Clearly Biased Experts*

Archishman Chakraborty† Rick Harbaugh‡

August 2008

Abstract

We consider the credibility, informativeness, and value of multidimensional cheap talk by an expert with transparent motives. Transparency ensures that the expert can credibly communicate information across dimensions and this information can be quite detailed. The expert always benefits from cheap talk if her preferences are quasiconvex, but is better off remaining silent if her preferences are quasiconcave. The model generates new results on the nature of persuasive advertising, the revenue gains from auction disclosure, the informational efficiency of voting rules, and the tradeoffs between cheap talk and delegation.

JEL Classification: D82, L15, C72, D72.

Key Words: multidimensional cheap talk, transparent motives, media bias, advertising, voting, auctions, delegation.

---

*For helpful comments we thank Curtis Eberwein, Sid Gordon, Michael Grubb, Navin Kartik, Hervé Moulin, Marco Ottaviani, Michael Rauh, Joel Sobel, and seminar participants at CUNY-Baruch, HEC, Ohio State University, the University of Cincinnati, the University of Oregon, the IIOC, the Midwest Theory Conference, and the MEDS Workshop on Communication, Game Theory, and Language. An early version of the paper was circulated under the title "Multidimensional cheap talk with transparent motives."

†Baruch College-CUNY & York University, achakraborty@schulich.yorku.ca

‡Kelley School of Business, Indiana University, riharbau@indiana.edu

1 Introduction

Experts are often biased. Political analysts are biased toward certain candidates, salespeople receive different commissions on different products, and media outlets benefit from emphasizing particular issues. If an expert's biases are known it is often argued that a decision maker can "see through" them and still obtain useful information. This idea that transparency facilitates communication has motivated a wide range of reforms, including requirements that investment advisors reveal any conflicts of interest, that lobbyists reveal their clients, and that political advertisements contain statements of approval by candidates.[1] As the Federal Communications Commission has argued, "... the public is entitled to know by whom it is persuaded" (Coase, 1979).

Despite the popularity of these efforts, it is not obvious that transparency should improve communication. If an expert is biased toward pushing the decision maker in a certain direction, knowledge of this bias allows the decision maker to adjust accordingly. But then the expert has an incentive to exaggerate even further, so there is no assurance that communication improves in equilibrium. Indeed, in the canonical Crawford and Sobel (1982) model of cheap talk communication in a single dimension, knowledge of the expert's bias often precludes communication that would otherwise be possible. Morgan and Stocken (2003) and Dimitrakas and Sarafidis (2005) show that revelation of the expert's bias can hurt communication when the size of the possible bias is uncertain, while Li and Madarasz (2008) show that revelation of the expert's bias always hurts communication when the direction of the bias is uncertain.[2]

To better understand the intuition that transparency can facilitate communication we use a stronger notion of transparency than used in these papers. We say that the expert's preferences over the decision maker's actions are transparent when they are common knowledge and therefore independent of the state of the world that the expert is privately informed about. Preferences are not transparent in this sense in the Crawford and Sobel model since

[1] These requirements follow respectively from the Securities Exchange Act of 1934, the Lobbying Disclosure Act of 1995, and the Bipartisan Campaign Reform Act of 2002.

[2] Uncertain expert bias is also analyzed in related models in Sobel (1985), Benabou and Laroque (1992), and Morris (2001), as well as in Gordon (2006) where the bias is allowed to be state-dependent. Lyon and Maxwell (2004) analyze transparency of the expert's bias in an environment with costly signals, Inderst and Ottaviani (2007) consider the transparency of a salesperson's endogenous commission structure, and the central banking literature examines the transparency of central bank policy, e.g., Moscaroni (2007).

even when the expert's bias is known the expert may still prefer a higher or lower action depending on the unknown state. Nevertheless, transparency is a natural assumption in many environments, e.g., when it is common knowledge that a seller wishes to maximize the expected price of a good regardless of its actual quality. Preferences of this form are standard in signaling games (e.g., Spence, 1973), screening games (e.g., Stiglitz, 1975), and disclosure/persuasion games (e.g., Milgrom, 1981, and Glazer and Rubinstein, 2004).

