

Cheap Talk

Vijay Krishna and John Morgan

June 2005

1 Introduction

In the context of games of incomplete information, the term “cheap talk” refers to direct and costless communication among players. Cheap-talk models should be contrasted with more standard signalling models. In the latter, informed agents communicate private information indirectly via their choices—concerning, say, levels of education attained—and these choices are costly. Indeed, signalling is credible precisely because choices are differentially costly—for instance, high productivity workers may distinguish themselves from low productivity workers by acquiring levels of education that would be too costly for the latter.

The central question addressed in cheap-talk models is the following: How much information, if any, can be credibly transmitted when communication is direct and costless? Interest in this question stems from the fact that with cheap talk, there is always a “babbling” equilibrium in which the participants deem all communication to be meaningless—after all it has no direct payoff consequences—and as a result, no one has any incentive to communicate anything meaningful. It is then natural to ask if there are also equilibria in which communication is meaningful and informative.

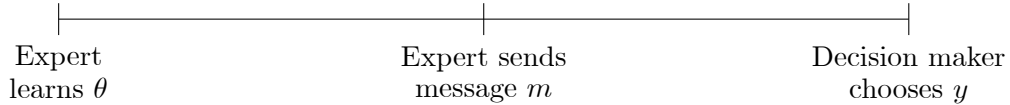
We begin by examining the question posed above in the simplest possible setting: there is a single informed party—an expert—who offers information to a single uninformed decision maker. This simple model forms the basis of much work on cheap talk and was introduced in a now classic paper by Crawford and Sobel (1982). In what follows, we first outline the main finding of this paper—while there are informative equilibria, these entail a significant loss of information. We then examine various remedies that have been proposed to solve (or at least alleviate) the “information problem.”

2 The Information Problem

We begin by considering the leading case in the model of Crawford and Sobel, henceforth CS. A decision maker must choose some decision y . Her payoff depends on y and on an unknown state of the world θ , which is distributed uniformly on the

unit interval. The decision maker can base her decision on the costless message m sent by an expert who knows the precise value of θ . The decision maker’s payoff is $U(y, \theta) = -(y - \theta)^2$, and the expert’s payoff is $V(y, \theta, b) = -(y - (\theta + b))^2$, where $b \geq 0$ is a “bias” parameter that measures how closely aligned the preferences of the two are. Because of the tractability of the “uniform-quadratic” specification, this paper, and indeed much of the cheap talk literature, restricts attention to this case.

The sequence of play is as follows:



What can be said about (Bayesian-perfect) equilibria of this game? As noted above, there is always an equilibrium in which no information is conveyed—even in the case where preferences are perfectly aligned (i.e. $b = 0$). In such a “babbling” equilibrium, the decision maker believes (correctly it turns out) that there is no information content in the expert’s message and hence chooses her decision only on the basis of her prior information. Given this, the expert has no incentive to be convey any information—he may as well send random, uninformative messages—and hence the expert indeed “babbles.” This reasoning is independent of any of the details of the model other than the fact that the expert’s message is “cheap talk.”

Are there equilibria in which all information is conveyed? When there is any misalignment of preferences, the answer turns out to be no. Specifically,

Proposition 1 *If the expert is even slightly biased, all equilibria entail some information loss.*

The proposition follows from the fact that if the expert’s message always revealed the true state and the decision maker believed him, then the expert would have the incentive to exaggerate the state—in some states θ , he would report $\theta + b$.

Are there equilibria in which some but not all information is shared? Suppose that following message m , the decision maker holds posterior beliefs given by distribution function G . The action y is chosen to maximize her payoffs given G . Because payoffs are quadratic, this amounts to choosing a y satisfying:

$$y(m) = E[\theta|m] \tag{1}$$

Suppose that the expert faces a choice between sending a message m that induces action y or an alternative message, m' , that induces an action $y' > y$. Suppose further that in state θ' the expert prefers y' to y and vice-versa in state $\theta < \theta'$. Since the preferences satisfy the *single-crossing* condition, $V_{y\theta} > 0$, the expert would prefer y' to y in all states higher than θ' . This implies that there is a unique state a , satisfying $\theta < a < \theta'$, in which the expert is indifferent between the two actions. Equivalently,

the distance between y and the expert’s “bliss” (ideal) action in state a is equal to the distance between action y' and the expert’s bliss action in state a . Hence,

$$a + b - y = y' - (a + b) \tag{2}$$

Thus, message m is sent for all states $\theta < a$ and message m' for all states $\theta > a$.

