Indeed, we can write condition (1) as

$$x(\alpha_1, \ldots, \alpha_I|\omega) = \left[\prod_{i \notin \hat{I}} \mathbb{I}\{\sigma_i(t_i) = \alpha_i\}\right] \sum_{t \in T} \pi_S(t|\omega) \prod_{i \in \hat{I}} \mathbb{I}\{\sigma_i(t_{N_i}) = \alpha_i\}.$$

There are two ways to proceed. One is to combine robust obedience as in Definition 4 with some constraint on x that captures these statistical properties; the other is to modify the notion of feasible outcomes and obedience, which may offer additional insights.

For the first way to be useful, the additional constraint on outcomes should be *linear*, so that together with obedience we continue to have a linear problem. One example of such a constraint is as follows: For any $\alpha_{-\hat{I}}, \alpha'_{-\hat{I}} \in \boldsymbol{x}_{-\hat{I}}$,

$$\sum_{\substack{i \notin \hat{I} \\ \omega \in \Omega \\ \alpha_{\hat{I}} \in \boldsymbol{x}_{\hat{I}}}} \sum_{a_i \in A_i} (\alpha'_i(a_i) - \alpha_i(a_i))^2 x(\alpha'_{-\hat{I}}, \alpha_{\hat{I}}|\omega)\mu(\omega) = 0.$$

Clearly, this holds if and only if, for every $i \notin \hat{I}$ and $\alpha_i, \alpha'_i \in \boldsymbol{x}_i$, we have $\alpha_i(a_i) = \alpha'_i(a_i)$ for all $a_i \in A_i$. But this means that x "recommends" to every isolated player the same action independently of the state and the recommendations to others.

For the second approach, we can describe outcomes as a profile of mixed actions for the isolated players and a recommendation mechanism restricted to the non-isolated players. That is, a feasible outcome for (G, N, S) is a pair $(\alpha_{-\hat{I}}, x_{\hat{I}})$, where $x_{\hat{I}}$ satisfies

$$x_{\hat{I}}(\alpha_{\hat{I}}|\omega) = \sum_{t \in T} \pi_S(t|\omega) \prod_{i \in \hat{I}} \mathbb{I}\{\sigma_i(t_{N_i}) = \alpha_i\}$$

for some π_S and $(\alpha_{-\hat{I}}, \sigma_{\hat{I}}) \in \mathrm{BNE}(G, f_N(\pi_S))$. We can then say that $(\alpha_{-\hat{I}}, x_{\hat{I}})$ is spillover-robust obedient for (G, N) if for all $i \in \hat{I}$, $\alpha_{N_i} \in \boldsymbol{x}_{N_i}$, and $a_i \in A_i$,

$$\sum_{\substack{\omega\in\Omega \\ \alpha_{-N_i}\in\boldsymbol{x}_{-N_i}}} \Big(u_i(\alpha_i, \alpha_{\hat{I}\setminus\{i\}}, \alpha_{-\hat{I}}; \omega) - u_i(a_i, \alpha_{\hat{I}\setminus\{i\}}, \alpha_{-\hat{I}}; \omega)\Big) x_{\hat{I}}(\alpha_i, \alpha_{\hat{I}\setminus\{i\}}|\omega)\mu(\omega) \geq 0, \tag{7}$$

and for all $i \notin \hat{I}$ and $a_i \in A_i$,

$$\sum_{\substack{\omega\in\Omega \\ \alpha_{\hat{I}}\in\boldsymbol{x}_{\hat{I}}}} \Big(u_i(\alpha_i, \alpha_{\hat{I}}, \alpha_{-\hat{I}\setminus\{i\}}; \omega) - u_i(a_i, \alpha_{\hat{I}}, \alpha_{-\hat{I}\setminus\{i\}}; \omega)\Big) x_{\hat{I}}(\alpha_{\hat{I}}|\omega)\mu(\omega) \geq 0.$$

In words, each non-isolated player finds it optimal to follow the recommendations from $x_{\hat{I}}$, conditional on knowing his sources' recommendations (by spillovers) and the isolated players' behavior (by correct equilibrium conjectures) and on any inference about the other non-isolated players' behavior and the state (through $x_{\hat{I}}$). Each isolated player finds it optimal to implement his mixed action, conditional on knowing the other isolated players' behavior (by correct equilibrium conjectures) and the joint distribution of actions and states induced by $x_{\hat{I}}$. This formulation of obedience shows that now feasibility involves a fixed-point condition between $x_{\hat{I}}$ and the isolated players' behavior.

