

Bayesian Persuasion and Information Design

Emir Kamenica

Booth School of Business, University of Chicago, Chicago, Illinois 60637, USA;
email: emir.kamenica@chicagobooth.edu

**ANNUAL
REVIEWS CONNECT**

www.annualreviews.org

- Download figures
- Navigate cited references
- Keyword search
- Explore related articles
- Share via email or social media

Annu. Rev. Econ. 2019. 11:249–72

First published as a Review in Advance on
May 13, 2019

The *Annual Review of Economics* is online at
economics.annualreviews.org

<https://doi.org/10.1146/annurev-economics-080218-025739>

Copyright © 2019 by Annual Reviews.
All rights reserved

JEL codes: D02, D82

Keywords

communication, manipulation of beliefs, rating systems

Abstract

A school may improve its students' job outcomes if it issues only coarse grades. Google can reduce congestion on roads by giving drivers noisy information about the state of traffic. A social planner might raise everyone's welfare by providing only partial information about solvency of banks. All of this can happen even when everyone is fully rational and understands the data-generating process. Each of these examples raises questions of what is the (socially or privately) optimal information that should be revealed. In this article, I review the literature that answers such questions.

1. WHAT IS BAYESIAN PERSUASION?

An implicit premise in most of economics is that behavior is driven by three factors: preferences, technology, and information. Consequently, if we wish to influence economic outcomes, then there are three broad ways of doing so. The most straightforward is to change the (induced) preferences over actions via incentives, e.g., contingent payments, threat of violence, or supply of complementary goods. A second way to engender an outcome is to make it easier for a decision maker to achieve it, i.e., to improve the relevant technology.¹ This article is about the third path—persuasion—which we can define as influencing behavior via provision of information (Kamenica & Gentzkow 2011).² Throughout, the focus is on standard decision makers who understand how information is generated and react to information in a rational (Bayesian) manner: thus the term Bayesian persuasion. Bayesian persuasion is also referred to as information design,³ and a comparison with mechanism design is instructive (Bergemann & Morris 2016b, Taneva 2016). In mechanism design problems, the allocation of information (i.e., who knows what) is given, and the designer influences the outcome by selecting the game that the agents will play. In information design problems, the game that the agents play is given, and the designer influences the outcome by specifying the allocation of information.

Bayesian persuasion can alternatively be seen as a communication protocol, in the tradition of cheap talk (Crawford & Sobel 1982), verifiable message (Grossman 1981, Milgrom 1981), and signaling games (Spence 1973). Relative to these other models of communication, Bayesian persuasion endows the sender with more commitment power. In the most common formulation, Bayesian persuasion allows the sender to commit to sending any distribution of messages as a function of the state of the world. This full commitment formulation, however, yields equilibrium outcomes that are identical to those that arise in an alternative model where the sender publicly chooses how much information he will privately observe and then strategically decides how much of this private information to reveal via verifiable messages (Gentzkow & Kamenica 2017c). This equivalence result may not be especially important, however, since in most applications, the full commitment formulation corresponds more closely to the real-world institution being analyzed.

In this review, I focus exclusively on environments where the information designer is motivated by the desire to influence the actions of those who observe the signal realization. A separate line of research, surveyed by Bergemann & Bonatti (2019), studies markets where a seller designs information to sell it. Bergemann et al. (2018a), for example, consider a seller who offers a menu of signals to a buyer with unknown private information, while Kastl et al. (2018) analyze the sale of information to competitive firms about their suppliers' marginal costs.

Bayesian persuasion is also closely related to the literature on Bayes correlated equilibria (Bergemann & Morris 2013). Bayes correlated equilibria take as given a basic game (a set of players, a set of feasible actions for each player, and players' payoffs as a function of the state of the

¹Kamenica (2012) argues that some of the methods used in nudging and choice architecture (Thaler & Sunstein 2008) can be seen as technological interventions. For instance, teaching drivers to open their car door with their right hand when they are about to exit the car (which forces them to swivel and see whether there is a bicycle coming) is a technological innovation that allows drivers to reduce the chance that they cause an accident.

²Of course, there is no reason why these three methods of influencing behavior must be used in isolation. A promising area for future research might be to explore how to best combine incentives, choice architecture, and information provision. Li (2017) adds transfers to the simplest two-action, two-state model of persuasion. Lewis & Sappington (1994) study a monopolist who chooses both what price to charge and what information to provide to consumers about their valuations for the firm's product.

³I use the two terms interchangeably, but the former is probably used more when the designer is one of the players in the game and there is a single receiver, while the latter is used more when the designer is a social planner or there are multiple, interacting receivers.

world and the actions taken) and describe the set of all possible outcomes that could arise (as Bayes Nash equilibria) regardless of what each player knows (about the state and about what the other players know). One benefit of deriving this set is that it provides a prediction about the outcome of the basic game that is robust to the uncertainty about what the players engaged in the game know, but the set of Bayes correlated equilibria, by definition, also coincides with the set of outcomes that can be attained through information design. Thus, a Bayesian persuasion problem is equivalent to a problem of selecting an optimal Bayes correlated equilibrium given an objective function. The relationship between these two literatures is also discussed in a complementary survey by Bergemann & Morris (2019).

The research on Bayesian persuasion has developed in two main directions. One strand of research, which is the primary focus of this article, is more abstract and seeks to extend the basic model in various dimensions and/or develop new approaches for solving the designer's optimization problem. The other strand is more applied and aims to understand or improve real-world institutions via information design. Research in this second strand includes applications to (in no particular order) financial sector stress tests (Goldstein & Leitner 2018, Inostroza & Pavan 2018, Orlov et al. 2018b), grading in schools (Boleslavsky & Cotton 2015, Ostrovsky & Schwarz 2010), employee feedback (Habibi 2018, Smolin 2017), law enforcement deployment (Hernandez & Neeman 2017, Lazear 2006, Rabinovich et al. 2015), censorship (Gehlbach & Sonin 2014), entertainment (Ely et al. 2015), financial over-the-counter markets (Duffie et al. 2017), voter coalition formation (Alonso & Camara 2016b), research procurement (Yoder 2018), contests (Feng & Lu 2016, Zhang & Zhou 2016), medical testing (Schweizer & Szech 2019), medical research (Kolotilin 2015), matching platforms (Romanyuk & Smolin 2019), price discrimination (Bergemann et al. 2015), financing (Szydlowski 2016), insurance (Garcia & Tsur 2018), transparency in organizations (Jehiel 2015), and routing software (Das et al. 2017, Kremer et al. 2014). The breadth of these topics reveals the wide applicability of Bayesian persuasion models.

2. THE MODEL AND ITS INTERPRETATIONS

2.1. The Basic Model

The basic Bayesian persuasion model (of the static variety with a single sender and a single receiver) takes the following form. A player called Receiver (she) has a utility function $u(a, \omega)$ that depends on her action $a \in A$ and the state of the world $\omega \in \Omega$. Another player, Sender (he), aka the information designer, has a utility function $v(a, \omega)$ that depends on Receiver's action and the state of the world. In some applications, we might think of Sender not as a player but rather as the social planner with social welfare function v . Sender and Receiver share a common prior μ_0 on Ω .

The key object in Bayesian persuasion models is the thing that Sender chooses, which goes by many names, including signal, signal structure, information structure, experiment, Blackwell experiment, or data-generating process. I refer to it as signal. Let S be some sufficiently large set of signal realizations.⁴ In the basic model that we consider, it suffices to assume that $|S| \geq \min\{|A|, |\Omega|\}$, i.e., that the signal realization is not smaller than both the state space and the action space.⁵ The notation that I use further assumes that S is finite. A signal is a map from the state to the

⁴An alternative to taking S as given is to let Sender choose it in the process of selecting his signal. Doing so, however, creates an economically irrelevant, if philosophically delightful, complication that prevents us, because of Russell's paradox, from saying that Sender optimizes over the set of all signals.

⁵With multiple receivers, things are a little more tricky. Suppose that there are n receivers, each with some action space A_i . If we consider the private signals environment (see the discussion in Section 4.1), and if equilibrium selection is resolved in favor of the designer, then it suffices to set $|S| \geq |A_1 \times \dots \times A_n|$ since, by the

distribution over signal realizations, $\pi : \Omega \rightarrow \Delta(S)$.⁶ In other words, a signal specifies the statistical relationship between truth ($\omega \in \Omega$) and data ($s \in S$). Another, equivalent, way to define a signal is as a joint distribution over states and signal realizations, $\pi \in \Delta(\Omega \times S)$, with the requirement that the marginal distribution over Ω coincides with the prior. Let Π denote the set of all signals. The timing is as follows:

1. Sender chooses a signal π .
2. Receiver observes which signal was chosen.
3. Nature chooses ω according to μ_0 .
4. Nature chooses s according to $\pi(\omega)$.
5. Receiver observes the realized s .
6. Receiver takes action a .

Receiver's behavior is mechanical. Given her knowledge of π , she uses Bayes' rule to update her belief from the prior μ_0 to the posterior $\mu_\pi(\omega|s) = [\pi(s|\omega)\mu_0(\omega)]/[\sum_{\omega'} \pi(s|\omega')\mu_0(\omega')]$, and then she simply selects an action $a^*(\mu_\pi(\cdot|s))$ that maximizes $\mathbb{E}_{\omega \sim \mu_\pi(\cdot|s)} u(a, \omega)$.⁷ Given this behavior by Receiver, Sender solves

$$\max_{\pi \in \Pi} \mathbb{E}_{\omega \sim \mu_0} \mathbb{E}_{s \sim \pi(\omega)} v(a^*(\mu_\pi(\cdot|s)), \omega). \quad 1.$$

The optimization problem in Equation 1 looks somewhat daunting for a couple of reasons. First, Π is a pretty large set. Second, the choice of π influences Sender's payoff both by changing the distribution of signal realizations and by changing the action induced by a given signal realization. Much of the progress in the Bayesian persuasion literature has relied on recasting Sender's optimization problem in a more approachable way. I discuss those reformulations and environments where they are applicable in Section 3.

2.2. Interpretations

The motivating example of Kamenica & Gentzkow (2011) considers a courtroom setting where Sender is a prosecutor and Receiver is a judge. The state of the world is the guilt of the defendant. We can think of the choice of the signal as consisting of forensic tests, questions asked to witnesses, etc. The prosecutor can ask for a DNA test but does not have to; he can call an expert witness but does not have to; etc. The assumption that the judge necessarily sees all of the evidence uncovered by the prosecution might seem problematic, but anything unfavorable to the accused the prosecutor will willingly share (the prosecutor prefers conviction), and any exculpatory evidence he is required by law to reveal. [In *Brady v. Maryland* (1963), the Supreme Court of the United States ruled that a prosecutor violates the Due Process Clause of the Fourteenth Amendment when he

revelation principle–type argument, we can restrict attention to signal realizations that are action recommendations. However, analogously to the fact that the revelation principle no longer applies in mechanism design once we drop the assumption that equilibrium selection favors the designer (e.g., Maskin 1999), an S consisting of action recommendations no longer necessarily suffices in Bayesian persuasion problems with multiple receivers and adversarial equilibrium selection (see also Mathevet et al. 2018). Another variety of models with multiple receivers requires that all receivers observe the same signal realization. In that case, it suffices to set $|S| \geq |\Omega|$.