To comprise an equilibrium where exactly two actions are induced, one would need to find values for a , y , and y' that simultaneously satisfy equations (1) and (2). Since m is sent in all states $\theta < a$, from equation (1), $y = \frac{a}{2}$. Similarly, $y' = \frac{1+a}{2}$. Inserting these expressions into equation (2) yields

$$a = \frac{1}{2} - 2b \tag{3}$$

Equation (3) has several interesting properties. First, notice that a is uniquely determined for a given bias. Second, notice that when the bias gets large ($b > \frac{1}{4}$), there is no feasible value of a , so no information is conveyed in any equilibrium. Finally, notice that when the expert is unbiased ($b = 0$), there exists an equilibrium where the state space is equally divided into “high” ($\theta > \frac{1}{2}$) and “low” ($\theta < \frac{1}{2}$) regions and the optimal actions respond accordingly. As the bias increases, the low region shrinks in size while the high region grows; thus, less information is conveyed the higher is the bias.

For all $b < \frac{1}{4}$, we constructed an equilibrium that partitions the state space into two intervals. As the bias decreases, equilibria exist that partition the state space into more than two intervals. Indeed, Crawford and Sobel (1982) showed that:

Proposition 2 *All equilibria partition the state space into a finite number of intervals. The information conveyed in the most informative equilibrium is decreasing in the bias of the expert.*

If the expert were able to commit to fully reveal what he knows, *both* parties would be better off than in any equilibrium of the game described above. With full revelation, the decision maker would choose $y = \theta$ and earn a payoff of 0 while the expert would earn a payoff of $-b^2$. It is easily verified that in any equilibrium the payoffs of both parties are lower than this. The overall message of the CS model is that, absent any commitment possibilities, cheap talk inevitably leads to information loss, which is increasing in the bias of the expert. The remainder of the paper studies various “remedies” for the information loss problem: more extensive communication, delegation, contracts, and multiple experts.


3 Remedies

3.1 Extensive communication

In the CS model, the form of the communication between the two parties was one-sided—the expert simply offered a report to the decision maker who then acted on

this. Of course, communication can be much richer than this, and it is natural to ask if its form affects information transmission. One might think that it would not. First, one-sided communication where the expert speaks two or more times is no better than having him speak once since any information the expert might convey in many messages can be encoded in a single message. Now, suppose the communication is two-sided—it is a conversation—so the decision maker also speaks. Since she has no information of her own to contribute, all she can do is to send random messages and at first glance, this seems to add little. As we will show, however, random messages improve information transmission by acting as *coordinating devices*.

To see this, suppose the expert has bias $b = \frac{1}{12}$. As we previously showed, when only he speaks, the best equilibrium is where the expert reveals whether the state is above or below $\frac{1}{3}$. Suppose instead that we allow for *face-to-face* conversation—a simultaneous exchange of messages—and that the sequence of play is:



The following strategies constitute an equilibrium: The expert reveals some information at the face-to-face meeting, but there is also some randomness in what transpires. Depending on how the conversation goes, the meeting is deemed by both parties to be a “success” or a “failure.” After the meeting, and depending on its outcome, the expert may send an additional “written report” to the decision maker.

During the meeting, the expert reveals whether θ is above or below $\frac{1}{6}$; he also sends some additional messages that affect the success or failure of the meeting. If he reveals that $\theta \leq \frac{1}{6}$, the meeting is adjourned, no more communication takes place, and the decision maker chooses a low action $y_L = \frac{1}{12}$ that is optimal given the information that $\theta \leq \frac{1}{6}$.

