Information Design in Operations

Ozan Candogan
University of Chicago, Booth School of Business, ozan.candogan@chicagobooth.edu

Abstract Consider a set of agents (receivers) whose payoffs depend on an underlying state of the world as well as each other's actions. Suppose that a designer (sender) commits to a signaling mechanism which reveals payoff-relevant signals to agents when the state is realized. The availability of such signals influences the agents' actions, and by choosing the signaling mechanism appropriately the designer can induce a desired outcome. Information design studies signaling mechanisms that maximize the payoff of the designer. In this paper, we first present the classical information design framework and discuss different approaches for characterizing the optimal information structures. We then discuss various applications in the recent operations literature. The applications include signaling (i) content/product quality in networked systems, (ii) product availability in revenue management settings, and (iii) seller quality in two-sided markets. Finally, we present recent work that discusses the design of optimal information structures when some of the key assumptions in the classical information design problems (which may not hold in operational settings of interest) are relaxed.

Keywords Information design, persuasion, persuasion in networks.

---

1. Introduction

Information design studies how a designer (or sender) can influence the actions taken by agents (or receivers) by committing to a mechanism that reveals signals about a payoff relevant state once the state is realized. Since the seminal papers [10, 30, 38], information design has become a very active research area.

Information is a natural lever in many important operational settings. Review systems provide useful information to consumers about various service operations, ranging from restaurants to hospitals. Ride-sharing platforms signal rider demand to their drivers and on-demand labor platforms signal freelancer quality and supply to potential employers. Many applications provide real-time traffic information to their users, which impacts the routes individuals take in transportation networks. E-commerce platforms often ask the buyers to rate third-party sellers they transact with, and make these ratings available to future buyers. Recent work has leveraged ideas from information design to study how information can be used to induce desired outcomes in different operational settings.

The purpose of this paper is three-fold. First, we present the key approaches from the literature that can be used to characterize the optimal information structures in a variety of information design problems. Second, we cover some of the recent work from operations that applies information design ideas in different operational settings. Third, we highlight some assumptions made in the classical information design formulations, which may not hold in important operational settings. We then discuss the recent work that sheds light on the design of optimal information structures when these assumptions are relaxed.

To understand the gist of information design, let us start with a simple example adapted from the "courtroom" example of [30]. A firm faces a single buyer who decides whether to adopt its new product. The product may be a good match with the buyer or not (which corresponds to the state of the world). The payoffs are normalized so that if the buyer takes the "correct" action (i.e., purchases the product when the match is good, and does not purchase when the match is poor) her payoff is one. Otherwise, her payoff is zero. The firm gets a payoff of one when the product is sold and zero otherwise.

A priori the quality of the match is unknown to the firm and the buyer. They share a common prior belief which suggests that the match is good with probability 0.3. The firm can choose a (signaling) mechanism that provides informative signals to the buyer about the product. For instance, it can provide free trials (where the buyer can explore a limited subset of the product capabilities) and/or highlight some features of the product with factual advertising.

If the firm does not reveal any information about the product, then the buyer is better off not purchasing it, which yields a payoff of zero to the firm. If the firm chooses to reveal the information about the product fully (e.g., through a long trial of the product with full capabilities) then, the buyer purchases the product with probability 0.3 – whenever the match is good. Suppose that the firm uses a trial/advertising campaign which reveals limited information about the match and suggests that the match is good (i) with probability 1 when it is indeed good, and (ii) with probability 3/7 when it is not. A mechanism of this sort makes the buyer indifferent between purchasing and not purchasing the product whenever she receives a signal indicating that the match is good. This guarantees that the product will be sold with probability 0.6 – which is the maximum achievable for the setting described here. Thus, the firm can strictly improve its expected payoff by designing the mechanism appropriately.

Information design tries to address how to design the information structures in a way that maximizes the designer's payoffs in general environments. Often, we focus on problems where the designer's and the agents' payoffs depend on the receivers' actions and an underlying state. The designer, before observing the state, commits to a mechanism (or information structure) that reveals signals once the state is realized. The optimal mechanism is the one that maximizes the designer's payoff.

Here, the commitment assumption plays an important role and makes the environment different from the cheap talk setting. It removes ambiguity on how the receivers interpret the designer's messages, and in a way simplifies the design problem. While commitment may be a strong assumption in some settings, it is reasonable in other settings. For instance, in some settings (e.g., limited product trials) the designer's choice of the mechanism as well as the realizations of the signals (e.g., how much the receiver enjoys the features included in the trial) are readily observable to the receiver. Reputation effects, which are relevant in many operational settings, could also play a role. A platform revealing information about the (possibly low) qualities of sellers to the buyers, considers not only the short term loss of sales, but also its long term reputation. [6, 36] formalize some of these ideas and illustrate how repeated interaction of a long-lived sender with receivers can restore the commitment payoffs. In many platforms and review systems, choosing an algorithm (which is only infrequently modified) that reveals signals to users can be viewed as commitment to a mechanism as well. Yet in other settings, there are legal or contractual constraints that make commitment feasible (see, e.g., [30]).

As highlighted in [5] information design perspective provides a useful benchmark even when there is not a literal information designer. In particular, characterizing the set of outcomes an information designer can induce, equivalently characterizes the set of all outcomes that emerge under *some* information structure. For instance, consider a setting where the agents have access to some exogenous signals about the state and focus on the maximum payoff that will be induced under these signals. The best payoff a designer can achieve (using the optimal information structure) upper bounds the aforementioned payoff. Thus, the characterization of optimal information structures and the corresponding payoffs yields an important benchmark that is relevant even in the absence of an information designer.

The recent literature has studied a variety of operational problems through the lens of information design. The applications of information design ideas have emerged in (i) classical queueing settings, (ii) revenue management, (iii) social networks, (iv) platform operations, among others. In one of the earliest applications of information design in operations, [35] studies how a designer can reveal informative signals on the queue length to influence customers' decisions on whether to join a queue. The objective of the designer is to maximize the revenue collected from the customers who join. [2] builds on this model to study the welfare maximization problem in a setting where different types of customers who have different needs decide whether to join the queue. [17, 33, 34] study complementary revenue management settings. [17] and [34] explore how a seller can signal product availability to persuade buyers to purchase the product earlier (and at a higher price). [33] explores how a seller can disclose information about historical sales to convince customers who arrive over time to purchase the product. [16, 12] study the applications of information design ideas in the context of social and economic networks. [16] considers engagement decisions of agents with the available content on social media, and designs signaling mechanisms that reveal informative signals about the accuracy of the content. [12] studies how appropriate public reviews/signaling mechanisms could shape agents' adoption decisions of a product that exhibits local network externalities. Different aspects of platform operations have been studied by using tools from information design. [37] studies how platforms can induce exploration by provisioning information appropriately to their users. In [39], motivated by ride-sharing problems, the authors study how a platform can improve the outcome of a spatial resource competition and increase welfare by using appropriate public or private signaling mechanisms. [29] explores how a platform managing a two-sided market can signal the quality of the sellers to increase the total transaction value in the market. [7] studies a setting where both sides of the market endogenously decide whether to participate, and the authors shed light on the supply-side benefits of information provision by the platform managing this market. [26] studies the relationship between a platform and a third-party seller that does not know consumer demand for its product but can use dynamic pricing to learn from consumer purchase decisions. To induce the seller to set platform-preferred prices, the platform commits to an initial information structure that discloses (some) information about demand and then takes costly actions to constrain the seller's ability to learn from purchase decisions. The paper explores how the seller's ability to learn after the initial information disclosure impacts the optimal information structure. Applications in other new domains are also explored by the recent literature. One interesting example is [1], where the authors explore how a public health agency can signal the severity of a pandemic to induce agents to undertake costly measures.

Thus far, many papers in the literature focused on stylized theoretical models and developed guiding principles for the design of information provision schemes that are relevant in different operational settings. They also illustrated how some insights on the value of operational levers (such as pricing) may change when information is used as another lever (see the related discussion in Section 3.2). Going forward, the information design ideas are likely to find fruitful applications in many other operational settings. We comment on some future directions in our concluding remarks.

In this paper, we first focus on an abstract setting and present the information design problem as well as various approaches that can be used for obtaining optimal information structures (Section 2). Then, in Section 3, we discuss a few papers that develop applications of information design in the recent operations literature. In Section 4, we discuss some key assumptions that are imposed in the classical information design settings and present recent work that sheds light on how the design of optimal mechanisms changes when these assumptions do not hold. We conclude in Section 5.

2. Different Approaches to Information Design

In this section, we review different ideas from the literature that are useful for the solutions of information design problems. We start by introducing the notation that will be used in the remainder of the paper.

The state of the world belongs to a set $\mathcal{T} \subset \mathbb{R}$. We use T to denote the random state and t to denote either the realization of the state or a dummy state $t \in \mathcal{T}$. There is a set of receivers (agents) V. The payoff of receiver i, denoted by $u_i(a_i, a_{-i}, t)$, is a function of her action a_i (which belongs to a set of possible actions A_i), the remaining receivers' actions $a_{-i} = \{a_j\}_{j\in V\setminus\{i\}}$, and the state realization $T = t$. The receivers do not observe the state realization prior to choosing their actions. An information designer's payoff $\nu(a_i, a_{-i}, t)$ depends on the actions chosen by the receivers and possibly on the state. The designer and the agents share a common prior on the state given by the (cumulative) distribution function $F : \mathcal{T} \to [0, 1]$. We let[1] $F^{-1}(q) = \inf\{t \in \mathcal{T} | F(t) \geq q\}$.

Before the state is realized the designer commits to a mechanism (or information structure) π which maps each state realization (possibly after randomization) to a signal realization. We denote by S_i the random variable that represents receiver i's signal. Once the state is realized, the designer's mechanism shares the realization of S_i with receiver i. Then, the receivers take actions to maximize their expected payoffs (conditional on the observed signal). Formally, given the designer's mechanism, the agents play a game of incomplete information. In this game, agent i's strategy x_i is a mapping from the possible realizations of S_i to an action in A_i. We focus on the Bayesian Nash equilibria of this game, and denote the set of these equilibria by $\mathcal{Q}$. If there are multiple Bayesian Nash equilibria we break ties in favor of the equilibrium that maximizes the designer's payoff – or focus on the sender-preferred equilibrium. By designing her mechanism appropriately, the designer can influence the equilibria of the game among the agents and improve her own payoff.

Let Π denote the set of feasible mechanisms of the designer. There are two important cases to consider. In the first one the mechanisms in Π are not restricted, and only mild measurability assumptions are made on the mechanisms (i.e., for all $i \in V$ signal S_i needs to be measurable). In the second one, attention is restricted to public signaling mechanisms, i.e., in addition to measurability we require $S_i = S$ for all $i \in V$. The designer's problem is to choose $\pi \in \Pi$ that maximizes her expected payoff. This problem can mathematically be stated as follows: $\max_{\pi\in\Pi} \max_{(x_i,x_{-i})\in\mathcal{Q}} \mathbb{E}[\nu(x_i(S_i), x_{-i}(S_{-i}), T)]$.