This second approach also helps to extend the comparative statics results of Section 4.1. Fix I and consider the information systems (N, S) and (N', S') that lead to the same set of non-isolated players $\hat{I}$. With regard to Lemma 1, given $\alpha_{-\hat{I}}$ equation (7) implies that if (N, S) is *less* connected than (N', S'), then the set of feasible outcomes among the players in $\hat{I}$ under (N, S) contains that under (N', S') for every G. Thus, overall $X(G, N', S') \subseteq X(G, N, S)$. The converse also holds (by the same argument as the baseline case) if we require the ranking of feasible sets to hold for all G.

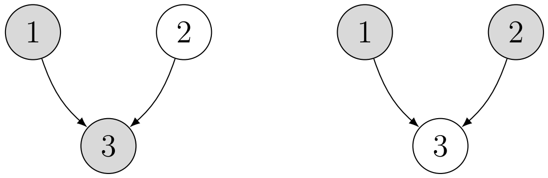

Figure 12: Influence and Disconnected Information Systems

Similarly, we can extend the analysis of Section 4.1.2 to settings where (N, S) is not connected. Fix N, S', and S''. Define $I' = \{i \in I : N_i \cap S' \neq \varnothing\}$ and $I'' = \{i \in I : N_i \cap S'' \neq \varnothing\}$ and assume that either $I' \subsetneq I$ or $I'' \subsetneq I$. There are two cases worth considering. Suppose that $I' = I''$. Then, S' is more influential than S'' if and only if $(N, S'') \trianglerighteq (N, S')$. This is a direct consequence of what we said in the previous paragraph and can be seen as an extension of Corollary 1. Alternatively, suppose that $N^{S'} = N^{S''}$. Then, S' is more influential than S'' if $I' \supseteq I''$. This latter case cannot arise under

Assumption 2 (which implies $I' = I'' = I$) and it is therefore conceptually distinct. We provide an example of this in Figure 12. In such example, let $S' = \{1, 2\}$ and $S'' = \{1, 3\}$. Note that $I' = I$ and $I'' = S''$. Thus, $I' \supseteq I''$. Moreover, $N^{S'} = N^{S''} = N$. Therefore, we have $X(G, N, S') = X(G, N^{S'}, S') = X(G, N^{S''}, I) \supseteq X(G, N, S'')$ for any G. We conclude that S' is more influential than S''.

D.3 More-connected Systems and Information Aggregation

We briefly explain the relation between more-connected information systems and their ability to aggregate information. For every $i \in I$, let $\Delta^{\pi_I}(\Omega \times A_i)$ be the set of distributions such that

$$Pr(\omega, a_i) = \sum_{t \in T} \zeta(a_i|t_i)\pi_I(t_i, t_{-i}|\omega)\mu(\omega), \quad (\omega, a_i) \in \Omega \times A_i,$$

for some $\zeta : T_i \to \Delta(A_i)$. In the spirit of Blackwell (1951), we say that π_I is more informative than π'_I for player i if $\Delta^{\pi'_I}(\Omega \times A_i) \subseteq \Delta^{\pi_I}(\Omega \times A_i)$. Moreover, we say that (N, I) aggregates information better than (N', I) if, for all π_I, $f_N(\pi_I)$ is more informative than $f_{N'}(\pi_I)$ for every player. We have the following characterization.

