

Giving and Receiving Advice*

Joel Sobel[†]

May 5, 2011

Abstract

This paper reviews literature on communication between informed experts and uninformed decision makers. The research provides some insight into what constitutes a persuasive statement and under what conditions a decision maker will benefit from consulting an expert. I classify the literature along four dimensions: strategic, technological, institutional, and cultural. To the extent that decision makers and experts have different preferences, communication creates strategic problems. Technological considerations describe the domain of uncertainty, the cost of acquiring information, and the cost of manipulating information. The institution determines who has responsibility for making decisions and the rules that govern communication. Cultural factors describe the way in which agents interpret language.

Keywords: communication, signaling.

JEL Classification Numbers: D02, D03, D44.

*I thank Nageeb Ali, Marco Battaglini, Andreas Blume, Antonio Cabrales, Ying Chen, Vincent Crawford, Stefano Demichelis, Wouter Dessein, Eddie Dekel, Navin Kartik, Kohei Kawamura, Chulyoung Kim, Young-Gwan Kim, Frédéric Koessler, David Miller, Roger Myerson, Philip Neary, Lucas Siga, Ran Spiegler, Joel Watson, and Kathryn Woolard for encouragement and wisdom. I am grateful to the NSF for financial support.

[†]Department of Economics, University of California, San Diego, La Jolla, CA 92093, U.S.A. E-mail: jsobel@ucsd.edu

1 Introduction

Models of signaling, adverse selection, and moral hazard make it clear that differences in information can lead to inefficiency. Workers may invest in non-productive education in order to convince an employer they have skills. Markets may fail when sellers know more about the quality of their item than buyers. Risk-neutral principals may fail to offer perfect insurance to risk-averse workers when they cannot directly observe the workers' choice of effort.

If asymmetric information leads to inefficiency, why can't agents improve outcomes by direct, costless communication? This question motivates the study of cheap-talk models. It has a straightforward answer. In simple economic environments it is not in the interest of one agent to reveal private information to another. Naive and honest sellers who accurately reveal the quality of their item may lose all bargaining power. Credulous employers who believe a worker's claims about productivity encourage workers to make exaggerated claims about their quality. Adding realistic complications to these situations creates situations in which this kind of cheap talk can be beneficial.

In this essay, I discuss some concepts central to the study of communication. I illustrate most of the ideas using a model in which there is an informed agent who has the ability to communicate with an uninformed decision maker. The informed Sender gives advice. The uninformed Receiver decides how to use it. The fundamental questions in this environment are: How credible can communication be? What factors lead to no communication? What conditions are consistent with full communication? How does one organize environments so that communication is most effective?

The next section outlines models in which one can formulate these questions. Section 3 describes how the differences in preferences between the Sender and Receiver may limit the ability to communicate. The Sender may have strategic reasons for distorting information and the Receiver may have strategic reasons to be skeptical. The theory formalizes the intuition that communication is more successful if the interests of the parties are more closely aligned.

Interesting communication problems must be associated with multiple equilibria. In Section 4, I argue that non-trivial communication can be guaranteed only if players use exogenous factors that induce common understanding of messages. Theory does not permit the modeler to make precise pre-

dictions without making assumptions about how agents interpret language. The cultural context of interaction may place restrictions on the relationship between the Sender's private information and the messages she uses to describe it.

Section 5 studies the communication problem when there are many informed agents. Adding additional sources of information typically increases the opportunities for information exchange in standard models even when the informed agents have access to identical information.

In a simple environment with a single informed agent and a single decision maker, direct one-shot communication is one of many ways to structure the communication problem. To gain an insight into how best to communicate, it is essential to look at the problem of institution design. Section 6 examines models that give some actor commitment power and describes how to structure communication to benefit the uninformed agent.

Section 7 looks at the institution design problem from the perspective of the informed agent. When the Sender has commitment power, the Receiver can be more credulous as the Sender can convincingly promise not to distort her information when it is in her strategic interest to do so. Section 8 shows that related results are possible when information is verifiable so that what the Sender can say is constrained by what the Sender knows.

Section 9 points out some connections between economic models of communication and linguistics.

Section 10 and 11 discuss two significant variations on the basic model that call attention to assumptions on the communication technology. Section 10 relaxes the assumption that information is exogenously given. The actors have different opinions not only about the final decision, but about how much to invest in collecting information. There is typically a tension between factors that provide incentives to gather information and those that provide incentive to communicate the information accurately.

Section 11 points out that the limiting factor in communication may not be incentive problems, but the complexity of information itself. It describes some issues that arise when information is costly to prepare and costly to interpret.

Section 12 is a conclusion.

2 Framework

I use a simple model to describe communication. The essential features are two agents, an informed Sender (S) and an uninformed Receiver (R). These agents have preferences that depend on an action $a \in A$ and a state of the world $\theta \in \Theta$. Denote these preferences by $U_i(\cdot)$ for $i = R, S$. Nature selects θ according to a common-knowledge distribution, $p(\cdot)$. For most of the discussion, I assume that the Sender learns θ without cost.

Starting from these basic assumptions, there are several possible ways to complete the description of the strategic interaction. I introduce some variations in this section. In a cheap-talk game (Crawford and Sobel [21] and Green and Stokey [41]) there is an abstract set of messages M , a set of states (or types) Θ , and a set of actions, A . S learns θ , selects a message $m \in M$ and R selects an action $a \in A$. Formally, a (mixed) strategy for the Receiver is denoted by α , where $\alpha(a | m)$ is the probability that the Receiver takes action $a \in A$ given message $m \in M$. A (mixed) strategy for the Sender is denoted by σ , where $\sigma(m | \theta)$ is the probability that the Sender sends message m when her type is $\theta \in \Theta$. An equilibrium consists of a strategy profile (α^*, σ^*) and a belief function μ^* ($\mu^*(\theta | m)$ is the probability that the Receiver believes that the Sender's type is θ given message m), such that α^* is a best response to μ^* ; σ^* is a best response to α^* ; and μ^* is consistent with the prior and σ^* in the sense that Bayes's Rule determines $\mu^*(\cdot | m)$ for all m such that $\sum_{\theta \in \Theta} p(\theta) \sigma^*(m | \theta) > 0$.¹ I will discuss cheap-talk models in more detail in Section 3 and Section 4. Section 5 discusses models with many Senders.

In cheap-talk games, the set of messages available to the Sender does not depend on the state of the world. Disclosure games arise when what the Sender knows constrains what the Sender can say. Grossman [43] and Milgrom [66] introduce models of hard, or verifiable, information. When information is verifiable, the set of messages M is the set of all subsets of Θ and S 's strategies are constrained so that the Sender of type θ can only use a message m if $\theta \in m$. That is, the Sender can withhold information, but cannot misrepresent information. I discuss these ideas in Section 8.

The basic model of cheap talk assumes that players lack the ability to make commitments. To answer the question of how to design organizations

¹Jung [51] comments on the appropriate definition of equilibrium when strategy spaces are large.

to improve communication, it is important to consider the possibility that one or the other agent has the ability to commit to a strategy or that the players can use a third party to facilitate the interaction. Section 6 and Section 7 discuss environments in which one player has commitment power.

This paper neglects many topics including career concern or dynamics, monetary transfers, and privately informed Receivers.

3 Strategic Considerations

If the Receiver has a best response to the prior distribution (an assumption that I maintain throughout this essay), then cheap-talk games have a “babbling” equilibrium in which there is no communication. Assume that the uninformed agent takes the same action independent of the signal (this action should be a best response to the prior). In this case all Sender types are indifferent between all signals. In particular, it is a best response for all of them to say the same thing.² This equilibrium outcome is similar to the pooling outcome of classical signaling models. In labor-market signaling, however, the response to out-of-equilibrium messages typically must be different than the response to the message used on the path of the pooling equilibrium (in particular, if there is pooling at a positive level of education, then the response to “no education” must be a lower salary than the equilibrium salary, otherwise no one would get education). In cheap-talk games all responses can be the same (and, in some examples, they *must* be the same) as the response on the equilibrium path.

The observation that cheap-talk games always have babbling equilibria motivates two important questions one can ask about communication games. Under what conditions is babbling the only equilibrium outcome? Under what conditions is there a fully revealing equilibrium?

When the only equilibria involve babbling, the opportunity to communicate freely does not enlarge the set of predictions. In this case, cheap talk cannot ameliorate problems caused by asymmetries in information. When

²One can always take a game and form a new game in which players can engage in costless pre-play communication. When the underlying game has an equilibrium, the augmented game will have a babbling equilibrium in which all messages are ignored. Seidman [75] provides an example of a game with pre-play communication that has a unique informative equilibrium outcome but no babbling equilibrium. This is possible because the game without communication does not have an equilibrium.

there is a fully revealing equilibrium, cheap talk can eliminate all informational asymmetries. The presence of a revealing equilibrium is not sufficient to predict that the Receiver will learn everything that the Sender knows, however, because the babbling equilibrium always exists. Still, characterizing when a fully revealing equilibrium exists identifies cases in which full communication is possible.

Communication can change beliefs, actions, and payoffs. Communication is *informative* if $\mu^*(\cdot | m)$ is not constant on the equilibrium path.³ Communication is *influential* if $\alpha^*(\cdot | m)$ is not constant on the equilibrium path.⁴ Communication is *payoff relevant* for the Receiver if $E[U_R(\alpha^*(\sigma^*), \theta)] > \max_{a \in A} E[U_R(a, \theta)]$. In words, an informative equilibrium is one in which the Receiver changes his beliefs after some message. An influential equilibrium is one in which the Receiver does not always take the same action. Communication is payoff relevant for the Receiver if it permits him to improve his payoff relative to the babbling equilibrium.⁵ An informative equilibrium will fail to be influential if R does not use the information contained in a message to change his action.⁶ In order for communication to be payoff relevant for the Receiver it must be both informative and influential. The converse fails in general, but holds in non-degenerate situations.