Note that the set of feasible mechanisms is very rich, and in general the dependence of the optimal mechanism on the prior belief and the designer's/agents' payoffs is quite nontrivial. Partly due to these features, information design problems can become algorithmically challenging. We next outline different ideas for the solution of information design problems. In Section 2.1, we discuss the concavification idea, which provides an elegant geometric approach for understanding when the designer can use information as a lever to improve her payoff. Section 2.2 provides a formulation in terms of Bayes correlated equilibria (BCE), which paves the way for natural optimization formulations for the characterization of optimal information structures in the presence of multiple receivers. The approaches of these two subsections are useful when the designer's/agents' payoffs have a general dependence on the posterior distributions induced by the signals. On the other hand, the optimization formulations obtained from the BCE concept involve a decision variable for each state and profile of agents' actions which may lead to algorithmic challenges. In Sections 2.3 – 2.5 we focus on settings where the designer's payoff depends only on the induced posterior mean (as opposed to the entire posterior distribution). This dependence enables alternative formulations that exploit the aforementioned structure. Section 2.3 formulates the information design problem as an optimization problem over a set of convex functions, whereas Section 2.4 provides an infinite-dimensional optimization formulation and leverages ideas from duality. Section 2.5 studies settings where the designer's payoff is a step function of the induced posterior mean, and provides a tractable (finite-dimensional) convex optimization formulation for obtaining the optimal information structure. We illustrate the last approach in Section 3.1.

[1] Observe that for $q \in [0, 1]$ and strictly increasing $F(\cdot)$, this definition is consistent with the inverse of the cumulative distribution function.

2.1. Concavification

The seminal paper [30] presents the Bayesian Persuasion framework. The baseline model focuses on a setting with a single receiver and finitely many states (hence F is atomic) – though the results are then extended to richer environments. The designer chooses a mechanism with finitely many possible signal realizations to influence the receiver's decisions.

Given the mechanism of the designer, the receiver chooses an expected payoff maximizing strategy. It is without loss of optimality to restrict attention to *straightforward* mechanisms where the designer recommends the receiver to take an action, and the mechanism is designed so that it is optimal for the receiver to follow this recommendation. The idea behind straightforwardness is similar to the revelation principle in mechanism design.

Recall that F denotes the prior belief about the state. Each signal realization $S = s$ induces a posterior belief in $\Delta(\mathcal{T})$. The mechanism chosen by the designer, therefore, induces a distribution of posterior beliefs $\tau \in \Delta(\Delta(\mathcal{T}))$. Observe that the latter distribution needs to be consistent with prior beliefs. In particular, it needs to be *Bayes Plausible*:

$$\sum_{\mu \in Supp(\tau)} \mu\tau(\mu) = F,$$

where $Supp(\tau)$ is the (finite) set of beliefs in the support of τ. For a given belief $\mu \in \Delta(\mathcal{T})$, denote the payoff of the designer by

$$\hat{\nu}(\mu) = \mathbb{E}_{\mu}[\nu(a(\mu), T)],$$

where with some abuse of notation we let $a(\mu)$ denote the action that maximizes the receiver's payoff when her belief is μ (and if there are multiple such actions once again we break ties in favor of the designer).

A fundamental question in information design is when the designer benefits from an appropriate choice of the information structure. [30] answers this question through the idea of concavification. Let $\hat{\nu}_c$ denote the concave closure of $\hat{\nu}$, i.e.,

$$\hat{\nu}_c(\mu) = \sup\{z | (\mu, z) \in co(\hat{\nu})\},$$

where $co(\hat{\nu})$ denotes the convex hull of the graph of $\hat{\nu}$. [30] establishes that the designer can improve her payoff by choosing an appropriate information structure (or benefits from persuasion) if and only if:

$$\hat{\nu}_c(F) > \hat{\nu}(F).$$

That is, to see whether the designer benefits from persuasion it suffices to check whether the payoff function $\hat{\nu}$ and its closure $\hat{\nu}_c$ match at the prior belief F. Note that this immediately implies that if $\hat{\nu}$ is concave, the designer does not benefit from persuasion for any prior. Conversely, when $\hat{\nu}$ is convex and not concave, the designer benefits from persuasion for every (non-degenerate) prior.

It can also be shown that the payoff of the designer under the optimal mechanism is exactly $\hat{\nu}_c(F)$. Moreover, in some special cases, it is possible to characterize $\hat{\nu}_c(\cdot)$ explicitly, and use this characterization to derive the beliefs the designer's mechanism needs to induce to maximize her payoff.

For instance, consider the example from the introduction. In this example, there are two possible state realizations, and the distributions can be equivalently represented in terms of the probability weight of one of them. Moreover, the buyer finds it optimal to adopt the product if she believes that the probability of the product being a good match is greater than 0.5. Thus, using the aforementioned alternative representation of the distributions, the designer's payoff can be expressed as a step function of the belief (with a cutoff at 0.5). The concavification approach establishes that when the prior probability of the product being a good match is less than 0.5, the designer can improve her payoff by employing an appropriately chosen information structure. Moreover, given its simple structure, in this example it is possible to explicitly compute the concave closure of the designer's payoff and the optimal information structure (see [30] for details).

In an important class of persuasion problems, the designer's payoff depends on the expected state. That is, there is a function $\tilde{\nu}$ such that $\tilde{\nu}(\mathbb{E}_\mu[T]) = \hat{\nu}(\mu)$. The concavification idea applies in this setting as well, and it can be shown that the designer benefits from persuasion if and only if $\tilde{\nu}_c(\mathbb{E}_F[T]) > \tilde{\nu}(\mathbb{E}_F[T])$, where $\tilde{\nu}_c$ is the concave closure of $\tilde{\nu}$. It is natural to expect that $\tilde{\nu}_c(\mathbb{E}_F[T])$ is the designer's payoff under an optimal mechanism. But this is no longer the case, and the aforementioned quantity is just an upper bound.

2.2. Bayes Correlated Equilibrium

An alternative approach to information design was provided in [4, 5]. Consider once again a setting with finitely many states, receivers, and actions. Suppose that the designer commits to a *decision rule* σ, which maps each possible state $t \in \mathcal{T}$ to a distribution of action profiles (i.e., a member of $\Delta(A)$, where $A = \times_{i \in V} A_i$). Given the realization of the state, the designer makes a recommendation to each agent consistent with the decision rule. The decision rule is required to be *obedient*, i.e., given her action recommendation a_i, agent i finds it optimal to take action i. Thus, given an obedient decision rule of the designer, after only observing the designer's action recommendation for her, each receiver follows this recommendation. A decision rule satisfying obedience is referred to as a *Bayes correlated equilibrium (BCE).* Using a revelation principle argument (and similar to straightforwardness above), it can be shown that there exists a mechanism that gives rise to a decision rule in a Bayesian Nash equilibrium if and only if the decision rule is obedient (and hence is a BCE).

The set of all possible obedient decision rules admits a characterization in terms of a set of linear inequalities. To see this, first recall the concept of correlated equilibrium from game theory. Consider a setting where there is no state that impacts agents' payoffs. A distribution $\sigma \in \Delta(A)$ is a correlated equilibrium if for each i and $a_i \in A_i$, we have

$$\sum_{a_{-i} \in A_{-i}} u_i(a_i, a_{-i})\sigma(a_i, a_{-i}) \geq \sum_{a_{-i} \in A_{-i}} u_i(a'_i, a_{-i})\sigma(a_i, a_{-i}), \quad \forall a'_i \in A_i.$$

Here we suppress the dependence of the payoffs on t, since there is no state. Intuitively, if the designer randomly draws an action profile $a = (a_i, a_{-i})$ (from distribution σ), and shares with each agent the corresponding component of a, then agent i maximizes her payoff by taking action a_i. Note that in this setting, the uncertainty for agent i stems from the random draw of the action profile (and in particular she does not observe the actions a_{-i} prior to choosing her action). Suppose next that a state impacts agents' payoffs, and let $f(t)$ denote the probability that the state is t (i.e., f is the probability mass function associated with the cumulative distribution function F). Let σ be the decision rule and $\sigma(a_i, a_{-i}|t)$ denote the probability that the draw of the strategy profile is (a_i, a_{-i}) when the state is t. The Bayes correlated equilibria (or the set of obedient decision rules) are characterized in a similar fashion in terms of the following inequalities:

$$\begin{aligned} \sum_{a_{-i} \in A_{-i}, t \in \mathcal{T}} & u_i(a_i, a_{-i}, t)\sigma(a_i, a_{-i}|t)f(t) \geq \\ & \sum_{a_{-i} \in A_{-i}, t \in \mathcal{T}} u_i(a'_i, a_{-i}, t)\sigma(a_i, a_{-i}|t)f(t), \quad \forall a'_i \in A_i. \end{aligned} \tag{1}$$

Also observe that the designer's expected payoff under this decision rule is given by:

$$\sum_{a_i \in A_i, a_{-i} \in A_{-i}, t \in \mathcal{T}} \nu(a_i, a_{-i}, t)\sigma(a_i, a_{-i}, t)f(t). \tag{2}$$

Thus, it follows that an optimal information structure can be obtained by searching over decision rules σ that satisfy (1) and maximize (2).

A few points about this approach are useful to highlight. First, it establishes that to obtain an optimal information structure, it suffices to solve a linear program (with decision variables $\{\sigma(a|t)\}_{a\in A,t\in\mathcal{T}}$). For instance, the optimal information structure for the motivating example from the introduction could be easily obtained by solving such a linear program. On the other hand, the number of decision variables is exponential in the number of agents and this could limit the applicability of the approach. Moreover, if the state belongs to a continuum, then this will be an infinite-dimensional linear program. Finally, imposing side constraints is not straightforward. For instance, suppose that the designer is restricted to using public signals. Then, each agent knows the information the others have as well. Thus, (1) no longer characterizes the set of obedient decision rules and a different characterization is needed. That said, in addition to its conceptual simplicity, the formulation in this subsection is also a computationally useful tool when there is a small number of agents, the state belongs to a finite set, and the designer can send different signals to different agents.

2.3. Rothschild–Stiglitz Approach

In an important class of information design problems, the receivers' payoffs depend on the posterior mean induced by the designer's (public) signal (as opposed to the entire distribution) and the designer's payoff depends only on the actions chosen by the receivers. In this case, the designer's payoff for a signal realization can be expressed as a function of only the corresponding posterior mean. Moreover, the designer's expected payoff can be explicitly characterized in terms of the distribution of the posterior means induced by her mechanism. [24] focuses on such settings, and assumes that the state belongs to $[0,1]$ and there is a single receiver who chooses her action from a finite set.

Given a mechanism π, let G_π denote the distribution of the posterior means induced by this mechanism. Associate a function $c_\pi(x) = \int_0^x G_\pi(t)dt$ with each such mechanism. There are two extreme cases to consider: $\pi = \underline{\pi}$ is the uninformative mechanism and $\pi = \bar{\pi}$ is the mechanism that completely reveals the state. [24] argues that the function c_π associated with any mechanism π is convex and satisfies

$$c_{\bar{\pi}}(x) \geq c_\pi(x) \geq c_{\underline{\pi}}(x), \quad x \in [0,1]. \tag{3}$$

Conversely, for any convex function c_π satisfying (3), there is a mechanism π that induces it. However, these observations allow for an alternative formulation of the designer's problem. The designer can optimize over the convex functions c_π that satisfy (3) and then she can construct a mechanism that supports this optimal solution.

When the action space is relatively simple (e.g., binary or small number of actions), this approach yields an elegant geometric way of characterizing the optimal information structures. When the action space is more complicated, it becomes less clear how to find an optimal mechanism by searching over convex functions that satisfy (3).

