

A Survey of Strategic Communication and Persuasion

Ayça Özdoğan*

TOBB University of Economics and Technology

Abstract

This survey paper presents an extended literature review of strategic communication and persuasion. The goal is to provide a systematic categorization of the different approaches in the literature, accompanied with the discussions of the pioneer and leading papers as well as the frontier research in each of the categories.

Keywords: persuasion games, cheap talk, Bayesian persuasion, information disclosure.

JEL Classification: C70, D82, D83.

Stratejik İletişim ve İkna Üzerine Bir İnceleme

Özet

Bu inceleme çalışması, stratejik iletişim ve ikna etme üzerine detaylı bir literatür taraması sunmaktadır. Araştırmanın amacı literatürdeki farklı yaklaşımların sistematik bir sınıflandırmasını yaparken, her bir kategoride öne çıkan çalışmaları, belirleyici ve öncü makaleleri tartışmaları ile ortaya koymaktır.

Anahtar Kelimeler: ikna oyunları, ucuz konuşma, Bayesian ikna, bilgi ifşası.

JEL Sınıflaması: C70, D82, D83.

With its widespread applications to the issues such as advertising, lobbying, political campaigning, judicial decision making, gathering information from possibly biased experts, convincing an investor to invest in a project among many others, the question of when or if a person (i.e. sender/speaker) who is better informed about the state of the world can persuade an uninformed decision maker (i.e. receiver/listener) to act in a particular way is of great economic importance. If the receiver knows with what intention the sender is presenting information, could the sender still convince the receiver to change her opinion? Under what conditions (such as verifiability of information, cost of untruthfulness, alignment of preferences) can the receiver be persuaded by the sender, and if such conditions exist, what

* Ayça Özdoğan is an Assistant Professor in the Department of Economics at TOBB University of Economics and Technology, 06560, Söğütözü, Ankara, Turkey. E-mail: aozdogan@etu.edu.tr

could be an optimal persuasion strategy or mechanism? On the other hand, how much information is revealed or concealed by the sender and is the transmitted information credible (in cases when the sender cannot warrant that she is truthful or that the information cannot be verified are crucial from the perspective of the receiver in making correct decisions. These and related questions have been investigated in different environments including settings that extend the benchmark models to multi-sender situations, incomplete-information structures that incorporate behavioral types, and dynamic strategic communication models.

There is an extensive literature on the theoretical underpinnings of strategic communication and effective persuasion for a wide range of environments examined under various assumptions. This survey has emerged from the need to categorize the diversified literature on strategic communication and persuasion with an attempt to allow the reader to unify the results obtained in different approaches. In pursuing that, the current survey discusses the pioneer and leading papers in each category as well as the frontier research on the subject. This paper hopes to shed light on the big picture and enable identification of the questions that are still open.

The studies that are the focus of this review analyze variations of situations where a privately informed and possibly biased agent transmits this information to an uninformed decision maker who takes an action based on the revealed information that determines the welfare of both. There are a couple of aspects that influence the approach used in the analysis. The first one is related to the verifiability of the sender's information through some hard evidence. Naturally, one expects that the tools that are used to examine information disclosure with verifiable information to be different than those with unverifiable information, since the latter brings about issues related to the credibility of the sender's messages and the unreliability of the transmitted information. The second aspect concerns whether the sender's information is exogenous and the sender chooses optimally how much of it to reveal or withhold, or whether the information is endogenous when the sender designs the information technology (signaling mechanism) before transmitting the realization of it. Depending on these two aspects, this survey offers three broad categories that include studies which assume that

- i. *the sender's information is exogenous and unverifiable,*
- ii. *the sender's information is exogenous and verifiable,*
- iii. *the sender's information is endogenous^[1]*

The papers belonging to the first category, pioneered by Crawford and Sobel (1982) and Green and Stokey (2007) (first version 1981), assume that information is unverifiable and the sender transmits messages or signals that are not necessarily truthful and that incur no direct costs. The games describing those situations are called *cheap talk games*. The papers in the second category, initiated by Grossman (1981) and Milgrom (1981), assume

^[1] Most of the literature where the sender endogenously acquires information or designs a signaling mechanism assumes that the sender commits to that technology whose realization can be publicly observed or truthfully transmitted. The Bayesian receiver takes an action after she updates her prior belief about the state of the world based on the signal designed and transmitted by the sender.

that the information is verifiable and that the senders may conceal information but that they cannot lie. These games are referred to as *persuasion* or *verifiable disclosure games*. Finally, the third and more recent category is referred to as *Bayesian persuasion games*.

There are many important papers that span these categories or lie in the intersection of these approaches. However, for the purpose of this survey, the focus will be given to the work that fits in one of these categories. The rest of the paper is organized as follows: Section 2 presents the first approach, i.e., *cheap talk games* and their variations. Section 3 reviews the literature on *verifiable* or *partially verifiable information disclosure* which will be further categorized as *persuasion games* and *persuasion mechanisms*. Section 4 provides the literature on *Bayesian persuasion games*. Finally, Section 5 concludes the paper.

Cheap Talk (Exogenous and Unverifiable Information)

A non-negligible portion of the literature dealing with strategic communication examines situations where an informed sender tries to convey this information to an uninformed receiver, who then makes a decision that affects the payoff of both through costless messages or possibly noisy signals. Some real life applications include doctor-patient, lawyer-client, taxi driver-customer interactions (among many others) in which one party has information about the state of the world (such as the severity of the patient's sickness, the difficulty of the lawsuit, the congestion of alternative routes etc.) that may be hard to verify by the other party, who relies on the information transmitted by the informed agent through costless messages before making her decision that affects the welfare of both. The alignment of the parties' interests can be seen as the key feature behind how much information is truthfully revealed by the informed party and whether the uninformed agent's decision obeys the informed party's recommendation in these examples.

Three attributes related to direct and costless communication are often the main concerns: (i) credibility (whether the messages transmitted and the information revealed by the sender is truthful or not), (ii) persuasiveness (whether the receiver takes the sender's preferred action or not), and (iii) informativeness (how much of the private information possessed by the sender is revealed, if any, or withheld).^[2]

One of the pioneer papers on *cheap talk* is the seminal work of Crawford and Sobel (1982), which shows that equilibrium signaling always takes a simple form of partitioning the scalar variable that represents the private information of the sender. The sender reports only which partition his private information belongs to, and which adds noise into his actual signal by providing an interval around it. This is the outcome of the tradeoff between optimally including enough information in the signal to induce

^[2] The full-revelation of information refers to the existence of a fully separating equilibrium where each piece of information is perceived to be a type for the sender and different types of the sender transmit different messages. However, in order to gain more, a sender's different types may optimally pool each side's behavior (i.e. send the same message) which makes the information vague, noisy and partially informative. Moreover, there may be situations in which the transmitted message has no informational value at all.

the receiver to respond to it and concealing enough information to benefit more. Crawford and Sobel (1982) establish that equilibrium signaling is more informative when the agents' preferences are more aligned; and, in a game of perfectly conflicting interests, the sender's signal is completely uninformative and thus the receiver ignores the messages transmitted by the sender.^[3] In fact, with *cheap talk*, there is always an equilibrium where the message is uninformative and the receiver ignores the message. The partially revealing (semi-separating) partition equilibrium exists only when the preferences or interests are somewhat aligned.