⁶Given a set X , $\Delta(X)$ denotes the set of all probability distributions on X .

⁷If Receiver does not have duplicate actions that yield the same utility in each state, then she can only be indifferent across multiple actions at a set of beliefs that has dimensionality lower than $\Delta(\Omega)$. In any equilibrium that leads her to hold such beliefs with positive probability, it will necessarily be the case that she breaks her indifference in Sender's favor.

fails to disclose evidence favorable to the accused.] The courtroom example might not seem to match the model because an actual prosecutor can conduct additional investigations (i.e., generate another signal) once he observes the realization of the initial signal. It turns out, however, that allowing for such additional signals does not change the game since any contingent information gathering can be fully baked into the choice of the initial signal. Finally, it is implausible that a prosecutor could generate arbitrarily informative signals, i.e., discover the whole truth. This concern, again, does not diminish the applicability of the model; since we can redefine the state to be the realization of the most informative signal that Sender can generate, the assumption that a fully informative signal is feasible is vacuous. Thus, the Bayesian persuasion approach can be useful even when the environment does not match the literal description of the model very well.

The justifications underlying the commitment assumption vary substantially across applications. Consider grades in schools. There is a large population of students with varying ability. I denote the ability of a particular student by ω . The distribution of ability in the student population, μ_0 , is known to everyone. It may be the case—in contrast with the literal description of the model—that each student’s ability is directly observed by the school, perhaps even before the school chooses its grading policy (i.e., its signal). This will not matter.⁸ What is important is that the way in which the school assigns grades (the school’s grading policy) is publicly known. A grading policy maps each ability level ω into a (potentially randomized) grade. The idea that the mapping from ability to grades can be stochastic might reflect a policy that takes attributes orthogonal to ability, e.g., attendance, into consideration when assigning grades.⁹ The set of possible grades in this case is S , and $\pi(s|\omega)$ is the probability that a student of ability ω obtains grade s . Receiver is the labor market at large, and the placement of a student whose ability is perceived to be μ is some $a^*(\mu)$. The school values placement a of a student of ability ω at $v(a, \omega)$; Equation 1 then implies that the school chooses its grading policy to maximize its average valuation of the placement of its students.¹⁰

For yet another interpretation of what constitutes a signal, consider deployment of law enforcement. Lazear (2006) introduces the following model. There are Z miles of road. A driver can either speed or obey the speed limit on each mile. Speeding generates utility V per mile, and the fine for speeding if caught is $K > V$. There are $G < Z$ police, and each policeman can patrol one mile of road. The police wish to minimize the number of miles over which drivers speed. To map this environment to the Bayesian persuasion model, let $\omega \in \Omega = \{0, 1\}$ denote the presence of a policeman on a given mile; thus, we have $\mu_0 = G/Z$. The set of signal realizations corresponds to the miles of the road: $S = \{1, \dots, Z\}$. The police is Sender and the driver is Receiver. A signal in this case represents the consistency or predictability of the patrolling strategy. A patrolling strategy induces a joint distribution over Ω and S , i.e., over the presence of a policeman and a mile of the road. The case where police randomly choose where to set up their speed traps each day corresponds to the completely uninformative signal $\underline{\pi}$ (with ω and s uncorrelated) and $\mu_{\underline{\pi}}(\cdot|s) = G/Z$ on every mile of the road s . This policy induces drivers to speed everywhere if $V > (GK)/Z$ and to speed nowhere otherwise. The case where the police always patrol the exact same locations corresponds to the completely informative signal $\bar{\pi}$ (with ω and s perfectly correlated), and $\mu_{\bar{\pi}}(\cdot|s)$

⁸If the grading policy is chosen after seeing the abilities, it is important that the student population is large, so there is no uncertainty about the realized distribution of ability, μ_0 .

⁹If Ω is uncountable, and A is finite, then Sender is not harmed by restricting Π to deterministic, i.e., partitional, signals, $\pi : \Omega \rightarrow S$.

¹⁰Of course, the impact of grades on the labor market is more complex—placement of a student with a given apparent ability might depend on the grading policy at other schools, schools might differ in the quality of their education, etc. Papers that focus on applications of information design to grading develop extensions that take these issues into account (e.g., Boleslavsky & Cotton 2015, Ostrovsky & Schwarz 2010).

is either zero or one: one on the G/Z share of miles that are consistently patrolled and zero on the remainder of the road. When speeding is sufficiently appealing (i.e., $V > (GK)/Z$), neither of these policies will be optimal—a partially informative signal induced by an imperfect consistency in the location of the speed traps will be the best.

These three disparate examples are meant to illustrate the variety of situations where some party has control over the information that another party will observe. Other applications involve yet other interpretations of how a signal is generated. Best & Quigley (2017), for instance, identify the circumstances under which reputation-building motives (e.g., of a long-run Sender facing a sequence of short-run receivers) allow Sender to generate arbitrary signals.

3. SENDER'S OPTIMIZATION PROBLEM

3.1. Concavification

Let us return to the optimization problem in Equation 1. We aim to reformulate this problem in a way that will, at least in some circumstances, be more manageable than brute force optimization over all possible signals. Our reformulation relies on two steps.

First, note that, whatever signal Sender chose, his expected payoff is fully determined by Receiver's posterior. In particular, if Receiver holds belief μ , then Sender's expected utility is

$$\hat{v}(\mu) = \mathbb{E}_{\omega \sim \mu} v(a^*(\mu), \omega).$$

Note that μ enters \hat{v} in two ways—it both affects Receiver's action and impacts Sender's expected utility from that action. The latter channel relies on the fact that Receiver's beliefs, being formed by Bayes' rule, are expected to be well calibrated: If we look at all of the cases in which Receiver had some belief μ , then we should expect to find that the state was ω in $\mu(\omega)$ share of those cases.

The second step draws on the relationship between signals and distributions of beliefs. When Sender chooses some signal π , each signal realization s leads to some posterior $\mu_\pi(\cdot|s)$. From the ex ante perspective, however, before the realization of the signal, we can think of the choice of π as inducing a distribution of posteriors. We use notation $\tau = \langle \pi \rangle$ to indicate that a distribution of posteriors τ is induced by signal π .¹¹ We say that a distribution of posteriors τ is Bayes plausible if it equals the prior in expectation, i.e., $\mathbb{E}_{\mu \sim \tau} \mu = \mu_0$. By the law of iterated expectations, we know that every distribution of posteriors induced by a signal is Bayes plausible. Moreover, Bayes plausibility is the only restriction on induced distributions of posteriors: For every Bayes-plausible τ , there is a $\pi \in \Pi$ such that $\tau = \langle \pi \rangle$ (Kamenica & Gentzkow 2011). The proof of this observation is constructive and allows us to easily find a signal that induces any given Bayes-plausible τ .¹²

Combining these two observations allows us to reformulate Sender's problem as

$$\max_{\tau} \mathbb{E}_{\mu \sim \tau} \hat{v}(\mu), \tag{2}$$

subject to $\mathbb{E}_{\mu \sim \tau} \mu = \mu_0$.

This formulation has a nice geometric interpretation. Consider **Figure 1**. Sender's payoff from a signal that induces some beliefs $\{\mu_l, \mu_b\}$ is the height of the point at the intersection of the line segment connecting $\hat{v}(\mu_l)$ and $\hat{v}(\mu_b)$ and the vertical line from μ_0 . This observation makes it clear that the optimal binary signal is the one that induces the distribution of posteriors with the support

¹¹Algebraically, a signal π induces a distribution of posteriors τ if $\tau(\mu) = \sum_{s: \mu_\pi(\cdot|s) = \mu} \sum_{\omega' \in \Omega} \pi(s|\omega') \mu_0(\omega')$.

¹²Specifically, associate every μ in the support of τ with some signal realization s , and let $\pi(s|\omega) = [\mu(\omega)\tau(\mu)]/\mu_0(\omega)$.




Figure 1

Sender's value function and its concavification (*thick red line*).

on $\{\mu_l^*, \mu_h^*\}$. Moreover, it is easy to see that a signal with more than two realizations cannot improve Sender's payoff.

We can generalize this line of reasoning with the notion of concavification. A concavification of \hat{v} is the smallest concave function everywhere greater than \hat{v} . It is indicated as the thick red line in **Figure 1**. The concavification of \hat{v} evaluated at μ_0 equals $\max\{z | (\mu_0, z) \in \text{co}(\hat{v})\}$, where $\text{co}(\hat{v})$ denotes the convex hull of the graph of \hat{v} . Therefore, since the set of Sender's payoffs across all signals is $\{z | (\mu_0, z) \in \text{co}(\hat{v})\}$,¹³ Sender's payoff under the optimal signal is precisely the concavification of \hat{v} evaluated at the prior.

When there are only two or three states of the world, we can plot \hat{v} , and the concavification approach then allows us to simply read off the optimal distribution of posteriors. It is then straightforward to identify a signal that induces this distribution of posteriors. When the state space is larger, the concavification approach can still be used to derive some qualitative features of the optimal signal, but it does not immediately deliver the solution to Sender's problem. Where I discuss the various extensions of the basic model in Section 4, I note the cases where the concavification approach remains useful.

A few words on the intellectual history of information design and the concavification approach might be in order. At the height of the Cold War, from 1966 to 1968, a group of economists was retained by the United States Arms Control and Disarmament Agency to study game-theoretic aspects of arms control and disarmament. This group included Robert Aumann, Michael Maschler, and Richard Stearns. These three authors wrote a number of reports that focused on a particular concern, namely "that the negotiating strategy used by the Americans in a series of arms control conferences might implicitly send signals to the Russians about the nature of the US arsenal" (Aumann & Maschler 1995, p. xiii).¹⁴ To study this issue, Aumann & Maschler (1966) consider the following model. There are two players, called Informed (player I) and Uninformed (player U). There are two zero-sum games, G_A and G_B , with identical action spaces. With probability μ_0 , the players will repeatedly play G_A ad infinitum, and with the complementary probability, they will repeatedly play G_B ad infinitum. Before they start playing, player I learns which game they will be playing. After each period, player U observes the action of player I, but she does not observe her payoff, nor which game they had played (and will continue to play). Both players seek to

¹³Note that $\{z | (\mu_0, z) \in \text{co}(\hat{v})\} = \{\mathbb{E}_{\mu \sim \tau} \hat{v}(\mu) | \mathbb{E}_{\mu \sim \tau} \mu = \mu_0\}$.