2.4. Ideas from Infinite-dimensional Optimization

[21] also focuses on settings where the designer's payoff can be expressed as a function of the posterior mean her signals induce. The authors do not explicitly specify a receiver (hence they do not make assumptions such as a finite action set for the receiver). Instead they focus on an abstract setting where the designer's payoff is $\tilde{\nu}(x)$ when the induced posterior mean is x, for some function $\tilde{\nu}$. Given such a payoff function, the authors formulate the designer's problem as

$$\max_G \int_0^1 \tilde{\nu}(x)dG(x) \tag{4}$$

subject to the constraint that the prior F is a mean-preserving spread of G.

The authors provide an interpretation of this problem as characterizing the Walrasian equilibria of a "persuasion economy". Note that without any further restriction, the problem above can be viewed as an infinite-dimensional optimization problem. The optimal solution admits a dual characterization, which is used to obtain the welfare theorems for the aforementioned economy. The equilibrium conditions (or dual characterizations) can be used to verify if a feasible solution in the designer's problem is indeed optimal. This is especially helpful in special cases (e.g., settings with binary actions), where it is easy to guess the structure of the optimal mechanism. However, it is less clear how to obtain the optimal mechanism in a tractable way in general environments.

2.5. A Reduced-form Approach

In this section, we assume that the designer's payoff is a *step function* of the induced posterior mean, and present an approach due to [11] and [14], which enables a tractable characterization of optimal mechanisms through the solutions of finite-dimensional convex programs. In many practical settings (including some of the examples discussed in Section 3) the designer's payoff has the aforementioned structure, making the approach broadly applicable. This approach also reveals important structural properties of the optimal mechanisms. In particular, it establishes that the signal realizations in an optimal mechanism correspond to the steps of the underlying step function. Moreover, the optimal mechanism partitions the set of states and associates a partition element with each signal realization (or step of the payoff function). Each partition element is a union of at most two subintervals of the set of states. Once the state is realized the mechanism simply reveals to which partition element the state belongs.

To see why a step function structure is important, consider a natural setting where the receivers' payoffs are affine in the state, and the sets of actions $\{A_i\}$ are finite. Assume that the state is absolutely continuous. Suppose that the designer is restricted to using public signaling mechanisms and the designer's payoff depends only on the profile of actions $a \in A = \times_{i\in V} A_i$ chosen by the receivers. Hence, we suppress the dependence on the state and denote the designer's payoff by $\nu(a)$. It is not difficult to show that in such problems, when the posterior mean of the state induced by the designer's mechanism belongs to certain subintervals of $\mathcal{T}$, the receivers always take the same action (see Section 3.1 for an illustration of this point in network games). The following result is adapted from [14]:

Lemma 1. *Suppose that A is finite and $u_i(a_i, a_{-i}, t)$ is affine in t for all $i \in V$, $a \in A$. Suppose further that for $t \in \mathcal{T}$, if $a, a' \in A$ are pure Nash equilibria of the normal form game with payoffs $\{u_i(\cdot, t)\}$, then there exists $t' \in \mathcal{T}$ such that $t' \neq t$, and for t'' that is in between t and t' (i.e., t'' such that $t \le t'' \le t'$ or $t' \le t'' \le t$) a and a' continue to be Nash equilibria of the normal form game with payoffs $\{u_i(\cdot, t'')\}$.*

Consider a public signaling mechanism with signal S. There exist cutoffs $\{b_k\}_{k=0}^K$ and action profiles $\{a^k \in A\}_{k=1}^K$ such that at the induced sender-preferred equilibrium:

(i) $b_{k-1} < b_k$ *for* $k > 0$, $b_0 = \inf \mathcal{T}$, $b_K = \sup \mathcal{T}$,
(ii) $a^{k-1} \neq a^k$ *for* $k > 0$,
(iii) the receivers play action profile a^k *when* $\mathbb{E}[T|S=s] \in (b_{k-1}, b_k)$,
(iv) the receivers play action profile $a' \in \arg\max_{a\in\{a^k, a^{k+1}\}} \nu(a)$ *when* $\mathbb{E}[T|S=s] = b_k$ *for* $0 < k < K$.

Note that the assumption on pure Nash equilibria $a, a' \in A$ is made for expositional simplicity, and the result can be extended by explicitly handling the cases where this assumption fails. If this assumption is relaxed, at the sender-preferred equilbiria the receivers can play an action profile for a single posterior mean level (as opposed to an interval of posterior mean levels). This assumption simply rules out such cases.

[14] established this result under the assumption that there is a single receiver, but the approach immediately generalizes to settings where there are multiple receivers. Because the payoffs are affine in the state the receivers' actions only depend on the induced posterior mean. As the posterior mean increases the induced equilibrium outcome can change. In the lemma, $\{b_k\}$ capture precisely the posterior mean levels at which the equilibrium outcome changes, and a^k captures the equilibrium outcome when the posterior mean is strictly between b_{k-1} and b_k. Note that when the posterior mean is b_k or b_{k-1} there are multiple equilibrium outcomes that can emerge in the game among the receivers, and we break ties in favor of the sender-preferred one. We denote by $\mathcal{B}_k$ the interval where the equilibrium outcome a^k emerges, and note that it can be left- or right-closed and satisfies $(b_{k-1}, b_k) \subset \mathcal{B}_k \subset [b_{k-1}, b_k]$. This discussion implies that the designer's payoff is a step function $\tilde{\nu}$ of the induced posterior mean where

$$\tilde{\nu}(x) = r_k = \nu(a^k) \text{ for } x \in \mathcal{B}_k, \tag{5}$$

and $r_k \in \mathbb{R}$ denotes the reward associated with the kth step.

In the remainder of this section, we focus on such payoffs. Specifically, we assume that there are parameters $\{r_k\}_{k=1}^K$ and subsets of states $\{\mathcal{B}_k\}_{k=1}^K$ such that when the posterior mean is in $\mathcal{B}_k$, the designer's payoff is r_k. Furthermore, there is an increasing sequence of cutoffs $\{b_k\}_{k=0}^K$ with $b_0 = \inf \mathcal{T}$, $b_K = \sup \mathcal{T}$ such that $(b_{k-1}, b_k) \subset \mathcal{B}_k \subset [b_{k-1}, b_k]$ for all $k \in [K] := \{1, \ldots, K\}$.

How can the step function structure be exploited to tractably characterize the optimal mechanisms? [14] refers to a (public signaling) mechanism π as a level mechanism if (i) the set of signal realizations is $[K]$, and (ii) signal realization k induces a posterior mean in $\mathcal{B}_k$ and a reward of r_k. It is straightforward to establish that in the designer's problem it is without loss of optimality to restrict attention to level mechanisms. Intuitively this holds because given a (public) mechanism, a new mechanism can be defined such that when the signal realization in the original mechanism generates a posterior mean in $\mathcal{B}_k$ the signal realization of the new mechanism is k. It can be readily checked that the latter mechanism is a level mechanism and is also payoff-equivalent to the initial mechanism.

Denote the set of level mechanisms by Π_L. [14] formulates the designer's problem over the level mechanisms as follows:

$$\begin{aligned} \max_{\pi \in \Pi_L} \quad & \sum_{k \in [K]} r_k \mathbb{P}(S = k) \\ s.t. \quad & \mathbb{E}[T | S = k] \in \mathcal{B}_k \quad \text{for } k \in [K] \text{ such that } \mathbb{P}(S = k) > 0, \end{aligned} \tag{6}$$

where S denotes the signal of π.

Observe that a level mechanism π induces a distribution of posterior means: $\{\mathbb{E}[T|S = k], \mathbb{P}(S = k)\}$. The discussion above implies that the designer's problem can equivalently be formulated as an optimization problem over posterior mean distributions. However, not all posterior mean distributions are relevant, and it suffices to optimize over (atomic) posterior mean distributions with at most one atom in each $\mathcal{B}_k$. We proceed by providing an equivalent characterization of the relevant posterior mean distributions which yields a natural convex optimization formulation for the designer's problem.

Given a level mechanism with signal S, define a tuple $\{p_k, z_k\}_{k \in [K]}$ such that:

(C1) $p_k = \mathbb{P}(S = k)$
(C2) $z_k = \mathbb{E}[T \cdot \mathbf{1}\{S = k\}]$

for all $k \in [K]$. Observe that (C2) implies that $z_k / p_k = \mathbb{E}[T|S = k]$ for $p_k > 0$. Thus, by the definition of level mechanisms we have:

(C3) $z_k / p_k \in \mathcal{B}_k$ for k such that $p_k > 0$.

Following [14], we let $\mathcal{D}$ denote the set of $\{p_k, z_k\}_{k\in[K]}$ tuples that are *consistent* with a level mechanism; i.e., for $\{p_k, z_k\}_{k\in[K]} \in \mathcal{D}$, there exists a level mechanism (with signal S) such that (C1)–(C3) hold.

[14] establishes a useful representation of the tuples that are consistent with a level mechanism:

Theorem 1 ([14]). *$\{p_k, z_k\} \in \mathcal{D}$ if and only if*

(C1') $p_k \geq 0$ for $k \in [K]$ and $\sum_{k\in[K]} p_k = 1$,
(C2') $z_k/p_k \in \mathcal{B}_k$ for $k \in [K]$ such that $p_k > 0$, and $z_k = 0$ for $k \in [K]$ such that $p_k = 0$,
(C3') For $\ell \in [K]$, we have

$$\sum_{k\geq \ell} z_k \leq \int_{1-\sum_{k\geq \ell} p_k}^{1} F^{-1}(x)dx, \tag{7}$$

where the inequality holds with equality for $\ell = 1$.

This representation readily gives rise to a two-step *reduced form approach* for characterizing optimal mechanisms. In this approach, as opposed to optimizing directly over the mechanisms (e.g., deciding on which signal to send at each state), the designer optimizes over $\{p_k, z_k\}$ tuples that satisfy the conditions (C1')-(C3'). Then, she uses the optimal tuple to construct a consistent mechanism.

Specifically, in the first step the designer solves the following optimization problem.

$$\begin{aligned}
\max_{\{p_k, z_k\}_{k\in[K]}} \quad & \sum_{k=1}^{K} p_k r_k \\
s.t. \quad & \sum_{k\geq \ell} z_k \leq \int_{1-\sum_{k\geq \ell} p_k}^{1} F^{-1}(x)dx \quad \text{for } \ell \in [K]\setminus\{1\}, \\
& \sum_{k} z_k = \int_0^1 F^{-1}(x)dx, \\
& p_k b_{k-1} \leq z_k \leq p_k b_k \quad \text{for } k \in [K], \\
& \sum_k p_k = 1, \\
& p_k \geq 0 \quad \text{for } k \in [K].
\end{aligned} \tag{OPT}$$

The constraints of this optimization problem correspond to (C1')-(C3'). The only difference is that for k such that $p_k > 0$ the constraint $z_k/p_k \in \mathcal{B}_k$ is replaced with $p_k b_{k-1} \leq z_k \leq p_k b_k$. Recalling that $(b_{k-1}, b_k) \subset \mathcal{B}_k \subset [b_{k-1}, b_k]$, it can be seen that the latter constraint is effectively a relaxation. This relaxation turns out to be immaterial, and optimal solutions always satisfy (C2'). Given a tuple $\{p_k, z_k\}$ consistent with a level mechanism, the objective of this problem is the expected payoff of the designer under this mechanism.