A necessary condition for the existence of influential communication is that the Sender and the Receiver have some interests in common. More precisely, there must exist non-empty sets of types T_1 and T_2 , and distinct actions a_1 and a_2 such that

- i. a_i is a best response (for the Receiver) to $\theta \in T_i$.
- ii. If $\theta \in T_i$ the type θ Sender prefers a_i to a_j .

These conditions identify shared interest between Sender and Receiver because conditional on $\theta \in T_i$, the Sender and Receiver have the same preferences over $\{a_1, a_2\}$.

³A message m is on the equilibrium path if it is sent with positive probability ($\sum_{\theta \in \Theta} p(\theta) \sigma^*(m | \theta) > 0$).

⁴Chakraborty and Harbaugh [15] use the term influential equilibrium to describe what I call informative equilibrium.

⁵Payoff-relevant communication for a type θ Sender changes her payoff relative to a babbling equilibrium. Relative to babbling, payoff-relevant communication must increase the Receiver's expected utility but may make the Sender worse off.

⁶An influential equilibrium can fail to be informative in degenerate cases where there are multiple best replies to the prior distribution.

There are many interesting situations in which cheap talk may be influential. Crawford and Sobel [21] study a one-dimensional model in which the Sender has a uniform bias (the Sender's ideal action is strictly greater than the Receiver's action for each θ). Her most preferred action for each state is strictly higher than that of the Receiver. If the bias is sufficiently small, different Sender types will have different preferences and there are partially informative equilibria in this model.

In some circumstances, it is natural to assume that the Sender's preferences are independent of type. If the Receiver's action is a wage payment to the Sender and the Sender's utility is strictly increasing in money, then one would not expect informative equilibria. If the game is a persuasion game in which the Receiver must decide whether to accept or reject a proposal and the Sender prefers to maximize the probability of acceptance, then preferences again are independent of type. In these natural situations, cheap talk is not influential.

Chakraborty and Harbaugh [14] show how common interests between Sender and Receiver arise naturally when A is multidimensional. Consider first a one-dimensional setting in which the Sender wants the Receiver to take the highest possible action independent of type. This environment corresponds to many applications in which the Receiver's action is the Sender's payment. I have argued that cheap talk is not influential in this setting. Now imagine that both the state space and the action space are two dimensional. Further assume that the Sender has linear and increasing preferences over the Receiver's actions, but that the marginal rate of substitution depends on the state and the components of the state are independently distributed. Under these conditions, Chakraborty and Harbaugh [14] demonstrate how equilibria in which the Sender communicates categorical information – which component is higher – can exist when A is multidimensional.

They extend this insight in Chakraborty and Harbaugh [15], which shows that informative equilibria exist in multidimensional models when the Sender's preferences are state independent. Notice that in this environment all Sender types must be indifferent between all messages used in equilibrium, so Condition ii cannot hold with strict preference in an informative equilibrium. The construction in Chakraborty and Harbaugh [15] is simple. Since the Sender has state-independent preferences, U_S depends only on action. An informative equilibrium exists if it is possible to partition the state space into two non-empty regions, Θ_1 and Θ_2 , such that if a_i is the Receiver's best response given that $\theta \in \Theta_i$, then $U_S(a_1) = U_S(a_2)$. In this case, the Sender would

be indifferent between the two actions and it would be a best response to communicate which Θ_i contains her type. Take a point θ_0 in the interior of the state space and imagine a hyperplane through θ_0 with normal p . This hyperplane divides the state space into two non-empty regions. Let $a_i(p)$ be the best response of the Receiver given that the state is in Region i . There will be an informative equilibrium in which the Sender's message identifies her region if and only if $U_S(a_1(p)) - U_S(a_2(p)) = 0$. Chakraborty and Harbaugh point out that it follows from the Borsuk-Ulam theorem that there exists a p such that $U_S(a_1(p)) - U_S(a_2(p)) = 0$.

Chakraborty and Harbaugh [15] make another useful observation. In all cheap-talk models, the Receiver (weakly) prefers an informative equilibrium to the babbling outcome. The Sender's ex ante preferences are less clear. Chakraborty and Harbaugh point out that relative to the babbling equilibrium, the Sender prefers a communicative equilibrium if her utility function is quasiconvex. The observation follows (in their model) because increasing information leads to mean-preserving spreads in the distribution of actions induced and increases in risk are beneficial under quasiconvexity.

Chakraborty and Harbaugh's arguments do not guarantee the existence of informative equilibria in all multidimensional models. In fact, the clear intuition from one-dimensional models that large differences in preferences make effective communication impossible still holds. Levy and Razin [61] demonstrate that even in multi-dimensional models, informative equilibria typically do not exist when there is sufficiently strong conflict of interest between the Sender and Receiver. Imagine a two-dimensional model in which the Receiver wishes to minimize the distance between the state and the action and the Sender has a bias of the form $(b, 0)$, so that she agrees with the Receiver's preferences along the second dimension, but has a bias in the first component. The preceding analysis suggests that the Receiver could ask the Sender to report the second component honestly. This argument works if the two components of θ are independently distributed. Denoting the average of θ_1 by $\bar{\theta}_1$, the game has a type-action equilibrium distribution in which the Receiver takes the action $(\bar{\theta}_1, \theta_2)$ when the state is θ . When the components are not independently distributed, this argument breaks down. If the Receiver expects the Sender to report θ_2 accurately, then he sets the first component of his action equal to the conditional expectation θ_1 given θ_2 . Faced with this, the Sender is unlikely to report honestly. Indeed, Levy and Razin [61] provide conditions under which large biases rule out informative equilibria even in multidimensional models.

Battaglini [8] observes that in n -dimensional models with quadratic preferences, the Sender and Receiver have common interests along $n - 1$ dimensional subsets, which means that there will exist equilibria in which infinitely many actions are induced even when there is a conflict of interest between Sender and Receiver. Section 5 describes the intuition behind this result. Battaglini points out that this observation is not general and the literature contains no precise statement on how expanding the domain of preferences influences opportunities to communicate.

In addition to the Sender and Receiver having some interests in common, in typical situations, influential communication requires a conflict of interest between Sender types. If Condition ii holds strictly for some Sender type, then it will be the case that there exists a pair of Sender types and a pair of actions such that the Sender types have different preferences over the actions. Alternatively, suppose that all types of Sender have identical, strict preferences over the Receiver's actions. The Receiver can take only one action in equilibrium. To see this note that the Receiver's strategy α makes available a set of actions, $A_\alpha = \{\alpha(m) : m \in M\}$.⁷ By assumption, the Sender's favorite action in A_α is unique and does not depend on type, so she must send a message that induces this action. While it is possible that the Sender's message is informative (there may be two different messages that lead to the same action), the Receiver does not use the information to determine his action.

When the set of actions is not linearly ordered, there will be conflicts between Sender types even in the simplest economic examples. Consider the following formal modification of the standard labor-market signaling model. The student learns her type and makes a cheap-talk message to the "market" of potential employers. The market responds by associating with each message a contract in the form of a wage and education pair.⁸ That is, the worker writes a cheap letter of application to an employer and receives in response a binding contract to work for a particular salary (conditional on providing a certain level of education). The set of actions available to the Receiver is two dimensional (wage, education pairs) and different Sender types have different preferences (because they have different marginal rates of substitution

⁷In general, $\alpha(\cdot)$ can be a probability distribution over A . In applications one typically rules this out by assuming that the Receiver has a unique best response to each distribution over types.

⁸To treat this as a two-player game, assume that the Receiver wants to minimize the difference between the wage offered and the expected productivity of the worker.

between wages and education). Standard arguments from signaling games show that informative equilibria exist in this model.

Conflict between Sender types is a necessary condition for payoff-relevant communication, but I have argued that this condition holds in interesting environments. Under what conditions does a fully revealing equilibrium exist? A necessary and sufficient condition for fully revealing equilibria is a strong form of common interests: Assume that $\max U_R(a, \theta)$ has a unique solution for each θ . Denote the solution by a^R . Fully revealing equilibria exist if and only if $U_S(a^R(\theta), \theta) \geq U_S(a^R(\theta'), \theta)$ for all θ and θ' .

4 Equilibrium Selection

Effective communication requires coordination. The Sender must send informative messages and the Receiver must understand them. This creates a multiple-equilibrium problem that is distinct from the kind that affects standard signaling models in which preferences satisfy a single-crossing condition.

There are three kinds of indeterminacy in cheap-talk models: multiple off-the-path responses, multiple meanings of messages, and multiple equilibrium associations between types and actions. The first indeterminacy is familiar in dynamic games. There may be many equilibria that give rise to the same outcome. This kind of indeterminacy is the result of different possible specifications of behavior off the path of equilibrium. The second and third kinds of indeterminacy involve multiple equilibrium outcomes.⁹ The possibility of multiple equilibrium outcomes is the focus of attention in standard signaling games. For these games, selection arguments based on Kohlberg and Mertens [57]’s notion of strategic stability frequently focus attention on separating outcomes that are efficient from the point of view of the Sender.¹⁰ In cheap-talk games, this type of multiplicity is inevitable.

The second kind of indeterminacy, multiple meanings of messages, arises in any non-babbling outcome. One can take any equilibrium outcome and form a new equilibrium outcome by permuting the messages. Consequently, when messages have no intrinsic meaning, one cannot deduce the meaning of a message from equilibrium analysis alone. If there is an equilibrium in which type θ sends message m and induces action a and type θ' sends message m'

⁹An equilibrium outcome is the distribution over types, messages, and actions induced by the equilibrium strategies.

¹⁰See Banks and Sobel [7] and Cho and Kreps [19].

and induces action a' , then there must be an equilibrium in which type θ sends message m' and induces action a and type θ' sends message m and induces action a' . Hence statements of the form “in all equilibria, type θ sends message m ” or “in all equilibria, message m induces action a ” cannot be true.

The third indeterminacy is that in equilibrium there will generally be many associations of types to actions.¹¹ Since babbling equilibria always exist, if there exists an equilibrium with influential communication, there must be multiple equilibrium type-action distributions.