It is then natural to ask under what situations a fully informative equilibrium can be attained through direct and costless communication if the interests of the parties are not perfectly aligned or if they have conflicting interests. In Crawford and Sobel's (1982) paper, the information is unverifiable and the sender can lie without any direct costs. However, the sender may bear a cost of lying because of various reasons.^[4] On top of that, some actors may be acting non-strategically -- i.e., the sender and/or receiver may have behavioral types that act in a certain manner such as always truthfully or naively.^[5]

Another aspect that can be considered is about the dimensions of the issues that the sender is informed about. For instance, an expert often has information on multiple aspects of a decision-problem, e.g. a salesperson has information about multiple attributes of a single product he wants to sell. Moreover, in some situations, there are multiple experts having possibly different biases and information who try to affect the decision maker to take their own preferred action. Lastly, the communication may have a back and forth (dynamic) and gradual (the sender reveals information step by step) structure. The theoretical studies that are to be discussed below accommodate direct lying costs, the possibility of boundedly rational agents, the multi-dimensionality of information, multi-senders and dynamic communication into the benchmark *cheap talk* model of Crawford and Sobel (1982); they also analyze if one can attain fully revealing, credible or persuasive equilibria by incorporating these additional features.

Cheap Talk with Lying Costs

As mentioned earlier, the informed party may incur a cost of lying while transmitting a message untruthfully due to legal or moral constraints; for instance, a financial analyst

^[3] Farrell (1993) calls these as babbling equilibria.

^[4] There may be a cost associated with cooking the books, possibility of getting caught, reputation concerns etc. And also, recent experimental work suggests that people have an intrinsic aversion to lying. For a partial list of these studies, see Gneezy (2005), Hurkens and Kartik (2009), Sanchez-Pages and Vorsatz (2007), Sanchez-Pages and Vorsatz (2009), Gurdal et al. (2014).

^[5] Dickhaut et al. (1995) and Cai and Wang (2006) provide laboratory experiments demonstrating that receivers may react non-strategically and differently to the same message. In particular Cai and Wang (2006) test the Crawford and Sobel (1982) model. Although their findings support some results of the theory (such as less information being transmitted when preferences are less aligned), they also show that players overcommunicate (in the sense that the senders' messages are more informative and that the receivers rely more on the senders' messages before choosing their actions), compared with what the theory admits in the most informative equilibrium. Gurdal et al. (2014) show that this overcommunication phenomenon is robust to the existence of a behavioral type for the sender who is honest.

may bear a legal cost for providing systematically wrong recommendations of the stock price (also see footnote 4). Building on Crawford and Sobel (1982)'s *cheap talk* model, one of the leading papers on costly talk is by Kartik et al. (2007), who study a model of strategic communication by an informed and upwardly biased sender to one or more receivers in which it is costly for the sender to misrepresent information and the receivers may naively believe the senders' message. In contrast to the predictions obtained in the benchmark *cheap talk* model, their model admits a fully separating equilibrium where the sender's message completely reveals her private information, provided that the state space is unbounded above. In the equilibria, the message sent by the sender is higher than the true state; i.e., the message is inflated (as the sender is upwardly biased). However, a strategic receiver correctly infers the true state. These findings differ from one of the main results of Crawford and Sobel (1982) which suggests that given that the sender's and the receiver's preferences diverge, every equilibrium is partially revealing or uninformative. However, one should emphasize that the full-revelation of information in Kartik et al. (2007) holds under the crucial assumption of unboundedness of the state and message spaces.

In a related work, Kartik (2009) analyzes information revelation and language (mapping from types to messages) with lying costs for a bounded state and message spaces as in Crawford and Sobel (1982). The equilibria show that low sender types separate whereas high types pool into some segments, regardless of the extent of the lying costs. Like the previous study, the equilibria exhibits language inflation; i.e., the sender claims to be a higher type than he truly is even though he incurs a cost for lying. The degree of language inflation and how much information is revealed depend upon the intensity of the lying cost. However, unlike Kartik et al. (2007), Kartik (2009) show that fully separating equilibria do not exist since the upper bound on the type space prevents full separation given that the messages must be inflated.

Cheap Talk with Boundedly-Rational Agents

One of the first studies that accommodates behavioral (nonstrategic) types for the agents is by Crawford (2003) and Hendricks and McAfee (2006) who study misrepresentation of intention by an agent to his/her competitor with conflicting interests. In order to explain the occurrence of deception observed in real life examples with conflicts of interests (and those which cannot be induced by the benchmark model of Crawford and Sobel (1982)), Crawford (2003) considers a one-shot zero-sum game preceded by a *cheap talk* stage with behavioral types for both players. The introduction of behavioral types generates equilibria in which the strategic sender lies and deceives the strategic receiver.^[6] Hendricks and McAfee (2006) also model deception as a one-shot zero-sum attacker-defender game in which the attacker's actions are imperfectly signaled. An equilibrium with deception can emerge under some conditions on the noisy signaling technology.

^[6] In motivating the misrepresentation of intentions to enemies and possibility of deception with behavioral types, Crawford (2003) starts with the real life example of the Allies' decision to feint at Calais and attack Normandy, which resulted in the successful deception of the German troops. In the benchmark cheap talk equilibrium, the Germans should have ignored the fake messages (attack at Calais) sent by the allied forces.

Ottaviani and Squintani (2006) introduce the possibility of a naive receiver in a game of information transmission where the parties have partially aligned interests. The construction of an equilibrium exhibits full-revelation in a low range of the state space whereas partial-revelation is in the top; the language is inflated, and more precise information (compared with that in the benchmark fully strategic model of Crawford and Sobel (1982)) is transmitted. Moreover, the equilibrium converges to the fully-revealing equilibrium as the state space becomes unbounded.

Chen (2011) studies strategic communication both with a possibly honest sender and a naive receiver and characterizes message monotone equilibria (i.e., the senders' reporting strategy is nondecreasing in the state). The properties of the equilibria indicate sender exaggeration, receiver skepticism and message clustering. Assuming that the strategic sender has an upward bias, he pools his message with the honest sender (who has a higher observation in expectation) resulting in sender exaggeration. While the messages at the low range are sent only by the honest sender and thus they are credible, the model generates receiver skepticism for a sufficiently high end. The other aspect of the game is the clustering of messages, i.e., only the messages at the top of the message space are sent in large measure by the dishonest sender. The limit equilibrium approaches the most-informative equilibrium of the canonical Crawford and Sobel (1982) model as the probabilities of the behavioral types approach zero.