To see how transparent preferences can permit a nontrivial role for cheap talk, we depart from the standard one-dimensional model and consider environments where the expert's information is a multidimensional variable and the decision maker takes an action (or actions) based on his estimate of the true value of the variable. With multidimensional information, even if the expert has a strict preference ordering over the decision maker's estimate in each dimension, e.g., always prefers a higher estimate, the expert will not have a strict ordering over all estimate vectors if preferences are continuous. Therefore an expert might be able to credibly use cheap talk to trade off more favorable estimates on some dimensions for less favorable estimates on other dimensions.

Regardless of any biases or other asymmetries across dimensions, we show that influential and informative cheap talk equilibria exist as long as the expert's preferences are transparent and continuous. Transparency ensures that the decision maker knows the tradeoffs that the expert faces and can adjust for any biases that might make the expert misreport the information. Continuity ensures that the tradeoffs can exactly offset any incentive to lie. For instance, if a stock analyst is known to gain more from pushing one stock than another, in equilibrium an investor takes this into account and is less influenced on the favored stock. Or if a media outlet is known to favor one candidate over another, in equilibrium voters can still obtain useful information about the candidates.

Informative communication induces a mean preserving spread in the decision-maker's estimates, implying that an expert with transparent preferences is better off from communication via cheap talk if her preferences are quasiconvex. Quasiconvex preferences arise naturally in many environments where the literature previously has not recognized a role for cheap talk and we examine several in detail. First, in a voting model, we show that a defense attorney can credibly acknowledge her client's bad behavior in one dimension in the hope of persuading some jurors that her client does not warrant conviction based on his better behavior in other dimensions. Such cheap talk uniformly lowers the probability of conviction under the unanimity rule. Second, in an auction model, we show that by indicating to buyers the relative strengths of a product, the seller induces a better match of the product with the buyer who values it most. The seller captures part of this gain in allocational efficiency and raises her expected revenue if and only if there are sufficiently many buyers. Third, in a model of horizontal and vertical product differentiation, we show that an entrant profits from cheap talk advertising that emphasizes either the uniqueness of the product or its quality relative to the existing product.

In all of these situations the expert does not need to commit to a policy of revealing information. Instead the expert benefits from simple cheap talk that is credible because of the transparency of preferences and the multidimensional nature of communication. The information revealed is favorable on some dimensions and unfavorable on others, but overall the expert always benefits regardless of the realization of the information. Therefore, as long as the decision maker accounts for the expert's preferences correctly, the expert faces no temptation to deviate from the communication strategy ex post depending on the information that will be disclosed. This distinguishes our cheap talk approach from other models of communication that emphasize one-dimensional uncertainty and assume that the expert can commit to reveal even unfavorable information.[3]

How much information can be revealed by cheap talk? Even if the expert has a strong incentive to exaggerate within each of N dimensions, there is always an $N-1$ dimensional subspace over which the expert has no incentive to deceive the decision maker. Therefore, full revelation on these $N-1$ "dimensions of agreement" may be possible if the expert's preferences are known to the decision maker.[4] With linear preferences we show that this limit is always attainable even with arbitrary distributions and arbitrary biases across dimensions through a series of increasingly detailed statements. This result is of particular relevance since we show that standard Euclidean preferences converge to linear preferences as the bias within each dimension becomes large. For preferences that are strictly quasi-convex, similar detailed cheap talk using mixed strategies is also possible, with the expert's payoff strictly increasing the more detailed is her speech.

[3] For instance, commitment to a disclosure policy is the assumption in most of the literature on seller communication in auctions following Milgrom and Weber (1982), and in most of the literature on information sharing between firms as surveyed in Vives (2001).

[4] The idea of a dimension of agreement is developed by Battaglini (2002) in the context of state-dependent preferences.

We test the robustness of our existence results to small deviations from full transparency in three ways. First, if the probability that the expert's preferences are different from the expected type is sufficiently low, then under mild regularity conditions an equilibrium still exists in which the low probability types are correctly anticipated to always recommend the same action. Second, if there are only a finite number of different types of expert preferences, then we show that an informative equilibrium exists as long as there is a larger number of dimensions. Third, if state-dependent preferences converge uniformly to state-independent preferences, then there is an epsilon cheap talk equilibrium in which any incentive to deviate from the cheap talk equilibrium goes to zero. This is the case with Euclidean preferences converging uniformly to linear preferences as the biases in some or all dimensions increase.