Lemma 9. *(N, I) aggregates information better than (N', I) if and only if $(N, I) \trianglerighteq (N', I)$.*

Proof. **Part 1 ($\Leftarrow$)**: Suppose $(N, I) \trianglerighteq (N', I)$ and so $N_i \supseteq N'_i$, for every i. We want to show that, for every π_I, $\Delta^{f_{N'}(\pi_I)}(\Omega \times A_i) \subseteq \Delta^{f_N(\pi_I)}(\Omega \times A_i)$ for all i. Fix a payer i and (π_I, T). Let $(\hat{T}, f_N(\pi_I))$ and $(\hat{T}', f_{N'}(\pi_I))$ be the induced information structure for N and N'. Fix $y \in \Delta^{f_{N'}(\pi_I)}(\Omega \times A_i)$. We will show that $y \in \Delta^{f_N(\pi_I)}(\Omega \times A_i)$. Since $y \in \Delta^{f_{N'}(\pi_I)}(\Omega \times A_i)$, there exists $\gamma' : \hat{T}'_i \to \Delta(A_i)$ (where $\hat{T}'_i = \times_{j \in N'_i} T_j$) such that

$$y(\omega, a_i) = \sum_{t} \gamma'(a_i|t_{N'_i})\pi_I(t|\omega)\mu(\omega), \quad (\omega, a_i) \in \Omega \times A_i.$$

Now define $\gamma : \hat{T}_i \to \Delta(A_i)$ (where $\hat{T}_i = \times_{j \in N_i} T_j$) from γ' by letting

$$\gamma(a_i|t_{N'_i}, t_{N_i \setminus N'_i}) = \gamma'(a_i|t_{N'_i})$$

for all $t_{N_i \setminus N'_i}$. That is, γ depends only on the components in N_i that also belong to N'_i and in the same way that γ' depends on them. This is well defined because $N_i \supseteq N'_i$. Clearly, γ leads to the distribution y induced by γ', which therefore belongs to $\Delta^{f_N(\pi_I)}(\Omega \times A_i)$.

Part 2 ($\Rightarrow$): We prove this by contrapositive. The only relevant case to consider is that $(N, I) \ntrianglerighteq (N', I)$ and $(N, I) \ntrianglelefteq (N', I)$. This implies that for some i, there exists a k such that $k \in N'_i$ and $k \notin N_i$, and for some j (possibly $i = j$), there exists m such that $m \in N_j$ and $m \notin N'_j$. It follows that there exists a player i_k such that $i_k \neq k$ and there is a direct link from k to i_k in N' but not in N, and there exists a player i_m such that $i_m \neq m$ there is a direct link from m to i_m in N but not in N'.

First, take an information structure $\pi_{\{1\}}$ such that k gets full information and all other players always get fully uninformative signals. Then, under $f_{N'}(\pi_{\{1\}})$ player i_k gets full information, while under $f_N(\pi_{\{1\}})$ he still gets no information. Therefore, (N, I) does not aggregate more information than (N', I). Now, take an information structure $\pi_{\{2\}}$ such that m gets full information and all other players always get fully uninformative signals. Then, under $f_N(\pi_{\{2\}})$ player i_m gets full information, while under $f_{N'}(\pi_{\{2\}})$ he still gets no information. Therefore, (N', I) does not aggregate information better than (N, I). ☐

D.4 Pure-Strategy Recommendations in Example 4

In Section C.2, we restricted attention to pure-action recommendations. We now show that this is without loss of generality by showing that x^* solves the designer's problem even if we allow for mixed recommendations. The argument follows the duality approach of Galperti and Perego (2018). The first step is to construct a solution to the dual of the design problem restricted to pure-recommendation mechanisms and show that the value of this dual is the same as the value of the primal under x^*. Given this, note that for the unrestricted problem that allows for mixed recommendations, the value of the primal must be weakly larger than under x^* and the value of the dual must be weakly smaller than under the solution of the restricted-problem dual. Since by weak duality the value of the primal has to be smaller than the value of the dual in the unrestricted problem, it follows that x^* also solves the unrestricted problem by strong duality.