Observe that (OPT) is a convex optimization problem. It can be readily seen that the objective as well as the constraints other than the first one are linear. The first one is convex in the decision variables, since the c.d.f. (as well as its inverse) are non-decreasing functions.

Given an optimal solution to this problem, the next step is to construct a mechanism that is consistent with the optimal solution. To this end [14] introduces the notion of a laminar interval partition of the set of states and establishes that if a $\{p_k, z_k\}$ tuple satisfies conditions similar to (C1')–(C3'), then it is possible to obtain a laminar interval partition of states such that p_k corresponds to the probability with which the state belongs to the kth partition element, and z_k/p_k yields the corresponding posterior mean. The formal definition of laminar interval partitions, and a partition lemma which plays a key role in the construction of an optimal mechanism, are presented next.

Definition 1. A collection of sets $\{I_k\}_{k\in\mathcal{A}}$ constitutes a *laminar family* if for any $k,\ell\in\mathcal{A}$ either I_k and I_ℓ do not intersect (i.e., $I_k\cap I_\ell=\emptyset$) or one contains the other (i.e., $I_k\subset I_\ell$ or $I_\ell\subset I_k$). If, in addition, each $I_k\subset\mathbb{R}$ is an interval, then we refer to $\{I_k\}_{k\in\mathcal{A}}$ as a laminar interval family. A partition $\cup_{k\in\mathcal{A}}\mathcal{T}_k=\mathcal{T}'$ is referred to as a laminar interval partition of $\mathcal{T}'$ if $\mathcal{T}_k=I_k\setminus\cup_{\ell\in\mathcal{A}|\ell>k}I_\ell$ for all $k\in\mathcal{A}$ and some laminar interval family $\{I_k\}_{k\in\mathcal{A}}$.

Lemma 2 ([14], Partition Lemma). *Fix a finite collection $\{p_k,z_k\}_{k\in\mathcal{A}}$, where $\mathcal{A}\subset\mathbb{N}_{++}$ and real numbers $q_0,q_1\in[0,1]$ satisfying $q_0<q_1$. Suppose that (i) $p_k>0$ for $k\in\mathcal{A}$, and $\sum_{k\in\mathcal{A}}p_k=q_1-q_0$, (ii) z_k/p_k is strictly increasing in $k\in\mathcal{A}$, and (iii) for all $\ell\in\mathcal{A}$ we have*

$$\sum_{k\in\mathcal{A}|k\ge\ell} z_k \le \int_{q_1-\sum_{k\in\mathcal{A}|k\ge\ell}p_k}^{q_1} F^{-1}(x)dx, \tag{8}$$

where the inequality holds with equality only for $\ell=\min\mathcal{A}$.

There exists a laminar interval family $\{I_k\}_{k\in\mathcal{A}}$ and sets $\mathcal{T}_k=I_k\setminus\cup_{\ell\in\mathcal{A}|\ell>k}I_\ell$ for all $k\in\mathcal{A}$ such that

(a) $I_{(\min\mathcal{A})}=\cup_{k\in\mathcal{A}}I_k$, and if $|\mathcal{A}|>1$, then the end points of the interval $F(I_{(\max\mathcal{A})})$ are strictly in between those of $F(I_{(\min\mathcal{A})})$. Moreover, $\cup_{k\in\mathcal{A}}\mathcal{T}_k=[F^{-1}(q_0),F^{-1}(q_1)]$.
(b) $\mathbb{P}(T\in\mathcal{T}_k)=p_k$ for all $k\in\mathcal{A}$.
(c) $\mathbb{E}[T|T\in\mathcal{T}_k]=z_k/p_k$ for all $k\in\mathcal{A}$.

In [14], the partition lemma is established by following an inductive approach. When set $\mathcal{A}$ has cardinality two, the result follows from the intermediate value theorem. In this case, by solving a single-parameter equation it is possible to obtain an interval for each element of $\mathcal{A}$ such that the interval associated with the largest element of this set is contained in the interval associated with the smallest one, and taking their differences produces a partition of the set of states $[F^{-1}(q_0),F^{-1}(q_1)]$ that is consistent with the $\{p_k,z_k\}_{k\in\mathcal{A}}$ tuple. When $\mathcal{A}$ has larger cardinality, we group the elements of $\mathcal{A}$ other than $\max\mathcal{A}$ to construct another problem instance where the aforementioned set has cardinality two. Then, we use the partition lemma for this problem instance to obtain an interval for $\max\mathcal{A}$ and another for the elements of $\mathcal{A}\setminus\{\max\mathcal{A}\}$. This reduces the problem to obtaining a consistent partition of a subset of states to the elements of $\mathcal{A}\setminus\{\max\mathcal{A}\}$. Proceeding recursively, a partition to the elements of $\mathcal{A}$ is obtained. This approach also yields a simple recursive algorithm (which relies on repeatedly using the result for the case where $\mathcal{A}$ has cardinality two) for constructing a partition that is consistent with a given tuple $\{p_k,z_k\}$. We refer the reader to [14] for the details of the algorithm.

Now consider an optimal solution $\{p_k^\star,z_k^\star\}$ to (OPT). Let $\{\ell_1,\dots,\ell_m\}$ denote the set of ℓ for which the first constraint of (OPT) is binding and $p_\ell^\star>0$. Label the elements of this set such that $\ell_1<\ell_2<\cdots<\ell_m$. For $i\in[m]$ define $\bar{L}_i=\{k\in\{\ell_i,\dots,\ell_{i+1}-1\}|p_k>0\}$ and $q_i=1-\sum_{k\ge\ell_{i+1}}p_k^\star$, where $\ell_{m+1}=K+1$ by convention. The fact that (7) holds with equality for $\ell\in\{\ell_1,\dots,\ell_m\}$ implies that for $\ell'\in\bar{L}_i$ and $i\in[m]$, we have

$$\sum_{k\in\bar{L}_i|k\ge\ell'} z_k^\star \le \int_{q_i-\sum_{k\in\bar{L}_i|k\ge\ell'}p_k^\star}^{q_i} F^{-1}(x)dx, \tag{9}$$

where the inequality holds with equality only for $\ell'=\min\bar{L}_i$.

This observation provides a way of constructing a level mechanism consistent with $\{p_k^\star,z_k^\star\}$. First, we identify $\ell\in[K]$ for which the first constraint of (OPT) is binding and construct the sets $\{\bar{L}_i\}_{i\in[m]}$. Then, by using the algorithm outlined above for each set $\bar{L}_i$, a partition of $[F^{-1}(q_{i-1}),F^{-1}(q_i)]$ which contains a set for each element of $\bar{L}_i$ is obtained. Collectively these partitions yield a partition $\{\mathcal{T}_k\}_{k\in[K]}$ of $\mathcal{T}$ to the elements of $[K]$. A mechanism which sends signal k when the state belongs to $\mathcal{T}_k$ is consistent with $\{p_k^\star,z_k^\star\}$, and hence is optimal.

Summarizing, by following the approach outlined in this section the designer's problem boils down to solving a simple convex optimization problem (OPT). Studying which constraints are binding at the optimal solution and using a simple recursive algorithm, a mechanism consistent with the optimal solution is obtained. This approach has three additional benefits. First, given the convex optimization formulation it is straightforward to do comparative statics by leveraging tools from sensitivity analysis in convex optimization. Second, the approach is flexible, and allows for incorporating side constraints. We illustrate this in Section 4.2 where we allow a receiver to have private information, which necessitates the mechanism to satisfy additional incentive compatibility constraints. Third, by studying optimality conditions in (OPT), it is possible to shed light on the structure of an optimal mechanism. For instance, in [14], it is established that there exist optimal solutions to this problem such that each set $\bar{L}_i$ has cardinality at most two. In this case, the final partition is obtained by associating an interval with each of the two elements of $\bar{L}_i$ and taking their set differences. Consequently, it follows that each element of the partition $\{\mathcal{T}_k\}$ is a union of at most two intervals. [11, 14] refer to this as the double-interval structure and establish its optimality in a wide range of information design problems (see Section 3.1 for an illustration of this structure). When the designer's problem has side constraints (e.g., incentive compatibility constraints), $\bar{L}_i$ has larger cardinality (which can again be characterized by studying the optimality conditions in (OPT)). In these cases, richer laminar interval partitions are required to satisfy the side constraints, and more complicated mechanisms that are based on this partition structure achieve optimality. We revisit this point in Section 4.2.

3. Applications

We next discuss a few recent papers that study various operational problems through the lens of information design. In these applications, we consider a designer who tries to influence the actions of a group of agents. In some of these settings (e.g., when private signals are employed), the information received by an agent could spill over to others, thereby complicating the designer's problem. We ignore such spillovers until Section 4.1, where we discuss the impact of informational spillovers.

3.1. Persuasion in Networks

The literature on social and economic networks sheds light on how the structure of an underlying network impacts the outcome of the strategic interactions among the agents in the network. Building on the models from this literature, a number of papers explore how a firm can make use of the available social network data to improve its pricing/advertising decisions. For instance, suppose that the firm offers a product that exhibits local network externalities, i.e., an agent enjoys the product more if her peers also use the same product. Knowing this a seller can target some agents with appropriate discounts to harvest the network effects and improve her revenues (see, e.g., [8, 15]). At a high level, these papers use pricing as a lever to induce a desired outcome at the end of a network game. Information is another important (and arguably more natural) lever to induce a desired outcome in social networks. We next discuss two related papers that explore different applications in operations.

In [16], the authors focus on online social networks, and the agents' decisions to engage with the available content (e.g., news articles). The content may contain inaccuracies, and the social networking platform can send informative signals about the content quality to influence agents' engagement decisions. Moreover, agents' engagement decisions exhibit local strategic complementarities. That is, if her friends decide to engage with the available content, agent i receives a larger payoff from engaging with the content. Specifically, the paper focuses on the following payoff structure:

$$u_i(a_i, a_{-i}, t) = a_i(v - bt + \sum_j g_{ij} a_j),$$

where $a_i \in \{0,1\}$ captures the engagement decision of agent i. The state realization t measures how inaccurate the content is: the larger t represents larger content inaccuracy. The state is assumed to be uniformly distributed in $[0,\alpha]$.

In this model, agent i's payoff is normalized to zero if she decides not to engage with the content. Otherwise, her payoff consists of three terms: (i) a constant payoff term v corresponding to the satisfaction that she (unilaterally) derives from engaging with the content, (ii) a negative term that captures the disutility that the agent incurs due to engaging with content that has inaccuracies, and (iii) a network externality term that captures the additional utility that she derives from engaging with content with which her friends also engage. The second term depends on the b parameter, which captures the importance of the state on the agents' payoff. The last term depends on $\{g_{ij}\}$ parameters which correspond to the entries of the adjacency matrix. We assume that the underlying network is unweighted and $g_{ij} \in \{0,1\}$.

The platform commits to a mechanism (e.g., chooses an algorithm that evaluates the available content) before the content is realized. Once the content is realized this mechanism reveals informative signals about the quality of the content. In this setting, the platform has two natural goals. The first one is to maximize engagement, i.e., the number of agents who take action 1. The second one is to minimize misinformation, i.e., $t\sum_i a_i$. If t is interpreted as the probability that the content is fake or contains errors, then the latter quantity focuses on minimizing the total engagement with such content.