Cheap Talk with Multi-Dimensional Information

The informed and biased sender may sometimes have differential information about multiple attributes of the same issue and may present differential and comparative statements (messages) about these attributes to enhance his gain in his interaction with a decision maker who takes the action that affects the payoff of both. For instance, a salesperson has information about a couple of properties of the same product that he wants to sell regardless of the state (i.e., there is a strongly biased preference). This salesperson may reveal differential information about the different aspects of the same product. The question is whether the recommendations of the salesperson are persuasive and whether a fully-revealing equilibrium is possible in a pure *cheap talk* game (no lying costs) with multi-dimensional information.

One of the leading papers, Chakraborty and Harbaugh (2007) on multi-dimensional *cheap talk*, demonstrates that simple complementarity conditions guarantee an expert with private information about multiple issues will credibly rank the issues for a decision maker. The communication gains can be substantial as the complete ranking is asymptotically equivalent to full-revelation of information when the number of issues grows large. In a related work, Chakraborty and Harbaugh (2010) analyze the credibility, persuasiveness, and informativeness of multidimensional *cheap talk* by an expert who is biased towards the same action regardless of the state of the world. They show that both the expert and the decision maker may benefit from the multi-dimensional structure. This is because the expert can always make credible comparative statements which allow for more information that balances her incentive to cheat on each dimension (given that her preferences are state-independent) and the decision maker can accommodate for the

experts biases and receive the credible content of the messages. Communication benefits the expert (*cheap talk* is persuasive) if her preferences are quasiconvex. The intuition is as follows: the messages make the decision makers updated estimates of the true state more extreme in comparison to ex-ante expectations, and thus the expert benefits from such communication when she prefers extreme estimates on some dimensions to average ones on all dimensions. Hence, multi-dimensional information and state-independent preferences are sufficient to provide credibility and persuasiveness.

Levy and Razin (2008) also extend the *cheap talk* model of Crawford and Sobel (1982) to a multidimensional state and message space with a focus on information transmission when the agents' interests diverge greatly. They show that although it is possible to construct equilibria with information revelation even for unboundedly large conflicts, it requires knife-edge assumptions. Thus, adding more dimensions cannot suggest an improvement in information revelation when interests are too divergent.

Cheap Talk with Multi-Senders

There is a growing literature on *cheap talk* with multiple senders. In many instances, decision makers consult multiple experts with possibly different biases. For instance, a patient may see multiple doctors before deciding on getting an operation; a policy maker often seeks the opinion of multiple advisors in attaining her final decision on an issue; a judge listens to the members of a jury before giving her final call about a lawsuit; or, a CEO decides whether to extend the functions of the firm after consulting with the managers of different divisions in the company, etc.

Early work on *cheap talk* with multiple senders include Gilligan and Krehbiel (1989), Austin-Smith (1993), Epstein (1998) and Khrisna and Morgan (2001a) who study two experts with totally different biases simultaneously communicating with a decision maker. In particular, Khrisna and Morgan (2001a) shows that full revelation of information can be attained as an equilibrium.^[7] Khrisna and Morgan (2001b), on the other hand, examine a more general model where the two experts with alike or opposing biases communicate sequentially. They find that while consulting two experts is useless for the decision maker when experts have like biases; consulting with both experts is always more informative when the experts have opposing biases. A common feature of these papers is that the state space is one-dimensional and both senders have the same information.^[8]

Battaglini (2002) provides a comprehensive analysis of two-sender *cheap talk* with a multidimensional Euclidean state space. He shows that a generically fully-revealing perfect Bayesian equilibrium that survives any refinement puts restrictions on off the equilibrium beliefs. The construction is as follows: each sender conveys information only in directions along which her interest coincides with that of the receiver. These

^[7] However, this result is not robust since truth telling is a weakly dominated strategy. Battaglini (2002) defines an equilibrium refinement by allowing players to make mistakes. But, this refinement eliminates all equilibria with full-revelation.

^[8] Austin-Smith (1993) is an exception with senders receiving correlated signals regarding the true state of the world. Austin-Smith (1993) also finds that multiple informed and biased senders are beneficial in a political economy setting.

directions of common interest generically span the whole state space, which allows the decision maker to perfectly identify the true state by combining the messages.

Ambrus and Takahashi (2008) differs from the Battaglini (2002) analysis in allowing for a multidimensional state space that is a closed subset of the Euclidean space. This enables the analysis of multidimensional cheap talk to be more comparable to the canonical *cheap talk*, where the one-dimensional state space is assumed to be bounded. They show that if biases are large enough in opposite directions, then no fully-revealing perfect Bayesian equilibrium exists, which suggests that boundedness, not dimensionality, plays an important role in determining the transmission of information in *cheap talk* games. They also give necessary and sufficient conditions for the existence of fully-revealing equilibria for a compact state space and find that the senders must have similar biases for arbitrarily large biases to achieve full-revelation of information in equilibria.

In Battaglini (2002) and Ambrus and Takahashi (2008), each sender observes the same information in all dimensions. In a recent paper, McGee and Yang (2013) study a cheap talk model with two senders having partial and nonoverlapping private information about the state of the world that takes two dimensions. Senders who have different biases simultaneously communicate their information with an uninformed receiver. The equilibria are shown to be partition equilibrium in which each sender indicates only which interval the realized state that he has observed belongs to as in the standard Crawford and Sobel (1982) cheap talk model. Focusing on the most informative equilibria, they demonstrate that more information transmitted by one sender induces the other sender to transmit more information. Li et al. (2014) also examine a cheap talk model with two senders. Each sender in their model is responsible for a single project and only observes the realization of the return on that project. Exactly one of the two projects is going to be implemented by the decision maker. Both senders have some common interest with the receiver, but each has her own project biases. All equilibria are shown to be partially revealing (partition equilibria as before). But each agent transmits more information in a symmetric equilibrium as the number of agents increases.

Dynamic Cheap Talk

In some situations, we observe that agents interact repeatedly (sometimes with a predetermined terminal date and sometimes not). For instance, a car owner (uninformed agent) may go to the same mechanic (informed agent) to have a repair, a patient may see the same doctor, a bank may work with the same client. In repeated settings, reputational concerns of the informed agents may play an important role in the analysis.^[9] One of the first studies on dynamic model of credibility is by Sobel (1985) who shows that one agent must trust another who has precise information about the state of the world but also has unknown intentions. He shows that it pays to build a reputation by repeatedly providing the accurate information and then to cash in on this reputation.

^[9] In a static model, Ottaviani and Sorensen (2006) also study a reputational cheap talk model where the privately informed sender cares about his reputation as a provider of accurate information. They show that senders can credibly reveal only partial information in equilibrium.