¹⁴This quote is from the preface of the book that collected and reorganized the original technical reports.

maximize their undiscounted average payoff. While this situation might seem quite distinct from the Bayesian persuasion model, an important step in characterizing the equilibria of these repeated games of incomplete information involves solving an optimization problem that is analogous to the one in Equation 2. Aumann & Maschler (1966) developed the concavification approach for solving such problems. For over half a century, however, Aumann & Maschler's contribution played a part in the analysis of repeated games but did not spur the development of the applications mentioned in Section 1. Formulating a model with an explicit information design step seems to have been important for stimulating further research on the topic.

Another early, although more recent, contribution to information design is that of Brocas & Carrillo (2007).¹⁵ They analyze an environment that fits squarely within the framework of the model in Section 2.1, with one exception: Sender, rather than choosing any signal whatsoever from Π , selects how many independent and identically distributed (IID) draws of a fixed (binary) signal to generate. Even though it might seem easier to optimize over this one-dimensional set, the fact that the concavification approach is only applicable when Sender can choose from all possible signals has led to the unrestricted model being more widely used.

3.2. An Important Special Case

As mentioned above, the concavification approach delivers a visual solution to the information design problem only when the state space is small, with two or three elements. Another special case that has received much attention is the case where the state space is large, in fact uncountable, but Sender's payoff depends only on the mean of Receiver's posterior.

Suppose that $\Omega = [0, 1]$ and that $a^*(\mu) = f(\mathbb{E}_\mu \omega)$ for some function f .¹⁶ Then, there exists a function \tilde{v} that captures Sender's payoff as a function of Receiver's posterior mean, i.e., $\tilde{v}(\mathbb{E}_\mu \omega) = \hat{v}(\mu)$. Even though μ is infinite dimensional, \tilde{v} can be plotted on a piece of paper. We might hope that the concavification of \tilde{v} would yield the solution to Sender's optimization problem, but unfortunately this turns out not to be the case.¹⁷ The problem is that, even though we can induce every distribution of posteriors whose average value is the prior, this is not the case for every distribution of posterior means whose average value the prior mean.

Thus, we cannot solve these cases via concavification, but some progress has been made using other routes. Gentzkow & Kamenica (2016) derive a way to solve these problems when the action space is small. The first step is to transform each signal into a convex function. Let G_π denote the cumulative distribution function of the posterior means induced by π . Then, let c_π be the integral of G_π , i.e., $c_\pi(x) = \int_0^x G_\pi(t) dt$. Note that c_π is an increasing convex function. Let me illustrate the construction of c_π with some examples. Suppose that the prior μ_0 is uniform on $[0, 1]$. First, consider a completely uninformative signal $\underline{\pi}$. This signal induced a degenerate distribution of posterior means always equal to $1/2$, so $G_{\underline{\pi}}$ is a step function equal to 0 below $1/2$ and equal to 1 above. Thus, $c_{\underline{\pi}}$ is flat on $[0, 1/2]$ and then linearly increasing on $[1/2, 1]$ with a slope of 1. Second, at the other extreme, consider a fully informative signal $\bar{\pi}$. Under this signal, the posterior mean equals the true state, so the distribution of posterior means is uniform, $G_{\bar{\pi}}$ is linear, and $c_{\bar{\pi}}$ is quadratic: $c_{\bar{\pi}}(x) = x^2/2$. Finally, consider a partial signal \mathcal{P} that gives a distinct signal

¹⁵In work contemporaneous with that of Kamenica & Gentzkow (2011), Rayo & Segal (2010) analyze a specific case of the baseline model from Section 2.1 with an added twist that Receiver has some private information about her preferences.

¹⁶It only matters that Ω is a compact subset of \mathbb{R} . I assume that it is a unit interval to simplify notation.

¹⁷The concavification is not entirely useless in this case. If the concavification of \tilde{v} is strictly above \tilde{v} at the prior mean, then we can at least conclude that providing no information is strictly suboptimal (Kamenica & Gentzkow 2011).

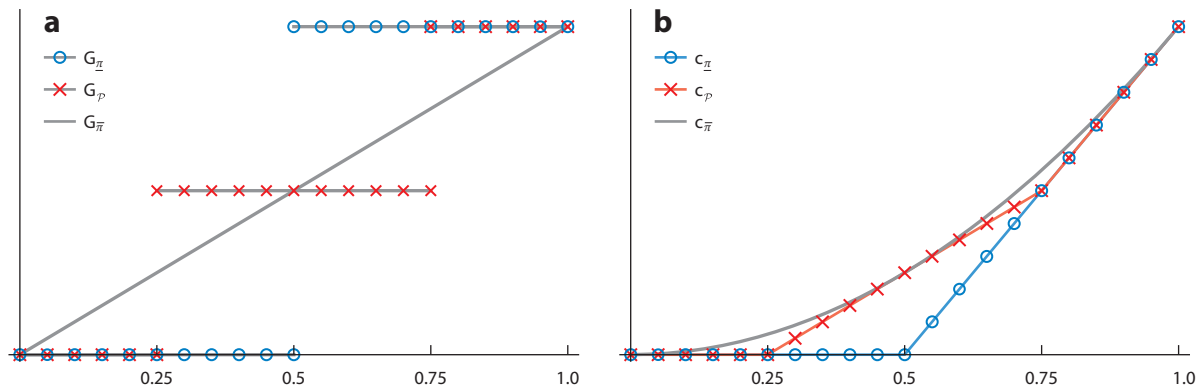


Figure 2

Signals as convex functions. (a) Cumulative distribution functions of posterior means. (b) Functions induced by signals.

realization depending on whether the state is above or below $1/2$. Then, G_p is a step function, and c_p is piecewise linear. **Figure 2** depicts these functions. Note that signal π is more informative¹⁸ than π' if and only if $c_\pi \geq c_{\pi'}$ (Blackwell 1951, Blackwell & Girshick 1954). This implies that, for any signal π , we have $c_{\bar{\pi}} \geq c_\pi \geq c_{\underline{\pi}}$. In fact, every convex function that is sandwiched between $c_{\bar{\pi}}$ and $c_{\underline{\pi}}$ is induced by some signal (Gentzkow & Kamenica 2016). Thus, we can represent the set of all signals Π as the set of convex functions. If Receiver's action space is finite (so that \bar{v} is a step function), then we can solve for the optimal signal by working inside this new space. For a simple example, suppose that prior is uniform, and \bar{v} is a step function equal to 0 below some cutoff γ and 1 above it.¹⁹ Then, Sender's payoff from π is the likelihood that Receiver's induced posterior mean is γ or above, which corresponds to $1 - G_\pi(\gamma)$ or $1 - c'_\pi(\gamma^-)$.²⁰ Thus, Sender wants to minimize the left derivative of c_π at γ , which is attained by c_{π^*} , indicated in **Figure 3**.

The special case where Sender's payoff depends only on the mean of Receiver's posterior is also studied by Ivanov (2015), Kolotilin (2018), and Dworzak & Martini (2019).²¹ Ivanov (2015) considers an extension of the basic model that allows for Sender's payoff to also depend on the rank of the realized posterior mean among the posterior means that might be generated. Kolotilin (2018) and Dworzak & Martini (2019) draw on linear programming methods and represent Sender's problem as a consumer-choice situation where Sender purchases posterior means using the prior as her endowment.²² Kolotilin (2018) assumes that the action space is binary but allows for Receiver to have private information about his preferences. Dworzak & Martini (2019) allow for a general action space, and thus for a \bar{v} that can take any shape, and derive a simple way to verify whether a given signal is optimal.²³

¹⁸Throughout this article, when I say more informative, less informative, or comparable, I mean in terms of Blackwell order.

¹⁹This extremely simple case can also be solved directly, but I use it rather than a case that cannot be solved using other methods for the simplest illustration of how this method works.

²⁰The expression $c'_\pi(\gamma^-)$ denotes the left derivative of c_π at γ .

²¹Ostrovsky & Schwarz (2010) similarly consider this special case but in the setting with multiple senders discussed in Section 4.2.

²²Moreover, Dworzak & Martini (2019) show that this reformulation does not require the assumption that Sender's preferences depend only on the posterior mean.

²³These and other papers also deliver results on circumstances under which optimal signals fall into a particular class. Ivanov (2015) and Dworzak & Martini (2019) provide necessary and sufficient conditions for optimality of monotone partitional signals. Kolotilin (2018) provides necessary and sufficient conditions for optimality of




Figure 3

Deriving the optimal signal. The set of convex functions in the shaded region represents the set of distributions of means that can be induced by a signal.

3.3. Computational Methods

Despite the aforementioned progress in techniques for deriving optimal signals in various settings, many information design problems of applied interest may not be amenable to analytic solutions. This naturally leads to a question of whether we might be able to employ computational methods instead. A recent literature in computer science and algorithmic game theory delivers a number of results on this question. Dughmi (2017) provides an excellent survey.

Dughmi & Xu (2016) analyze algorithmic approaches to the Bayesian persuasion problem with a single Receiver. They deliver two positive results and one negative one. The first positive result concerns environments with strong symmetry across actions and states. In particular, given any action $a \in A$, we can think of both $u(a, \omega)$ and $v(a, \omega)$ as real-valued random variables. Suppose that, for any pair of distinct actions a and a' , we have that $u(a, \omega)$ and $u(a', \omega)$ are IID, and that $v(a, \omega)$ and $v(a', \omega)$ are also IID. In this case, there is a polynomial time algorithm for computing the optimal signal. The second positive result concerns approximately optimal signals and approximately rational behavior by Receiver. If we only require that $\mathbb{E}_\mu u(a^*(\mu), \omega) \geq \mathbb{E}_\mu u(a', \omega) - \epsilon$ for all $a' \in A$, then we can find a signal that delivers Sender a payoff within ϵ of the maximal one using an algorithm that is polynomial in $|A|$ and $\frac{1}{\epsilon}$. The negative result states that, without the aforementioned simplifications, the general problem of computing Sender's maximal payoff is $\#P$ -hard (which basically means it is very hard).²⁴ The literature has also explored computation

a signal that reveals moderate types and hides extreme types. Mensch (2018) establishes necessary and sufficient conditions for optimality of monotone partitions without assuming that Sender's payoff depends only on the mean of Receiver's posterior. Guo & Shmaya (2017) show that when the action space is binary, and Receiver has private information about the state, the optimal signal has a particular structure that they term a nested interval.

²⁴I suspect that economists are likely to have encountered complexity classes P and NP but might be less familiar with $\#P$. Roughly, $\#P$ is the set of counting problems associated with decision problems in NP . For example, asking whether a traveling salesman can visit a set of cities and return home while traversing less than X miles is in class NP , so asking how many distinct paths allow the salesman to visit a set of cities and return home while traversing less than X miles is in class $\#P$. Every problem in $\#P$ is at least as hard as the corresponding problem in NP .

of Bayesian persuasion problems with multiple receivers, but I postpone discussion of those until Section 4.1.