The first kind of indeterminacy is familiar and does not prevent the modeler from making definite predictions about outcomes. Possibly the second kind of indeterminacy is an advantage of the modeling approach. One obtains predictions without making exogenous assumptions about the meaning of messages. If the analyst (like the players) does not care about the actual choice of messages but only about the payoff-relevant association of actions to types, this kind of multiplicity of equilibrium is not a problem. One cannot dismiss the existence of multiple equilibrium type-action distributions, however, and some approaches to the problem involve making assumptions that assign particular meanings to specific messages.

Farrell [31] attacks the multiple-equilibrium problem by requiring that some out-of-equilibrium messages have unambiguous interpretations. He assumes that there always exist unused messages and argues that given an equilibrium profile (α^*, σ^*) the Receiver would interpret a new message to mean “My type is an element of K . Take action a^* .” if there exists an action $a^* \in A$ and a set $K \subset \Theta$ such that

$$a^* \text{ solves } \max_{\theta \in K} \sum_{\theta \in K} U_R(a, \theta) p(\theta) \quad (1)$$

and

$$K = \{\theta : U_S(a^*, \theta) > E[U_S(\alpha^*(\sigma^*(\theta)), \theta)]\}. \quad (2)$$

Condition (1) states that a^* is a best response to the prior belief conditional on $\theta \in K$. Equation (2) states that precisely those types in K prefer the payoff generated by inducing a^* to the equilibrium payoff. If there exists

¹¹Given strategies (α^*, σ^*) the distribution induced on actions by each type, $\gamma(a | \theta)$, is

$$\gamma(a | \theta) = \sum_m \alpha^*(a | m) \sigma^*(m | \theta).$$

a^* and K satisfying (1) and (2), then Farrell says that a *credible neologism* exists. When a credible neologism exists, the Sender can send a message that will induce an action attractive to a non-empty subset of types. For this reason, Farrell proposes refining equilibria by restricting attention to those in which no credible neologism exists.¹²

The assumption that there exist unused messages is difficult to formulate in the context of an equilibrium theory because associated with any equilibrium type-action distribution there exists an equivalent distribution in which all messages are used with positive probability. On the other hand, the existence of unused messages – or the ability of agents to create new ones – is in the spirit of the theory. Farrell’s formulation refines away all equilibria in many settings (including the Crawford-Sobel model). While the literature contains variations of Farrell’s idea (notably that of Matthews, Okuno-Fujiwara, and Postlewaite [64]), no paper provides clear guidance about selecting equilibria, either because it lacks general existence properties or, in the non-equilibrium models of Rabin [72] and Zapater [80], because they fail to rule out equilibrium outcomes in interesting settings with multiple equilibria.¹³

Evolutionary arguments provide another approach to the problem of equilibrium selection. The basic idea is due to Robson [74]. Imagine an environment in which large populations (one group of Senders and one group of Receivers) play the basic game in pairs and in which, after each period, successful strategies become more prevalent, less successful strategies die out, and novel strategies may invade. Take a game with a unique Pareto-efficient payoff, so the players’ interests are perfectly aligned. Fix an equilibrium that fails to arrive at an efficient outcome. The population can drift to a configuration in which some messages are used with zero probability, but the existing strategies do not punish a strategy that uses one of these messages. At this point, there can be an invasion of a strategy that uses the novel message. When the invading strategy meets a member of the existing population, it does not lose. When it meets another individual playing the invading strategy, the invaders can use the novel message as a “secret

¹²Grossman and Perry [44] introduce a variation of this idea and apply it to a more general class of games.

¹³Myerson [71] presents an axiomatic characterization of a solution concept related to Farrell’s that has general existence properties when the informed player has commitment power and correlation is feasible. de Groot Ruiz, Offerman, and Onderstal [23] develop a modification of Farrell’s ideas to organize some experimental results on communication.

handshake” that permits them to coordinate on the efficient outcome. The invading strategy therefore earns a strictly higher payoff than the existing strategies and will move the population towards the efficient outcome. The population coordinates on the efficient equilibrium outcome. This argument applies to repeated games, communication about intentions, and communication games.¹⁴ The evolutionary arguments resolve the third kind of indeterminacy (in a restricted class of games), but do not predict the meaning of particular messages, so the second kind of indeterminacy remains. Since the modeler does not make any assumptions about meaning initially, one cannot hold that a particular message takes on a particular meaning in equilibrium. A weakness of the evolutionary approach is that the results require strong assumptions about common interests.¹⁵

More recent work uses the existence of behavioral types or payoff perturbations to select equilibria in variations of the standard one-dimensional cheap-talk model. Chen [17] assumes that with positive probability the Sender reports her type honestly and with positive probability the Receiver interprets a message literally. She then demonstrates that only one of the equilibrium type-action distributions in the standard cheap-talk model is the limit of equilibria in the perturbed game as the probability of behavioral types goes to zero.¹⁶ In independent work, Kartik [53] assumes that the Sender faces a cost of lying and, like Chen, gives conditions under which only one of the equilibrium outcomes in the standard cheap-talk model is the limit of refined equilibria of the perturbed game. Chen, Kartik, and Sobel [18] point out that the limit equilibria in Chen’s paper satisfies a condition that, if imposed directly, selects equilibria in the standard cheap-talk model. The same selection arises in Kartik [53]’s model of costly lying and in Gordon [40]’s dynamic analysis. Because of the perturbations, messages have an exogenous meaning in Chen and Kartik’s models. The message space cannot be arbitrary. It must be linked to the set of Sender types in order to define honesty. As a consequence, the analysis selects an equilibrium outcome and not only an equilibrium type-action distribution. In particular, in an environment in

¹⁴Blume, Kim, and Sobel [12], Kim and Sobel [56], Demichelis and Weibull [24], Gilboa and Matsui [35], Sobel [76], and Wärneryd [79] provide details.

¹⁵Blume, Kim, and Sobel [12] establish selection in communication models under a restrictive partial common interest assumption. The other referenced papers essentially require that there is a unique Pareto efficient feasible payoff.

¹⁶Chen requires that the equilibria in the perturbed game satisfy a monotonicity refinement.

which the Sender would like to convince the Receiver that her type is higher than it really is, in equilibrium the Sender tends to exaggerate by sending a message that would be honest for a higher type.

While Farrell’s approach assumes that some messages have natural meanings that will be believed if they do not violate strategic motivations of agents, it does not directly associate specific messages to specific types. For Farrell it is important that there is a message that means “my type is θ ,” but there is no reason why the message that means “my type is θ' ” should be related to “my type is θ ” even when θ and θ' are close. Chen and Kartik’s models do preserve a connection between messages and types.¹⁷ It appears that simply connecting messages to types or actions through lying costs or behavioral types is not sufficient to produce a unique correspondence between meanings and messages. Both Chen and Kartik add an equilibrium refinement – a condition that guarantees monotonicity of strategies as functions of messages – in order to select a specific equilibrium type-action distribution. It is not clear whether arguments using only monotonicity restrictions (and not strategy perturbations) are sufficient to select equilibria in simple cheap-talk games.

Cheap-talk models treat messages abstractly. Because messages have no intrinsic meaning, games have a layer of strategic complexity, the need to coordinate on a common language, that does not exist in other signaling games. One cannot make confident predictions about the effectiveness of communication without making assumptions about the extent to which players share a common language. It should be a goal of theory to model messages with exogenous meanings and the extent to which these meanings are manipulated by strategic actors. I am convinced that a productive approach to this problem is to impose restrictions on the use of messages within a game that capture the way that messages are used outside of strategic interactions. A natural conjecture is that communication would be more effective the greater the players have a common “culture” that leads to shared interpretation of available messages and places constraints on the ways messages are used and interpreted. Laboratory evidence may clarify this conjecture.¹⁸

¹⁷Lo [63] combines restrictions on strategy spaces and dominance arguments to study cheap-talk outcomes.

¹⁸Crawford [20] reviews some experimental work on cheap-talk communication in games including direct tests of the Crawford and Sobel [21] model.

5 Many Senders

It is natural to seek advice from more than one person. The literature identifies two ways in which having multiple informants may make an uninformed decision maker better off. Having multiple Senders allows the Receiver to check facts. If the Senders have access to identical information, the Receiver expects identical messages. If the set of actions is sufficiently rich, then the Receiver can respond to conflicting messages with an action that punishes both Senders, inducing informative reporting in equilibrium. Alternatively, while the interests of the Receiver may differ from those of the Senders, an individual Sender may be willing to provide information along a dimension of common interest. Multiple Senders supplying complementary information may permit the Receiver to reconstruct the state.

Motivated by legislative decision making, Gilligan and Krehbiel [36] introduce a cheap-talk model with multiple informed agents and a one-dimensional decision. Krishna and Morgan [58] show that when Senders have symmetrically biased, quadratic utilities, and the set of possible states is sufficiently large, then there exists a fully revealing equilibrium.¹⁹ In this simple setting, the construction of the equilibrium is straightforward. The Receiver's strategy when he receives two identical messages is to treat the common message as an honest report of the state. When the messages are different, the Receiver takes an extreme action. If the set of rationalizable actions available to the Receiver is sufficiently large, there will be an extreme action that supports the equilibrium.