Benabou and Laroque (1992) provide a model of repeated strategic communication with an agent (an insider trader, a journalist or a guru) who has noisy private information about the value of an asset and aims to manipulate asset prices. Unlike Sobel (1985), since the agent's information is noisy, he can engage in manipulation repeatedly without being fully detected. Morris (2001) adapts Benabou and Laroque's (1992) model to a setting where the sender is either a friend or a bad type. He shows that the friend type may have an incentive to lie just as the bad type sometimes has an incentive to tell the truth to improve her reputation. In an application of reputation models in financial markets, Mathis et al. (2009) examine whether the reputation concerns are sufficient to discipline the credit rating agencies. The CRA can be either an honest type who it is committed to always telling the truth or an opportunist. They show that there is a unique Markov perfect equilibrium consisting of several phases: the reputation building phase, the cashing on reputation phase and the crisis of confidence where the opportunist CRA is detected and its reputation falls.

Besides the reputational papers that make use of repeated games, there are dynamic communication models in which the sender sequentially transmits information through back and forth communication and the decision maker takes an action once at the end of the communication process. In that regard, Ivanov (2015b) investigates a dynamic model of strategic communication between a principal and an expert with conflicting preferences. In each stage, the uninformed decision maker can select the precision of the expert's information about an unknown state without observing the realization of the information itself. He shows that the decision maker can elicit perfect information from the expert about the state in two stages only if the expert's preference bias is not too large. If the state space is unbounded, full-information revelation is possible for an arbitrarily large preference conflict. In a related paper, Ivanov (2015a) studies a multi-stage version of Crawford and Sobel (1982). The receiver may influence the precision of the sender's information through assigning tests about the unknown state that is privately observed by the sender, and request a report from her in each period. The sender conducts each test and sends a costless message about the outcome of the test. Then the receiver selects a test in the next period. After a finite number of periods, the receiver makes a decision. The receiver can design a sequence of tests that allow for learning information with any precision over an interval of states which converges on the state space as the preference bias disappears.

Information Disclosure with Verifiable Information

Some applications fit more to situations where the information disclosed by the informed agent can be fully or partially verified by the uninformed one. For instance, a customer who wants to buy a product may ask for a warranty or instructions from the salesperson, which show at least some properties of the product that reflect the quality. The studies concerning strategic communication with verifiable information will be classified into two according to the tools used: (i) persuasion games, and (ii) persuasion mechanisms. The major seminal and frontier papers in each class are discussed in the following sections.

Game Theoretical Approach: Persuasion Games

The leading benchmark model of strategic communication with verifiable messages in which senders can withhold information but not fully lie (referred as persuasion games) was initiated by Grossman (1981) and Milgrom (1981). They show complete unravelling of information in equilibrium. Shin (1994), on the other hand, shows that unravelling of information disappears when the decision maker does not know the precision of the sender's information. Verrecchia (1983) adds cost to information disclosure and again finds that it prevents complete unravelling of information.

Milgrom and Roberts (1986), on the other hand, study whether full information disclosure is possible when multiple experts who observe the same information compete to influence a decision maker by presenting verifiable information as recommendations on decisions. They show that competition may not be necessary for full revelation of information if the decision maker is strategically sophisticated, and well-informed about some relevant variable and about the preferences of the experts. On the other hand, if the decision maker is not well-informed or not sophisticated, then competition may be insufficient for a full revelation of information. But, if the interested parties' interests are sufficiently opposed, then competition can diminish the decision maker's need for prior information about the relevant variables or strategic sophistication in achieving full revelation of information.

In a more recent study, Dziuda (2011) analyzes a game between an uninformed decision maker and a possibly biased expert who helps the decision maker to choose between two alternatives. The expert privately observes a set of arguments each of which favors one of two alternatives. Even though he can credibly disclose each argument, he cannot prove whether he has disclosed all the arguments. The expert may be an honest type who transmits all of the arguments, or he may be biased in convincing the decision maker to choose a particular alternative. The decision maker prefers the alternative which is favored by a sufficient fraction of arguments, but she lacks information about the number of arguments. It is found that the biased expert sends messages containing arguments both for and against his preferred alternative in all equilibria. The paper also shows that full disclosure of information is not an equilibrium and the biased expert is able to persuade the decision maker only if he has sufficient favorable arguments.

In a related paper, Honryo (2011) develops a game theoretic model of dynamic persuasion. A sender with a fixed number of hard evidence tries to persuade a decision maker by providing enough favorable evidence. The sender can sequentially reveal at most one piece at a time with a cost. The decision maker, on the other hand, chooses whether to receive another piece of evidence or not, which delays decision making, which is also costly. So, at each period, the sender chooses when to give up the persuasion and the decision maker decides when to make the decision. In equilibrium, each time the sender communicates a piece of evidence, the decision maker updates her belief, and she accepts the proposal with a strictly positive probability. As the game proceeds, the decision maker accumulates more information and the probability of making the wrong decision decreases. The decision maker typically accepts the sender's proposal even when she knows that the sender's evidence may be unfavorable. However, in a Pareto efficient equilibrium, she does not refuse a proposal with favorable evidence

unless delay costs are very large. While there is a lower bound to the probability of immediate acceptance every time the sender communicates a piece of evidence, silence never welcomes immediate acceptance, which suggests that only the costly message is persuasive.

Mechanism Design Approach: Persuasion Mechanisms

In this section, we discuss some of the leading papers that use a mechanism design approach to analyze persuasion and information transmission with hard evidence or (partially) verifiable information.^[10]

Glazer and Rubinstein (2004) study a situation in which one agent (the speaker) wishes to persuade another agent (the listener) to take a particular action among two alternatives. The receiver's decision depends on both the sender's messages and the hard evidence that she can collect. Specifically, there are two aspects that determine the preferred action for the listener. The speaker has private information about both of the aspects, whereas the listener can check the realization of at most one of them. Based on the speaker's message and hard evidence, the listener is either persuaded to take the speaker's preferred action or not. A mechanism in this setting is a three part set of messages from which the speaker can choose, a) function that specifies which aspect is to be checked depending on the speaker's message, b) the action the listener takes as a function of the message sent by the speaker, and c) the obtained evidence. Glazer and Rubinstein (2004) aim to characterize mechanisms that maximize the probability that the listener chooses the correct action given that the speaker maximizes the probability that his preferred action is going to be chosen by the listener. An optimal mechanism is found by solving a linear programming problem for which the set of constraints is derived from what is called the L-principle.

In a related paper by Glazer and Rubinstein (2006), a speaker wishes to persuade a listener who has a binary action space. The speaker who can support his claim through hard evidence is restricted by the amount of evidence he can present. They study persuasion rules that maximize the probability that the listener accepts the recommendation of the speaker only when it is just, given that the speaker maximizes the probability that his request is accepted. The results of their paper show that the optimal persuasion rule features two properties: (i) there always exists an optimal persuasion rule which is deterministic (no need for randomization); and (ii) every optimal persuasion rule has a credible implementation which is an equilibrium of the persuasion game without commitment (commitment has no value).