This final convergence result provides a close link between the idea of an expert being biased towards a higher action as developed in the Crawford-Sobel model, and the idea of an expert being biased across dimensions as emphasized in this paper. In particular we show that as the expert's bias toward a higher action in some or all dimensions increases, Euclidean preferences converge uniformly to linear preferences with biases across dimensions (the expert's "slant") equal to the ratios of the biases within dimensions. Therefore, a simple model of linear preferences with different weights on different dimensions can capture how equilibrium communication is affected by different biases within dimensions when such biases are large and common knowledge.

Our analysis in this paper is related to that in Chakraborty and Harbaugh (2007) who consider a multidimensional version of the Crawford-Sobel model in which the expert's preferences are state-dependent in that she wants a higher action in a dimension when the state is higher in that dimension. When the environment is sufficiently symmetric, the expert's and decision maker's incentives are sufficiently aligned across dimensions that the expert can disclose complete and partial rankings of the variables, e.g., a stock analyst can provide a complete ranking of different stocks or a categorization of stocks into "buy", "hold" and "sell" groupings. The main difference in this paper is the transparency assumption that the expert's preferences are state-independent and common knowledge. With such preferences, we show that influential cheap talk equilibria exist even in arbitrarily asymmetric environments, e.g., ones where the complete ranking is common knowledge ex ante. For more than two dimensions, the papers also differ in the nature of communication they focus on. Specifically, in this paper communication is informative about the value of linear combinations of the variables, e.g., recommendations about stock portfolios instead of rankings of different stocks.

We show that transparency of the expert's preferences ensures that communication is possible, but the equilibria we consider do not fully reveal the expert's information so there remains room for stronger policy measures that attempt to modify or eliminate biases. Similarly there remains room for other factors that we have not modeled to affect communication, e.g., multiple periods (Sobel, 1985; Benabou and Laroque, 1992; Morris, 2001; Gentzkow and Shapiro, 2006; Ottaviani and Sørenson, 2006) and multiple competing experts (Gilligan and Krehbiel, 1989; Austen-Smith, 1993; Krishna and Morgan, 2001; Battaglini, 2002; Mullainathan and Shleifer, 2006; Ambrus and Takahashi, 2008; Gick, 2006; Visser and Swank, 2007).

In Section 2 we set up our model, demonstrate the existence of informative equilibria, identify conditions under which communication benefits the expert, and show when extended communication can reveal increasingly fine information. Section 3 considers the robustness of our results to small deviations from the assumption of complete transparency. We provide a number of examples that illustrate the scope our results in Section 4 and then Section 5 concludes. The Appendix contains all proofs.

2 Cheap Talk under Transparency

An expert is privately informed about the ideal actions $\theta \in \mathbf{\Theta}$ of a decision maker where $\mathbf{\Theta}$ is a compact convex subset of $\mathbb{R}^N$ with a non-empty interior and $N \geq 2$. She sends advice in the form of a costless[5] unverifiable message m from an arbitrary set $\mathbf{M}$ to an uninformed decision maker whose prior beliefs about θ are summarized by a joint distribution F with density f that has full support on $\mathbf{\Theta}$. The decision maker then chooses actions $a \in \mathbf{A} \equiv \mathbf{\Theta}$ equal to $E[\theta|m]$, the expected value of θ given his priors and the expert's message m.[6] The expert's preferences over the decision makers' actions a are described by a continuous utility

[5] In practice the expert might have to pay to send a message (e.g., an advertising fee), or might receive a payment for sending a message (e.g., a subscription fee). This has no effect on equilibrium behavior as long as the amount paid or received does not vary with the message.

[6] This standard behavioral assumption for the decision maker reflects underlying preferences for estimating the state θ as precisely as possible. We do not explicitly specify the preferences of the decision maker so that we can capture situations where multiple decision makers play a game as a function of their common estimate a of the content of the expert's message.

function $U(a)$ that does not depend on the state θ and is common knowledge.[7]

A communication strategy for the expert specifies a probability distribution over messages in $\mathbf{M}$ as a function of the state θ. A communication strategy is partitional if the inverse image of messages used with positive probability constitutes a partition of $\mathbf{\Theta}$, e.g., any pure strategy. Such a partition is convex if each element of the partition is a convex set, i.e., the intersection of half-spaces created by hyperplanes and $\mathbf{\Theta}$. In this paper we focus, for the most part, on convex partitional communication strategies for the expert.[8]

Given the specification of decision maker behavior, a (perfect Bayesian) equilibrium of the cheap talk game is fully specified by a communication strategy for the expert.[9] We say that the expert induces an action $a = E[\theta|m]$ in equilibrium if the decision maker chooses the action a after the message m. In any equilibrium, every action induced by the expert must maximize $U(\cdot)$. An equilibrium is influential if there are at least two different actions chosen by the decision maker with strictly positive probability, i.e., the expert uses at least two different messages with distinct (equilibrium) meanings. We use the term k-message equilibrium to refer to the case where the decision maker chooses k different actions, i.e., the expert uses messages with k distinct equilibrium meanings. We are now ready to state our first result.