Regarding the first step, the dual of the problem restricted to pure-recommendation mechanisms involves choosing $p \in \mathbb{R}^{\Omega}$,[26] $\lambda_1 \in \mathbb{R}^{A_1 \times A_1}$, and $\lambda_2 \in \mathbb{R}^{A_2 \times A}$ so as to minimize

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

subject to $\lambda_1(a_1'|a_1) \geq 0$ for all $a_1, a_1' \in A_1$, $\lambda_2(a_2'|a) \geq 0$ for all $a_2 \in A_2$ and $a \in A$, and for all $(a, \omega) \in A \times \Omega$,

$$\begin{aligned} p(\omega) &\geq w(a_1, a_2, \omega, \lambda_1, \lambda_2) \\ &\equiv \mathbb{I}\{a_1 = y_1\} + \mathbb{I}\{a_2 = y_2\} + \Big(u_1(a_1, a_2, \omega) - u_1(a_1', a_2, \omega)\Big)\lambda_1(a_1'|a_1) \\ &\quad + \Big(u_1(a_1, a_2, \omega) - u_1(a_1, a_2', \omega)\Big)\lambda_2(a_2'|a). \end{aligned}$$

---

[26] The dual variable $p(\omega)$ corresponds to the constraint $\sum_{a \in A} x(a|\omega) = 1$ for all ω.

When $\omega = b$, we have

$$w(a_1, a_2, b, \lambda) = \begin{cases} 2 - [1-\varepsilon]\lambda_1(n_1|y_1) - [1-\varepsilon]\lambda_2(n_2|y_1, y_2) & \text{if } (a_1, a_2) = (y_1, y_2) \\ 1 - \lambda_1(n_1|y_1) + [1-\varepsilon]\lambda_2(y_2|y_1, n_2) & \text{if } (a_1, a_2) = (y_1, n_2) \\ 1 + [1-\varepsilon]\lambda_1(y_1|n_1) - \lambda_2(n_2|n_1, y_2) & \text{if } (a_1, a_2) = (n_1, y_2) \\ \lambda_1(y_1|n_1) + \lambda_2(y_2|n_1, n_2) & \text{if } (a_1, a_2) = (n_1, n_2). \end{cases}$$

Similarly, when $\omega = g$, we have

$$w(a_1, a_2, g, \lambda) = \begin{cases} 2 + [\gamma_1+\varepsilon]\lambda_1(n_1|y_1) + [\gamma_2+\varepsilon]\lambda_2(n_2|y_1, y_2) & \text{if } (a_1, a_2) = (y_1, y_2) \\ 1 + \gamma_1\lambda_1(n_1|y_1) - [\gamma_2+\varepsilon]\lambda_2(y_2|y_1, n_2) & \text{if } (a_1, a_2) = (y_1, n_2) \\ 1 - [\gamma_1+\varepsilon]\lambda_1(y_1|n_1) + \gamma_2\lambda_2(n_2|n_1, y_2) & \text{if } (a_1, a_2) = (n_1, y_2) \\ -\gamma_1\lambda_1(y_1|n_1) - \gamma_2\lambda_2(y_2|n_1, n_2) & \text{if } (a_1, a_2) = (n_1, n_2). \end{cases}$$

In Galperti and Perego (2018), we show how to derive the solution of the dual (and primal) following a few simple steps. Here, we only conjecture a solution and apply duality. Let

$$\lambda_1(n_1|y_1) = 1, \quad \lambda_2(n_2|y_1, y_1) = \frac{1+\varepsilon}{1-\varepsilon}, \quad \text{and} \quad \lambda_2(n_2|n_1, y_2) = 1;$$

for the other cases, set both λ_1 and λ_2 equal to zero. Since $|\varepsilon| \approx 0$ and $\gamma_2 < \gamma_1 < 1$, we have that $p(b) = 0$ and

$$p(g) = 2 + [\gamma_1+\varepsilon]\lambda_1(n_1|y_1) + [\gamma_2+\varepsilon]\lambda_2(n_2|y_1, y_2).$$

It follows that the value of the dual is

$$\frac{1}{2}p(g) + \frac{1}{2}p(b) = 1 + \frac{\gamma_1+\varepsilon}{2} + \frac{\gamma_2+\varepsilon}{2}\left[\frac{1+\varepsilon}{1-\varepsilon}\right] = \chi_1^* + \chi_2^*,$$

where the latter is the value of the restricted primal under x^*.