In a follow-up paper, Sher (2011) points out that the credibility and determinism results of Glazer and Rubinstein (2006) do not hold in more general spaces (than binary action spaces) without further assumptions. He shows that concavity is the essential assumption for both results (no value to commitment and no value to randomization) to carry over more general action spaces. I.e., when the listener's utility function is a

^[10] A general theory of mechanism design with state verification or verifiable evidence is studied by Bull and Watson (2007), Kartik and Tercieux (2012), Deneckere and Severinov (2008) among many others.

concave transformation of the speaker's utility function, both credibility and determinism are restored.

Sher (2014), on the other hand, compares static and dynamic persuasion. In static persuasion, the speaker presents evidence and the listener responds with a decision. In dynamic persuasion, back-and-forth communication is allowed. In particular, the speaker makes an opening argument with a cheap talk. The listener then requests some of the evidence that the speaker claims to have. The speaker responds by presenting either the requested evidence or some other evidence. Finally, the listener makes his decision.^[11] A persuasion rule determines which evidence is persuasive. Dynamic persuasion may outperform static persuasion when it is infeasible for the speaker to present all the evidence due to associated costs. He characterizes the relation between optimal dynamic and static rules by presenting a single linear program whose solution corresponds to both the listener's optimal dynamic rule (with no further constraints) and the optimal static rule with additional constraints. He proposes a condition, called foresight, under which the optimal persuasion problem is reduced to the classical maximum flow problem that makes the optimal dynamic and static persuasion rules coincide.

Bayesian Persuasion

This section provides a discussion of the papers which study that the senders endogenously design the informational environment, i.e., signaling structures (mechanisms) which induce the Bayesian decision maker (receiver) to update her prior beliefs so that the probability of choosing the sender's preferred action by the decision maker increases on average.^[12] As an application of *Bayesian persuasion*, Gentzkow and Kamenica (2011) start with a motivating example of a prosecutor (sender) who tries to convince a rational and Bayesian judge (receiver) that a defendant is guilty and to persuade her to convict the defendant (which is the sender's preferred action). The question is whether the prosecutor structures her arguments (through selection of signals, evidence or revealing facts suitable for her purpose, etc.) so as to increase the probability of conviction on average. By a neat design of signaling technology, they show that the prosecutor can manipulate the posterior beliefs so that she can benefit from them.

As one of leading papers Rayo and Segal (2010) study the optimal disclosure rule in a setting where a sender first chooses and commits to an information disclosure rule about a prospect that is associated with a profitability for the sender and a value for the receiver.^[13] Next, a prospect is drawn at random by the sender and a signal is transmitted to the receiver according to the chosen rule. The receiver then makes a Bayesian inference

^[11] As an example, Sher (2014) gives a time-constrained interview: a job candidate must decide which qualifications to mention. If she emphasizes her leadership capacity, the employer may ask to provide evidence of her past leadership, whereas, if she emphasizes her analytical ability, she may be asked to solve a puzzle.

^[12] For a list of studies where the receiver is not perfect Bayesian, we refer to Mullainathan and Shleifer (2008) and Ettinger and Jehiel (2010).

^[13] As an economic application, one can think of bond rating agencies as senders who choose what information to disclose to investors about bond issuers, from whom they also get their payments for the rating. Each rating (prospect) has a profitability for the rating agency and a value for the investors. The purpose is to characterize the optimal (profit-maximizing) disclosure rule and the properties of it.

about the prospects value and chooses whether to accept the prospect or not. Rayo and Segal (2010) characterize the optimal disclosure rule under particular assumptions about the preferences and provide several properties that the optimal disclosure rule must satisfy.^[14]

The other pioneer paper is by Gentzkow and Kamenica (2011) who point out that a sender can persuade a Bayesian receiver (decision maker) to pick his favorite action (with a higher probability on average compared to no communication) by optimally designing a signal structure. The two agents share a common prior assessment on the state of the world. The sender publicly generates a signal by choosing a persuasion rule that is a map from the true state of the world for a distribution over some signal space. The realization of the signals are truthfully revealed to the receiver. Gentzkow and Kamenica (2011) characterize the necessary and sufficient conditions for the existence of signals that strictly benefit the sender. Their results are due to simplifying the sender's problem with the help of two observations. The first one considers the sender's payoff as a value function over the posterior belief induced by the signal realization; the second observation points out that there exists a signal that generates any distribution of posteriors (whose expectation equals to the prior). These enable to derive a value function over beliefs and determine the existence of an optimal signal that can benefit the sender from the concavification of that value function. They show that the sender benefits from designing signals if the receiver does not take the sender's preferred action by default and the receiver's action is constant in some neighborhood around the prior belief. They find that the concavity or convexity of the sender's payoff in the receiver's belief determines whether the *Bayesian persuasion* is beneficial to the sender or not.

In Gentzkow and Kamenica (2014), they extend their benchmark model to accommodate a costly signaling environment. The problem that arises with costly information generation is that the sender's payoff is not a value function over the posterior belief alone anymore; it also depends on the signal because of its cost. Since one cannot express the sender's payoff as a value function over beliefs, the concavification approach cannot be used. This paper introduces a family of cost functions that is compatible with the concavification approach in attaining the optimal signal. One example in this family is the cost of a signal being proportional to the expected reduction in uncertainty.

In a related work, Gentzkow and Kamenica (2016a) study situations where a number of senders with common prior and the same set of signals choose what information to communicate to a receiver. They first argue that the effect of competition on information

^[14]In particular, this assumption deals with the receiver's private reservation value (opportunity cost of accepting a prospect), which says that the receiver's reservation value is drawn from a uniform distribution in $[0, 1]$. This makes the probability of the receiver accepting a prospect equal to the expectation of its value. Moreover, the distribution from which prospects are drawn is assumed to be of finite value, and the randomization of signals is allowed. Under these assumptions, the optimal rule must satisfy several properties that include: (i) It may be optimal to pool two prospects (send the same signal for each of them with a positive probability) when one has a higher profit and a lower value than the other. On the other hand, it is never optimal to pool them when one prospect dominates the other in both profit and value. (ii) The prospects that are pooled with each other with a positive probability must lie on a straight line in the profit-value space. For the generic case, no three prospects are on the same line, which implies that a signal may pool at most two prospects.

revelation is ambiguous. This paper identifies a necessary and sufficient condition (which they call Blackwell-connectedness) on the information environment under which the competition cannot produce less information revelation regardless of the preferences of the agents.