4. EXTENSIONS

There are three main extensions of the basic model: (a) multiple receivers, (b) multiple senders, and (c) dynamic environments.²⁵ Before discussing each of these three, I briefly mention a few other generalizations that have been considered to date.²⁶

A natural and easy extension is to allow for the possibility that Receiver might have some private information.²⁷ This information might be about her own preferences (e.g., Kolotilin 2018, Rayo & Segal 2010) or about the state of the world (e.g., Guo & Shmaya 2017).^{28,29} In both cases, the literature has also examined the possibility that Sender elicits information from Receiver prior to releasing the signal (Kolotilin et al. 2017, Li & Shi 2017).³⁰ Matyskova (2018) analyzes situations where Receiver has no private information at the outset but can gather additional costly information after observing the realization of Sender's signal. She shows that it is without loss of generality to focus on cases where Receiver never actually gathers information on the equilibrium path. The threat of additional information gathering, of course, weakly harms Sender and can be beneficial or harmful for Receiver.

Alonso & Camara (2016a) analyze Bayesian persuasion in situations where Sender and Receiver have heterogeneous priors.^{31,32} They establish a striking result that, as long as Receiver's action

²⁵Of course, these extensions are sometimes combined, both with each other and with the other extensions discussed below. Koessler et al. (2018) examine basic properties of games with multiple senders and multiple receivers. Ely (2017) considers a dynamic model with multiple receivers and derives insights about information policies that reduce the likelihood of bank runs.

²⁶There are also other extensions that space constraints prevent me from discussing. For example, Au & Li (2018) add reciprocity to Receiver's preferences, and Tsakas et al. (2017) provide some commitment power to Receiver by allowing her to reduce her utility from certain actions.

²⁷The literature on Bayes correlated equilibria typically assumes that receivers obtain some exogenous information about the state, but this information is not private: Sender can condition his signal on the realizations of the exogenous signal (see also the discussion in Section 4.1).

²⁸This distinction might seem incoherent at first glance. Since Receiver's utility is $u(a, \omega)$, uncertainty about the state is the same thing as uncertainty about preferences over actions. What we mean when we say that Receiver has private information about her preferences is that the actual state space is some $\Theta = \Omega \times T$, the prior over Ω is independent of the prior over T , and Receiver's utility $u(a, \theta)$ depends on the full state $\theta \in \Theta$, but Sender's signal is a function only of $\omega \in \Omega$, i.e., it cannot depend on $t \in T$.

²⁹Both varieties of the models allow a reformulation to the concavification-friendly form of Equation 2; we simply redefine \hat{v} by integrating private information out (see also Kamenica & Gentzkow 2011).

³⁰Kolotilin et al. (2017) consider the case where Receiver's private information is about her preferences, and Sender observes Receiver's reported type. Li & Shi (2017) consider the case where Receiver (a buyer) has private information about the state, while Sender (a seller) conditions his signal on Receiver's report but does not observe the report prior to setting the price. Kolotilin et al. (2017) show that Sender does not benefit from the ability to elicit Receiver's private information. Li & Shi (2017) show that discriminatory disclosure (sending different signals for different reported types) dominates full disclosure.

³¹Their model allows for the concavification approach, since there is a function that maps Receiver's posterior into Sender's posterior that does not depend on the signal or the signal realization that caused the updating of the beliefs.

³²Kosterina (2018) considers a model with heterogeneous priors and an ambiguity-averse Sender. Specifically, Sender is uncertain about Receiver's prior and thinks that, whatever signal he chooses, Receiver's prior will turn out to be the one that minimizes his payoff. Kosterina characterizes the optimal signal and shows that it has features that are qualitatively different from the optimum in the standard setup. In particular, under

depends only on her posterior mean,³³ Sender generically benefits from his ability to generate a signal. This result is particularly surprising because we might think that, in the case where Receiver’s prior differs from Sender’s in the direction that benefits Sender, the last thing that Sender would want to do is to generate information and thus sober up Receiver, who is about to (mistakenly from Sender’s perspective) take an action that Sender likes. This intuition indeed helps us understand why Sender will not want to generate a fully informative signal, but Π is rich enough that, even when the audience is mistaken in a favorable direction, Sender can benefit from some manipulation of beliefs.

Gentzkow & Kamenica (2014) extend the basic model by making the signals potentially costly for Sender; if Sender generates signal π , then his overall payoff is $v(a, \omega) - c(\pi)$ for some cost function π .³⁴ This model of costly persuasion provides a bridge between the literatures on Bayesian persuasion and rational inattention (e.g., Sims 2003). Both of these literatures consider special cases of costly persuasion: Bayesian persuasion assumes $c(\pi) = 0$, while rational inattention assumes $u = v$. The assumption that $c(\pi)$ has a posterior-separable form, i.e., that it can be expressed as $c(\pi) = \mathbb{E}_{\mu \sim (\pi)}[H(\mu_0) - H(\mu)]$ for some function H , is nearly universal in the literature on rational inattention (where H is typically assumed to be entropy; for discussion, see Caplin et al. 2017, Frankel & Kamenica 2018) and has important implications for the tractability of the costly persuasion model.³⁵ When c is posterior separable, Sender’s optimization problem, which is difficult if expressed analogously to Equation 1,

$$\max_{\pi \in \Pi} (\mathbb{E}_{\omega \sim \mu_0} \mathbb{E}_{s \sim \pi(\omega)} v(a^*(\mu_\pi(\cdot|s)), \omega)) - c(\pi),$$

can be reformulated³⁶ analogously to Equation 2:

$$\max_{\tau} \mathbb{E}_{\mu \sim \tau} [\hat{v}(\mu) + H(\mu)],$$

subject to $\mathbb{E}_{\mu \sim \tau} \mu = \mu_0$. Thus, the concavification approach can be used to derive the optimal signal. When c is not posterior separable, such reformulation is not possible.

ambiguity aversion, it is no longer the case that, when Receiver takes an action that Sender dislikes, she is certain of the state (see also Kamenica & Gentzkow 2011, proposition 4). The ambiguity-averse model can also be reinterpreted as a model in which Sender is trying to convince a group of receivers with heterogeneous priors to unanimously approve his proposal (see also Section 4.1).

³³Note that this is a weaker condition than the one assumed in Section 3.2, since Sender’s preferences are allowed to depend on the state.

³⁴We can also think of situations where Sender’s choice of signals is restricted to some strict subset of Π (e.g., Brocas & Carrillo 2007) as special cases of the costly model, but in these cases, we typically lose the natural monotonicity of the cost function, i.e., the property that, if π is more informative than π' , then $c(\pi) \geq c(\pi')$. Ichihashi (2017) explores how restrictions of Sender’s signal impact Receiver’s welfare. For the case of binary actions, he fully characterizes which Sender’s choice set maximizes Receiver’s utility. Tsakas & Tsakas (2018) explore limitations on signals that arise from exogenous noise in signal realizations. Perez-Richet & Skreta (2018) analyze an environment with a given signal π where Sender chooses a manipulation strategy, i.e., the probability with which, in state ω , signal realizations will be generated according to $\pi(\cdot|\omega')$ rather than $\pi(\cdot|\omega)$. The manipulation strategy is observable to Receiver, so this model is equivalent to a modification where Sender has access to a smaller set of signals; Sender cannot generate every signal that is less informative than π , however, so the set of available signals has a nontrivial structure.

³⁵Function H is referred to as a measure of uncertainty, so posterior-separable functions assume that the cost of a signal is the expected reduction in uncertainty. Gentzkow & Kamenica (2014) show that the baseline level of uncertainty need not be $H(\mu_0)$, but instead can be taken to be the uncertainty at any benchmark belief. Posterior-separable functions satisfy monotonicity (see also Footnote 34) as long as H is concave. All measures of uncertainty with a decision-theoretic foundation are indeed concave.

³⁶I suppress the constant term $-H(\mu_0)$, since it does not affect the optimization problem.

Finally,³⁷ an important class of extensions try to weaken the assumption of Sender's commitment.³⁸ Min (2017), Frechette et al. (2017), and Lipnowski et al. (2018) all consider a model where, after publicly choosing the signal, Sender can, with some probability, change the signal realization.³⁹ Lipnowski et al. (2018) provide the strongest results for this setting.⁴⁰ They show that, as long as Sender's preferences are state independent, a geometric approach similar to concavification can be used to derive the equilibrium.⁴¹ They observe that Receiver might be better off when Sender becomes less credible (i.e., has a higher chance of altering the signal realization).⁴² Sender is obviously always harmed by a reduction in his credibility. Moreover, there is generically a key level of credibility at which Sender's payoff changes discontinuously. Under an additional assumption,⁴³ however, full credibility is not such a threshold; this means that a small departure from the baseline model (i.e., a small chance that Sender can manipulate the signal realization) can only lead to a small reduction in Sender's payoff.

4.1. Multiple Receivers

By far the most important extension of the basic model is to allow for multiple receivers. Space constraints prevent me from giving this case as much attention as it deserves, but the consequences of this omission are lessened by the excellent survey by Bergemann & Morris (2019) that focuses on this case.

There are two classes of multiple receiver environments that are as easy to analyze as the single receiver case. One is when Sender can only send public signals observed by all receivers. In this case, we simply need to reinterpret $a^*(\mu)$ as the vector of (potentially mixed) equilibrium actions when receivers share the posterior μ ; the analysis then proceeds as with a single receiver.⁴⁴ The other case is when Sender can send separate signals to each receiver, each receiver cares only

³⁷A curious pair of papers (Danilov & Lambert-Mogiliansky 2018a,b) modifies the Bayesian persuasion setting by replacing the standard formulation of uncertainty with quantum uncertainty. My knowledge of quantum mechanics is insufficient to understand these papers, but from what I can gather, the idea is not to expand Sender's toolkit (say, by allowing him to quantum entangle some signal realizations), but rather to assume a different (as far as I can tell, an irrational) model of Receiver's belief formation process based on an analog to quantum systems.

³⁸Perez-Richet (2014) and Hedlund (2017) consider settings where Sender chooses the signal after observing some information about the state.

³⁹These models are related to the literature on lying costs (e.g., Guo & Shmaya 2017, Kartik 2009).

⁴⁰Min (2017) shows that commitment helps both Sender and Receiver in Crawford & Sobel's (1982) uniform-quadratic setting. Frechette et al. (2017) conduct laboratory experiments and analyze the extent to which subjects' behavior corresponds to equilibrium predictions. Other explorations of Bayesian persuasion models in the lab include those of Nguyen (2017) and Au & Li (2018).

⁴¹Lipnowski & Ravid (2017) show that, in a model of pure cheap talk, if Sender's preferences do not depend on the state, then his maximal equilibrium payoff is determined by the quasiconcavification of \hat{v} . (Recall that concavification is the smallest concave function above \hat{v} . Quasiconcavification is the smallest quasiconcave function above \hat{v} .) Lipnowski et al. (2018) extend this result to Bayesian persuasion with limited commitment by showing that Sender's payoff can be characterized by an object that combines the concavification and the quasiconcavification. (Alas, it is not a simple convex combination of the two.)