D.5 Omitted Proofs

Proof of Lemma 2

Step 1. Consider any finite-support (T, π_I) and let σ be the designer-preferred equilibrium of $(G, f_N(\pi_I))$. For every i, every t_{N_i} determines a non-empty subset of optimal actions:

$$A_i(t_{N_i}) = \arg\max_{a_i \in A_i} \mathbb{E}_{\pi_I,\sigma}\Big(u_i(a_i, a_{-i}, \omega) \mid t_{N_i}\Big).$$

Since A_i is finite, every (π, σ) can determine at most finitely many subsets $A_i(t_{N_i})$ for every player i. This requires no more than $|2^{A_i}|$ signals for player i. Therefore, every (π_I, σ) can determine at most finitely many profiles of optimal-action sets of the form $A(t) = \times_i A_i(t_{N_i})$. We conclude that if we are interested in only such profiles, it is enough to consider information structures that satisfy $|T_i| = |2^{A_i}|$ 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 outcome, hence satisfies obedience. That is, for every i, $\alpha_{N_i} \in \boldsymbol{x}_{N_i}$, and $a_i, a_i' \in A_i$,

$$\sum_{\omega,\alpha_{-N_i}} \Big(\sum_{a_{-i}} \Big(u_i(a_i, a_{-i}, \omega) - u_i(a_i', a_{-i}, \omega) \Big) \alpha_{N_i}(a_{N_i}) \alpha_{-N_i}(a_{-N_i}) \Big) x(\alpha_{N_i}, \alpha_{-N_i}|\omega)\mu(\omega) \geq 0,$$

where $\alpha_{N_i}(a_{N_i}) = (\alpha_j(a_j))_{j\in N_i}$ and $\alpha_{-N_i}(a_{-N_i}) = (\alpha_j(a_j))_{j\notin N_i}$. We want to construct an alternative x' that is also feasible 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 $A^x(\alpha) = \times_{i\in N} A_i^x(\alpha_{N_i})$, 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 for every state: For every a and ω,

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

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

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

Now, recall that for every $\alpha \in A^x$, we have that the set of optimal actions for player i conditional on α_{N_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_{N_i}. Therefore, the entire sum must be non-negative. This shows that x' satisfies obedience for player i and every $\alpha'_{N_i} \in \boldsymbol{x}'_{N_i}$. By the same argument, x' satisfies obedience for all players. ☐

Proof of Lemma 3 Since $V^*(G, N, I) = \sup_{x \in X(G,N,I)} \sum_{\omega,\alpha} v(\alpha, \omega) x(\alpha|\omega)\mu(\omega)$, for every $n \geq 1$ there exists x^n that satisfies

$$V^*(G, N, I) \geq \sum_{\omega,\alpha} v(\alpha, \omega) x^n(\alpha|\omega)\mu(\omega) \geq V^*(G, N, I) - \frac{1}{n}.$$

Moreover, by Lemma 2 we can choose x^n so as to satisfy $|\boldsymbol{x}_i| \leq |2^{A_i}|$ for every i. Let $K_i = |2^{A_i}|$ for every i. To every such x^n there correspond finite subsets $A^n_i \subset \Delta(A_i)$ such that $|A^n_i| = K_i$ for all i, which define a grid in $\times_{i \in I} \Delta(A_i)$ over which we can restrict the construction of x^n itself. (Note that the support of x^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 $\alpha^{k_i,n}_i \in \Delta(A_i)$ where $\alpha^{k_i,n}_i \in A^n_i$ is an element of the grid of player i with (fixed) K_i elements to construct x^n. Also, for each ω and every $(k_1, \ldots, k_I)$ where $k_i = 1, \ldots, K_i$ for every i, we have a sequence of elements $x^n(\alpha^{k_1,n}_1, \ldots, \alpha^{k_I,n}_I, \omega) \in [0, 1]$. 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:

$$\lim_{\tilde{n}\to\infty} \alpha_i^{k_i,\tilde{n}} = \hat{\alpha}_i^{k_i} \in \Delta(A_i), \quad k_i = 1,\dots,K_i, i \in I;$$
$$\lim_{\tilde{n}\to\infty} x^{\tilde{n}}(\alpha_1^{k_1,\tilde{n}},\dots,\alpha_I^{k_I,\tilde{n}}|\omega) = \hat{x}(\hat{\alpha}_1^{k_1},\dots,\hat{\alpha}_I^{k_I}|\omega), \quad k_i = 1,\dots,K_i, \omega \in \Omega.$$

Since $x^{\tilde{n}} \in X(G,N,I)$ for all $\tilde{n}$ by assumption, it is easy to see that $\hat{x} \in X(G,N,I)$ by continuity of the linear constraints that define $X(G,N,I)$. Finally, we have that

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

Therefore, $\sum_{\omega,\hat{\alpha}} v(\hat{\alpha},\omega)\hat{x}(\hat{\alpha}|\omega)\mu(\omega) = V^*(G,N,I)$, which proves the result. $\square$

D.6 News Consumption and Elections: Analysis

This section explains the outcomes x that lead to the passage probabilities in Figure 10.(a) for the four sets of seeds. At the end, it also presents other parametrizations of Figure 10.

Broad Outlet Alone

Suppose we shut down both one-sided outlets: $\hat{S} = \{i_B\}$. Let x_B be the outcome that maximizes the passage probability. To understand x_B, it suffices to examine its support. We numerically calculate x_B and report its support, which does not change in the feasible space of paramteters c and θ, in Table 2. The profile (a_M, a_m) indicates that x_B recommends $a_M \in \{y,n,h\}$ to all majority voters and $a_m \in \{y,n,h\}$ to all minority voters. The structure of x_B is equivalent to a standard Bayesian-persuasion problem with one receiver: the majority voters. If the state is favorable to them ($\omega = -1$), x_B always recommends them to vote y; if the state is unfavorable to them ($\omega = 1$), with some probability x_B recommends them to vote n—thus revealing the state—and with the remaining probability it recommends them to vote y. These probabilities are such that recommendation y renders the majority voters indifferent between y and h. Since such recommendations are public, the information they convey must induce the minority voters to choose opposite actions to those of the majority voters.[27]

---

[27] The cases with only the majority or only the minority outlet have similar solutions—except that in the latter case, for the minority to carry the bill, the majority must abstain under no news and so c must be large enough.

		m voters		
		y	h	n
	y	0	0	•
M voters	h	0	0	0
	n	0	0	0

$\omega = -1$

		m voters		
		y	h	n
	y	0	0	•
M voters	h	0	0	0
	n	•	0	0

$\omega = 1$

Table 2: Support of Optimal x_B.

Broad Outlet and Majority Outlet

Suppose we shut down only the minority outlet: $\hat{S} = \{i_B, i_M\}$. Let x_{BM} be the outcome that maximizes the passage probability. Similarly to before, Table 3 represents the numerically calculated support of x_{BM}, which again does not change across feasible (c, θ). To see how x_{BM} leads to a higher passage probability than under x_B, note that x_{BM} sometimes induces the minority to vote y and the majority to choose h. This is perhaps surprising, as the majority always knows the minority's information, which comes from i_B. The key is that the recommendation y to the minority is only partially informative. We can interpret how x_{BM} is determined by news coming from the two outlets as follows. The news from broad outlet induces the minority to vote y or n. In addition, when they vote n, the news from the majority outlet induces its audience to vote y; when the minority votes y, instead, the news from the majority outlet induces its audience to sometimes vote n and sometimes abstain. This illustrates the effect of the majority outlet's ability to convey news privately to its audience.

		m voters		
		y	h	n
	y	0	0	•
M voters	h	•	0	0
	n	0	0	0

$\omega = -1$

		m voters		
		y	h	n
	y	0	0	•
M voters	h	•	0	0
	n	•	0	0

$\omega = 1$

Table 3: Support of Optimal x_{BM}.