Gentzkow and Kamenica (2016b) extend their analysis in their basic model to incorporate large state spaces. They assume that the receiver's optimal action depends on the expectation of a random variable that represents the state, and sender's preferences over the action chosen by the receiver are independent of the state. In their earlier paper, Gentzkow and Kamenica (2011) show that the concavification of that value function over the posterior beliefs pins down whether the sender can benefit from generating information; however, it is not sufficient to determine the optimal signal. The problem is that it is not possible to create every distribution of posterior means whose expectation is the prior mean. In the most recent paper, they derive the characterization of all feasible distributions of the posterior mean by using insights from Blackwell (1953) and Rothschild and Stiglitz (1970).^[15]

Kolotilin (2014), on the other hand, examines an environment similar to Gentzkow and Kamenica (2016b) with the same assumptions on preferences but also by allowing the receiver to have private information. The receiver decides whether to act or not to act. The sender prefers the receiver to act. The receiver's utility depends both on his action and on his information. The receiver takes an action that maximizes his expected utility, given his private information and the information transmitted by the sender. Before the sender obtains her private information, the sender can commit to a disclosure rule. Kolotilin (2014) shows that neither the sender's nor the receiver's payoff is necessarily monotone in the precision of the receiver's private information: the sender may find it easier to influence a more informed receiver, and the receiver may suffer from having more precise private information.

Similarly, Kolotilin et al. (2014) consider the Bayesian environment of Gentzkow and Kamenica (2011) with the receiver having private tastes and types.^[16] The receiver reports a private type first, then the sender generates information and chooses a mechanism that sends messages based on the type reported by the receiver. The sender can screen different types of receivers by targeting different information disclosure to the receiver's reports, which they call private persuasion. They compare private persuasion with the benchmark model (public persuasion) in which the sender designs a public signal about the state that is independent of the receiver's report. They show that private and public persuasion are equivalent from the perspective of the sender and that sender never strictly benefits by allowing for private persuasion. They characterize optimal persuasion rules in connection with the standard linear mechanism design with transfers. The optimal persuasion rule features upper censorship (i.e., the public signal sends the message equal to the realized state below some threshold, but sends no message above it) if the

^[15] In particular, they borrowed the characterization of second-order stochastic dominance from Rothschild and Stiglitz (1970).

^[16] As stated by Kolotilin et al. (2014), many applications fit this environment including public educational campaigns, media censorship regulations and, informational lobbying among others.

probability density of the receiver's types is logconcave. Symmetrically, lower censorship is optimal if the probability density of the receiver's type is logconvex.

Alonso and Camara (2015) studies a *Bayesian persuasion* model where the sender and the receiver hold heterogeneous prior beliefs about the state of the world. The sender tries to affect the receiver's behavior by designing the informativeness of a public signal (experiment). They solve for the senders' optimal signal by characterizing the set of distributions of posterior beliefs that can be generated by the signal. Then they provide a necessary and sufficient condition for a sender to benefit from information design. One of the main results establishes that the martingale property of the senders' beliefs and the perspective transformation (which is an expression that depends on the likelihood ratios of prior beliefs) characterize the set of distributions of posterior beliefs that are consistent with Bayesian rationality. The perspective transformation condition implies that the set of joint distributions of players' posterior beliefs under common priors and heterogeneous priors form a bijection, which suggests that belief disagreement does not allow the sender to generate more ways to persuade the receiver. The necessary and sufficient condition is a geometric condition for the value of persuasion to be zero; that is, the sender does not benefit from experimentation if and only if the value admits a supporting hyperplane at the sender's prior belief.

Wang (2015) examines a setting where a biased sender adopts a *Bayesian persuasion* mechanism to provide a group of uninformed receivers with signals. They compare public persuasion (all receivers observe the sender's mechanism choice and the generated signals) to private persuasion (each receiver receives a separate signal even though the sender's mechanism is commonly known). They find that the sender always obtains the concave closure of the set of possible expected payoffs under public persuasion, regardless of the number of signal realizations, and that the sender is weakly worse off under private persuasion.

Li and Norman (2015) consider a class of multi-sender persuasion games that accommodate both sequential and simultaneous moves. The senders generate signals that can be correlated or independent. They also analyze the effect of incorporating a new sender in the original game on information revelation. They find that there is no information loss if the signals are arbitrarily correlated and the new senders are added at the beginning of the game. They also present situations where the addition of a sender results in information loss.

Brocas and Carrillo (2007) analyze the benefit a sender obtains from controlling the flow of public information. The sender who has free access to information decides how much public evidence to collect. Conditional on the content of the news collected, the sender decides whether to keep adding evidence or to stop the learning process. Then, a decision maker with conflicting preferences chooses an action that affects the payoff of both. They characterize the gain of the sender as a result of his control of the generation of public information by deriving the sender's optimal stopping rule for the generation of information. They show that the ability of the sender to induce the receiver to choose his preferred alternative increases as the sender's preferred state is more likely and as the preferences become more aligned.

Horner and Skrzypacz (2014), on the other hand, study a dynamic model of variable information transmission in which a seller can transmit information gradually or partially as the buyer makes payment for it. The sender puts an effort into collecting information, and then transmits this information to a decision maker. The sender only cares about how much payment he can receive for the information he acquires. The problem is not trivial since the buyer is reluctant to pay for information whose value he does not know and the seller cannot credibly reveal how valuable information could be for the decision maker without disclosing the information itself, after which he cannot get any payment. They characterize the properties of equilibria that maximize incentives to acquire information. Basically, they show that the seller can gain a substantial fraction of the value through gradual revelation and the entire value can be extracted with the help of an intermediary. The properties of equilibria include: (i) an equilibria in which the receiver pays the sender; (ii) a particular type of sender gains from using partial evidence and releasing information gradually; (iii) the best equilibrium (for that particular type of sender) involves an initial release of information given away for free, followed by gradual information sales; (iv) the highest sender payoff is independent of the receiver's prior belief and the payoff is less than the full value of information; and, (v) an intermediary allows the sender to get the full value of the information.

Concluding Remarks

This survey presents a partial list of studies in strategic communication and persuasion by providing three categories of the relevant literature, namely *cheap talk* games (strategic communication with unverifiable information), *persuasion games* and *mechanisms* (persuasion with verifiable information disclosure) and *Bayesian persuasion* (information transmission with endogenous sender information). The focus has been mainly on the pioneer and frontier papers that fit one of these categories. The aim has been to unify the results obtained in this diversified literature and identify the questions that are open for future research.