In the baseline model, [16] allows for private signaling mechanisms where different agents receive different signals. Leveraging the straightforwardness idea discussed earlier, the authors establish that for both objectives mentioned above (as well as their convex combinations) optimal mechanisms always take a simple recommendation structure: if the state realization is below a threshold t_i, the platform recommends agent i to engage with the content and otherwise she recommends the agent not to engage. Moreover, these thresholds are chosen so that the agents always find it optimal to follow their recommendation. Note that this imposes a nontrivial restriction on the thresholds. Suppose that agent i receives a recommendation to engage. She can infer that the state is in $[0,t_i]$ and conditional on agent i's signal her connection j receives a recommendation to engage (and follows it) with probability $\min\{t_i,t_j\}/t_i$. Thus, the agent i's expected payoff from engaging is given by: $v - \frac{b}{2}t_i + \frac{1}{t_i}\sum_j g_{ij}\min\{t_i,t_j\}$. This quantity should be nonnegative for agent i to engage, which after rearranging terms can be stated as

$$t_i^2 \le \frac{2}{b}vt_i + \frac{2}{b}\sum_{j\in V} g_{ij}\min\{t_i,t_j\}.$$

A similar constraint can also be derived to ensure that the thresholds are such that agent i finds it optimal not to engage when the recommendation is not to engage. Imposing these constraints, we obtain a a simple convex optimization problem whose solution yields the optimal thresholds:

$$\begin{aligned}
\max_{\{t_i\}} \quad & \frac{1}{\alpha}\sum_{i\in V} t_i \\
s.t. \quad & t_i^2 \le \frac{2}{b}vt_i + \frac{2}{b}\sum_{j\in V} g_{ij}\min\{t_i,t_j\} \quad \text{for } i\in V, && (10)\\
& \alpha^2 - t_i^2 \ge \frac{2}{b}v(\alpha - t_i) + \frac{2}{b}\sum_{j\in V} g_{ij}\max\{t_i,t_j\} - \frac{2}{b}d_i t_i \quad \text{for } i\in V, && (11)\\
& 0 \le t_i \le \alpha \quad \text{for } i\in V, && (12)
\end{aligned}$$

where d_i is the number of connections agent i has.

How are the optimal thresholds related to agents' network positions? Under mild conditions on the primitives, the first constraint in this problem is binding at the optimal solution. This allows for an equivalent fixed-point characterization of the optimal thresholds:

$$t_i = \frac{2v}{b} + \frac{2}{b}\sum_{j\in V} g_{ij}\min\left\{1, \frac{t_j}{t_i}\right\}. \tag{13}$$

This fixed-point equation closely mimics the fixed-point equation that characterizes agents' Bonacich centralities $\{\kappa_i\}$ in the network (see, e.g., [3]). The latter equation is given by:

$$\kappa_i = 1 + \gamma \sum_{j\in V} g_{ij}\kappa_j, \tag{14}$$

where γ is a fixed constant. Intuitively, this fixed-point equation implies that an agent is more central, if she is connected to more central agents in the network. This feature is also present in (13). However, in (13), the "contribution" of j to its neighbor i's centrality is capped at one and scaled down by t_i. Thus, the contribution of an agent to the centrality of a very central agent is smaller than to that of a less central one. Motivated by these observations, [16] uses $\{t_i\}$ that solve (13) to define agents' *engagement centrality* and concludes that an engagement maximizing platform chooses larger thresholds for more central agents. This in turn, implies that on average these agents engage with content that is less accurate.

The structure of the optimal mechanism is illustrated in Figure 1 (adapted from [16]). As can be seen from this figure, the more central agents receive larger thresholds than the less central ones.

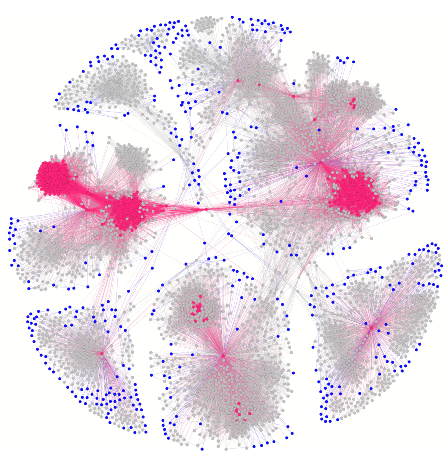

FIGURE 1. A Facebook subnetwork with 4,039 nodes. $b=200$, $v=20$, $\alpha=1$. Agents whose thresholds are in the top/bottom %10 are highlighted in red/blue respectively.

What happens, if the platform tries to minimize misinformation? In this case, the structure of the optimal thresholds is completely different. In particular, it turns out that the platform chooses identical thresholds for all agents. Moreover, these thresholds are independent of the network structure. To understand this intuitively, suppose that the platform chooses thresholds by ignoring the network effects – which yields identical thresholds for all agents. Suppose that agents other than i follow the recommendation of the induced mechanism and they do not engage when the state is above the aforementioned threshold. Thus, from agent i's point of view the network effects do not play a role when the state is larger than the threshold. Hence, she also finds it optimal to follow the recommendation of the platform. In other words, the outcome induced by these thresholds in the initial network is identical to the one that would emerge in the absence of network effects. On the other hand, when network effects are absent the agents engage with content less, and this results in lower misinformation. Hence, the misinformation obtained in such settings is a lower bound on what can be achieved in the initial network. This implies that the aforementioned mechanism, obtained by ignoring network effects, is in fact optimal for the initial network. Moreover, similar results extend to the settings where the platform maximizes a weighted combination of the engagement and misinformation objectives. In particular, when the weight of misinformation is large, the network effects can be ignored and in the other extreme the centralities of agents play a key role in the design of optimal thresholds.

Our discussion so far focused on mechanisms where the designer can send different signals to different agents. What if she is restricted to using a public signaling mechanism, which shares the same signal with all the agents? This question is addressed in [11, 12], which leads to fundamentally different mechanisms.

Here, the agents' payoffs have structure similar to that specified before. In particular, the payoff of agent i is given by:

$$u_i(a_i, a_{-i}, t) = a_i(t + \sum_j g_{ij} a_j).$$

The problem can be cast, as before, in the context of engaging with content that involves inaccuracies. Alternatively, we can think about a setting where agents decide whether to adopt a new product that exhibits local network externalities. The quality of the product (the state) is a priori unknown to the agents and the designer. The designer commits to a review system which reveals informative signals about the quality of the product once it is realized. The objective of the platform is to maximize the expected number of adopters. What are the optimal public signaling mechanisms?

Fix a mechanism π, and let S denote its public signal. Suppose that the designer's signal induces a posterior mean of $\mathbb{E}[T|S=s]$ and let k denote the smallest integer weakly larger than $-\mathbb{E}[T|S=s]$. Since all agents have the same posterior mean, it can be seen from the payoff structure that if there is a set of agents who have k or more connections within the set, they can guarantee nonzero payoff for all agents in the set by taking action 1. Furthermore, the maximal such set is the set of agents who take action 1 in the sender-preferred equilibrium. The aforementioned maximal set is actually an object that is familiar in graph theory, and it is often referred to as the *k-core* of the network. Let r_k denote the cardinality of the k-core.

These observations imply that if the posterior mean is in $[-k, -k+1)$ the agents in the k-core take action 1, and this yields a payoff of r_k to the designer. In other words, the designer's payoff is a step function of the induced posterior mean. Thus, using the framework in Section 2.5, an optimal mechanism can be obtained. In this mechanism, the designer partitions the states and associates a partition element with each step of the step function, or equivalently, with each (distinct) core of the graph. When the state belongs to the partition element associated with the k-core, the designer reveals this information and the agents in the k-core find it optimal to take action 1.

Consider the network from Figure 1. Suppose that the state is distributed uniformly in $[-50, 0]$. The optimal public mechanism exhibits a double-interval structure as discussed in Section 2.5, and is illustrated in Figure 2. In this figure, we assign a different color to each possible signal realization (which corresponds to a different core of the underlying network), and highlight the associated intervals with this color. The figure indicates that signal realizations that correspond to $k \in \{12, 33\}$ cores have two associated intervals. We also highlight the nodes of the network by assigning to each node the color of a signal realization. If a node has the color of the signal realization that induces the k-core to take action 1, then the relevant agent takes action 1 for any signal realization that triggers the

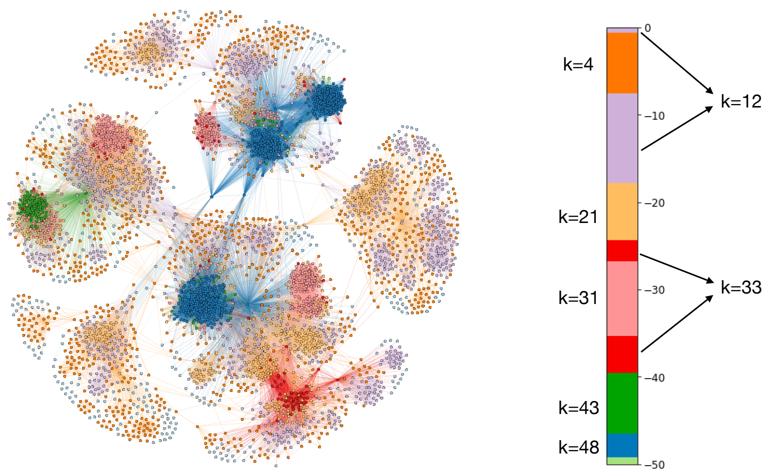

FIGURE 2. The legend on the right illustrates the partition of $\mathcal{T} = [-50, 0]$ into different signal realizations/cores. The unlabeled interval at the bottom corresponds to the k-core for $k = 49$.

k'-core to take action 1 for $k' \leq k$. For instance, red is the color of the signal realization that corresponds to the $k = 33$ core, and the nodes of the network that are colored in red take action 1 whenever the state belongs to the intervals associated with the k'-core for $k' \leq 33$. Observe that the set of states for which a node takes action 1 need not be convex (e.g., consider the nodes that belong to the $k = 31$ core but not to the $k = 33$ core). This counter-intuitive behavior is a byproduct of the public signaling restriction and was not present when we considered the optimal signaling mechanisms where this requirement was not imposed (the latter mechanisms exhibited a threshold structure).

The discussion so far assumes that the designer knows the network structure. Suppose that the designer does not know the exact network structure, but has access to limited information about it such as the degree distribution. Is it still possible to design a public signaling mechanism that improves the payoff of the designer? The answer, perhaps surprisingly, is that it is possible to obtain *asymptotically optimal mechanisms*.