If, however, he reveals that $\theta > \frac{1}{6}$, then the written report depends on whether the meeting was a success or a failure. If the meeting is a failure, no more communication takes place, and the decision maker chooses the “pooling” action $y_P = \frac{7}{12}$ that is optimal given that $\theta > \frac{1}{6}$. If the meeting is a success, however, the written report further divides the interval $[\frac{1}{6}, 1]$ into $[\frac{1}{6}, \frac{5}{12}]$ and $[\frac{5}{12}, 1]$. In the first subinterval, the medium action $y_M = \frac{7}{24}$ is taken and in the second subinterval the high action $y_H = \frac{17}{24}$ is taken. The actions taken in different states are depicted in Figure 1. The dotted line depicts the actions, $\theta + \frac{1}{12}$, that are “ideal” for the expert.

Notice that in state $\frac{1}{6}$, the expert prefers y_L to y_P (y_L is closer to the dotted line than is y_P) and prefers y_M to y_L . Thus if there were no uncertainty about the outcome of the meeting—for instance, if all meetings were “successes,” then the expert would not be willing to reveal whether the state is above or below $\frac{1}{6}$; for states $\theta = \frac{1}{6} - \varepsilon$, the expert would say $\theta \in [\frac{1}{6}, \frac{5}{12}]$, thereby inducing y_M instead of y_L . If all meetings were failures, then for states $\theta = \frac{1}{6} + \varepsilon$, the expert would say $\theta < \frac{1}{6}$, thereby




Figure 1: Equilibrium with Face-to-Face Meeting

inducing y_L instead of y_P .

There exists a probability $p = \frac{16}{21}$ such that when $\theta = \frac{1}{6}$ the expert is indifferent between y_L and a $(p, 1-p)$ lottery between y_M and y_P (whose certainty equivalent is labelled y_C in the figure). Also, when $\theta < \frac{1}{6}$, the expert prefers y_L to a $(p, 1-p)$ lottery between y_M and y_P , and when $\theta > \frac{1}{6}$, the expert prefers a $(p, 1-p)$ lottery between y_M and y_P to y_L .

It remains to specify a conversation such that the meeting is successful with probability $p = \frac{16}{21}$. Suppose the expert sends a message (*Low*, A_i) or (*High*, A_i) and the decision maker sends a message A_j , where $i, j \in \{1, 2, \dots, 21\}$. These messages are interpreted as follows: *Low* signals that $\theta \leq \frac{1}{6}$ and *High* signals that $\theta > \frac{1}{6}$. The A_i and A_j messages play the role of a coordinating device and determine whether the meeting is successful. The expert chooses A_i at random and each A_i is equally likely. Similarly, the decision maker chooses A_j at random. Given these choices, the meeting is a

$$\begin{aligned} \text{Success} & \quad \text{if} \quad 0 \leq i - j < 16 \text{ or } j - i > 5 \\ \text{Failure} & \quad \text{otherwise} \end{aligned}$$

For example, if the messages of the expert and the decision maker are (*High*, A_{17}) and A_5 , respectively, then it is inferred that $\theta > \frac{1}{6}$ and since $i - j = 12 < 16$, the meeting is a success. Observe that with these strategies, given any A_i or A_j , the probability that the meeting is a success is exactly $\frac{16}{21}$.

The equilibrium constructed above conveys more information than any equilibria

of the CS game. The remarkable fact about the equilibrium is that this improvement in information transmission is achieved by adding a stage in which the *uninformed* decision maker also participates. While the analysis above concerns itself with the case where $b = \frac{1}{12}$, informational improvement through via a “conversation” is a general phenomenon (Krishna and Morgan, 2004):

Proposition 3 *Multiple stages of communication together with active participation by the decision maker always improve information transmission.*

What happens if the two parties converse than once? Does every additional stage of communication lead to more information transmission? In a closely related setting, Aumann and Hart (2003) obtain a precise but abstract characterization of the set of equilibrium payoffs that emerge in sender-receiver games with a finite number of states and actions when the number of stages of communication is infinite. Because the CS model has a continuum of states and actions, their characterization does not directly apply. Nevertheless, it can be shown that even with an unlimited conversation, full revelation is impossible. A full characterization of the set of equilibrium payoffs with multiple stages remains an open question.