As laid out in Section 2, there is an extensive literature on variations of *cheap talk* games. Given that, there are still unexplored parts. For instance, there are a couple of studies on *cheap talk* with multiple receivers;^[17] yet, an analysis that incorporates both multiple receivers (who possibly have different biases) and a multi-dimensional issue space is missing. In that kind of a situation, would the sender prefer to transmit private/public messages about all or some dimensions to all or some receivers, or would it be optimal for the sender to conceal information about some dimensions from some receivers? A relevant real life application could be a salesperson who wants to sell a computer to two receivers who arrived at the same time and who possess different biases about different issues related to the quality of the computer. In fact, there are many other relevant questions on *cheap talk* games that could be answered by designing an experimental study. Although there is already a considerable amount of work in that

^[17]The reader is referred to Farrell and Gibbons (1989) and Goltsman and Pavlov (2011).

regard (see footnote 4 and 5) which suggests that subjects may act differently than the predictions of the theory (such as the overcommunication phenomenon), the experimental studies that incorporate multiple receivers and/or senders involving boundedly rational agents and a multi-dimensional state space are only a few.^[18] As the most recent piece of literature, the *Bayesian persuasion* (Section 4) involves many working papers and unresolved questions for future research. Some of these questions are related to the multiple sender's and receiver's environment. Even though there are a couple of papers working on *Bayesian persuasion* with multiple senders or receivers (presented in Section 4), when one looks into the details of the models, it can be seen that these models have concrete assumptions about the information and payoff structures. It would be interesting to examine a situation in which senders who have conflicting preferences act sequentially or simultaneously while aiming to undo the *Bayesian persuasion* done by the other. Which sender could gain in this kind of a situation and under which type of regulatory conditions on the signaling mechanisms? Can they generate any evidence they wish given that an evidence against their argument has been already presented before? What could be the best thing the incumbent and follower senders could do? Also, in this setting, the informational and payoff gain or loss to the receiver is not trivial. The receiver may prefer a specific ordering of senders to present evidence or design signals. Yet, another relevant question in *Bayesian persuasion* would concern the addition of boundedly rational agents in the basic model to see if the results would change. The experimental studies that test the predictions of these models would also constitute another line of future research.

Lastly, we would like to point out that there are many other influential studies that may not have been mentioned in the current survey either because they analyze more general strategic communication games that encompass sender-receiver games (such as Aumann and Hart, 2003 and Gerardi, 2004), or they lie in the intersection of these categories or are a hybrid of them (such as Forges and Koessler (2008)), or they offer mechanisms that are suited for different environments (such as Glazer and Rubinstein, 2012;^[19] Wolinsky, 2002; and Gerardi et al., 2009),^[20] or they examine situations that are not covered in the categories classified here.^[21]

^[18] See Lai et al., 2015 (for multi-dimensional issues with two senders); Gurdal et al, 2014 (for the addition of the limited rational sender) and Gurdal et al., 2013 (for two senders with simultaneous and sequential communication).

^[19] Glazer and Rubinstein (2012) model a persuasion situation where a listener first announces and commits to a set of conditions called codex to which the profile presented by the speaker must conform in order for the listener to be persuaded. The speaker then presents a profile that is not necessarily true. While the speaker tries to persuade the listener to choose his desired action regardless of his true profile, the listener wishes to be persuaded only if the speaker's profile satisfies certain conditions. The speaker's ability to come up with a persuasive profile is limited. The listener knows that the speaker may be providing false information that is not verifiable. The analysis provides a full characterization of the conditions under which the listener's objective can be implemented.

^[20] Wolinsky (2002) and Gerardi et al. (2009) take into account the possibility of communication or collusion among multiple-experts.

^[21] For instance, Galeotti et al. (2013) study multi-player communication in networks.

References

- Alonso, R. and Camara, O. (2015). “*Bayesian Persuasion with Heterogeneous Priors*,” Unpublished Manuscript.
- Ambrus, A. and Takahashi, S. (2008). “Multi-sender Cheap Talk with Restricted State Spaces,” *Theoretical Economics*, 3: 1-27.
- Aumann, R.J. and Hart, S. (2003). “Long Cheap Talk,” *Econometrica*, 71(6): 1619-1660.
- Austin-Smith, D. (1993). “Interested Experts and Policy Advice: Multiple Referrals under Open Rule,” *Games and Economic Behavior*, 5: 1-43.
- Battaglini, M. (2002). “Multiple Referrals and Multidimensional Cheap Talk,” *Econometrica*, 70: 1379-1401.
- Benabou, R. and Laroque, G. (1992). “Using Privileged Information to Manipulate Markets: Insiders, Gurus and Credibility,” *Quarterly Journal of Economics*, 107(3): 921-958.
- Blackwell, D. (1953). “Equivalent Comparisons of Experiments,” *The Annals of Mathematical Statistics*, 24(2): 265-272.
- Brocas, I. and Carrillo, J.D. (2007). “Influence through Ignorance,” *RAND Journal of Economics*, 38(4): 931-947.
- Bull, J. and Watson, J. (2007). “Hard Evidence and Mechanism Design,” *Games and Economic Behavior*, 58(1): 459-490.
- Cai, H. and Wang, J.T. (2006). “Overcommunication in Strategic Information Transmission Games,” *Games and Economic Behavior*, 56(1): 7-36.
- Chakraborty, A. and Harbaugh, R. (2007). “Comparative Cheap Talk,” *Journal of Economic Theory*, 132: 70-94.
- , (2010). “Persuasion by Cheap Talk,” *American Economic Review*, 100(5): 2361-2382.
- Chen, Y. (2011). “Perturbed Communication Games with Honest Senders and Naive Receivers,” *Journal of Economic Theory*, 146: 401-424.
- Crawford, V.P. (2003). “Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions,” *The American Economic Review*, 93(1): 133-149.
- Crawford, V.P. and Sobel, J. (1982). “Strategic Information Transmission,” *Econometrica*, 50(6): 1431-1451.
- Deneckere, R. and Severinov, S. (2008). “Mechanism Design with Partial State Verifiability,” *Games and Economic Behavior*, 64(2): 487-513.
- Dickhaut, J.W., McCabe, K.A., and Mukherjee, A. (1995). “An Experimental Study of Strategic Information Transmission,” *Economic Theory*, 6(3): 389-403.
- Dziuda, W. (2011). “Strategic Argumentation,” *Journal of Economic Theory*, 146: 1362-1397.
- Epstein, D. (1998). “Partisan and Bipartisan Signaling in Congress,” *Journal of Law, Economics, and Organization*, 14(2): 183-204.
- Ettinger, D. and Jehiel, P. (2010). “A Theory of Deception,” *American Economic Review: Microeconomics*, 2(1): 1-20.
- Farrell, J. (1993). “Meaning and Credibility in Cheap Talk Games,” *Games and Economic Behavior*, 5(4): 514-531.