This construction appears to depend on the assumption that the Senders receive identical signals, which allows the Receiver to construct off-the-equilibrium path punishments. It is natural to investigate whether fully revealing equilibria can be approximated by equilibria when Senders receive noisy information about the state of the world. Battaglini [8] demonstrates that, when Senders have large opposing biases, fully revealing equilibria in one-dimensional models fail to be robust in that they are not limits of equilibria of games where Senders receive noisy information about the state. The results in this area are subtle. Ambrus and Lu [3] construct equilibria in

¹⁹Formally, the model consists of one in which there are three players. Sender I has preferences $U_{S_I}(a, \theta, a) = -(\theta - a - b)^2$; Sender II has preferences $U_{S_{II}}(a, \theta) = -(\theta - a + b)^2$; and the Receiver has preferences $U_R(a, \theta) = -(\theta - a)^2$. The Senders simultaneously observe t and simultaneously make reports to the Receiver, who chooses an action based on the messages.

which the Receiver obtains approximately full information when Senders have noisy signals. Ambrus and Lu assume that the Senders learn the true state with positive probability strictly between zero and one and observe a state-independent random variable otherwise. They assume that the Sender and Receiver’s preferences are sufficiently similar that the type θ Sender strictly prefers $a^R(\theta)$ to extreme actions. They construct an equilibrium that approximates full revelation when the probability of the Senders learning the truth goes to one. Their construction works like this. Break the interval of types into n^2 intervals. Label the intervals (i, j) where $i = 0, \dots, n - 1$ and $j = 1, \dots, n$. Sender 1 reports $m_1 = j$ when she observes θ in interval (i, j) . Sender 2 reports $m_2 = j + i \bmod n$ when she observes θ in interval (i, j) . The Receiver responds optimally to the message. The signaling strategy has two important properties. When the noise is small, the Receiver will be confident that the state is in interval $(m_2 - m_1, m_1)$. Hence if the Senders have accurate information and follow the equilibrium strategy, the Receiver will also obtain accurate information. More subtly, if both Senders observe the true state and one sender follows the proposed strategy, then any deviation by the other Sender leads the Receiver to infer that the state is at least $T/(2n)$ units from the interval associated with the common signal. The coding of messages makes any deviation lead to a large change in the Receiver’s beliefs. If T is large and signals are accurate, deviations will therefore lead to “extreme” actions that are not attractive. Consequently the construction leads to an outcome that approximates full revelation.

Battaglini’s [8] analysis of multi-dimensional models demonstrates that the existence of fully revealing equilibria does not depend on the punishments used in one-dimensional models. When the set of actions has $n > 1$ dimensions and agents seek to minimize the distance between the action selected and an ideal point, then a Sender and the Receiver have a hyperplane on which their interests agree. Specifically, assume the players seek to minimize the distance between the action selected x and $b + \theta \in \mathbb{R}^n$ while the Receiver wishes to minimize the distance between x and θ . Here the Sender and Receiver have common interests on the hyperplane $A \cdot z = c$ where $A = b/\|b\|$. Battaglini shows that as long as the Senders have different biases, one can construct an equilibrium in which the Receiver obtains full information.²⁰ In

²⁰In a revealing equilibrium the Receiver can have a response function $a^R(\theta_1, \theta_2) = \theta_1$ when $\theta_1 = \theta_2$ and a point on the intersection of the hyperplanes $\{z : b_i z = (\|b_i\|/\|b_j\|)b_j \theta_j\}$ otherwise.

equilibrium each Sender is willing to project her information honestly onto a subspace of common interest and the Receiver can combine the two reports to obtain full information about the state of the world. This construction is more robust than the one-dimensional construction, although my summary ignores two important aspects of the construction. First, Battaglini’s arguments apply to more general preferences in which there is only a local dimension of common interest between Sender and Receiver. Battaglini’s argument is more subtle than my discussion suggests. Second, as pointed out by Ambrus and Takahashi [4], the construction depends on the geometry of the set of rationalizable actions. If the set of actions A is not all of \mathbb{R}^n , then the sets of common interest will not be hyperplanes, but intersections of hyperplanes with A . The constraint that the Receiver’s response must lie in A can (and will when biases are large) create a conflict of interest between the Sender and Receiver.

It is not hard to get more insight into the importance of role that A plays in the construction of fully revealing equilibria. Ambrus and Takahashi [4] identify necessary and sufficient conditions on preferences and action sets for the existence of fully revealing equilibria (Battaglini [8] earlier identified a special case of these conditions for a one-dimensional model). Intuitively, the smaller the set of feasible actions, the harder it is to support fully revealing equilibria. The reason for this is that the construction of equilibrium requires the Receiver to find “unattractive” actions as responses to unexpected pairs of messages. The larger the action space, the more likely such actions exist.²¹ Here is a sketch of the formal result in Ambrus and Takahashi. An argument due to Battaglini (and similar to standard proofs of the revelation principle) demonstrates that if there exists a fully revealing equilibrium, there exists a fully revealing equilibrium in which the message space is equal to the set of types and Senders honestly report their types. Hence a fully revealing equilibrium exists if and only if for all feasible reports (θ_1, θ_2) , $\theta_1 \neq \theta_2$, there is an action a such that Sender i prefers $a^R(\theta_j)$ to a for $i, j = 1$ and 2 and $j \neq i$.

Studies of models with multiple Senders have concentrated on extreme results. In contrast to the single-Sender models, fully revealing equilibria

²¹More precisely, in order for a fully revealing equilibrium to exist, the set of actions that can be a best response to some beliefs must be large. Ambrus and Takahashi assume that the set of types Θ is equal to the set of feasible actions A and that $a^R(\theta) = \theta$. They do not assume that A is convex. Hence in their paper the set of possible Receiver best responses in the convex hull of A .

exist under relatively weak conditions. Yet when domains are restricted and biases are large, non-trivial communication is not possible in equilibrium. These analyses stop short of full characterizations of equilibria and leave a gap between which outcomes are strategically possible and which are likely to arise. They also leave open the possibility that generalizations of refinement arguments that are effective in one-dimensional settings may not select fully informative equilibria in multiple Sender models. A simple example may illustrate the problem.

Imagine a setting in which the uninformed Receiver is a husband who must decide where to eat with his wife. He can go to Restaurant X or Restaurant Y or stay home. The husband's preferences depend on the state of the world – where an enemy is dining out. The husband prefers to eat at a restaurant, but only if he is confident that the enemy will not be there. He would prefer to eat at home to being in the same restaurant as the enemy. His prior is such that without further information, he prefers to stay home. Suppose that Sender (wife) knows where the enemy is dining. She prefers staying home to any other outcome no matter what her type is, but, if she goes to a restaurant, prefers to avoid the enemy. A fully revealing equilibrium exists, but it is not an obvious prediction because communication makes the wife worse off in both states. One can add another Sender to the example and show that there exists a babbling equilibrium preferred by the Senders and a separating equilibrium preferred by the Receiver. I am not aware of a refinement that selects the fully revealing equilibrium in this type of example and the pooling outcome is intuitively appealing.

There are two challenges for theoretical work. The first challenge is to provide conditions under which commonly observed forms of communication are important. For example, consulting more than two informed experts is obviously critical in some situations, but there is no incentive to do so if fully revealing equilibria exist with only two Senders. Using multiple informed experts is beneficial if different experts have access to different information, but it also may make revealing equilibria more robust. Sequential consultations and conferences between different Senders are also important in practice, but theory provides a limited understanding of the value of these ways to organize communication. The second challenge is to understand how natural language interacts with strategic considerations to determine outcomes when multiple equilibria exist. A challenge for future experimental work is to develop a body of evidence about what outcomes to expect in multi-Sender models.

6 Enhancing Communication

The simple non-cooperative model of communication concentrates on a plausible, but specific, extensive form. It is typically possible to improve decisions when more general communication protocols are available. The literature approaches this issue in two ways. One approach looks at the abstract design of institutions. These studies describe the set of outcomes that are feasible given incentive constraints. In particular, they identify the optimal outcomes from the point of view of the Receiver. The other approach is to look at different extensive forms. Even when there is only one informed player, it is possible to improve upon non-cooperative outcomes by permitting multiple rounds of communication.

A large part of the literature examines the special “uniform-quadratic” model in which the Sender and Receiver have quadratic preferences and the state is uniformly distributed on the unit interval. The quadratic case is important because it is tractable. Explicit solutions are possible for several models with quadratic preferences. Assuming quadratic preferences also simplifies the study of outcomes in which the distribution of actions associated with a particular type may be non-degenerate. In the simple one-dimensional model of cheap talk, an outcome associates with each type a deterministic action. The action space is linearly ordered and assumptions on preferences guarantee a single-crossing property: if a given type prefers one action to a lower one, then so will all higher types. This property makes it easier to characterize equilibrium behavior. In this section I discuss environments in which it is valuable for the Receiver to randomize. In this case, action distributions are not linearly ordered and there is generally no analog to the single-crossing property. When preferences are quadratic, however, the expected action is a sufficient statistic for the distribution and action distributions can be ranked according to their means. If a given type prefers one action to another action with a lower mean, then so will all higher types.

Goltsman, Hörner, Pavlov, and Squintani [39] compare different institutions in the uniform-quadratic model. They solve the problem under three different constraints on feasible institutions: arbitration, mediation, and negotiation. Arbitration is the most general institution. Here a disinterested third-party can ask the Sender for information and commit to a decision rule. One simple decision rule is for the Receiver to give the Sender the right to make a decision, but possibly to constrain the set of possible choices available to the Sender. That is, the Receiver could select a subset D of A and

permit the Sender to select her favorite element in D . It is straightforward to show that the solution to this delegation problem is to select $D = [0, d]$ for $d < 1$. That is, the arbitrator says to the Sender: “take any action you want, as long as it isn’t too big.” Because the Sender has an upward bias, the arbitrator rules out extreme high actions that would never be optimal for the Receiver. Goltsman, et. al. demonstrate that delegation in this way is actually the optimal arbitration scheme.²²

Under mediation, a third-party can collect information, but the Receiver retains decision-making authority and is constrained to make a best response given his information. There is no hope for effective communication when the preferences of the players are sufficiently different. However when preferences are similar, mediation generally can improve upon one round of cheap talk.²³ Mediation permits more informative communication by allowing the mapping from types to outcomes to be stochastic. To get an intuition for how this works, note that the equilibrium outcomes in the Crawford-Sobel model are discontinuous in the (difference between the) preferences of the players. The maximum number of actions induced is an integer-valued function of the preferences that increases as the bias decreases to zero. Loosely, mediation makes it possible to construct mediated outcomes with lower-variance than the k -step partition equilibrium for biases in which the $(k + 1)$ -step equilibrium does not exist.²⁴

In the quadratic case, the Receiver and Sender will have the same (ex ante) preferences over outcomes. To see this, denote the expected utility of Agent i by EU_i for $i = S$ and R . Note that

$$EU_S = -E(y - \theta - b)^2 = -E(y - \theta)^2 - 2bE(y - \theta) - Eb^2 = EU_R - b^2, \quad (3)$$

where the third equation follows because the optimal action of the Receiver is an unbiased estimate of the state in equilibrium. Hence when equilibrium considerations are taken into account, it is in the interest of both players to remove noise from the message. It may seem surprising, therefore, that stochastic mechanisms are valuable. But while (3) uses the property that

²²Alonso and Matouschek [2], Dessein [25], Holmström [46], Melumand and Shibano [65] describe the optimal deterministic mechanism in different levels of generality.