3.2 Delegation

A key tenet of organizational theory is the “delegation principle” which says that the power to make decisions should reside in the hands of those with the relevant information (Milgrom and Roberts, 1992). Thus, one approach to solving the information problem is simply to delegate the decision to the expert. However, the expert’s bias will distort the chosen action from the decision maker’s perspective. Delegation this leads to a trade-off between an optimal decision by an uninformed party and a biased decision by an informed party.

Is delegation worthwhile? Consider again an expert with bias $b = \frac{1}{12}$. The decision maker’s payoff from the most informative partition equilibrium is $-\frac{1}{36}$. Under delegation, the action chosen is $y = \theta + b$ and the payoff is $-b^2 = -\frac{1}{144}$. Thus delegation is preferred. Dessein (2001) shows that this is always true:

Proposition 4 *If the expert’s bias is not too large, delegation is better than all equilibria of the CS model.*

In fact, by exerting only slightly more control, the decision maker can do even better. As first pointed out by Holmström (1984), the optimal delegation scheme involves limiting the scope of actions from which the expert can choose. Under the uniform-quadratic specification, the decision maker should optimally limit the expert’s choice of actions to $y \in [0, 1 - b]$. When $b = \frac{1}{12}$, limiting actions in this way raises the decision maker’s payoff from $-\frac{1}{144}$ to $-\frac{1}{162}$.

Optimal delegation still leads to information loss. When the expert’s choice is “capped,” in high states the action is unresponsive to the state.

An application of the delegation principle arises in the US House of Representatives. Typically a specialized committee—analogueous to an informed expert—sends a bill to the floor of the House—the decision maker. How it may then be amended depends on the legislative rule under effect. Under the so-called *closed rule* the floor is limited in its ability to amend the bill while under the *open rule* the floor may freely amend the bill. Thus operating under a closed rule is similar to delegation while an open rule is similar to the CS model. The proposition above suggests, and Gilligan and Krehbiel (1987, 1989) have shown, that, in some circumstances, the floor may benefit by adopting a closed rule.

3.3 Contracts

Up until now we have assumed that the decision maker did not compensate the expert for his advice. Can compensation, via an incentive contract, solve the information problem? To examine this, we amend the model to allow for compensation and use mechanism design to find the optimal contract. Suppose that the payoffs are now given by

$$\begin{aligned} U(y, \theta, t) &= -(y - \theta)^2 - t \\ V(y, \theta, b, t) &= -(y - \theta - b)^2 + t \end{aligned}$$

where $t \geq 0$ is the amount of compensation.

Using the revelation principle, we can restrict attention to a direct mechanism where both t and y depend on the state θ reported by the expert. Notice that such mechanisms directly link the expert’s reports to payoffs—talk is no longer cheap.

Contracts are powerful instruments. A contract that leads to full information revelation and first-best actions is:

$$\begin{aligned} t(\hat{\theta}) &= 2b(1 - \hat{\theta}) \\ y(\hat{\theta}) &= \hat{\theta} \end{aligned}$$

where $\hat{\theta}$ is the state reported by the expert. Under this contract, the expert can do no better than to tell the truth, that is, to set $\hat{\theta} = \theta$, and, as a consequence, the action undertaken in this scheme is the “bliss” action for the decision maker. Full revelation is expensive however. When $b = \frac{1}{12}$, the decision maker’s payoff from this scheme is $-\frac{1}{12}$. Notice that this is worse than the payoff of $-\frac{1}{36}$ in the best CS equilibrium, which can be obtained with no contract at all. The costs of implementing the fully revealing contract outweigh the benefits.

In general, Krishna and Morgan (2005) show:

Proposition 5 *With contracts, full revelation is always feasible but never optimal.*




Figure 2: Optimal Contract with Perfect Commitment, $b \leq \frac{1}{3}$

The proposition above shows that full revelation is never optimal. No contract at all is also not optimal—delegation is preferable. What is the structure of the optimal contract? A typical optimal contract is depicted as the dark line in Figure 2. First, notice that even though the decision maker could induce his bliss action for some states, it is never optimal to do so. Instead, for low states ($\theta < b$), the decision maker implements a “compromise” action—an action that lies between θ and $\theta + b$. When $\theta > b$, the optimal contract simply consists of capped delegation.