⁴²This echoes the aforementioned results by Ichihashi (2017) on the potential benefit to Receiver of restricting the set of signals available to Sender.

⁴³The assumption is that Sender does not need to rule out any states to obtain his preferred action.

⁴⁴The fact that actions might be mixed means that we need to set $|S| \geq |\Omega|$ (see also the discussion in Footnote 5). This does mean that, when the state space is large, the problem can be difficult. Bhaskar et al. (2016) establish that, if two receivers engage in a zero-sum game and Sender wishes to maximize a weighted sum of the receivers' payoffs, then computing the optimal public signal is *NP*-hard. On the more positive side, Cheng et al. (2015) introduce a class of games where approximately optimal signals can be computed in polynomial time.

about her own action, and Sender's utility is separable across receivers' actions. Then, Sender can determine the optimal signal receiver by receiver and faces a set of independent problems of a single-receiver variety. If Sender can send separate signals to each receiver, and if either (a) a receiver's optimal action depends on what other receivers do or (b) Sender's utility is not separable across receiver's actions,⁴⁵ then the problem becomes significantly more difficult and cannot be expressed in the form equivalent to the single receiver case. Of course, the most general case allows for both of these possibilities.

One approach to Sender's optimization problem in the general version takes a two-step approach (Bergemann & Morris 2016b, Taneva 2016). The first step is to characterize the set of all outcomes (joint distributions over the state and receivers' actions) that can be attained by some signal.⁴⁶ This set of outcomes is referred to as the set of Bayes correlated equilibria (Bergemann & Morris 2013, 2016a).⁴⁷ Identifying the set of Bayes correlated equilibria is often of interest even in the absence of the second, optimization stage. This set describes the outcomes that we would think might arise in a game if we are agnostic about the information obtained by the players.⁴⁸ Moreover, it pinpoints the worst-case scenario for a given game and thus aids informationally robust mechanism design (see also Bergemann et al. 2017b, 2018b; Brooks & Du 2018).⁴⁹ To turn the analysis of Bayes correlated equilibria into information design, we simply select the best equilibrium given some objective function.⁵⁰ A different approach to the general class of information design problems with multiple receivers is provided by Mathevet et al. (2018). They propose a procedure where we first identify the optimal purely private signal for every prior μ .⁵¹ Denoting the payoff from such a signal by $\hat{v}(\mu)$ if the common prior were μ , we can compute the optimal public signal via the concavification of \hat{v} as in Equation 2. Combining the optimal public signal with the optimal private signals (which are contingent on the realization of the public signal) then yields the overall optimum. Unlike the approach via Bayes correlated equilibria, the method proposed by Mathevet et al. (2018) is applicable even when Sender is concerned that equilibrium selection may not be in his favor.

⁴⁵Babichenko & Barman (2016) provide results on the computation difficulty of the model where the receiver's optimal action does not depend on the actions of other receivers, but Sender's payoff is not separable across receivers' actions. Suppose that $\Omega = \{0, 1\}$, each receiver's action space $A_i = \{0, 1\}$, and Sender's payoff $V(Q)$ depend on the set of receivers Q who take action $a_i = 1$. This setup includes voting models such as that of Alonso & Camara (2016b). Babichenko & Barman (2016) show that, if Sender's utility is submodular [in the sense that $V(Q \cup \{i\}) - V(Q) \geq V(Q' \cup \{i\}) - V(Q')$ for every $Q' \subset Q$ and every receiver i], then computing the optimal signal is *NP*-hard, but computing a signal that yields a payoff of at least $(1 - 1/e)$ times the maximal one can be done in polynomial time.

⁴⁶By the logic analogous to the revelation principle, this involves deriving a set of linear constraints, each of which requires that a receiver prefers an action recommended by her signal realization over all other actions.

⁴⁷The literature on Bayes correlated equilibria also allows for receivers to obtain some exogenous information about the state that is not private, i.e., that is observable by Sender.

⁴⁸Bergemann et al. (2015) establish a striking, beautiful result that characterizes what combinations of firm profits and consumer surplus might arise as we span all possible signals that the firm might observe about the consumer's valuation. Bergemann et al. (2017a) characterize combinations of revenue and bidder surplus that might arise in a first-price auction as we span all possible signals that each bidder might observe about her and other bidders' valuations.

⁴⁹By informationally robust mechanism design, I mean choosing the mechanism that yields the best possible outcome under the worst-case signal.

⁵⁰This decomposition requires the assumption that equilibrium selection works in Sender's favor. Otherwise, not every Bayes correlated equilibrium is attainable.

⁵¹Computing this signal requires optimizing over the set of minimal consistent distributions over belief hierarchies. A consistent distribution over hierarchies is one where receivers' beliefs arise from a common prior. A consistent distribution b is minimal if there is no consistent distribution b' such that the support of b' is a strict subset of the support of b .

Much of the research on information design with multiple receivers has focused on specific applications.⁵² For instance, several papers analyze the problem of persuading voters. Alonso & Camara (2016b) and Kosterina (2018) restrict Sender to public signals.⁵³ Wang (2015) restricts Sender to IID signals. Arieli & Babichenko (2016), Bardhi & Guo (2018), and Chan et al. (2018) allow for arbitrary private signals.⁵⁴ While each of these papers makes a nice contribution on its own, taken together, they deliver a laundry list of results rather than a coherent picture of how to persuade voters. One of the challenges facing the Bayesian persuasion literature going forward will be to synthesize existing results, rather than simply add to the list of cases that have been considered.⁵⁵

4.2. Multiple Senders

Bayesian persuasion models with multiple senders have proved especially useful in engaging issues of competition in the marketplace of ideas, an important topic that far predates the idea of information design.⁵⁶ A long tradition in political and legal thought places special emphasis on competition in information provision. A widely held view that competition increases information revelation has motivated protection of freedom of speech and freedom of the press, media ownership regulation, the adversarial judicial system, and many other policies (Gentzkow & Kamenica 2017b, Gentzkow & Shapiro 2008).

Several papers in information economics provide support for this view (e.g., Battaglini 2002, Milgrom & Roberts 1986, Shin 1998), but Bayesian persuasion-style models have proved especially tractable for exploring the issue. Gentzkow & Kamenica (2017a) consider a model that closely follows the basic model, with the only change being that, instead of a Sender with utility $v(a, \omega)$, they have a set of senders indexed by i , each with utility $v_i(a, \omega)$. All senders simultaneously choose a signal.⁵⁷ Receiver observes the signal realizations from all senders prior to choosing her action. This description of the model, however, leaves an important issue unspecified: If sender i chooses some signal π , and sender j also chooses π , then how much information does Receiver obtain? In particular, are the two signals redundant, or does Receiver get two IID draws from π ?

⁵²Das et al. (2017) apply information design to reducing congestion. They provide a simple example where a suitably designed IID signal about the state of traffic can fully eliminate congestion externalities. Congestion games seem to be a particularly fertile ground for future work on information design. If crucial data about traffic are user generated, then network externalities might lead to the rise of a monopolist provider of routing software that could afford to withhold some information from its users without the fear of being displaced by a competitor. Such a monopolist could raise everyone's welfare by providing drivers with suitably designed partial information about the route that is optimal for them. This information design problem is unlikely to be amenable to analytic solutions, but development and deployment of welfare-maximizing, dynamic information-provision algorithms might greatly reduce time wasted in traffic.

⁵³Alonso & Camara (2016b) assume that voters vary in their preferences. Kosterina (2018) assumes that they vary in their priors.

⁵⁴A receiver's payoff is assumed to depend on her own action (Arieli & Babichenko 2016), the collective action (Bardhi & Guo 2018), or both (Chan et al. 2018).

⁵⁵This article fails in this respect, as it has turned out to be more of a survey and less of a synthesis than one might have hoped.

⁵⁶I limit my discussion to papers that focus on the impact of competition on information revelation. Papers with multiple senders that focus on other topics include those of Ostrovsky & Schwarz (2010), Brocas et al. (2012), Gul & Pesendorfer (2012), and Boleslavsky et al. (2017).

⁵⁷Li & Norman (2018a,b) analyze the impact of competition under rich signal spaces when senders move sequentially. They show that adding an additional sender who moves first cannot result in a strictly less informative equilibrium, but adding a sender to the middle or the end of the lineup can. They also show that an equilibrium of the sequential move game can never be strictly less informative than the equilibrium of the simultaneous move game with the same set of senders.




Figure 4

Signals as partitions of $\Omega \times [0, 1]$.

To specify a Bayesian persuasion model with multiple senders, we need to reformulate our definition of the set of all signals in a way that will eliminate such ambiguity. The following formulation does the job.⁵⁸

The set of signal realizations S is the set of (Lebesgue-measurable)⁵⁹ subsets of $\Omega \times [0, 1]$. Thus, a signal realization $s \in S$ is a subset of $\Omega \times [0, 1]$. A signal π is a partition of $\Omega \times [0, 1]$, with each element of the partition in S . The interpretation is that a random variable drawn uniformly from $[0, 1]$ determines the signal realization conditional on the state; the probability of observing $s \in \pi$ when the state is ω is the probability that this uniform random variable lands inside $\{x \in [0, 1] | (\omega, x) \in s\}$. **Figure 4** illustrates this formalism. In this example, we have $\Omega = \{L, R\}$ and $\pi_1 = \{l, r\}$ where $l = (L, [0.07]) \cup (R, [0, 0.3])$ and $r = (L, [0.7, 1]) \cup (R, [0.3, 1])$. Signal π_1 is a partition of $\Omega \times [0, 1]$, and the state-specific likelihood of signal realizations is $Pr(l|L) = Pr(r|R) = 0.7$. This formulation immediately yields the joint informational content of multiple signals. The set of partitions has a lattice structure, and we can define the \vee operator that denotes the coarsest common refinement of two signals. In **Figure 4**, $\pi_1 \vee \pi_2$ is the signal that yields the same information as observing both signal π_1 and signal π_2 . In this sense, the operator \vee adds signals together. Given a vector of signals $\pi = (\pi_1, \dots, \pi_n)$, we write $\vee \pi$ for $\pi_1 \vee \pi_2 \dots \vee \pi_n$. Note that the set of all signals is quite rich in the sense that every sender can choose any signal whatsoever, including one the realizations of which are arbitrarily correlated with signal realizations of other senders. This richness assumption is in the spirit of the basic model, which does not impose any restrictions on signals that could be generated.⁶⁰

We now have an easy way of describing how each sender's payoff depends on the vector of signals chosen by the senders. As above, we can denote sender i 's payoff when Receiver's posterior is μ by $\hat{v}_i(\mu)$, and we can denote the distribution of posteriors induced by signal π by $\langle \pi \rangle$. Thus, if the senders' strategy profile is $\pi = (\pi_1, \dots, \pi_n)$, then sender i 's payoff is $\mathbb{E}_{\mu \sim (\vee \pi)} \hat{v}_i(\mu)$. This means

⁵⁸Green & Stokey (1978) first proposed this definition of a signal. Gentzkow & Kamenica (2017a) employed it to study Bayesian persuasion with multiple senders.