Broad Outlet and Minority Outlet

Suppose we shut down only the majority outlet: $\hat{S} = \{i_B, i_m\}$. Let x_{Bm} be the outcome that maximizes the passage probability. Similarly to before, Table 4 and 5 represent the numerically calculated support of x_{Bm}, which now changes across feasible (c, θ). When c (respectively, θ) is sufficiently low (high) $x_{Bm} = x_B$. This is because the majority votes n under the prior and it is informationally too costly to induce it to abstain. By contrast, when c (respectively, θ) is sufficiently high (low), x_{Bm} differs from x_B.[28] To understand how x_{Bm} leads to a higher passage probability than under x_B, note that with some probability x_{Bm} induces the minority to vote y while the majority abstains (h). We can interpret how the outlets' news induce this as follows. If c is high, the majority abstains under no news. Exploiting this, now the broad outlet conveys news that either induces the majority to vote y or to abstain. If the latter happens, the news from the minority outlet either reveals that the state is $\omega = -1$—leading the minority to vote n—or induces the minority to vote y.

m voters

		y	h	n
M voters	y	0	0	●
	h	0	0	0
	n	0	0	0

$\omega = -1$

m voters

		y	h	n
M voters	y	0	0	●
	h	0	0	0
	n	●	0	0

$\omega = 1$

Table 4: Support of Optimal x_{Bm} for Low c (High θ).

m voters

		y	h	n
M voters	y	0	●	●
	h	●	0	●
	n	0	0	0

$\omega = -1$

m voters

		y	h	n
M voters	y	0	●	0
	h	●	0	0
	n	0	0	0

$\omega = 1$

Table 5: Support of Optimal x_{Bm} for High c (Low θ).

---

[28] In this case, x_{Bm} sometimes induces the majority to vote y, which renders the minority irrelevant. Since, unlike the majority, the minority also receives news from the minority outlet, such news can induce the minority to choose different actions while the majority votes y. This gives rise to multiple solutions, of which we report one kind.

All Outlets

Finally, suppose we keep all outlets: $\hat{S} = \{i_B, i_M, i_m\}$. Let x_{BMm} be the outcome that maximizes the passage probability. Similarly to before, Table 6 and 7 represent the numerically calculated support of x_{BMm}, which again changes across feasible (c, θ). As under x_{BM} and x_{Bm}, under x_{BMm} sometimes the majority passes the bill, and sometimes the minority passes the bill while the majority abstains. In contrast to the previous cases, however, x_{BMm} sometimes induces unanimous support for the bill *in state* $\omega = 1$. This hinges on news that each one-sided outlet conveys privately to its audience. When c is low (respectively, θ is high), under x_{BMm} the only scenario when the bill does *not* pass involves the majority's opposition and the minority's support. When c is high (respectively, θ is high), the only scenario where the bill does not pass involves the majority's abstention and the minority's opposition. This regime change explains the kink in Figure 10. Overall, these points illustrate the effect of both one-sided outlets' conveying news privately to their audiences.

		m voters		
		y	h	n
M voters	y	0	0	•
	h	•	0	0
	n	0	0	0

$\omega = -1$

		m voters		
		y	h	n
M voters	y	•	0	0
	h	•	0	0
	n	•	0	0

$\omega = 1$

Table 6: Support of Optimal x_{BMm} for Low c (High θ).

		m voters		
		y	h	n
M voters	y	0	0	•
	h	•	0	•
	n	0	0	0

$\omega = -1$

		m voters		
		y	h	n
M voters	y	•	0	0
	h	•	0	0
	n	0	0	0

$\omega = 1$

Table 7: Support of Optimal x_{BMm} for High c (Low θ).

Alternative Parametrizations and Passage Probabilities

Figure 13 represents the passage probabilities for the four sets of seeds as in Figure 10, but for different choices of θ in panel (a) and c in panel (b). Figure 13 exhibits the same qualitative properties discussed in relation to Figure 10.