3.4 Multiple senders

Thus far we have focused attention on how a decision maker should consult a single expert. In many instances, decision makers consult multiple experts—often with similar information but differing ideologies (biases). Political leaders often form cabinets of advisors with overlapping expertise. How should a cabinet be constituted? Is a balanced cabinet—one with advisors with opposing ideologies—helpful? How should the decision maker structure the “debate” among her advisors?

To study these issues, we add a second expert having identical information to the CS model. To incorporate ideological differences, suppose the experts have differing biases. When both b_1 and b_2 are positive, the experts have *like bias*—both prefer higher actions compared to the decision maker. In contrast, if $b_1 > 0$ and $b_2 < 0$, then the experts have *opposing bias*—expert 1 prefers a higher action and expert 2 a lower action compared to the decision maker.


3.4.1 Simultaneous talk

When both experts report to the decision maker simultaneously, the information problem is apparently solved—full revelation is now an equilibrium. To see this, suppose the experts have like bias and consider the following strategy for the decision maker: choose the action that is the more “conservative” of the two recommendations. Precisely, if $m_1 < m_2$, choose action m_1 and vice-versa if $m_2 < m_1$. Under this strategy, each expert can do no better than to report θ honestly if the other does likewise. If expert 2 reports $m_2 = \theta$, then a report $m_1 > \theta$, has no effect on the action. However, reporting $m_1 < \theta$ changes the action to $y = m_1$, but this is worse for expert 1. Thus, expert 1 is content to simply tell the truth. Opposing bias requires a more complicated construction, but the effect is the same—full revelation is an equilibrium (see Krishna and Morgan, 2001b).

Notice that the above construction is fragile because truth-telling is a weakly dominated strategy. Each expert is at least as well off by reporting $m_i = \theta + b_i$ and strictly better off in some cases. Battaglini (2001) defines an equilibrium refinement for such games which, like the notion of perfect equilibrium in finite games, incorporates the usual idea that players may make mistakes. He then shows that such a refinement rules out all equilibria with full revelation regardless of the direction of the biases. While the set of equilibria satisfying the refinement is unknown, the fact that full revelation is ruled out means that simply adding a second expert does not solve the information problem satisfactorily.

3.4.2 Sequential talk

Finally, we turn to the case where the experts offer advice in sequence:



Suppose that the two experts have biases $b_1 = \frac{1}{18}$ and $b_2 = \frac{1}{12}$, respectively. It is easy to verify (using (2)) that if only expert 1 were consulted, then the most informative equilibrium entails his revealing that the state is either below $\frac{1}{9}$, between $\frac{1}{9}$ and $\frac{4}{9}$, or above $\frac{4}{9}$. If only expert 2 were consulted, then the most informative equilibrium is where he reveals whether the state is below or above $\frac{1}{3}$. If the decision maker were able to consult only one of the two experts, she would be better off consulting the more loyal expert 1.

But what happens if she consults both? It turns out that if both experts actively contribute information, then decision maker can do no better than the following equilibrium: Expert 1 speaks first and reveals whether or not the state is above or below $\frac{11}{27}$. If expert 1 reveals that the state is above $\frac{11}{27}$, expert 2 reveals nothing further. If, however, expert 1 reveals that the state is below $\frac{11}{27}$, then expert 2 reveals further

whether or not it is above or below $\frac{1}{27}$. That this is an equilibrium may be verified again by using (2) and recognizing that, in state $\frac{1}{27}$, expert 2 must be indifferent between the optimal action in the interval $[0, \frac{1}{27}]$ and the optimal action in $[\frac{1}{27}, \frac{11}{27}]$. In state $\frac{11}{27}$, expert 1 must be indifferent between the optimal action in $[\frac{1}{27}, \frac{11}{27}]$ and the optimal action in $[\frac{11}{27}, 1]$.