⁵⁹A set $X \subset [0, 1]$ is Lebesgue measurable if it is meaningful to ask what is the probability that a random variable uniformly distributed on $[0, 1]$ will fall inside X . A set $s \subset \Omega \times [0, 1]$ consists of a collection of subsets of $[0, 1]$, one for each state. The set s is Lebesgue measurable if each element of this collection is Lebesgue measurable.

⁶⁰Despite the richness, it is not clear whether, under this formalism, if we have n receivers, we can induce any coherent hierarchy of beliefs (see also Brandenburger & Dekel 1993, Mertens & Zamir 1985) for the receivers simply by having each of them observe a realization of some signal. This question is not relevant for this analysis, since we have a single Receiver who observes all signals, but the question seems interesting in its own right.

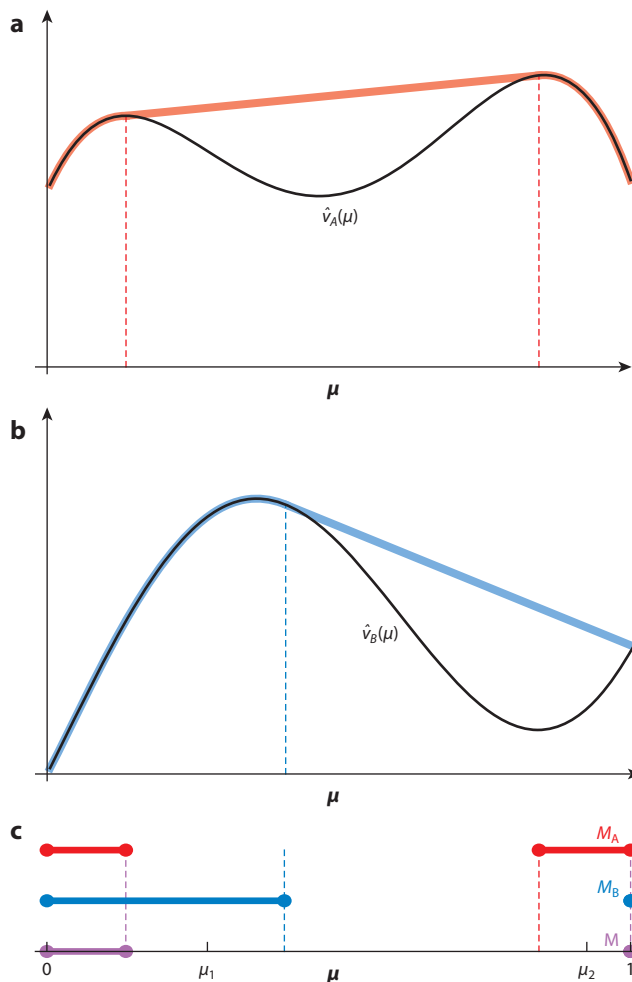


Figure 5

Characterizing equilibrium outcomes. (a) \hat{v} and its concavification for Sender *A* (thick red line). (b) \hat{v} and its concavification for Sender *B* (thick blue line). (c) Sets of beliefs where \hat{v} coincides with its concavification. The thick purple line indicates the intersection of the red set and the blue set.

that a profile π is a (pure strategy)⁶¹ equilibrium if $\mathbb{E}_{\mu \sim (\nu, \pi)} \hat{v}_i(\mu) \geq \mathbb{E}_{\mu \sim (\pi'_i, \nu, \pi_{-i})} \hat{v}_i(\mu)$ for all signals π'_i . If π is an equilibrium, then we say that $\tau = \langle \pi \rangle$ is an equilibrium outcome.

Characterizing the set of equilibrium outcomes turns out to be surprisingly easy. The approach again draws on concavification. As long as there are multiple senders, a Bayes-plausible distribution of posteriors τ is an equilibrium outcome if and only if, for every belief μ in the support of τ and for every sender i , \hat{v}_i coincides with its concavification at μ (Gentzkow & Kamenica 2017a). **Figure 5** illustrates how this result can be used in practice. Suppose that there are two senders *A* and *B* with value functions \hat{v}_A and \hat{v}_B . Then, if we denote by M_i the set of beliefs where \hat{v}_i coincides with its concavification, we know that a distribution of posteriors τ is an equilibrium outcome if and only if its support lies in $M = M_A \cap M_B$.

⁶¹For now, I focus exclusively on pure strategy equilibrium. This is a substantive restriction for reasons that I discuss below.

Let us return to the question from the beginning of this section, namely whether competition increases the amount of information revealed. There are a few things that we might mean by that question, but one natural interpretation would be to compare a collusive outcome, which would arise if senders jointly maximized their welfare, with the competitive, equilibrium outcome.⁶² Specifically, suppose that τ^c is the unique distribution of posteriors that maximizes the sum of sender's utilities, i.e., $\tau^c = \operatorname{argmax}_{\tau} \mathbb{E}_{\mu \sim \tau} \sum_i \hat{v}_i(\mu)$.⁶³ Suppose that τ^* is some equilibrium outcome. Can we conclude that τ^* is more informative than τ^c ?⁶⁴ This strong version of the conjecture that competition increases information revelation turns out not to be true: It could happen that τ^c and τ^* are not comparable, so that, for some utility functions u , Receiver is better off under collusion, and for other utility functions, she is better off under competition. We can establish a somewhat weaker form of the conjecture, however. In particular, if τ^c and τ^* are comparable, then it must be the case that τ^* is more informative than τ^c . In other words, no matter what preferences senders have, it can never be the case that the collusive outcome is strictly more informative than an equilibrium outcome. This analysis thus provides some qualified support for the common view about competition and information that has played such an important role in shaping public policy.

Of course, the rich signal space assumed in the analysis above is rather special, so we would like to know whether the aforementioned conclusions generalize. Suppose that each sender i has access to some set of signals Π_i . We refer to $\mathbf{\Pi} = (\Pi_1, \dots, \Pi_n)$ as the informational environment. Other than relaxing the assumption that each Π_i is the rich set of signals, we keep the model the same as above. Can we still conclude that the collusive outcome can never be strictly more informative than any equilibrium? The answer turns out to be: It depends.

Without the rich signal space, we no longer have the \vee operator, but we can still talk about the distribution of posteriors induced by any set (or vector) of signals. We denote by $\tau = \langle \cup \pi_i \rangle$ the distribution of posteriors induced by the set $\{\pi_1, \dots, \pi_n\}$. We say that an outcome τ is feasible if there exist $(\pi_1, \dots, \pi_n) \in \mathbf{\Pi}$ such that $\tau = \langle \cup \pi_i \rangle$. We say that the informational environment is Blackwell connected if, for any feasible τ , any sender i , and any $\pi_{-i} \in \mathbf{\Pi}_{-i}$ such that τ is more informative than $\langle \cup_{j \neq i} \pi_j \rangle$, there exists a $\pi \in \Pi_i$ such that $\tau = \langle \pi \cup (\cup_{j \neq i} \pi_j) \rangle$. In other words, an information environment is Blackwell connected if, given any strategy profile, each sender can unilaterally deviate to induce any feasible outcome that is more informative. Importantly, if each Π_i is the rich set of signals (as in the model considered above), then the environment is Blackwell connected. This turns out to be the key feature of the rich signal space that leads to our conclusions. Gentzkow & Kamenica (2017b) show that the collusive outcome cannot be strictly more informative than an equilibrium (regardless of senders' preferences) if and only if the information environment is Blackwell connected.⁶⁵ This result further sharpens our understanding of the circumstances under which we can safely assume that competition cannot reduce the amount of information revealed.⁶⁶

⁶²Gentzkow & Kamenica (2017a,b) also analyze the impact of introducing additional senders or increasing the misalignment of senders' preferences.

⁶³This argmax will generically be unique. If we allow for multiple collusive outcomes, then analogous results can be stated using orders on sets (see also Gentzkow & Kamenica 2017a,b).

⁶⁴The Blackwell order is usually defined as an order over signals, not over distribution of posteriors. I abuse the terminology somewhat and say that τ is more informative than τ' if a signal that induces τ is Blackwell more informative than a signal that induces τ' , i.e., if τ is a mean-preserving spread of τ' .

⁶⁵Gentzkow & Kamenica (2017b) establish this result in a slightly more general setting where senders can have an arbitrary utility over induced distribution of posteriors, not necessarily one that takes the form of $\mathbb{E}_{\mu \sim \tau} \hat{v}_i(\mu)$. Their proof of the only-if direction utilizes preferences that do not conform to the $\mathbb{E}_{\mu \sim \tau} \hat{v}_i(\mu)$ formulation. It is an open question whether there is a weaker condition on the information environment that suffices for the result to hold when each sender's utility is some $\mathbb{E}_{\mu \sim \tau} \hat{v}_i(\mu)$.

⁶⁶This result also explains why it is important to focus on pure strategy equilibria. Once senders use mixed strategies, the environment immediately fails to be Blackwell connected. Consequently, even with the rich

Board & Lu (2018) analyze a closely related question in a search setting. Senders are sellers who all offer the same product.⁶⁷ There is a binary state of the world that determines whether the buyer is better off buying the product; price is taken to be exogenous. At the prior, the buyer prefers not to buy. Sellers wish to sell the product regardless of the state. The buyer approaches sellers one by one, paying a search cost $c > 0$ to visit each additional seller. When a buyer visits a particular seller, the seller chooses a signal about the state to reveal to the buyer. The buyer then decides whether to purchase the product from this seller, visit another seller, or exit the market. Board & Lu (2018) show that, if each seller has access to the rich signal space (and can thus correlate with signal realizations of other sellers), and the buyer's belief is observable when she visits a seller,⁶⁸ then competition will not have any bite: The unique equilibrium is one where each seller generates the same signal that he would generate if he were the only seller in town. However, if sellers cannot perfectly correlate their signal realizations, and the buyer's belief is not observable, then the unique equilibrium as c converges to zero leads to full revelation of the state. It is quite interesting that, in the static setting, senders' access to the rich signal space suffices to ensure that competition induces information revelation, whereas in the search setting, senders' ability to correlate signals reduces the impact of competition.⁶⁹

4.3. Dynamics

Let us add time to the basic model. If the state of the world and the set of available actions do not change, and the players are patient (so the payoffs do not depend on when the action is taken), then dynamics do not matter. Any sequence of signals (including those where signal sent in period t depends on the signal realization in period $t - 1$) in the end induces some Bayes-plausible distribution of posteriors τ , so Sender might as well just induce his preferred τ at the outset.