To establish this result, [12] considers a natural random graph model, which we summarize next. We say that a sequence of nonnegative integers $\{d_i\}_{i=1}^n$ is *graphical*, if there exists a simple network whose degree sequence matches this sequence.[2,3] Consider the set of all networks with n nodes and a given graphical degree sequence $\{d_i(n)\}_{i=1}^n$. In all networks that belong to this set, node $i \in [n] = \{1, \ldots, n\}$ has degree $d_i(n)$, but the connection structure among nodes may be different for different elements of this set. We focus on the uniform distribution on this set of networks, and use G_n to denote the random network obtained after a draw from this distribution. We are interested in providing results for large networks, and focus on settings where $n \to \infty$. As we conduct our asymptotic analysis, we require the degree sequence $\{d_i(n)\}_{i=1}^n$ to be consistent with an underlying degree distribution $\{\rho_l\}_{l=0}^\infty$, i.e., $|\{i | d_i(n) = l\}|/n \to \rho_l$ for every $l \in \mathbb{Z}_+$ as $n \to \infty$. Under this assumption, and additional mild regularity conditions on the degree sequence (many of which are always satisfied when the degrees are bounded, see [12]), a result of [28] implies that there exists $\{\bar{r}_k\}$ such that the fraction of nodes that belong to the k-core converges to $\bar{r}_k$ in probability as n goes

[2] A network is simple if it does not involve any self-loops or multiple edges between any pair of nodes.

[3] It is possible to check whether a given sequence is graphical efficiently, e.g., by using the Erdös-Gallai theorem or the Havel-Hakimi theorem – see [9].

to infinity. Moreover, $\bar{r}_k$ can be obtained by solving an equation whose coefficients depend on the underlying degree distribution (for settings with bounded degrees this is a simple polynomial equation).

For network G_n, let $\pi(n)$ denote a corresponding optimal public mechanism. [12] considers a variant of (OPT) where some constraints are relaxed. In this problem, attention is restricted to k such that $\bar{r}_k$ is nonzero, and for such k the r_k parameter in the objective is replaced with $\bar{r}_k$ (and the constraints are appropriately adjusted). Consider an optimal solution to this problem, and let $\pi^\star$ denote a mechanism consistent with this solution (obtained by following the approach outlined in Section 2.5). Let $\mathcal{A}(\pi, G)$ denote the designer's objective under mechanism π for network G. [12] establishes that

$$\frac{\mathcal{A}(\pi^\star, G_n)}{\mathcal{A}(\pi(n), G_n)} \xrightarrow{p} 1,$$

i.e., the ratio of the performance under the optimal mechanism and the mechanism $\pi^\star$ obtained using only the limiting degree distribution information, converges (in probability) to 1. In other words, it is possible to construct asymptotically optimal mechanisms by using only the limiting degree distribution. Furthermore, even when the number of nodes n is relatively small (e.g., in the order of thousands), [12] illustrates that the gap between the optimal mechanism and $\pi^\star$ is quite small. These observations suggest that even without the precise knowledge of the network structure, the designer can substantially improve her payoff by using appropriate mechanisms.

3.2. Signaling Product Availability

In many revenue management settings, the availability of the product offered by a seller has a first-order impact on the prices as well as the revenues. A strand of the recent literature (see, e.g., [17] and [34]) focuses on understanding how a seller can signal product availability in order to improve her revenues. In this subsection, we discuss some of the relevant contributions in the literature by focusing on [17].

A product is offered by a seller over two periods to a unit mass of buyers. The buyers' values follow a known cumulative distribution $G(\cdot)$ and for a given price p, the total demand of the buyers is given by $\bar{G}(p)$, where $\bar{G}$ denotes the complementary c.d.f. We assume that G is absolutely continuous and has a non-decreasing hazard rate.[4] The seller does not a priori know the available inventory, which is q_H (q_L) with probability $\mathbb{P}(H)$ ($\mathbb{P}(L)$). Inventory becomes available in the beginning of the selling season and cannot be replenished. We assume that $q_H = 1 > q_L$ so that there is no shortage in the "high" state but there is a possibility of shortage in the "low" state. Before the inventory is realized the designer commits to (i) prices p_1, p_2 respectively for periods 1 and 2, and (ii) a mechanism that reveals informative signals to buyers about the available inventory once it is realized. Each buyer's strategy maps the observed signal realization to an action in $\{0, 1, 2\}$, where 0 denotes the decision of not buying and 1 (2) denotes the decision of buying the product in period 1 (2) if it is available. If demand in a given period exceeds the available inventory, then the product is rationed (uniformly) among all buyers who demand it. Denote by $A_\ell(v)$ the event that the agent with value v receives the product if she demands it in period ℓ. The agent's expected payoff from taking action a conditional on signal realization s is given by:

$$\mathbf{1}(a \in \{1, 2\})(v - p_a)\mathbb{P}(A_a(v)|S = s)$$

Suppose that the designer does not reveal any information about the available inventory. Due to inventory uncertainty, the designer can benefit from charging a higher price in the first period and a lower one in the second period so that in the first period the agents with

[4] That is, $g(p)/(1 - G(p))$ is non-decreasing in p, where g is the associated p.d.f.

higher values purchase the product (without exposing themselves to stockout risk). In the second period the agents with lower values purchase the product when there is sufficient inventory. Can the designer further improve the revenues by choosing an appropriate signaling mechanism? If the designer is restricted to using public signals, the answer, surprisingly, turns out to be negative:

Theorem 2 ([17]). *Let $p^m = \arg\max_p \bar{G}(p)p$ and suppose that the designer is restricted to using public signals. Providing no information and setting $p_1 = \mathbb{P}(L)\max\{\bar{G}^{-1}(q_L), p^m\} + \mathbb{P}(H)p^m$ and $p_2 = p^m$ is an optimal mechanism.*

It is worth noting that the optimal mechanism is not unique. In fact, [17] establishes that another optimal mechanism can be obtained by revealing information about the state fully, but using different prices.

The conclusions completely change if the designer is allowed to use private signaling mechanisms. In this case, [17] establishes that the designer can substantially increase her revenue. The optimal policy also admits a simple structure: The designer identifies two prices $p_1 \geq p_2$. When the inventory realization is low, she recommends agents whose values are above p_1 to buy (in the first period) and she recommends the remaining agents not to buy. When the inventory realization is high, she recommends agents whose values are above p_1 to buy (in the first period) with some probability (independently). Note that the probability of each recommendation depends on the type/value of the buyer, and some of these agents may receive a recommendation to purchase the product with probability one whereas others may receive the same recommendation with strictly lower probability. Agents whose values are between p_1 and p_2 receive a recommendation to wait. The remaining agents are recommended not to buy the product (in any of the periods). The aforementioned randomization is chosen carefully, in a way that makes it incentive compatible for the agent to follow the recommendation they receive.

An important qualitative takeaway of this paper is that by discriminating agents appropriately in terms of the information they have access to, the designer can effectively persuade them to trade at different prices (in expectation). This form of price discrimination, in turn, yields a revenue improvement. The extensions of this insight to richer settings, as well as the interplay between personalized information and prices, are active areas of research.

3.3. Optimal Information Structures in Two-sided Markets

Recent work in operations has also explored the applications of information design ideas to two-sided markets. Here, we focus on [29] which investigates how a platform where buyers and sellers trade, can reveal information about the quality of the sellers to influence the trading outcome. The aforementioned work considers different models of trade between the buyers and sellers. For brevity, we restrict attention to one of these models where the platform chooses the prices, and the sellers decide on the quantity to supply to the market.

Assume that there is a unit mass of buyers and sellers. Each seller's quality belongs to an interval $\mathcal{T}_s \subset \mathbb{R}_+$ and is distributed according to a known distribution F_s. Similarly, each buyer's type belongs to an interval $\mathcal{T}_b \subset \mathbb{R}_+$, and is distributed according to F_b. The buyers do not know the sellers' quality levels.

If a type $m \in \mathcal{T}_b$ buyer trades with a type $q \in \mathcal{T}_s$ seller at price p, this yields a payoff of $mq - p$ to the buyer. Each seller can supply goods to multiple buyers and decides on the quantity h to supply to the market. If a type q seller manages to produce and sell h units at price p, then this results in a payoff of

$$U(q,h,p) = hp - k(q)\frac{h^{\alpha+1}}{\alpha+1}.$$

Here, the first term is the seller's revenue and the second one is her production cost. The function k captures how the sellers' production costs depend on their quality levels. It is also assumed that production costs are convex in the supplied quantity, i.e., $\alpha > 0$.

The platform has partial information about the sellers' quality levels. Specifically, it has access to a partition $\{A_k\}$ of $\mathcal{T}_s$ and knows to which of the partition elements each seller belongs. The platform can coarsen this information and share it with the buyers. Specifically, she can choose a disjoint collection of sets $\{B_k\}$ such that each element of this collection is a union of a subset of the elements in $\{A_k\}$. $\{B_k\}$ is referred to as an information structure. Intuitively, as opposed to providing refined information about the sellers, the platform can pool some types of sellers together and provide coarser information on them. Importantly, the union of $\{B_k\}$ need not be $\mathcal{T}_s$. This corresponds to removing sellers $\mathcal{T}_s \setminus \cup_k B_k$ from the platform altogether.

The platform's problem is to find an information structure $\{B_k\}$ and corresponding prices $\{p(B_k)\}$ that maximize the total transaction value – which might be relevant, for instance, when the platform charges for each transaction a commission that is equal to a fraction of the transaction value. The paper focuses on understanding when simple information structures that involve (i) banning a group of low quality sellers from the platform altogether and (ii) not providing any information on the remaining sellers, are optimal. Such information structures are referred to as 1-separating information structures.

Buyers and sellers take $\{B_k\}$ and $\{p(B_k)\}$ as given, and at equilibrium they choose optimal supply/demand levels. Namely, (a) each seller whose quality belongs to B_k decides on the optimal amount to supply to the market (assuming she can sell each supplied unit at price $p(B_k)$), (b) each buyer forms beliefs about the sellers' expected qualities (given the sellers' decisions) and decides on the group of sellers $\{B_k\}$ to transact with (if any) to maximize her expected payoff, and (c) the market clears.

Note that the information design problems we have discussed so far involve an uncertain state of the world, and a designer who commits to sending informative signals about this state to the receivers. The design problem in [29] has a different flavor. There is no (single-dimensional) state and given that there is a continuum of sellers the buyers know the mass of sellers present in the market whose qualities belong to some given set. On the other hand, they do not know which seller belongs to which quality level, and the platform commits to providing informative signals about this. The payoff of the platform depends on the sets of sellers that are pooled together. This is qualitatively similar to our discussion in the previous sections: the designer pools certain states together, obtains a partition of states, and reveals to which partition element the realization of the state belongs. The designer's payoff function effectively maps each chosen partition to a payoff level.

A key observation in [29] is that a given $\{B_k\}$ and $\{p(B_k)\}$ pair yields a menu of price-expected quality pairs. Thus, the designer's problem can equivalently be formulated as an optimization problem over such menus that are consistent with an equilibrium induced by some information structure and corresponding prices. This reformulation does not directly allow for solving for the optimal information structures/prices since the set of aforementioned pairs in general need not admit a crisp characterization. That said, it is still conceptually useful as the problem reduces to optimizing over menus of prices/quality levels and it is possible to shed light on the optimal solutions of such problems. Specifically, by leveraging such a connection the paper establishes that when the buyers' type distribution is such that $F_b(m)m$ is convex in m and satisfies an additional technical condition (which is characterized in the paper) the platform finds it optimal to use a 1-separating information structure.

Note that when the aforementioned convexity assumption does not hold, 1-separating information structures need not be optimal. It is an interesting research direction to characterize optimal information structures in broader settings and shed light on their structure.

4. Important Extensions

In this section, we revisit some of the key assumptions made in many information design problems and present recent research related to understanding the challenges that emerge once these assumptions are relaxed.