Theorem 1 *An influential cheap talk equilibrium exists for all U and F.*

We prove Theorem 1 by considering a partitional communication strategy created by a single hyperplane h passing through a point c in the interior of $\mathbf{\Theta}$, i.e., the expert discloses which "side" of the hyperplane the true state θ lies in. Figure 1(a) shows an example with $N = 2$ where a news network knows the seriousness of two scandals represented by the random variables θ_1 and θ_2. The audience knows the ex ante distributions are uniform i.i.d. on $[0, 1]$ and forms updated estimates given a message m as represented by actions $(a_1, a_2) = (E[\theta_1|m], E[\theta_2|m])$. To boost ratings the network wants to promote the seriousness of either scandal, and for partisan reasons the network wants to exaggerate the first scandal in particular, as seen by the indifference curves in the figure for $U = 4a_1 + a_2$.

[7] This is where we depart from the Crawford-Sobel model with state-dependent preferences. See more on this difference in Section 3.

[8] Lemma 1 in Crawford and Sobel (1982) shows that all equilibria involve partitional communication strategies, a fact that does not extend to our setup.

[9] We assume that all messages in $\mathbf{M}$ are used in equilibrium, and accordingly avoid specifying off-equilibrium-path beliefs. This is without loss of generality in a cheap talk game.

The hyperplane h divides $\mathbf{\Theta}$ into two non-empty convex regions, $\mathbf{R}^+$ and $\mathbf{R}^-$ with associated actions $a^+ = E[\theta|\theta \in \mathbf{R}^+]$ and $a^- = E[\theta|\theta \in \mathbf{R}^-]$. As the hyperplane is spun around $c = (\frac{1}{2}, \frac{1}{2})$ from vertical to vertical again, the actions for the two regions trace out the "circular" path shown in the figure. Since the actions reverse themselves as the hyperplane is rotated continuously, and since preferences are continuous, at some intermediate point the actions must fall on the same indifference curve. For such a hyperplane the messages balance out the network's incentives on each dimension so it is indifferent between the induced actions and therefore has no incentive to lie.

Theorem 1 uses the Borsuk-Ulam theorem to show that this argument is completely general and does not depend on the specific choice of preferences and priors nor on the number of dimensions $N \geq 2$. In general the hyperplane h through c is identified by its orientation vector s, a point on the unit sphere $\mathbb{S}^{N-1} \subset \mathbb{R}^N$. Since the difference $U(a^+) - U(a^-)$ is a continuous odd map as a function of the orientation s, the difference must be zero for some s^* (and associated hyperplane).[10] Since the regions $\mathbf{R}^+$ and $\mathbf{R}^-$ are convex, the actions are contained in each region and therefore distinct, $a^+ \neq a^-$ and the resulting equilibrium is influential. The expert's message can be interpreted to be disclosing in which direction (s^* or $-s^*$) the true state θ stands in relation to the reference point c.

The construction in Figure 1 illustrates one reason why communication is easier in a multidimensional state space. If $N = 1$ and $\mathbf{\Theta}$ is an interval, h would just be a point that divides the interval into two regions. If preferences are continuous then the intermediate value theorem implies that an influential equilibrium exists in the one-dimensional case if and only if the expert switches which of the two actions associated with each region is preferred as h varies over the interval. Such a "switching condition" cannot be satisfied, for example, by monotonic preferences.[11] In contrast, in our case with $N > 1$, the hyperplane can be spun around in a continuous manner so that the actions reverse themselves, implying that the switching condition is always satisfied if preferences are continuous, even when

[10]This is a direct application of the Borsuk-Ulam theorem: every continuous odd function g from $\mathbb{S}^{N-1}$ to $\mathbb{R}^{N-1}$ must have the origin in its image, i.e., $g(s^*) = 0$ for some $s^* \in \mathbb{S}^{N-1}$ (see, e.g., Matousek, 2003). A function $g(s)$ is odd if $g(-s) = -g(s)$ for all s. Throughout the paper, we think of h as the hyperplane passing through c that is parallel to the tangent to the unit sphere at s.