²³Myerson conjectured this when he discussed Crawford and Sobel [21] at the 1981 Summer Meetings of the Econometric Society.

²⁴An extensive literature exists that seeks condition under which mediated outcomes can be attained through direct communication. Ben-Porath [9], Forges [32], Gerardi [33], and Vida [78] are examples.

the Receiver’s expected action is equal to the expected state, it does not incorporate any other equilibrium constraints. I discuss below how stochastic actions weaken incentive constraints.

Blume, Board, and Kawamura [11] describe one way to implement the mediation mechanism that maximizes the Receiver’s expected payoff. Their construction illustrates why stochastic actions may be beneficial. Blume, Board, and Kawamura study direct communication in which, with positive probability, the Sender’s intended message does not reach the Receiver. Instead the Receiver hears a message drawn from an exogenously given distribution. Imagine a partition equilibrium with two steps. When the Sender’s type is below a given θ_0 , she will send the “low” message. When the Sender’s type is above θ_0 , she will send a “high” message. The marginal Sender type, θ_0 , is indifferent between the two messages. When there is no noise, the upward bias in the Sender’s preferences forces θ_0 to be less than one half. When there is a small probability that the Receiver does not receive a message, the Sender’s message will induce the action appropriate for her step most of the time and an “average” action otherwise. It follows that noise makes the expected responses to the two messages closer together, which in turn moves the marginal Sender type closer to one half, which makes the information partition more valuable to the Receiver.

There are other ways to implement the mediated outcome. Krishna and Morgan [59] describe a two-stage communication procedure that, like the noisy environment of Blume, Board, and Kawamura, can achieve the optimal mediated outcome without a mediator (for some parameter values). Ivanov [47] demonstrates that one can implement the mediation mechanism that maximizes the Receiver’s expected payoff using a strategic mediator whose preferences are determined by the Receiver.

Goltsman, et. al. [39] also study negotiated outcomes, which they define to be equilibria of games with a finite number of rounds of unmediated cheap talk followed by an action selected by the Receiver. The Krishna-Morgan procedure demonstrates that negotiation works as well as mediation for intermediate biases. When biases are either large or small, Goltsman, et. al. [39] show that mediation can give the Receiver a strictly higher expected payoff than negotiation.

It would be quite interesting to establish benchmark results like those of Goltsman, et. al. [39] in more general environments, either by relaxing the strong restrictions on preferences in one-dimensional, one-Sender problems, or by studying problems with higher dimensional action spaces or multiple

Senders. Quadratic preferences play a central role in most of the existing arguments. The ability to construct fully revealing equilibria in higher dimensional models trivializes the design problem for some applications, but characterizing the optimal mechanism remains an open problem either when fully revealing equilibria fail to exist or when the Senders have different information.²⁵

7 Persuasion

The previous section describes how to design communication from the standpoint of the Receiver (although in the quadratic example, preferences of the two agents are aligned *ex ante*). There are natural settings in which the Sender may have control over the communication environment. What is the best way for the Sender to package her advice?

When the Sender has commitment power, the strategic complexity connected with coordinating on a common language disappears. By committing to a strategy, the Sender establishes the meaning of messages (each message corresponds to a posterior distribution over states). Hence while the choice of actual messages may be arbitrary, there will be no coordination failures due to the Receiver failing to interpret messages correctly in equilibrium. In addition, commitment power frees the Sender from the requirement that she best respond to the Receiver's strategy. Consequently more outcomes are possible.

Kamenica and Gentzkow [52] analyze a general version of the problem.²⁶ The players have preferences given by $U_S(\cdot)$ and $U_R(\cdot)$. S knows θ ; R knows the prior distribution p . To avoid technicalities, assume that there are only finitely many states. The Sender sends an arbitrary message to the Receiver. If $q(m | \theta)$ is the probability that the Sender uses the message m when the state is θ , then

$$\mu(\theta | m) = \frac{p(\theta)q(m | \theta)}{\sum_{\theta'} p(\theta')q(m | \theta')} \quad (4)$$

²⁵Alonso, Dessein, and Matouschek [1] provide one such analysis in a model that investigates the trade offs between centralized and decentralized control when there are multiple senders who wish to coordinate their actions. Kawamura [54] and Morgan and Stocken [69] study how the quality of information transmitted through polls varies with the number of Senders.

²⁶Brocas and Carillo [13] and Rayo and Segal [73] study related models.

is the induced posterior. The strategy $q(\cdot)$ induces the posterior $\mu(\cdot | m)$ with probability $P(m) = \sum q(m | \theta)p(\theta)$. A rational Receiver will respond to each message with a best response to the induced beliefs. Let $a^R(\mu)$ denote the Receiver's best response to μ (if there are multiple best responses, assume that a^R is an action that is preferred by the Sender). The maximum expected payoff for the Sender when she induces the beliefs μ is

$$\hat{U}_S(\mu) = \sum_{\theta \in \Theta} U_S(a^R(\mu), \theta)\mu(\theta). \quad (5)$$

It is straightforward to show that the distribution over posteriors induced by the Sender's strategy must satisfy a "Bayesian plausibility" condition:

$$\sum_{m \in M} \mu(\theta | m)P(m) \equiv p(\theta). \quad (6)$$

Kamenica and Gentzkow demonstrate that (6) is the only restriction placed on the Sender's mechanism. That is, the optimal mechanism from the Sender's point of view solves²⁷

$$\max \int \hat{U}_S(\mu(m))P(m)dm \quad (7)$$

subject to (6).

In this setting, one should not expect more communication when the preferences of the Sender and Receiver are more closely aligned. Consider a setting in which the Sender is a district attorney, who learns the guilt or innocence of a suspect. Assume that the Sender always wishes to convict the suspect, while the Receiver (judge) will convict only if he believes that the defendant is guilty with probability at least r . The lower is r , the more closely aligned are the preferences of the Receiver and Sender. In this example, however, the Sender's optimal mechanism is to commit to saying nothing when the prior probability that the defendant is guilty exceeds r , and to send two signals, a "guilty" signal that indicates that the defendant is guilty with probability r and an "innocent" signal that indicates that the defendant is innocent with probability one otherwise.²⁸ The Receiver obtains the most

²⁷Aumann and Maschler [6] contains the basic mathematical result in their analysis of repeated games with incomplete information.

²⁸To do this, the Sender always sends the guilty signal when the defendant is guilty and randomizes in the appropriate way when the defendant is innocent.

precise information when the standard of proof is highest. Kamenica and Gentzkow [52] provide conditions under which the Sender's value is greater when the Receiver's preferences are closer to those of the Sender.

The example illustrates that the Receiver may wish to commit to a decision rule to improve the information content of the Sender's signal. Specifically, if the Receiver could commit to convicting only when he was certain that the defendant was guilty, then the Sender would be honest and the Receiver would make perfect decisions. This behavior is similar to the skeptical behavior observed in equilibria of models with verifiable information that I discuss in Section 8.

The requirement that the Sender's communication strategy is Bayesian plausible depends, of course, on the assumption that the Receiver has a correct prior and draws rational inferences using Bayes's Rule. When one relaxes the rationality assumption, a clever Sender may be able to exploit the Receiver's biases. Ettinger and Jehiel [30], Jehiel and Koessler [50] and Mullainathan, Schwartzstein, and Shleifer [70] present examples of this kind of approach. Systematic evidence of behavioral biases will motivate different ways in which opportunistic Senders can relax the Bayesian plausibility restriction and take advantage of biased Receivers. It is not necessary that a cognitive bias will make the Receiver worse off. It might be interesting to investigate circumstances in which behavioral biases are not costly. When biases are not costly, they would presumably survive evolutionary arguments designed to eliminate non-optimizing decision rules.

8 Verifiable Information

Several studies focus on the important special case of the persuasion problem in which the Sender prefers the same action independent of the state of the world. When the Sender has no commitment power and talk is cheap, there is no meaningful communication in this setting. Kamenica and Gentzkow [52] and Rayo and Segal [73] demonstrate that communication is possible when the Sender has commitment power. These problems are also interesting when there are exogenous restrictions on the messages available to the Sender.

Especially sharp results are available when information is verifiable. Grossman [43], Milgrom [66], and Milgrom and Roberts [68] study simple disclosure games in which the message space consists of all subsets of Θ and the Sender's messages are constrained so that the Sender of type θ can only send messages

m that contain θ . In a simple disclosure game, the Sender learns her type, sends a feasible message to the Receiver, who takes an action. Assume states and actions belong to finite sets and are ordered so that the Sender wants the Receiver to take higher actions. Further assume that first-order stochastically dominating shifts in the distribution of states induces the Receiver to take higher actions. In applications, higher states correspond to higher quality, which makes the Receiver willing to pay more. In this environment, equilibrium behavior involves maximal skepticism and full revelation. That is, the Receiver believes that message m indicates that θ is the lowest element in m (skepticism) and the Sender's strategy will allow R to infer the true state. To see this, note that the highest type of Sender is able to reveal herself and that doing so induces the Receiver to take her most preferred action. Hence the highest type Sender will separate from other types in equilibrium (otherwise, she would receive a lower payoff). Iterating this argument confirms that all of the information will "unravel" from the top.