4.1. Informational Spillovers

In many information design problems with multiple receivers it is often assumed that the designer uses public mechanisms and shares the same signal with all the receivers. An alternative and also common assumption is that the designer can target different receivers with different signals but the receivers do not communicate with each other the signals that they have received. For instance, this assumption was implicitly made in one of the information design problems studied in Section 3.1. This assumption may be strong for some applications. How can a designer obtain her mechanisms if there are informational spillovers among the receivers? [13] provides an answer to this question by characterizing spillover structures that yield tractable information design problems and by presenting algorithmic approaches for their characterization. On the other hand, the aforementioned paper also establishes that the presence of agents who follow multiple information sources can render the designer's problem intractable.

Specifically, [13] focuses on a setting where the (continuous) state belongs to an interval $\mathcal{T}$ of $\mathbb{R}$, and the designer has access to a finite set of experiments $E = [0, \ldots, n, n+1]$. Experiment ℓ is associated with a threshold t_ℓ and reveals whether the state is above this threshold or not. We assume that $\inf \mathcal{T} = t_0 < t_1 < \cdots < t_n < t_{n+1} = \sup \mathcal{T}$ and note that the experiments $0, n+1$ almost surely reveal the same signal and hence are uninformative. The designer chooses which set of experiments $E_i \subset E$ to assign to each agent i. The same experiment can be assigned to multiple agents. We refer to a collection of experiment assignments $\{E_i\}$ as an information structure.

Agents do not observe the state once it is realized, but they observe the outcome of the experiments assigned to them. Moreover, there are informational spillovers. We denote by $G = (V, A)$ an underlying directed communication network, and assume that agents correspond to the nodes of this network. A directed arc $(i, j) \in A$ represents an informational spillover from i to j. In addition to the outcome of the experiments in E_i, agent i has access to the outcomes of all the experiments in $\bar{E}_i = \cup_{j \in U(i)} E_j \cup \{0, n+1\}$, where we denote by $U(i)$ the set of agents who have a directed path to i (and include i in this set by convention). We also include the dummy experiments $\{0, n+1\}$ in $\bar{E}_i$ by convention.

Due to the threshold structure, for each agent i a given information structure induces a partition of the set of states into intervals whose end points belong to $\{t_\ell\}_{\ell \in \bar{E}_i}$. Once the outcomes of the experiments are realized the agent can infer to which interval the state belongs. Denote by $\Gamma(S)$ the set of consecutive elements (ℓ_1, ℓ_2) of S, i.e., $\ell_1 < \ell_2$ such that $\ell_1, \ell_2 \in S$ and there does not exist $\ell' \in S$ such that $\ell_1 < \ell' < \ell_2$. It can be seen that agent i can infer whether the state belongs to $[t_{\ell_1}, t_{\ell_2}]$ but not to a smaller interval if and only if $(\ell_1, \ell_2) \in \Gamma(\bar{E}_i)$. We refer to such intervals as minimal intervals.

We assume that the designer's payoff is additive over the agents, and the minimal intervals available to each agent. Specifically, when agent i infers that the state is in a minimal interval $[t_{\ell_1}, t_{\ell_2}]$ for $\ell_1 < \ell_2$, this yields a payoff of $w_i(\ell_1, \ell_2)/p(\ell_1, \ell_2)$ to the designer, where $w_i : E \times E \to \mathbb{R}$ is an arbitrary function and $p(\ell_1, \ell_2) = \mathbb{P}(T \in [t_{\ell_1}, t_{\ell_2}])$. Thus, under information structure $\{E_j\}$, the designer's expected payoff from agent i is given by $\sum_{(\ell_1, \ell_2) \in \Gamma(\bar{E}_i)} w_i(\ell_1, \ell_2)$. With some abuse of notation we denote the total expected payoff of the designer by $\nu(\{E_j\})$. Note that this payoff can be expressed as follows:

$$\nu(\{E_j\}) = \sum_{i \in V} \sum_{(\ell_1, \ell_2) \in \Gamma(\bar{E}_i)} w_i(\ell_1, \ell_2). \tag{15}$$

It can be shown that any setting where (i) after observing the outcomes of the experiments, each agent i takes an action a_i to maximize her expected payoff (which depends on her action and the state), and (ii) the designer's payoff additively decomposes over the agents (and is such that the payoff from each agent depends on her action), exhibits this payoff structure.

The designer's problem is to choose an information structure that maximizes her total expected payoff:

$$\max_{\{E_j\}|E_j\subset E,\forall j} \nu(\{E_j\}). \tag{16}$$

Consider a strongly connected component of $G=(V,A)$. By the structure of spillovers, it follows that all agents in this connected component have access to the same set of experiments. Thus, the designer's problem can be simplified by replacing each such component with a single representative agent. Repeating this for all strongly connected components yields a directed acyclic graph which is known as the *condensation* of the underlying network G. Thus, the design problem for any network can be reduced to a design problem in an equivalent directed acyclic network.

Restrict attention to $G=(V,A)$ that is acyclic and directed. Suppose further that each agent "follows" at most one other agent, i.e., the in-degree of each node is at most one. In this setting, the undirected network that corresponds to the communication network is a tree. [13] establishes that in this setting the designer's problem admits a simple recursive characterization. The key idea is that conditional on the experiments available to a (representative) agent, the problem of the designer in the subtree that originates from that agent decouples into smaller problems. This suggests a simple dynamic programming recursion to obtain the optimal information structure in a tractable way.

Suppose that there is at least one agent who follows multiple information sources. Surprisingly, in this case, even when the communication network has a very simple structure (e.g., a star network where the center follows the leaves), the problem of the designer (or the corresponding decision problem) is NP-hard. Qualitatively, this result implies that following multiple information sources is what makes information design problems intractable in the presence of spillovers.

In settings where the payoffs exhibit special structure the problem may still remain tractable despite the presence of agents who follow multiple information sources. The paper explores one such setting and focuses on a "voting game" where receivers take binary actions, and each agent has incentive to take action 1 when the posterior mean of the state is larger than an associated threshold. The designer's payoff is equal to the number of receivers who take action 1. [13] sheds light on the optimal information structures in the voting game. It establishes that when the followers are more pessimistic (i.e., when the posterior mean requirements of downstream agents for taking action 1 are larger) the network structure does not play any role, and the designer can obtain the optimal information structure by simply ignoring spillovers. Moreover, the optimal information structure exhibits a certain monotonicity property (where among the agents on the same directed path those who have larger posterior mean requirements are assigned to experiments with larger thresholds). On the other hand, when the followers are more optimistic, the network structure impacts the optimal mechanisms and such monotone structures need not be optimal. That said, after restricting attention to monotone (straightforward) information structures, an optimal one within this class can be obtained in a tractable way as long as the underlying communication network has bounded treewidth. Treewidth captures how "tree-like" a given network is, and can be viewed as a measure of how simple a network is. This finding implies that for some simple payoff structures and communication networks the optimal monotone information structure can once again be obtained in a tractable way.

Information design with spillovers is an active research area. It is necessary to identify richer settings where the information design problem remains tractable despite the presence of spillovers. It is also of interest to understand how the possibility of spillovers impacts the optimal information structures in practically relevant environments. Finally, in settings where the designer uses private signals to influence the decisions of multiple receivers (as in Section 3.2) it is important to study the robustness of the findings when there is (possibly limited) information spillovers among agents.

4.2. Privately Informed Receivers

So far we have assumed that the information designer learns all the payoff relevant information once the state is realized. In many interesting settings, the receiver may have payoff relevant information as well (see, e.g., [25, 31]). In this section, we leverage the framework of [14] (which was presented in Section 2.5) to discuss how to obtain the optimal signaling mechanisms when the receiver has private information and shed light on the structure of the optimal mechanisms.

We assume that there is a single receiver, whose type comes from a finite set, denoted by Θ, and we denote the probability that the receiver is of type $\theta \in \Theta$ by $w_\theta > 0$. The receiver chooses her actions from a finite set $[K] = \{1, \ldots, K\}$, and her payoff when the kth action is chosen is given by $h_{\theta,k} + c_{\theta,k}(T - b_{\theta,k})$. For all $\theta \in \Theta$, $\{c_{\theta,k}\}_{k\in[K]}$ are parameters that are strictly increasing in k. Similarly, for all $\theta \in \Theta$, we set $b_{\theta,0} = \inf \mathcal{T}$, $b_{\theta,K} = \sup \mathcal{T}$ and assume that $\{b_{\theta,k}\}_{\{0\}\cup[K]}$ are strictly increasing in k. In addition, for all $\theta \in \Theta$ and $k < K$, the parameters $\{h_{\theta,k}\}$ satisfy

$$h_{\theta,k} = h_{\theta,k+1} + c_{\theta,k+1}(b_{\theta,k} - b_{\theta,k+1}). \tag{17}$$

Under condition (17), when the state is $T = b_{\theta,k}$, the type θ receiver is indifferent between actions k and $k+1$. Thus, it can be readily checked that under our assumptions, the type θ receiver finds it strictly optimal to take action k when the posterior mean belongs to $(b_{\theta,k-1}, b_{\theta,k})$.

When the receiver chooses action $k \in [K]$, and her type is $\theta \in \Theta$, the payoff of the designer is $r_{\theta,k} \in \mathbb{R}$. We assume that $r_{\theta,k} \neq r_{\theta,k+1}$. Under these assumptions for each θ the designer's payoff is a step function of the posterior mean her signals induce, with cutoffs at $\{b_{\theta,k}\}_k$ and reward levels $\{r_{\theta,k}\}_k$. We let $\mathcal{B}_{\theta,k} \subset \mathcal{T}$ denote the set of posterior means for which type θ receiver takes action k which in turn yields a reward of $r_{\theta,k}$ to the designer, and note that $(b_{\theta,k-1}, b_{\theta,k}) \subset \mathcal{B}_{\theta,k} \subset [b_{\theta,k-1}, b_{\theta,k}]$. We point out that the setup described here generalizes the one in Section 2.5 to settings with private information.

Before the state or the receiver's type is realized the designer commits to a direct mechanism π, which consists of a menu of signaling mechanisms $\{\pi_\theta\}$. We assume that for each $\theta \in \Theta$, π_θ is a level mechanism. The receiver reports her type and the designer shares the realization of the signal of the corresponding mechanism with her. We focus on incentive-compatible direct mechanisms, where the receiver finds it optimal to report her type truthfully. [14] establishes that it is without loss of optimality to focus on such mechanisms.