Sadly, by actively consulting both experts, the decision maker is worse off than if she simply ignored expert 2 and consulted only her more loyal advisor, expert 1. This result is quite general, as shown by Krishna and Morgan (2001a):

Proposition 6 *When experts have like biases, actively consulting the less loyal expert never helps the decision maker.*

The situation is quite different when experts have opposing biases; that is, when the cabinet is balanced. To see this, suppose that the cabinet is comprised of two equally loyal experts biases $b_1 = \frac{1}{12}$ and $b_2 = -\frac{1}{12}$. Consulting expert 1 alone leads to a partition $[0, \frac{1}{3}]$, $[\frac{1}{3}, 1]$ while consulting expert 2 alone leads to the partition $[0, \frac{2}{3}]$, $[\frac{2}{3}, 1]$. If instead, the decision maker asked both experts for advice, the following is an equilibrium: expert 1 reveals whether θ is above or below $\frac{2}{9}$. If he reveals that the state is below $\frac{2}{9}$, the discussion ends. If, however, expert 1 indicates that the state is above $\frac{2}{9}$, expert 2 is actively consulted and reveals further whether the state is above or below $\frac{7}{9}$. Based on this, the decision maker takes the appropriate action. One may readily verify that this is an improvement over consulting either expert alone. Once again the example readily generalizes:

Proposition 7 *When experts have opposing biases, actively consulting the both experts always helps the decision maker.*

Indeed, the decision maker can be more clever than this. One can show that with experts of opposing bias, there exist equilibria where a portion of the state space is *fully revealed*. By allowing for a “rebuttal” stage in the debate, there exists an equilibrium where *all* information is fully revealed.

References

- [1] Austen-Smith, D. 1993. Interested experts and policy advice: Multiple referrals under open rule. *Games and Economic Behavior* 5, 3-43.
- [2] Aumann, R. and Hart, S. 2003. Long cheap talk. *Econometrica* 71, 1619-1660.
- [3] Baron, D. 2000. Legislative organization with informational committees. *American Journal of Political Science* 44, 485-505.
- [4] Battaglini, M. 2002. Multiple referrals and multidimensional cheap talk. *Econometrica* 70, 1379-1401.

- [5] Crawford, V. and Sobel, J. 1982. Strategic information transmission. *Econometrica* 50, 1431–1451.
- [6] Dessein, W. 2002. Authority and communication in organizations. *Review of Economic Studies* 69, 811-838.
- [7] Gilligan, T. and Krehbiel, K. 1987. Collective decision-making and standing committees: An informational rationale for restrictive amendment procedures. *Journal of Law, Economics, and Organization* 3, 287-335.
- [8] Gilligan, T. and Krehbiel, K. 1989. Asymmetric information and legislative rules with a heterogeneous committee. *American Journal of Political Science* 33, 459-490.
- [9] Grossman, G. and Helpman, E. 2001. *Special Interest Politics*. Cambridge: MIT Press.
- [10] Holmström, B. 1984. On the theory of delegation. In M. Boyer and R. Kihlstrom (eds.), *Bayesian Models in Economic Theory*, Amsterdam: North-Holland.
- [11] Krishna, V. and Morgan, J. 2001a. A model of expertise. *Quarterly Journal of Economics* 116, 747-775.
- [12] Krishna, V. and Morgan, J. 2001b. Asymmetric information and legislative rules: Some amendments. *American Political Science Review* 95, 435-452.
- [13] Krishna, V. and Morgan, J. 2004. The art of conversation: Eliciting information from experts through multi-stage communication. *Journal of Economic Theory* 117, 147-179.
- [14] Krishna, V. and Morgan, J. 2005. *Contracting for Information under Imperfect Commitment*. Mimeo, Penn State and Berkeley.
- [15] Milgrom, P. and Roberts, J. 1992. *Economics, Organization and Management*, Englewood Cliffs: Prentice Hall.
- [16] Wolinsky, A. 2002. Eliciting information from multiple experts. *Games and Economic Behavior* 41, 141-160.