Dynamics become important⁷⁰ when state evolves over time, past behavior influences current opportunities, and/or Sender and Receiver have conflicting preferences on the optimal timing of Receiver's action.⁷¹ Ely (2017) derives the optimal information policy (a map from the history and the current state to a distribution over signal realizations) in environments where the state evolves through a Markov process and Receiver is myopic (in each period, she chooses an action that maximizes her contemporaneous utility given her current belief).⁷² For example, suppose that Receiver is an employee who is either ready for promotion ($\omega = 1$) or not ($\omega = 0$) and can either request the promotion ($a = 1$) or not ($a = 0$). Time is continuous. It is common knowledge that $\omega = 0$ at the outset and that the state transitions to $\omega = 1$ at Poisson rate λ . It is optimal

signal space, it is possible to have a collusive outcome that is strictly more informative than a mixed strategy equilibrium.

⁶⁷Board & Lu (2018) consider a more general setting. I focus on a special case of their model.

⁶⁸One interpretation of this would be that the buyer and sellers are in an online setting, and cookies in the buyer's browser reveal all of the information that she has received thus far.

⁶⁹Levy et al. (2018) consider a receiver who suffers from correlation neglect. A monopolist owner of multiple news outlets can, by suitably correlating signal realizations of the various outlets, manipulate Receiver's belief in a way that violates Bayes plausibility. Competition is assumed to reduce the ability to correlate signal realizations and thus benefits Receiver.

⁷⁰Space constraints prevent me from covering the entire literature on dynamic Bayesian persuasion. In addition to the papers discussed below, important contributions to this literature include those of Horner & Skrzypacz (2016) and Henry & Ottaviani (2019).

⁷¹Ely et al. (2015) put belief dynamics directly into the players' utility function. They postulate that entertainment utility stems from suspense (variance of next period's beliefs) and surprise (realized movement of beliefs) and analyze how to entertain a Bayesian audience. They apply this analysis to the design of mystery novels, engaging political primaries, casino gambling, game shows, charity auctions, and sports.

⁷²Renault et al. (2017) also analyze a case of this model.

for the worker to ask for the promotion if and only if her belief that she is ready is above some cutoff p . Sender is the firm that always prefers that employees not ask for promotion. Via its design of employee evaluations, feedback procedures, and transparency, the firm can implement any information policy. If the firm provides no information, then the employee's belief that she is ready will drift upward; at time t , she will think the probability that she is ready is $1 - e^{-\lambda t}$. Consequently, regardless of the true state, she will ask for promotion at $t = [-\ln(1 - p)]/\lambda$. If the firm is fully transparent, then the employee will ask for promotion exactly when she is ready for it, which on average happens at time $t = 1/\lambda$. This means that when the employee is very eager (p is low), it is better to be fully transparent; otherwise, providing no information is better. It turns out, however, that neither of these policies is ever optimal. The optimal policy is to reveal to the employee that she has become ready with some deterministic delay.⁷³ It is rather remarkable that the optimal policy is so simple.

Ely & Szydlowski (2019) examine a setting where Receiver's past decisions influence her current opportunities.⁷⁴ Suppose that Receiver is again an employee concerned about her promotion, but this time promotion is a function of employee's effort, not of her type. Specifically, time is continuous, at every moment the employee chooses whether to continue working (incurring a constant cost of effort), and she is eligible for promotion once the time that she spent working exceeds some unknown state ω . Sender is a firm that designs its information policy to maximize employee's effort (reaping benefits even if the employee continues working beyond the level required for promotion). In this case, the state is assumed to remain constant over time, but dynamics still matter for two reasons. First, the firm can use future information provision as an incentive to get the worker to exert current effort.⁷⁵ Second, after the employee has sunk some effort, provision of additional information can induce her to continue working; thus, the employee might exert more overall effort than she would have been willing to do at the outset. Ely & Szydlowski (2018) characterize the optimal policy and demonstrate that (for the two aforementioned reasons) dynamic information provision improves upon what could be achieved with static information design.

The model of Ely & Szydlowski (2019) can be seen as a special case of a broader class of environments.⁷⁶ Suppose that, in period 0, Sender chooses a signal π_0 , and Receiver, after observing its realization, selects an element of some partition P_0 of the action space A . Denote Receiver's choice by A_1 . The interpretation is that A_1 is the set of actions that will remain available to Receiver. In period 1, sender chooses a signal π_1 , and Receiver, after observing its realization, selects an element of some partition P_1 of A_1 , and so on. In the (discrete time version of) Ely & Szydlowski's (2019) model, the action set A is the total number of periods worked and $P_t = \{t, \{t + 1, t + 2, \dots\}\}$ until the employee quits. In their context, the nature of sunk costs determines this particular partition structure, but in other applications, one could imagine that specifying the partitions in each period is part of the design problem. For example, a college can control both how much information students get about their aptitude and the rules about when students have to commit to their major. The legal system can determine whether a prosecutor can provide additional information after the conviction but before the sentencing. Such combinations of designing narrowing paths of options alongside designing information seem like potentially fertile ground for future research.

⁷³As long as the delay is shorter than $[-\ln(1 - p)]/\lambda$, the employee will not ask for promotion until she has been informed that she is ready. Consequently, setting the delay to $[-\ln(1 - p)]/\lambda$ maximizes the firm's payoff.

⁷⁴Using a similar modeling approach, Smolin (2017) derives optimal feedback that employees should be given about their past performance.

⁷⁵The firm can fully commit to her information policy. Orlov et al. (2018a) and Bizzotto et al. (2018) analyze settings where Sender lacks commitment in the sense that he must choose a sequentially rational signal.

⁷⁶These observations come from a discussion with Andy Skrzypacz and Ilan Kremer.

DISCLOSURE STATEMENT

The author is not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

I am grateful to University of Chicago Booth School of Business for funding.

LITERATURE CITED

- Alonso R, Camara O. 2016a. Bayesian persuasion with heterogeneous priors. *J. Econ. Theory* 165:672–706
- Alonso R, Camara O. 2016b. Persuading voters. *Am. Econ. Rev.* 106(11):3590–605
- Arieli I, Babichenko Y. 2016. *Private Bayesian persuasion*. Work. Pap., Technion Isr. Inst. Technol., Haifa
- Au PH, Li KK. 2018. *Bayesian persuasion and reciprocity: theory and experiment*. Work. Pap., Nanyang Technol. Univ., Singapore
- Aumann RJ, Maschler MB. 1966. *Game theoretic aspects of gradual disarmament*. Rep. ST-80, Chapter 5, US Arms Control Disarm. Agency, Washington, DC
- Aumann RJ, Maschler MB. 1995. *Repeated Games with Incomplete Information*. Cambridge, MA: MIT Press
- Babichenko Y, Barman S. 2016. *Computational aspects of private Bayesian persuasion*. Work. Pap., Technion Isr. Inst. Technol., Haifa
- Bardhi A, Guo Y. 2018. Modes of persuasion toward unanimous consent. *Theor. Econ.* 13:1111–50
- Battaglini M. 2002. Multiple referrals and multidimensional cheap talk. *Econometrica* 70:1379–401
- Bergemann D, Bonatti A. 2019. Markets for information: an introduction. *Annu. Rev. Econ.* 11:85–107
- Bergemann D, Bonatti A, Smolin A. 2018a. The design and price of information. *Am. Econ. Rev.* 108(1):1–48
- Bergemann D, Brooks B, Morris S. 2015. The limits of price discrimination. *Am. Econ. Rev.* 105(3):921–57
- Bergemann D, Brooks B, Morris S. 2017a. First-price auctions with general information structures: implications for bidding and revenue. *Econometrica* 85(1):107–43
- Bergemann D, Brooks B, Morris S. 2017b. *Informationally robust optimal auction design*. Work. Pap., Yale Univ., New Haven, CT
- Bergemann D, Brooks B, Morris S. 2018b. *Revenue guarantee equivalence*. Work. Pap., Yale Univ., New Haven, CT
- Bergemann D, Morris S. 2013. Robust predictions in games with incomplete information. *Econometrica* 81(4):1251–308
- Bergemann D, Morris S. 2016a. Bayes correlated equilibrium and the comparison of information structures in games. *Theor. Econ.* 11:487–522
- Bergemann D, Morris S. 2016b. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *Am. Econ. Rev.* 106(5):586–91
- Bergemann D, Morris S. 2019. Information design: a unified perspective. *J. Econ. Lit.* 57:1–57
- Best J, Quigley D. 2017. *Persuasion for the long run*. Work. Pap., Oxford Univ., UK
- Bhaskar U, Cheng Y, Ko YK, Swamy C. 2016. Hardness results for signaling in Bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 479–96. New York: ACM
- Bizzotto J, Rudiger J, Vigier A. 2018. *Dynamic persuasion with outside information*. Work. Pap., Univ. Oslo, Nor.
- Blackwell D. 1951. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pp. 93–102. Berkeley, CA: Univ. Calif. Press
- Blackwell D, Girshick MA. 1954. *Theory of Games and Statistical Decisions*. New York: Wiley
- Board S, Lu J. 2018. Competitive information disclosure in search markets. *J. Political Econ.* 126(5):1965–2010
- Boleslavsky R, Cotton C. 2015. Grading standards and education quality. *Am. Econ. J. Microecon.* 7(2):248–79
- Boleslavsky R, Cotton C, Gurnani H. 2017. Demonstrations and price competition in new product release. *Manag. Sci.* 63(6):2016–26
- Brady v. Maryland*, 373 US 83, 87 (1963)