[11]In Crawford and Sobel (1982), this switching condition is generated by a sufficiently small conflict of interest (or bias) parameter, implying that for sufficiently low states the expert prefers a lower action to a higher action, and vice versa for sufficiently high states. Since preferences are state-dependent, a single-crossing condition is also required to ensure that non-indifferent types have the right incentives.

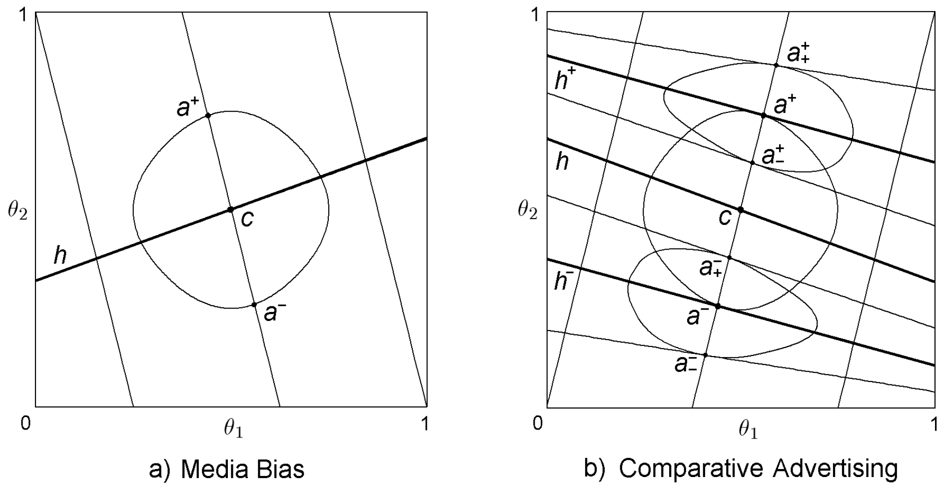

Figure 1: Equilibrium construction

the expert's preferences are monotonic as in Figure 1 (a), and even when the expert's utility is increasing in some dimensions and decreasing in others. Figure 1(b) shows such a case where $U = 4a_1 - a_2$, e.g., a political campaign wants to raise impressions of the quality of its candidate and to lower impressions of the quality of the competing candidate. In equilibrium if the campaign says something favorable about its own campaign then impressions of both candidates improve, while if the campaign says something negative about the competing candidate then impressions of both candidates fall, a pattern that has been observed empirically in consumer product advertising (Jain and Posovac, 2004) and political advertising (Lau et al., 1999).[12] The figure shows this tradeoff with actions a^+ and a^- for hyperplane h on the same rising indifference curve. The additional hyperplanes and actions in the figure will be used later to illustrate Theorems 3 and 4.

The value of cheap talk is limited if only one of the messages is typically ever sent, implying that the equilibrium is rarely informative. This is not an issue for these examples or more generally because, for any F, there exists an interior point (the centerpoint) such that for any hyperplane through the point the probability mass of each half-space is at least $1/(N+1)$ (Grunbaum, 1960). So if c is this centerpoint then each equilibrium action is chosen with at least probability 1/3 for $N = 2$.[13] Moreover, for $N > 2$, there are multiple

[12] Polborn and Yi (2006) model this phenomenon in a disclosure/persuasion game.

[13] For any logconcave density f and any hyperplane through $E[\theta]$, this lower bound on the probability of

possible equilibrium hyperplanes for any c in the interior of $\mathbf{\Theta}$, including at least one where both the differences $U(a^+) - U(a^-)$ and $\Pr[\theta \in \mathbf{R}^+] - \Pr[\theta \in \mathbf{R}^-]$ are equal to zero and each equilibrium action is chosen with ex ante probability of 1/2.[14]

Any (informative) communication strategy induces a mean-preserving spread in the decision maker's updated estimates, implying that the expert benefits from being informative, relative to remaining silent or "babbling", whenever her preferences are convex.[15] In any cheap talk *equilibrium* the weaker condition of quasiconvexity is sufficient for communication to be beneficial. Since the lower contour sets of a quasiconvex U are convex, and since the prior estimate $E[\theta]$ is a convex combination of the posterior estimates $E[\theta|\theta \in \mathbf{R}^+]$ and $E[\theta|\theta \in \mathbf{R}^-]$ that both lie on the same indifference curve in equilibrium, the prior estimate must lie on a lower indifference curve. Therefore the expert benefits from communication not just in expectation, but for every realization of her private information.