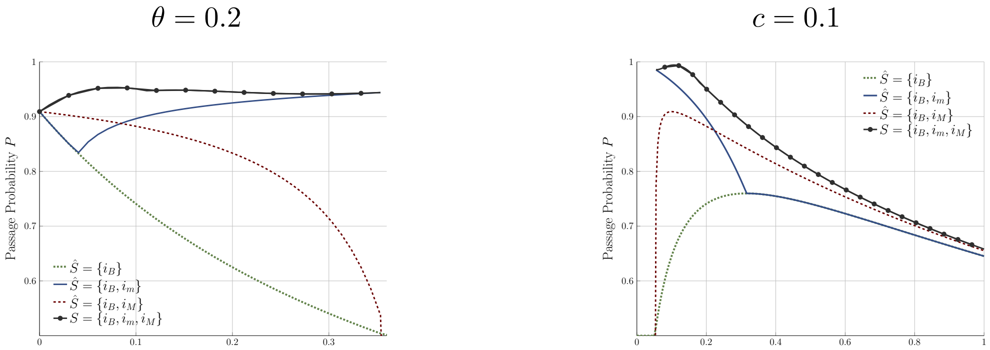
$\theta = 0.2$ $c = 0.1$

(a) Voting Cost c

(b) Polarization θ

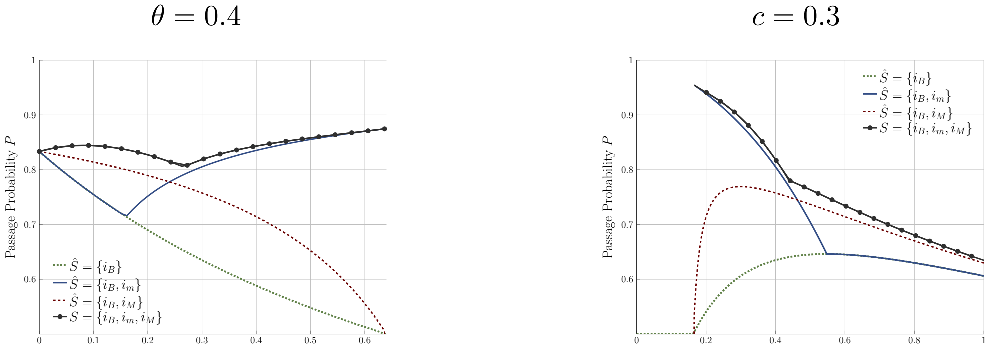
$\theta = 0.4$ $c = 0.3$

(a) Voting Cost c

(b) Polarization θ

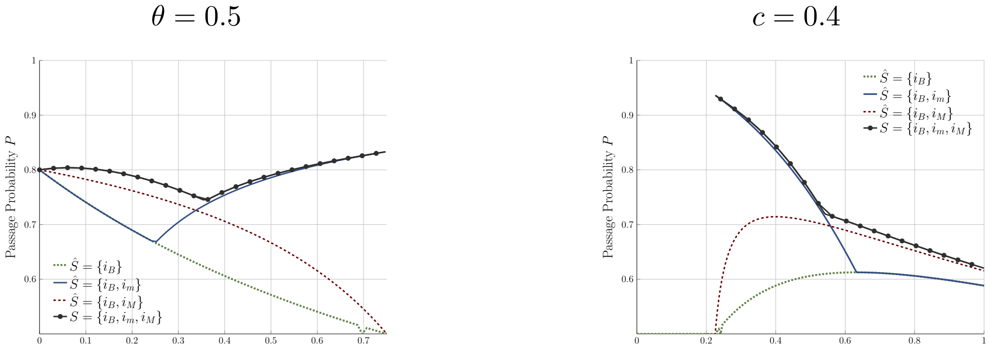
$\theta = 0.5$ $c = 0.4$

(a) Voting Cost c

(b) Polarization θ

Figure 13: Robust Passage Probability (Alternative Parametrizations)