- Brandenburger A, Dekel E. 1993. Hierarchies of beliefs and common knowledge. *J. Econ. Theory* 59:189–98
- Brocas I, Carrillo JD. 2007. Influence through ignorance. *RAND J. Econ.* 38(4):931–47
- Brocas I, Carrillo JD, Palfrey T. 2012. Information gatekeepers: theory and experimental evidence. *Econ. Theory* 51:649–76
- Brooks B, Du S. 2018. *Optimal auction design with common values: an informationally-robust approach*. Work. Pap., Univ. Chicago
- Caplin A, Dean M, Leahy J. 2017. *Rationally inattentive behavior: characterizing and generalizing Shannon entropy*. Work. Pap., New York Univ.
- Chan J, Gupta S, Li F, Wang Y. 2018. *Pivotal persuasion*. Work. Pap., New York Univ.
- Cheng Y, Cheung HY, Dughmi S, Emamjomeh-Zadeh E, Han L, Teng S-H. 2015. Mixture selection, mechanism design, and signaling. In *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science*, pp. 1426–45. Piscataway, NJ: IEEE
- Crawford V, Sobel J. 1982. Strategic information transmission. *Econometrica* 50(6):1431–51
- Danilov VI, Lambert-Mogiliansky A. 2018a. *Preparing a (quantum) belief system*. Work. Pap., Paris School Econ.
- Danilov VI, Lambert-Mogiliansky A. 2018b. Targeting in quantum persuasion problems. *J. Math. Econ.* 78:142–49
- Das S, Kamenica E, Mirka R. 2017. Reducing congestion through information design. In *Proceedings of the 55th Allerton Conference on Communication, Control, and Computing*, pp. 1279–84. Monticello, IL: Allerton
- Duffie D, Dworzak P, Zhu H. 2017. Benchmarks in search markets. *J. Finance* 72(5):1983–2044
- Dughmi S. 2017. Algorithmic information structure design. *ACM SIGecom Exch.* 15(2):2–24
- Dughmi S, Xu H. 2016. Algorithmic Bayesian persuasion. In *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC)*, pp. 412–25. New York: ACM
- Dworczak P, Martini G. 2019. The simple economics of optimal persuasion. *J. Political Econ.* In press
- Ely J. 2017. Beeps. *Am. Econ. Rev.* 107(1):31–53
- Ely J, Frankel A, Kamenica E. 2015. Suspense and surprise. *J. Political Econ.* 103(1):215–60
- Ely J, Szydlowski M. 2019. Moving the goalposts. *J. Political Econ.* In press
- Feng X, Lu J. 2016. The optimal disclosure policy in contests with stochastic entry: a Bayesian persuasion perspective. *Econ. Lett.* 147:103–7
- Frankel A, Kamenica E. 2018. *Quantifying information and uncertainty*. Work. Pap., Univ. Chicago
- Frechette G, Lizzeri A, Perego J. 2017. *Rules and commitment in communication*. Work. Pap., New York Univ.
- Garcia D, Tsur M. 2018. *Information design in competitive insurance markets*. Work. Pap., Univ. Vienna
- Gehlbach S, Sonin K. 2014. Government control of the media. *J. Public Econ.* 118:163–71
- Gentzkow M, Kamenica E. 2014. Costly persuasion. *Am. Econ. Rev.* 104(5):457–62
- Gentzkow M, Kamenica E. 2016. A Rothschild-Stiglitz approach to Bayesian persuasion. *Am. Econ. Rev.* 106(5):597–601
- Gentzkow M, Kamenica E. 2017a. Bayesian persuasion with multiple senders and rich signal spaces. *Games Econ. Behav.* 104:411–29
- Gentzkow M, Kamenica E. 2017b. Competition in persuasion. *Rev. Econ. Stud.* 84:300–22
- Gentzkow M, Kamenica E. 2017c. Disclosure of endogenous information. *Econ. Theory Bull.* 5(1):47–56
- Gentzkow M, Shapiro JM. 2008. Competition and truth in the market for news. *J. Econ. Perspect.* 22:133–54
- Goldstein I, Leitner Y. 2018. Stress tests and information disclosure. *J. Econ. Theory* 117:34–69
- Green J, Stokey N. 1978. *Two representations of information structures and their comparisons*. Tech. Rep. 271, Inst. Math. Stud. Soc. Sci., Stanford Univ., Stanford, CA
- Grossman S. 1981. The informational role of warranties and private disclosure about product quality. *J. Law Econ.* 24(3):461–83
- Gul F, Pesendorfer W. 2012. The war of information. *Rev. Econ. Stud.* 79:707–34
- Guo Y, Shmaya E. 2017. *The interval structure of optimal disclosure*. Work. Pap., Northwestern Univ., Evanston, IL
- Habibi A. 2018. *Motivation and information design*. Work. Pap., Univ. Coll. London
- Hedlund J. 2017. Bayesian persuasion by a privately informed sender. *J. Econ. Theory* 167:229–68
- Henry E, Ottaviani M. 2019. Research and approval process: the organization of persuasion. *Am. Econ. Rev.* 109:911–55

- Hernandez P, Neeman Z. 2017. *How Bayesian persuasion can help reduce illegal parking and other socially undesirable behavior*. Work. Pap., Univ. Valencia, Spain
- Horner J, Skrzypacz A. 2016. Selling information. *J. Political Econ.* 124(6):1515–62
- Ichihashi S. 2017. *Limiting sender's information in Bayesian persuasion*. Work. Pap., Stanford Univ., Stanford, CA
- Inostroza N, Pavan A. 2018. *Persuasion in global games with application to stress testing*. Work. Pap., Northwestern Univ., Evanston, IL
- Ivanov M. 2015. *Optimal signals in Bayesian persuasion mechanisms with ranking*. Work. Pap., McMaster Univ., Hamilton, Can.
- Jehiel P. 2015. On transparency in organizations. *Rev. Econ. Stud.* 82(2):736–61
- Kamenica E. 2012. Behavioral economics and psychology of incentives. *Annu. Rev. Econ.* 4:427–52
- Kamenica E, Gentzkow M. 2011. Bayesian persuasion. *Am. Econ. Rev.* 101(6):2590–615
- Kartik N. 2009. Strategic communication with lying costs. *Rev. Econ. Stud.* 76(4):1359–95
- Kastl J, Pagnozzi M, Piccolo S. 2018. Selling information to competitive firms. *RAND J. Econ.* 49(1):254–82
- Koessler F, Laclau M, Tomala T. 2018. *Interactive information design*. Work. Pap., Paris School Econ.
- Kolotilin A. 2015. Experimental design to persuade. *Games Econ. Behav.* 90:215–26
- Kolotilin A. 2018. Optimal information disclosure: a linear programming approach. *Theor. Econ.* 13:607–35
- Kolotilin A, Mylovannov T, Zapechelnyuk A, Li M. 2017. Persuasion of a privately informed receiver. *Econometrica* 85(6):1949–64
- Kosterina S. 2018. *Persuasion with unknown beliefs*. Work. Pap., Princeton Univ., Princeton, NJ
- Kremer I, Mansour Y, Perry M. 2014. Implementing the “wisdom of the crowd.” *J. Political Econ.* 122(5):988–1012
- Lazear E. 2006. Speeding, terrorism, and teaching to the test. *Q. J. Econ.* 121(3):1029–61
- Levy G, Moreno de Barreda I, Razin R. 2018. *Persuasion with correlation neglect: media power via correlation of news content*. Discuss. Pap. DP12640, Cent. Econ. Policy Res., Washington, DC
- Lewis TR, Sappington DEM. 1994. Supplying information to facilitate price discrimination. *Int. Econ. Rev.* 35(2):309–27
- Li C. 2017. A model of Bayesian persuasion with transfers. *Econ. Lett.* 161:93–95
- Li F, Norman P. 2018a. On Bayesian persuasion with multiple senders. *Econ. Lett.* 170:66–70
- Li F, Norman P. 2018b. *Sequential persuasion*. Work. Pap., Univ. North Carolina, Chapel Hill
- Li H, Shi X. 2017. Discriminatory information disclosure. *Am. Econ. Rev.* 107(11):3363–85
- Lipnowski E, Ravid D. 2017. *Cheap talk with a fixed agenda*. Work. Pap., Univ. Chicago
- Lipnowski E, Ravid D, Shishkin D. 2018. *Persuasion via weak institutions*. Work. Pap., Univ. Chicago
- Maskin E. 1999. Nash equilibrium and welfare optimality. *Rev. Econ. Stud.* 66(1):23–38
- Mathevet L, Perego J, Taneva I. 2018. *On information design in games*. Work. Pap., New York Univ.
- Matyskova L. 2018. *Bayesian persuasion with costly information acquisition*. Work. Pap., Cent. Econ. Res. Grad. Educ. Econ. Inst., Charles Univ., Prague
- Mensch J. 2018. *Monotone persuasion*. Work. Pap., Hebrew Univ., Jerusalem
- Mertens J-F, Zamir S. 1985. Formulation of Bayesian analysis for games with incomplete information. *Int. J. Game Theory* 14(1):1–29
- Milgrom P. 1981. Good news and bad news: representation theorems and applications. *Bell J. Econ.* 12(2):380–91
- Milgrom PR, Roberts J. 1986. Relying on the information of interested parties. *RAND J. Econ.* 17:18–32
- Min D. 2017. *Bayesian persuasion under partial commitment*. Work. Pap., Univ. Ariz., Tucson
- Nguyen Q. 2017. *Bayesian persuasion: evidence from the laboratory*. Work. Pap., Utah State Univ., Logan
- Orlov D, Skrzypacz A, Zryumov P. 2018a. *Persuading the principal to wait*. Work. Pap., Univ. Rochester, NY
- Orlov D, Zryumov P, Skrzypacz A. 2018b. *Design of macro-prudential stress tests*. Work. Pap., Univ. Rochester, NY
- Ostrovsky M, Schwarz M. 2010. Information disclosure and unraveling in matching markets. *Am. Econ. J. Microecon.* 2(2):34–63
- Perez-Richet E. 2014. Interim Bayesian persuasion: first steps. *Am. Econ. Rev.* 104(5):469–74
- Perez-Richet E, Skreta V. 2018. *Test design under falsification*. Work. Pap., Sci. Po, Paris

- Rabinovich Z, Jiang AX, Jain M, Xu H. 2015. Information disclosure as a means to security. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pp. 645–53. New York: ACM
- Rayo L, Segal I. 2010. Optimal information disclosure. *J. Political Econ.* 118(5):949–87
- Renault J, Solan E, Vieille N. 2017. Optimal dynamic information provision. *Games Econ. Behav.* 104:329–49
- Romanyuk G, Smolin A. 2019. Cream skimming and information design in matching markets. *Am. Econ. J. Microecon.* 11:250–76
- Schweizer N, Szech N. 2019. Optimal revelation of life-changing information. *Manag. Sci.* 64:4967–5460
- Shin HS. 1998. Adversarial and inquisitorial procedures in arbitration. *RAND J. Econ.* 29:378–405
- Sims C. 2003. Implications of rational inattention. *J. Monet. Econ.* 50:665–90
- Smolin A. 2017. *Dynamic evaluation design*. Work. Pap., Univ. Bonn, Switz.
- Spence M. 1973. Job market signaling. *Q. J. Econ.* 87(3):355–74
- Szydlowski M. 2016. *Optimal financing and disclosure*. Work. Pap., Univ. Minn., Minneapolis
- Taneva I. 2016. *Information design*. Work. Pap., Univ. Edinburgh, Scotl.
- Thaler RH, Sunstein CR. 2008. *Nudge: Improving Decisions About Health, Wealth, and Happiness*. New Haven, CT: Yale Univ. Press
- Tsakas E, Tsakas N. 2018. *Noisy persuasion*. Work. Pap., Maastricht Univ., Neth.
- Tsakas E, Tsakas N, Xefteris D. 2017. *Resisting persuasion*. Work. Pap., Maastricht Univ., Neth.
- Wang Y. 2015. *Bayesian persuasion with multiple receivers*. Work. Pap., Xiamen Univ., China
- Yoder N. 2018. *Designing incentives for heterogeneous researchers*. Work. Pap., Univ. Ga., Athens
- Zhang J, Zhou J. 2016. Information disclosure in contests: a Bayesian persuasion approach. *Econ. J.* 126(597):2197–217