The simple unraveling argument depends on the Receiver knowing that the Sender has information, the feasibility of conveying the information, and the rationality of the agents.²⁹

Glazer and Rubinstein [37] and [38] examine related simple models of persuasion. The Sender can make a statement about the state of the world. The Receiver does not know the state, but can make inferences about the state from what the Sender says. The Receiver makes a binary decision. The Sender always wants the Receiver to make a particular decision ("accept"). The Receiver wants to accept only if the state lies in a given set. Glazer and Rubinstein focus on communication regimes that minimize the probability that the Receiver makes a mistake. These are, at least for some specifications of preferences, Receiver optimal mechanisms. Glazer and Rubinstein assume that the Sender has hard information and that the Receiver has limited ability to evaluate the Sender's message.³⁰ Glazer and Rubinstein [38] note that the Receiver's optimal strategy necessarily involves drawing an inference. When the Sender reveals one piece of information, the Receiver infers that she did not have information that would be more more likely to convince him to take the Sender's favorite action. In particular, a message has a different meaning to the Receiver depending on whether it is provided by the Sender

²⁹Milgrom [67] reviews some variations.

³⁰In Glazer and Rubinstein [37] the Receiver can only partially evaluate the Sender's message. In Glazer and Rubinstein [38] the Sender can only reveal some of her information.

or randomly selected from available information by a disinterested party. It is not surprising to a game theorist to learn that how the Receiver interprets a message depends on the preferences of the Sender, but, as I discuss in the next section, the observation provides a way to think about questions important to linguists.

Dziuda [29] studies a model of strategic argumentation in which an informed Sender has hard evidence both for and against her desired outcome. The Sender may disclose or suppress these arguments. The Receiver does not know what evidence is available to the Sender and must make an inference about the state of the world based on what the Sender reveals. Dziuda shows that the standard unraveling intuition fails when the Sender is honest with positive probability. If the strategic Sender suppressed all negative arguments, then the Receiver would infer that negative evidence comes from an honest Sender. These beliefs would typically make it attractive for a strategic Sender to reveal negative information. The Receiver still draws inferences from the Sender's statements, but unlike the earlier treatment, a rational Receiver interprets statements using more than simple skepticism.³¹

9 Pragmatics

Pragmatics is a subfield of linguistics that studies how and why inferences that individuals draw from natural language statements differ from the actual semantic content of the statements. In this section I discuss two topics in pragmatics that connect naturally to other parts of the essay. There is potential for useful interaction between linguistics and the game theoretic study of language.³²

Linguistically it is a puzzle that when someone says “this room is stuffy,” a listener will react by opening a window. Even though the speaker did not say “open the window,” she probably induced a satisfactory action. Context permits people to draw inferences from statements that go beyond the abstract meaning of the words. Grice [42] explains how the Sender can say one thing and manage to convey something different. For Grice, the basis for the

³¹Le Quement [60] studies a related environment in which a strategic Sender reveals negative information in equilibrium.

³²For more examples, see the collection of Benz, Jäger, and Van Rooij [10]. Lewis [62]'s important philosophical study “Convention” contains detailed discussion of language and meaning.

extrapolation is the assumption that conversation is a cooperative enterprise. The Receiver can assume that the Sender will make honest statements that contain exactly the information that is relevant to the Receiver’s decision problem. This provides a framework in which the Receiver can draw inferences from the Sender’s messages. There are fundamental differences between Grice’s setting and the basic economic model of cheap-talk communication. First, in cheap-talk models statements have no intrinsic meaning. The Receiver draws sophisticated inferences in equilibrium, but these inferences do not depend on any exogenous interpretation of language. Second, in most economic applications, the Sender and Receiver have different preferences. It is naive for the Receiver to believe that the Sender will be honest.

Nevertheless, economic models do have something to say about how to draw inferences from messages. When, as in Section 8, information is verifiable, messages do have a natural meaning. The message space exogenously associates messages to subsets of types. But when the Sender says: “My type is in $T \subset \Theta$ ” in the Grossman and Milgrom models, the Receiver will draw the inference that “the Sender’s type is the least profitable element of T .” Similarly, in Glazer and Rubinstein’s [38] model, the Receiver infers that when the Sender sends one message it is because a more beneficial one (from the standpoint of the Sender) is not available. Hence the knowledge of the motives of the Sender enable the Receiver to draw inferences that go beyond the literal meaning of statements.

This observation is mathematically obvious, but still may be useful in forming a bridge between the two disciplines. Game theory does provide a framework in which to study implicature (how to make inferences from messages) and one can study the implications of messages without assuming that conversations are cooperative. The limit of the approach is that the strategy of being skeptical is not necessarily equilibrium behavior (as Dziuda’s [29] model demonstrates). Equilibria in disclosure games will always involve drawing inferences, but a principle that organizes the inferences will only be valid for subsets of games.

In models of cheap talk one can interpret a message in two ways: it can be a representation of information (“my type is in T ”) or it can be an instruction (“take action a ”). Linguists distinguish these uses. In fact, Jakobson [49] gives six basic functions of verbal communication. The *referential* function corresponds roughly to “my type is in T ” while the *conative* function cor-

responds to the instruction “take action a .”³³ Making a distinction between referential and conative functions is intuitive and natural in applications. One can imagine that there is a functional difference between an expert who reports the facts (communicates referentially) and one who gives advice (communicates conatively). In the most basic game-theoretic models of communication, there is no distinction between the conative and representative functions because in equilibrium a message describes both the set of types who send the message and the action the message induces. Distinctions do arise in more general models. When the Sender cannot predict the action a given message will induce (either because the Receiver randomizes or he has private information that determines how he responds to a message), it is more natural to interpret a message as referential. When the Receiver cannot infer the set of types from a message (for example, if the Sender randomizes or messages are sometimes unavailable), then it is more natural to interpret a message as conative. Perhaps ideas in Section 11 could form the basis for a game-theoretic distinction between different ways to interpret messages.

10 Information Acquisition

Until now I have assumed that what the Sender knows is exogenously determined. Rather than treat the Sender as an informed expert, it is often appropriate to view the Sender as an agent who has access to information, but the information is costly to acquire.

Che and Kartik [16] study a model in which the Sender must allocate effort to obtain information. The Sender’s effort increases the probability that the Sender will receive evidence, which takes the form of a noisy signal about the true state of the world. The Sender can then repress or transmit the information to the Receiver. Che and Kartik assume that the Sender and Receiver have identical preferences, but they have different prior beliefs. Different priors lead the players to interpret signals differently and hence the Sender and Receiver disagree on the utility-maximizing action in response to a given signal. The difference of opinion has two conflicting implications. On one hand, conditional on having information, a Sender is more likely to inform the Receiver the closer her prior is to the Receiver’s. Differences of

³³For the record, the other four functions are emotive, phatic, poetic, and metalingual. I am not prepared to make the case that these distinctions are important in economic models.

opinion limit information transmission. On the other hand, larger differences of opinion generate higher payoffs for information acquisition (because the Sender anticipates greater opportunities to change the Receiver’s action). Che and Kartik demonstrate that the second effect is large enough so that a Receiver with a (sufficiently rich) choice of Senders having different priors, would seek one with beliefs different from his. Unlike the simplest cheap-talk models, communication may be more effective when conflict increases.

Dewatripont and Tirole [27] present a model in which the Receiver must decide between a status quo and two novel policies. Information may exist in favor of either or both of the novel situations. Dewatripont and Tirole assume that the Receiver prefers a novel option if there is evidence that favors only it; if there is evidence favoring both novel options or no evidence, then he prefers the status quo. They assume that the Receiver must hire someone to look for evidence and that as in Che and Kartik, the expert (Sender) obtains hard evidence with a probability that increases with her effort.³⁴ Dewatripont and Tirole assume that the Receiver can pay the expert, but payment can depend only on the Receiver’s action (and not the information provided). They demonstrate that it is optimal for the Receiver to hire specialized experts, one to investigate each issue, rather than a single expert. This is because in order to get a single expert to investigate both options, she must be compensated for the possibility that evidence discovered for one option might destroy the benefit she receives for finding evidence in favor of the other option. Dewatripont and Tirole assume that experts do not care about the Receiver’s final decision. Kim [55] extends their model to include partisan advocates and makes an observation similar to Che and Kartik’s. It is advantageous to hire advocates who have an intrinsic preference for the option that they are investigating. In this environment, R can reduce the compensation to the expert. A partisan expert can be paid less to search for information because she gains utility if she can provide evidence that will convince the Receiver to take the expert’s preferred action.

In these models it is important that acquiring information permits the Sender to convince the Receiver to make the decision she prefers. Dur and Swank [28] analyze a model in which the uninformed decision maker is better off consulting a less biased expert than a partisan. In Dur and Swank’s model, the Receiver must make a binary decision, whether to stick with a status quo

³⁴The effort choice is binary in Dewatripont and Tirole and the expert can only receive information if she makes a costly investment.

($a = 0$) or change ($a = 1$). The best choice depends on the state. The Receiver has a prior on the state. The Sender can obtain better information at a cost. The Sender first decides how much to invest in improving her information, which determines the probability that her signal fully reveals the state (otherwise the signal contains no information). She then receives a signal (but cannot distinguish accurate information from noise). Given her signal, the Sender makes a cheap-talk recommendation to the Receiver, who makes his decision. The payoffs take the form $U_i(a, \theta) = a(b_i + \theta)$ for $i = S$ or R . b_i measures the bias of agent i . θ has mean zero, so an agent with $b_i = 0$ is unbiased. If the biases are sufficiently close, there exists an informative equilibrium in addition to the babbling outcome. When the informative equilibrium exists, R effectively permits S to select her preferred outcome. Receiver's equilibrium expected payoff depends on the quality of the Sender's information and the probability that the Sender and Receiver prefer the same action. Dur and Swank demonstrate that the quality of the Sender's information, which is determined by the Sender's choice of effort, decreases as the Sender's bias (measured by $|b_S|$) increases. This follows because the information is more valuable to the Sender when her bias is zero. From this observation, it follows that the Receiver prefers the Sender to be less extreme than he is. Starting from a situation in which the Sender and Receiver have the same bias, a reduction in the Receiver's bias leads to a first-order increase in the amount of information the Sender acquires, but a negligible effect on the quality of the Sender's message created by a small conflict of interest. Experts who are partisan (in the sense that their bias is greater in absolute value than the Receiver's) are not valuable to the Receiver in this setting both because the experts are less likely than less biased agents to recommend the Sender's preferred action and because they have less incentive to acquire information.