Theorem 2 *Relative to no communication, any informative cheap talk equilibrium (strictly) benefits the expert if U is (strictly) quasiconvex and (strictly) hurts the expert if U is (strictly) quasiconcave.*

Quasiconvex preferences for the expert arise in many environments of interest where the expert is communicating to multiple decision makers with different preferences. Several such examples are analyzed in Section 4 and pictured in Figure 2. The reader interested mainly in applications may want to jump directly to these examples without any loss in continuity. In the rest of this section and the next we characterize further the equilibrium set and perform some robustness tests.

A natural question is whether more information can be revealed than in the two-message cheap talk equilibrium of Theorem 1. Consider again Figure 1(b) where the hyperplane h is for a single comparative message corresponding to the upper and lower regions with equilibrium actions a^+ and a^-. Now suppose we follow the same procedure for the upper region to spin a new hyperplane h^+ around a^+, and also follow the same procedure for

each half-space rises to $1/e$ for any N (Caplin and Nalebuff, 1991).

[14]This is analogous to the standard example of the Borsuk-Ulam theorem: if both temperature and barometric pressure vary continuously over the earth's surface then there must be a pair of antipodal points with the same temperature and pressure. We exploit the power of extra dimensions given by the Borsuk-Ulam theorem further for later results.

[15]A babbling equilibrium in which the expert's messages convey no information exists in any standard cheap talk game.

the lower region to spin a new hyperplane h^- around a^-. The resulting actions for the top region are then a_+^+ and a_-^+, and for the lower region are a_+^- and a_-^-. Is the resulting partition of the entire space a four-message equilibrium? In this example with linear preferences the actions must all lie on the same indifference curve going through $E[\theta]$ so the expert has no incentive to lie and the partition is an equilibrium.

To better understand this potential for multiple messages to reveal more information, consider a general version of the linear specification used in Figure 1,

$$U(a) = \rho \cdot a \tag{1}$$

where the $\rho = (\rho_1, ..., \rho_N)$ is a vector of real numbers indicating the relative weights on each dimension.[16] One may think of ρ as measuring the expert's bias across dimensions, i.e., her "slant". For the following result we extend the argument from Figure 1(b) to show that, with linear preferences, it is possible to repeatedly apply Theorem 1 to obtain an arbitrarily fine partition. Indeed, we show that linearity allows us to obtain an even stronger result.

Theorem 3 *An informative cheap talk equilibrium revealing almost all information on $N-1$ dimensions exists if U is linear.*

To see this result for two dimensions, consider a 2^k-message equilibrium where each region is repeatedly subdivided by a hyperplane through the centerpoint of the region as in Figure 1(b), implying by the centerpoint theorem that each element has probability mass at most $\left(1 - \frac{1}{N+1}\right)^k$ which goes to zero as k increases. If $N = 2$, the full support assumption implies there must lie an action within any $\varepsilon > 0$ of any point on the equilibrium indifference line for k large enough. In this sense, the expert reveals almost all information in one dimension of the two-dimensional space as k becomes large. Now consider the case of $N = 3$ where Theorem 3 asserts that all actions must in the limit fill up a two-dimensional surface corresponding to the equilibrium hyperplane of the expert. We first "slice" $\mathbf{\Theta}$ by constructing a 2^k-message equilibrium where the equilibrium actions lie on a line in the indifference plane through $E[\theta]$. We can choose this line freely because of the extra degree of freedom given by the fact that $N > 2$. We then "dice" each slice with a 2^k-message equilibrium for the subregion where the equilibrium actions lie on lines that are, again

[16] Linear preferences are standard in the signaling and disclosure/persuasion game literatures, but are not generally used in the cheap talk literature since monotonic preferences preclude credible cheap talk in the one-dimensional case.

using the extra degree of freedom, chosen to be orthogonal to the original line. Since all these 2^{2k} actions lie on the expert's indifference plane through $E[\theta]$, they constitute a 2^{2k}-message equilibrium. For large k the action lines form an arbitrarily fine grid of the expert's equilibrium indifference hyperplane, allowing us to conclude by the centerpoint theorem arguments above that the equilibrium hyperplane is asymptotically "filled up" by the equilibrium actions as k increases. The same procedure of repeated slicing and dicing extends to all N.