- Farrell, J. and Gibbons, R. (1989). "Cheap Talk with Two Audiences," *The American Economic Review*, 79(5): 1214-1223.
- Forges, F. and Koessler, F. (2008). "Long Persuasion Games," *Journal of Economic Theory*, 143(1): 1-35.
- Galeotti, A., Ghiglino, C., and Squantini, F. (2013). "Strategic Information Transmission Networks," *Journal of Economic Theory*, 148(5): 1751-1769.
- Gentzkow, M. and Kamenica, E. (2011). "Bayesian Persuasion," *American Economic Review*, 101: 2590-2615.
- , (2014). "Costly Persuasion," *American Economic Review*, 104(5): 457-462.
- , (2016a). "Competition in Persuasion," *Unpublished Manuscript*.
- , (2016b). "A Rothschild-Stiglitz Approach to Bayesian Persuasion," *Unpublished Manuscript*.
- Gerardi, D. (2004). "Unmediated Communication in Games with Complete and Incomplete Information," *Journal of Economic Theory*, 114(1): 104-131.
- Gerardi, D., Postlewaite, R., and Squantini, A. (2009). "Aggregation of Expert Opinions," *Games and Economic Behavior*, 65(2): 339-371.
- Gilligan, T. and Krehbiel, K. (1989). "Asymmetric Information and Legislative Rules with a Heterogeneous Committee," *American Journal of Political Science*, 33: 459-490.
- Glazer, J. and Rubinstein, A. (2004). "On the Optimal Rules of Persuasion," *Econometrica*, 72: 1715-1736.
- , (2006). "A Study in the Pragmatics of Persuasion: A Game Theoretical Approach," *Theoretical Economics*, 1: 395-410.
- , (2012). "A Model of Persuasion with Boundedly Rational Agents," *Journal of Political Economy*, 120(6): 1057-1082.
- Gneezy, U. (2005). "Deception: the Role of Consequences," *American Economic Review*, 95: 384-394.
- Goltsman, M. and Pavlov, G. (2011). "How to Talk to Multiple Audiences," *Games and Economic Behavior*, 72(1): 100-122.
- Green, J. and Stokey, N. (2007). "A Two-person Game of Information Transmission," *Journal of Economic Theory*, 135: 90-104.
- Grossman, S.J. (1981). "The Informational Role of Warranties and Private Disclosure about Product Quality," *Journal of Law and Economics*, 24(3): 461-483.
- Gurdal, M.Y., Ozdogan, A., and Saglam, I. (2013). "Cheaptalk with Simultaneous versus Sequential Messages," *MPRA Paper 45727*, University Library of Munich, Germany.
- , (2014). "Truth-telling and Trust in Sender-receiver Games with Intervention: An Experimental Study," *Review of Economic Design*, 18(2): 83-103.
- Hendricks, K. and McAfee, R.P. (2006). "Feints," *Journal of Economics and Management Strategy*, 15(2): 431-456.
- Honryo, T. (2011). "Dynamic Persuasion," *Unpublished manuscript*.
- Horner, J. and Skrzypacz, A. (2014). "Selling Information," *Working paper*.

- Hurkens, S. and Kartik, N. (2009). "Would I Lie to You? On Social Preferences and Lying Aversion," *Experimental Economics*, 12: 180-192.
- Ivanov, M. (2015a). "Dynamic Information Revelation in Cheap Talk," *The B.E. Journal of Theoretical Economics*, 15: 251-275.
- , (2015b): "Dynamic Learning and Strategic Communication," *International Journal of Game Theory*. 1-27
- Kartik, N. (2009). "Strategic Communication with Lying Costs," *Review of Economic Studies*, 76: 1359-1395.
- Kartik, N., Ottaviani, M., and Squintani, F. (2007). "Credulity, Lies and Costly Talk," *Journal of Economic Theory*, 134: 93-116.
- Kartik, N. and Tercieux, O. (2012). "Implementation with Evidence," *Theoretical Economics*, 7(2): 323-355.
- Khrisna, V. and Morgan, J. (2001a). "Asymmetric Information and Legislative Rules: Some Amendments," *American Political Science Review*, 95: 435-452.
- , (2001b). "A Model of Expertise," *Quarterly Journal of Economics*, 116: 747-775.
- Kolotilin, A. (2014). "Optimal Information Disclosure: Quantity vs. Quality," *Unpublished Manuscript*.
- Kolotilin, A., Li, M., Mylovanov, T., and Zapechelnyuk, A. (2014). "Persuasion of a Privately Informed Receiver," *Unpublished Manuscript*.
- Lai, E.K., Lim, W., and Wang, J.T. (2015). "An Experimental Analysis of Multidimensional Cheap Talk," *Games and Economic Behavior*, 91: 114-144.
- Levy, G. and Razin, R. (2008). "Multidimensional Cheap Talk," *Unpublished manuscript*.
- Li, F. and Norman, P. (2015). "On Bayesian Persuasion with Multiple Senders," *Unpublished Manuscript*.
- Li, Z., Rantakari, H., and Yang, H. (2014). "Competitive Cheap Talk," *Unpublished Manuscript*.
- Mathis, J., McAndrews, J., and Rochet, J.C. (2009). "Rating the Raters: Are Reputational Concerns Powerful Enough to Discipline Rating Agencies?," *Journal of Monetary Economics*, 56: 657-674.
- McGee, A. and Yang, H. (2013). "Cheap Talk with Two Senders and Complementary Information," *Games and Economic Behavior*, 79: 181-191.
- Milgrom, P.R. and Roberts, J. (1986). "Relying on the Information of Interested Parties," *The RAND Journal of Economics*, 17(1): 18-32.
- Milgrom, P.R. (1981). "Good News and Bad News: Representation Theorems and Applications," *Bell Journal of Economics*, 12: 380-391.
- Morris, S. (2001). "Political Correctness," *Journal of Political Economy*, 109(2): 231-265.
- Mullainathan, S., Schwartzstein, J., and Shleifer, A. (2008). "Coarse Thinking and Persuasion," *Quarterly Journal of Economics*, 123(2): 577-619.
- Ottaviani, M. and Sorensen, P.N. (2006). "Reputational Cheap Talk," *RAND Journal of Economics*, 37: 155-175.
- Ottaviani, M. and Squintani, F. (2006). "Naive Audience and Communication Bias," *International Journal of Game Theory*, 35: 129-150.

- Rayo, L. and Segal, I. (2010). "Optimal Information Disclosure," *Journal of Political Economy*, 118(5): 949-987.
- Rothschild, M. and Stiglitz, J. (1970). "Increasing Risk: I. A definition," *Journal of Economic Theory*, 2: 225-243.
- Sanchez-Pages, S. and Vorsatz, M. (2007). "An Experimental Study of Truth-telling in a Sender-receiver Game," *Games and Economic Behavior*, 61: 86-112.
- , (2009). "Enjoy the Silence: An Experiment on Truth-Telling," *Experimental Economics*, 12: 220-241.
- Sher, I. (2011). "Credibility and Determinism in a Game of Persuasion," *Games and Economic Behavior*, 71(2): 409-419.
- Sher, I. (2014). "Persuasion and Dynamic Communication," *Theoretical Economics*, 9: 99-136.
- Shin, H. (1994). "The Burden of Proof in a Game of Persuasion," *Journal of Economic Theory*, 64: 253-264.
- Sobel, J. (1985). "A Theory of Credibility," *Review of Economic Studies*, 52(4): 557-573.
- Verrecchia, R.E. (1983). "Discretionary Disclosure," *Journal of Accounting and Economics*, 5: 179-194.
- Wang, Y. (2015). "Bayesian Persuasion with Multiple Receivers," *Unpublished Manuscript*.
- Wolinsky, A. (2002). "Eliciting Information from Multiple Experts," *Games and Economic Behavior*, 41(1): 141-160