Szalay [77] points out that the Receiver may benefit from the ability to commit to a restricted set of actions if doing so raises the value of information for the Sender. Restrictions take the form of ruling out moderate actions and arise even when the Sender and Receiver have the same preferences over actions. This result contrasts with the findings of Section 6. When the Receiver has commitment power and the Sender's information is exogenous, Holmström [46] shows that the Receiver wishes to rule out extreme actions.

The lesson from these studies is that when the Receiver wants the Sender to collect information, he should try to increase the value of information to the Sender. Filtered through a strategic interaction, it may be valuable for

the Receiver to limit his own actions or to seek an expert who has a strong interest in influencing his choice. On the surface, these results contradict the insight that communication improves when the Receiver's preferences are closer to those of the Sender. Stating things in this way is a bit misleading, however, because agents do have different preferences when search intensity is taken into account.

I am unaware of studies that characterize optimal institutions for acquisition and transmission of endogenously generated information for the basic model of Section 6.³⁵ This is a natural question for future research.

11 Describing and Interpreting Information

Decision makers are frequently at the mercy of experts who may have information relevant to decision making, but may fail to convey the information to them. The models that I have reviewed focus on incentive issues: The patient/Receiver should be skeptical of the doctor/Sender's claims that he needs surgery because the doctor may like to operate (or like to be paid for operating). The incentive issue is important, but there is another factor that interferes with successful information transfer. Suppose that the patient is sure that his doctor knows the truth about his medical condition and appropriate treatments, but does not think that she knows how he would trade off the costs and benefits of a particular treatment plan. The patient lacks the time and expertise to interpret the information that the doctor has in order to make an informed decision. The doctor lacks the time and expertise to describe medical options in a way that is useful to the patient.

Until now I have assumed that there is no direct cost associated with describing or interpreting information. Yet it takes skill and effort to describe information in a useful way and skill and effort to interpret information. There may be no conflict of interest between a doctor and her patient, but the better informed doctor may be unable to explain the treatment options in enough detail to permit the patient to make the best decision. Teaching requires careful preparation, but even the best course is ineffective unless students make an effort to learn the material. Modeling these problems requires a somewhat different model of information.

³⁵Gershkov and Szentes [34] analyze a mechanism-design problem with information acquisition and voting.

To illustrate the effect of assuming that it is costly to transfer information, consider a simple, one-dimensional cheap-talk game in which the Sender and Receiver have the same preferences. Assume that the state of the world is an element of the unit interval. Suppose that the message space consists of binary strings. The cost of communication to the Sender is increasing in the number of binary digits that she sends. The Receiver pays an amount that increases in the number of binary digits that he is able to decipher. If the Sender and Receiver decide simultaneously how much to invest in communication, the game has multiple equilibria.³⁶ In equilibrium, the Sender honestly sends the first n digits in the binary expansion of θ to the Receiver, the Receiver pays enough to understand precisely these digits. He draws the correct inference from the message and acts accordingly. Equilibria of this game share properties of the basic cheap-talk model: the Receiver takes a finite number of actions; the Sender's messages partition the state space into a finite set of intervals; and increasing the number of actions induced increases the utility of both agents. Because of complementarities, if there is an equilibrium in which the Sender sends n digits, there exists an equilibrium in which she sends k digits for $k < n$. While in the standard model, there is a sense in which communication is more effective when the preferences of the agents are more similar, in this case communication improves when the cost of communication decreases.

While there are models of noisy communication (for example, Blume, Board, and Kawamura [11]), for the most part economic models attribute a communication failure to strategic differences. Dewatripont and Tirole [26] is an exception. There is an informed Sender and an uninformed Receiver. The Receiver must make a binary decision between a status quo option that provides a certain payoff and a novel action with an uncertain payoff. The Sender prefers the novel action independent of her private information. The Receiver's best choice depends on the state of the world. Let β be the ex ante probability that the novel option is superior to the status quo for the Receiver and let β^* be the probability that leaves the Receiver indifferent between the two decisions. The authors assume communication is the outcome of two investments, $x_i \in [0, 1]$, $i = R, S$, that the Receiver and Sender select simultaneously. The paper interprets the investments as effort put into

³⁶Hertel and Smith [45] and Jäger, Koch-Metzger, and Riedel [48] analyze broadly related models in which the message space is finite and communication is costly for the Sender.

describing the information (x_S) or interpreting it (x_R). The probability that communication is successful is $x_R x_S$. If the communication is successful, the Receiver learns the state. Otherwise, the Receiver learns nothing.

Dewatripont and Tirole study two different modes of communication. Under supervisory decision making, the Receiver needs to approve a move from the status quo, but the Sender can implement it. The Receiver will permit the Sender to implement her favorite decision if he expects to gain by doing so. Under executive decision making, the Receiver must implement the new decision himself. In order to do this, the information transfer must be successful. Assume the Sender does not know whether the novel decision is beneficial to the Receiver.³⁷ Due to complementarities in the communication technology, there are typically multiple equilibrium effort levels; higher equilibrium effort from the Receiver is associated with higher effort from the Sender. Dewatripont and Tirole show that the maximum equilibrium effort level under executive decision making increases as β increases. This result follows because both the direct and indirect effects of increasing β lead to higher effort. Increasing β has the direct effect of inducing the Receiver to make a greater effort, because there is a greater probability that he will be able to use the information to implement a beneficial decision. Since efforts are complementary, the Receiver's increased effort will increase the Sender's equilibrium effort choice as well, which in turn will indirectly increase the Receiver's effort. This result is analogous to the familiar result that information transmission improves as the preferences of the agents move more closely together. The result does not hold under supervisory decision making. In this case, once $\beta > \beta^*$, S has no incentive to transmit information because if she sends no information, R will give her permission to take her favorite decision. Hence investing in communication is weakly dominated for the Sender.

Dewatripont and Tirole assume that the communication technology takes on a particular simple form. As in my variation of the cheap-talk game, they assume that there are complementarities in the production process. These complementarities give rise to multiple equilibria. There are situations in which it is more natural to model effort as substitutes. There are situations in which it is sensible to assume that the Receiver chooses effort after observing the Sender's effort (or some observable signal of her effort). For the

³⁷Dewatripont and Tirole also study the case in which the Sender knows the Receiver's value.

purposes of this essay, the specific assumptions on the technology of information transfer or the economic application of Dewatripont and Tirole is less important than the recognition that it requires effort to send and receive information. Modeling the technology of information exchange uses standard economic tools and promises insights into communication that are complementary to those derived from models that assume differences in preferences in order to describe communication failure. While some of these issues arise in studies that describe properties of optimal language or communication in organizations,³⁸ it would be valuable to devote more attention to the study of communication as a joint production process involving the Sender and Receiver.

12 Conclusion

The ideal conclusion would be a short list of deep general results that would give a casual reader a sense of where the literature is today and where it is headed. While the study of strategic communication supplies some powerful insights, the domain of plausible models is so rich that the most reliable intuition will fail in some simple environment. I warn the reader to be skeptical. The loose statements below summarize much of the state of the literature, but they are not mathematical truths.

The idea that when preferences are more similar, communication is more effective, is an excellent organizing principle, but this property fails in general (or would require such a restrictive definition of preference similarity as to be tautological).

Effective cheap talk requires some common interests. It is sometimes possible to find common interests by expanding the set of choices. Making preferences more similar is often good for everyone.

Effective communication is difficult. Strategic models demonstrate this by exhibiting multiple equilibria. Abstract refinements can provide default predictions, which are valuable without specific knowledge of the shared understanding of the meaning of messages that is not part of the standard description of the game. These predictions cannot substitute for incorporating restrictions on the use of language derived from experience outside a strategic interaction.

³⁸Arrow [5] provides a conceptual framework and more recent work, for example, Cremer, Garicano, and Prat [22] provides a formal analysis.

There are some convincing arguments that select against babbling outcomes when non-trivial equilibria exist. Arguments that select fully revealing equilibria over less informative outcomes are less convincing.

Commitment ability enables one agent to impose a common understanding of messages, reducing the multiple-equilibrium problems and the possibility of coordination failure that exist in models without commitment.

Economic models of communication have little to say about real conversations – dynamic exchanges in which people take turns.

Linguists are aware of strategic aspects of communication. Game theoretic models are a natural way to model the process by which agents draw inferences about information from messages. There are other issues of interest to linguists that are good candidates for strategic analysis.

Most economic models of communication attribute communication failures to strategic reasons – differences in preferences between the players. This approach is powerful, but misses the fact that it is difficult for an informed agent to summarize what she knows in a way that an uninformed agent can interpret and use. This difficulty remains even when agents have similar interests and a common language.

References

- [1] Ricardo Alonso, Wouter Dessein, and Niko Matouschek. When does coordination require centralization? *American Economic Review*, 98(1):145–179, March 2008.
- [2] Ricardo Alonso and Niko Matouschek. Optimal delegation. *Review of Economic Studies*, 75(1):259–293, 2008.
- [3] Attila Ambrus and Shi En Lu. Robust almost fully revealing equilibria in multi-sender cheap talk. Technical report, Harvard University, July 2010.
- [4] Attila Ambrus and Satoru Takahashi. Multi-sender cheap talk with restricted state spaces. *Theoretical Economics*, 3(1):1–27, 2008.
- [5] Kenneth J. Arrow. *The Limits of Organization*. Norton, New York, NY, 1974.