As in Section 2.5, we can characterize each level mechanism π_θ in terms of a corresponding tuple $\{p_{\theta,k}, z_{\theta,k}\}$. However, now the tuples for different $\theta \in \Theta$ need to be related in order to ensure incentive compatibility. To see this, given aforementioned tuples (associated with level mechanisms), for all $\theta, \theta' \in \Theta$, $k \in [K]$, define

$$s_{\theta,\theta',k} = \max_{k'\in[K]} h_{\theta,k'} p_{\theta',k} + c_{\theta,k'}(z_{\theta',k} - b_{\theta,k'} p_{\theta',k}). \tag{18}$$

Intuitively, $s_{\theta,\theta',k}/p_{\theta',k}$ captures the expected payoff of a type θ receiver from reporting her type as θ', subsequently receiving signal k, and taking the action that maximizes her payoff conditional on this signal. Thus, the expected payoff of a type θ agent from reporting her type as θ' is given by $\sum_k s_{\theta,\theta',k}$. Using this observation, [14] establishes that the optimal mechanism of the designer can be obtained by solving the following variant of (OPT):

$$
\begin{aligned}
\max_{\{p_{\theta,k}, z_{\theta,k}, s_{\theta,\theta',k}\}} \quad & \sum_{\theta\in\Theta} w_\theta \sum_{k\in[K]} p_{\theta,k} r_{\theta,k} \\
s.t. \quad & \sum_{k\geq \ell} z_{\theta,k} \leq \int_{1-\sum_{k\geq\ell} p_{\theta,k}}^{1} F^{-1}(x)dx \qquad \text{for } \theta\in\Theta,\ \ell\in[K]\setminus\{1\}\ , \\
& \sum_{k} z_{\theta,k} = \int_0^1 F^{-1}(x)dx \qquad \text{for } \theta\in\Theta, \\
& h_{\theta,k'} p_{\theta',k} + c_{\theta,k'}\left(z_{\theta',k} - b_{\theta,k'} p_{\theta',k}\right) \leq s_{\theta,\theta',k} \qquad \text{for } \theta,\theta'\in\Theta, k,k'\in[K], \\
& \sum_{k\in[K]} s_{\theta,\theta',k} \leq \sum_{k\in[K]} h_{\theta,k} p_{\theta,k} + c_{\theta,k}\left(z_{\theta,k} - b_{\theta,k} p_{\theta,k}\right), \qquad \text{for } \theta,\theta'\in\Theta, \\
& p_{\theta,k} b_{\theta,k-1} \leq z_{\theta,k} \leq p_{\theta,k} b_{\theta,k} \quad \text{for } \theta\in\Theta,\ k\in[K], \\
& \sum_{k\in[K]} p_{\theta,k} = 1 \qquad \text{for } \theta\in\Theta, \\
& p_{\theta,k} \geq 0 \qquad \text{for } \theta\in\Theta,\ k\in[K].
\end{aligned}
\tag{OPT2}
$$

Here, for each θ constraints other than the third and the fourth ensure that the chosen $\{p_{\theta,k}, z_{\theta,k}\}$ tuple is consistent with a level mechanism. The latter constraints ensure incentive compatibility, i.e., type θ agent maximizes her payoff by truthfully reporting her type as θ. [14] establishes that by solving this convex optimization problem and constructing a mechanism consistent with the optimal solution, an optimal mechanism can be obtained. The paper also outlines an algorithm for the construction of the optimal mechanism, and sheds light on the structure of this mechanism.

Theorem 3 ([14]). *Let* $\{p^\star_{\theta,k}, z^\star_{\theta,k}, s^\star_{\theta,\theta',k}\}_{\theta,\theta'\in\Theta,k\in\mathcal{S}}$ *denote an optimal solution of* (OPT2)*. There exists an optimal mechanism* $\pi = \{\pi_\theta\}$ *such that for* $\theta\in\Theta$*,* π_θ *is a level mechanism that is based on a laminar interval partition* $\{\mathcal{T}_{\theta,k}\}_{k\in[K]}$ *and that is consistent with* $\{p^\star_{\theta,k}, z^\star_{\theta,k}\}_k$*. Furthermore, there exists such an optimal mechanism where for* $\theta\in\Theta$*,* $k\in[K]$*,* $\mathcal{T}_{\theta,k}$ *is a union of at most* $|\Theta|+1$ *intervals.*

Note that in the absence of private information, $|\Theta| = 1$, and the optimal mechanism has a double interval structure described earlier. Intuitively, when the receiver has private information the partitions that support the optimal mechanism have a more intricate structure. This is necessary to ensure incentive compatibility. Moreover, it is easy to construct instances where the incentive compatibility constraints are binding and they have a drastic impact on the choice of the mechanism of the designer. We refer the reader to [14] for details.

5. Concluding Remarks

In the last decade, there has been substantial interest in information design. The existing literature has laid the foundation for systematically approaching information design problems. In addition, the applications of information design in many different settings have been explored. The purpose of this paper was to provide an overview of the different approaches that can be used to characterize optimal information structures and present various applications of information design in the recent operations management literature. In addition, we discussed key assumptions imposed in information design that need not hold in various operational settings, and reviewed the recent work that focuses on information design when such assumptions are violated.

The intersection of information design and operations presents many interesting future research opportunities. Four broad research directions are worth highlighting. First, as discussed in the introduction, information is a natural lever in many operational settings including service operations, ride-sharing platforms, transportation networks, and e-commerce platforms. This paper has reviewed some of the early work in the area. Going forward it is important to focus on other relevant settings, carefully model the operational details, and discuss how firms/decision makers can optimally use information to induce a desired outcome.

Second, other natural levers, such as pricing and admission control, have been widely studied by the operations community. However, exploring these levers jointly with informational levers has only recently started to receive attention. As briefly discussed in Section 3.2, their interplay can be nontrivial. When is choosing an appropriate information structure a substitute for using operational levers? In which settings do they complement each other? A thorough exploration of the interplay between different operational levers and information design presents many research opportunities.

Third, there are important algorithmic challenges in obtaining optimal information structures. Some of these challenges have received attention in the recent computer science and operations research literature (see, e.g., [18, 19, 20]). However, more research is needed especially in settings where the commonly made assumptions in the literature no longer hold. Some of these environments were discussed in Section 4, and some others (e.g., robust approaches to information design [22]) are active areas of research.

Lastly, in this paper we mainly focused on static information design environments. Information is an important lever in many dynamic settings as well. For instance, a firm can incentivize agents who arrive over time to experiment with different products (whose qualities are a priori unknown) by employing appropriately chosen information structures (such as review systems), which, e.g., reveal information about past decisions. Dynamic information design investigates how to choose information structures optimally in such settings and constitutes another active research area [23, 27, 32]. Applications of dynamic information design in different operational settings also remain to be an exciting avenue for future research.

Acknowledgments

I thank Kostas Bimpikis, Kimon Drakopoulos, Emir Kamenica, and Kris Iyer for their feedback.

References

[1] Saed Alizamir, Francis de Véricourt, and Shouqiang Wang. Warning against recurring risks: An information design approach. *Management Science*, 2020. https://doi.org/10.1287/mnsc.2019.3420.

[2] Jerry Anunrojwong, Krishnamurthy Iyer, and Vahideh Manshadi. Information design for congested social services: Optimal need-based persuasion. *arXiv preprint arXiv:2005.07253*, 2020.

[3] Coralio Ballester, Antoni Calvó-Armengol, and Yves Zenou. Who's who in networks. Wanted: The key player. *Econometrica*, 74(5):1403–1417, 2006.

[4] Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.

[5] Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, 2019.

[6] James Best and Daniel Quigley. Persuasion for the long run. *Available at SSRN 2908115*, 2017.

[7] Kostas Bimpikis, Yiangos Papanastasiou, and Wenchang Zhang. Information provision in two-sided platforms: Optimizing for supply. *Available at SSRN 3617351*, 2020.

[8] Francis Bloch and Nicolas Quérou. Pricing in social networks. *Games and Economic Behavior*, 80:243–261, 2013.

[9] John Adrian Bondy, Uppaluri Siva Ramachandra Murty, et al. *Graph Theory with Applications*, volume 290. Macmillan London, 1976.

[10] Isabelle Brocas and Juan D. Carrillo. Influence through ignorance. *RAND Journal of Economics*, 38(4):931–947, 2007.

[11] Ozan Candogan. Optimality of double intervals in persuasion: A convex programming framework. *Available at SSRN 3452145*, 2019.

[12] Ozan Candogan. Persuasion in networks: Public signals and k-cores. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pages 133–134. ACM, 2019.

[13] Ozan Candogan. On information design with spillovers. *Available at SSRN 3537289*, 2020.

[14] Ozan Candogan. Reduced form information design: Persuading a privately informed receiver. *Available at SSRN 3533682*, 2020.

[15] Ozan Candogan, Kostas Bimpikis, and Asuman Ozdaglar. Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905, 2012.

[16] Ozan Candogan and Kimon Drakopoulos. Optimal signaling of content accuracy: Engagement vs. misinformation. *Operations Research*, 68(2):497–515, 2020.

[17] Kimon Drakopoulos, Shobhit Jain, and Ramandeep S Randhawa. Persuading customers to buy early: The value of personalized information provisioning. *Available at SSRN 3191629*, 2018.

[18] Shaddin Dughmi. Algorithmic information structure design: A survey. *ACM SIGecom Exchanges*, 15(2):2–24, 2017.

[19] Shaddin Dughmi and Haifeng Xu. Algorithmic Bayesian persuasion. *In Proceedings of the 48th Annual ACM Symposium on Theory of Computing*, pages 412–425, 2016.

[20] Shaddin Dughmi and Haifeng Xu. Algorithmic persuasion with no externalities. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 351–368, 2017.

[21] Piotr Dworczak and Giorgio Martini. The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019.

[22] Piotr Dworczak and Alessandro Pavan. Preparing for the worst but hoping for the best: Robust (Bayesian) persuasion. 2020. Department of Economics, Northwestern University.

[23] Jeffrey C Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

[24] Matthew Gentzkow and Emir Kamenica. A Rothschild–Stiglitz approach to Bayesian persuasion. *American Economic Review*, 106(5):597–601, 2016.

[25] Yingni Guo and Eran Shmaya. The interval structure of optimal disclosure. *Econometrica*, 87(2):653–675, 2019.

[26] Yonatan Gur, Gregory Macnamara, and Daniela Saban. On the disclosure of promotion value in platforms with learning sellers. *arXiv preprint arXiv:1911.09256*, 2019.

[27] Johannes Hörner and Andrzej Skrzypacz. *Learning, Experimentation, and Information Design*, volume 1 of *Econometric Society Monographs*, pages 63 – 98. Cambridge University Press, 2017.

[28] Svante Janson and Malwina J Luczak. A simple solution to the k-core problem. *Random Structures & Algorithms*, 30(1–2):50–62, 2007.

[29] Ramesh Johari, Bar Light, and Gabriel Weintraub. Quality selection in two-sided markets: A constrained price discrimination approach. *arXiv preprint arXiv:1912.02251*, 2019.

[30] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[31] Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li. Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964, 2017.

[32] Ilan Kremer, Yishay Mansour, and Motty Perry. Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012, 2014.

[33] Can Küçükgül, Özalp Özer, and Shouqiang Wang. Engineering social learning: Information design of time-locked sales campaigns for online platforms. *Available at SSRN 3493744*, 2019.

[34] David Lingenbrink and Krishnamurthy Iyer. Signaling in online retail: Efficacy of public signals. *Available at SSRN 3179262*, 2018.

[35] David Lingenbrink and Krishnamurthy Iyer. Optimal signaling mechanisms in unobservable queues. *Operations Research*, 67(5):1397–1416, 2019.

[36] Laurent Mathevet, David Pearce, and Ennio Stacchetti. Reputation and information design, 2019. Working Paper, New York University.

[37] Yiangos Papanastasiou, Kostas Bimpikis, and Nicos Savva. Crowdsourcing exploration. *Management Science*, 64(4):1727–1746, 2018.

[38] Luis Rayo and Ilya Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, 2010.

[39] Pu Yang, Krishnamurthy Iyer, and Peter Frazier. Information design in spatial resource competition. *arXiv preprint arXiv:1909.12723*, 2019.