- [6] Robert J. Aumann and Michael B. Maschler. *Repeated Games with Incomplete Information*. MIT Press, 1995.
- [7] Jeffrey S. Banks and Joel Sobel. Equilibrium selection in signaling games. *Econometrica*, 55(3):647–661, May 1987.
- [8] Marco Battaglini. Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401, July 2002.
- [9] Elchanan Ben-Porath. Cheap talk in games with incomplete information. *Journal of Economic Theory*, 108:45–71, 2003.
- [10] Anton Benz, Gerhard Jäger, and Robert Van Rooij. *Game Theory and Pragmatics*. Palgrave Macmillan, 2005.
- [11] Andreas Blume, Oliver J. Board, and Kohei Kawamura. Noisy talk. *Theoretical Economics*, 2(4):395–440, 2007.
- [12] Andreas Blume, Yong-Gwan Kim, and Joel Sobel. Evolutionary stability in games of communication. *Games and Economic Behavior*, 5:547–575, 1993.
- [13] Isabelle Brocas and Juan D. Carrillo. Influence through ignorance. *RAND Journal of Economics*, 38:931–947, 2007.
- [14] Archishman Chakraborty and Rick Harbaugh. Comparative cheap talk. *Journal of Economic Theory*, 132(1):70–94, 2007.
- [15] Archishman Chakraborty and Rick Harbaugh. Persuasion by cheap talk. *American Economic Review*, 100(5):2361–82, 2010.
- [16] Yeon-Koo Che and Navin Kartik. Opinions as incentives. *Journal of Political Economy*, 117(5):815–860, October 2009.
- [17] Ying Chen. Perturbed communication games with honest senders and naive receivers. Technical report, Arizona State University, 2008.
- [18] Ying Chen, Navin Kartik, and Joel Sobel. On the robustness of informative cheap talk. *Econometrica*, 76(1):117–136, January 2008.
- [19] In-Koo Cho and David M. Kreps. Signaling games and stable equilibria. *The Quarterly Journal of Economics*, 102(2):179–221, 1987.

- [20] Vincent P. Crawford. A survey of experiments on communication via cheap talk. *Journal of Economic Theory*, 78(2):286–298, February 1998.
- [21] Vincent P. Crawford and Joel Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1451, November 1982.
- [22] Jacques Cremer, Luis Garicano, and Andrea Prat. Language and the theory of the firm. *The Quarterly Journal of Economics*, 122(1):373–407, 2007.
- [23] Adrian de Groot Ruiz, Theo Offerman, and Sander Onderstal. Equilibrium selection in cheap talk games: ACDC rocks when other criteria remain silent. Technical report, University of Amsterdam, December 2010.
- [24] Stefano Demichelis and Joërgen W. Weibull. Language, Meaning, and Games: A Model of Communication, Coordination, and Evolution. *The American Economic Review*, 98(4):1292–1311, 2008.
- [25] Wouter Dessein. Authority and Communication in Organizations. *Review of Economic Studies*, 69(4):811–838, 2002.
- [26] Mathias Dewatripont and Jean Tirole. Modes of Communication. *Journal of Political Economy*, 113(6):1217–1238, 2005.
- [27] Mattias Dewatripont and Jean Tirole. Advocates. *Journal of Political Economy*, 107(1):1–39, 1999.
- [28] Robert Dur and Otto H. Swank. Producing and manipulating information. *The Economic Journal*, 115(500):185–199, 2005.
- [29] Wioletta Dziuda. Strategic argumentation. Technical report, Northwestern University, May 2008.
- [30] David Ettinger and Philippe Jehiel. Theory of deception. *American Economic Journal: Microeconomics*, 2(1):1–20, February 2010.
- [31] Joseph Farrell. Meaning and credibility in cheap-talk games. *Games and Economic Behavior*, 5(4):514–531, October 1993.
- [32] François Forges. An approach to communication equilibria. *Econometrica*, 54(6):1375–1385, November 1986.

- [33] Dino Gerardi. Unmediated communication in games with complete and incomplete information. *Journal of Economic Theory*, 114:104–131, 2004.
- [34] Alex Gershkov and Balázs Szentes. Optimal voting schemes with costly information acquisition. *Journal of Economic Theory*, 144(1):36–68, 2009.
- [35] Itzhak Gilboa and Akihiko Matsui. Social stability and equilibrium. *Econometrica*, 59(3):859–867, 1991.
- [36] Thomas W. Gilligan and Keith Krehbiel. Asymmetric information and legislative rules with a heterogeneous committee. *American Journal of Political Science*, 33(2):459–490, May 1989.
- [37] Jacob Glazer and Ariel Rubinstein. On the optimal rules of persuasion. *Econometrica*, 72(6):1715–1736, 2004.
- [38] Jacob Glazer and Ariel Rubinstein. A study in the pragmatics of persuasion: A game theoretical approach. *Theoretical Economics*, 1:395–410, June 2006.
- [39] Maria Goltsman, Johannes Hörner, Gregory Pavlov, and Francesco Squintani. Mediation, arbitration, and negotiation. *Journal of Economic Theory*, 144(4):1397–1420, July 2009.
- [40] Sidartha Gordon. Iteratively stable cheap talk equilibrium. Technical report, Université de Montréal, December 2008.
- [41] Jerry R. Green and Nancy L. Stokey. A two-person game of information transmission. *Journal of Economic Theory*, 135(1):90–104, 2007.
- [42] H. Paul Grice. Logic and conversation. In *Studies in the Way of Words*, chapter 2, pages 22–40. Harvard University Press, Cambridge, 1989.
- [43] Sanford Grossman. The role of warranties and private disclosure about product quality. *Journal of Law and Economics*, 24:461–483, 1981.
- [44] Sanford Grossman and Motty Perry. Perfect sequential equilibria. *Journal of Economic Theory*, 39:97–119, 1987.

- [45] Johanna Hertel and John Smith. Not so cheap talk: Costly and discrete communication. Technical report, Rutgers University-Camden, June 2010.
- [46] Bengt Holmström. On the theory of delegation. Technical report, Northwestern University, 1980.
- [47] Maxim Ivanov. Communication via a strategic mediator. *Journal of Economic Theory*, 145(2):869–884, 2010.
- [48] Gerhard Jäger, Lars Koch-Metzger, and Frank Riedel. Voronoi languages. Technical report, Bielefeld University, 2009.
- [49] Roman Jakobson. Closing Statements: Linguistics and poetics. In Thomas A. Sebeok, editor, *Style in Language*, pages 350–377. Cambridge: MIT Press, 1960.
- [50] Philippe Jehiel and Frédéric Koessler. Revisiting games of incomplete information with analogy-based expectations. *Games and Economic Behavior*, 62(2):533–557, March 2008.
- [51] Hanjoon Michael Jung. Complete sequential equilibrium and its alternative. Technical report, Lahore University, 2009.
- [52] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, forthcoming.
- [53] Navin Kartik. Strategic communication with lying costs. *Review of Economic Studies*, 76(4):1359–1395, 2009.
- [54] Kohei Kawamura. A model of public consultation: Why is binary communication so common? *Economic Journal*, forthcoming.
- [55] Chulyoung Kim. Partisan advocates. Technical report, UCSD, May 2010.
- [56] Yong-Gwan Kim and Joel Sobel. An evolutionary approach to pre-play communication. *Econometrica*, 63(5):1181–1193, 1995.
- [57] Elon Kohlberg and Jean-François Mertens. On the strategic stability of equilibria. *Econometrica*, 54(5):1003–1037, September 1986.

- [58] Vijay Krishna and John Morgan. Asymmetric information and legislative rules: Some amendments. *American Political Science Review*, 95(2):435–452, June 2001.
- [59] Vijay Krishna and John Morgan. The art of conversation: eliciting information from experts through multi-stage communication. *Journal of Economic Theory*, 117(2):147–179, 2004.
- [60] Mark Le Quement. Persuasion and rhetorical moderateness. Technical report, University of Bonn, September 2010.
- [61] Gilat Levy and Ronny Razin. On the limits of communication in multi-dimensional cheap talk: a comment. *Econometrica*, 75(3):885–893, May 2007.
- [62] David Lewis. *Convention: A Philosophical Study*. Harvard University Press, Cambridge, 1969.
- [63] Pei-Yu Lo. Common knowledge of language and iterative admissibility in a sender-receiver game. Technical report, Brown University, October 2006.
- [64] Steven A. Mathews, Masahiro Okuno-Fujiwara, and Andrew Postlewaite. Refining cheap-talk equilibria. *Journal of Economic Theory*, 55(2):247–273, 1991.
- [65] Nahum D. Melumad and Toshiyuki Shibano. Communication in settings with no transfers. *The RAND Journal of Economics*, 22(2):173–198, 1991.
- [66] Paul R. Milgrom. Good news and bad news: Representation theorems and applications. *Bell Journal of Economics*, 21:380–391, 1981.
- [67] Paul R. Milgrom. What the seller won’t tell you: persuasion and disclosure in markets. *Journal of Economic Perspectives*, 22(2):115–131, 2008.
- [68] Paul R. Milgrom and John Roberts. Relying on the information of interested parties. *The Rand Journal of Economics*, 17(1):18–32, 1986.
- [69] John Morgan and Phillip C. Stocken. Information aggregation in polls. *American Economic Review*, 98(3):864–896, June 2008.

- [70] Sendhil Mullainathan, Joshua Schwartzstein, and Andrei Shleifer. Coarse thinking and persuasion. *Quarterly Journal of Economics*, 123(2):577–619, 2008.
- [71] Roger B. Myerson. Mechanism design by an informed principal. *Econometrica*, pages 1767–1797, 1983.
- [72] Matthew Rabin. Communication between rational agents. *Journal of Economic Theory*, 51:144–170, 1990.
- [73] Luis Rayo and Ilya R. Segal. Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987, October 2010.
- [74] Arthur J. Robson. Efficiency in evolutionary games: Darwin, Nash and the secret handshake. *Journal of Theoretical Biology*, 144:379–396, 1990.
- [75] Daniel J. Seidmann. Cheap talk games may have unique, informative equilibrium outcomes. *Games and Economic Behavior*, 4:422–425, 1992.
- [76] Joel Sobel. Evolutionary stability and efficiency. *Economics Letters*, 42:301–312, 1993.
- [77] Dezső Szalay. The economics of clear advice and extreme options. *Review of Economic Studies*, 72(4):1173–1198, 2005.
- [78] Péter Vida. *Long Pre-Play Communication in Games*. PhD thesis, Autonomous University of Barcelona, 2006.
- [79] Karl Wärneryd. Cheap talk, coordination, and evolutionary stability. *Games and Economic Behavior*, 5(4):532–546, 1993.
- [80] Iñigo Zapater. Credible proposals in communication games. *Journal of Economic Theory*, 72(1):173–197, January 1997.