Theorem 3 returns to an idea due to Battaglini (2002) that in two dimensions it can be possible to reveal full information in a one-dimensional subspace (the dimension of agreement) on which there is no conflict of interest. As Battaglini noted, for state-dependent preferences such revelation can only occur in special cases with special distributions.[17] We find that such revelation is possible generally for linear preferences that are transparent. Moreover, in Section 3 we show that linear preferences are of particular interest since they are the limiting case of standard Euclidean preferences as biases become large.

Now consider the potential for multiple messages when preferences are not linear. When we divide the initial two regions further as in Figure 1(b), we run into the problem that application of Theorem 1 to each region ensures that $U(a_+^+) = U(a_-^+)$ and $U(a_+^-) = U(a_-^-)$, but without linearity there is no assurance that all four actions are on the same indifference curve. The following Theorem shows that, when preferences are quasiconcave or quasiconvex, mixed messages can be used to equalize the payoffs and ensure that a cheap talk equilibrium with an arbitrarily large number of informative messages exists.

Theorem 4 *A 2^k-message informative cheap talk equilibrium, in which the expert's payoff is strictly increasing (decreasing) in k, exists for all strictly quasiconvex (quasiconcave) U and all $k \geq 1$.*

We prove the theorem by an inductive algorithm that creates a 2^k -message equilibrium from a 2^{k-1}-message equilibrium. We illustrate the logic here by creating a four-message equilibrium from a two-message one. Assume preferences are strictly quasiconvex and starting from a two message equilibrium with induced actions a^+, a^- we find a hyperplane for each of its halfspaces and four corresponding actions satisfying $U(a_+^+) = U(a_-^+)$ and $U(a_+^-) = U(a_-^-)$. If we think of the messages in the two-message equilibrium as m^+ and

[17] Battaglini discusses the difficulty of such revelation as a step toward understanding why multiple experts are needed to obtain full revelation in all dimensions.

m^-, we may think of the four-message equilibrium as one where each of these two messages is split into two, i.e., m^+ to m^+_+ and m^+_- and m^- to m^-_+ and m^-_-.

If, coincidentally, $U(a^+_+) = U(a^+_-) = U(a^-_+) = U(a^-_-)$, we have a four-message equilibrium. Suppose instead that $U(a^+_+) = U(a^+_-) > U(a^-_+) = U(a^-_-)$. Due to the strict quasiconvexity of U, we must also have $U(a^-_+) = U(a^-_-) > U(a^+) = U(a^-)$. Instead of disclosing the four-message partition above, consider now the possibility that the expert mixes over the messages m^+_+ and m^+_- when $\theta \in \mathbf{R}^+$. Mixing decreases the informativeness of the messages and moves the corresponding expectations a^+_+ and a^+_- closer together. In the extreme case where mixing by the expert conveys only the coarse information $\theta \in \mathbf{R}^+$, both a^+_+ and a^+_- equal a^+. By continuity, there must be some mix such that the expert's utility from mixing is the same as $U(a^-_+) = U(a^-_-)$, implying a four-message equilibrium exists in which two of the messages are mixed over. This algorithm can be extended to generate equilibria with 2^k messages for any $k > 2$, and for all $N \geq 2$. By strict quasiconvexity, the expert's payoff is strictly increasing in k. A similar algorithm also exists if preferences are strictly quasiconcave, though in this case the extra communication hurts the expert.

Notice that by construction such an equilibrium communication strategy typically requires the expert to mix between messages and it is not partitional, although it has a partition of $\mathbf{\Theta}$ associated with it. Also note that since all the induced actions lie on the same indifference curve, this mixed message equilibrium is also a sequential cheap talk equilibrium, in the sense that we can think of the expert making successive statements corresponding to the successive stages of the algorithm. In this sense the "longer" that cheap talk continues, the more information that is revealed.[18]

3 Robustness: Near Transparency

Our results so far concern the case where the expert's preferences $U(a)$ over the decision maker's actions are transparent, i.e., common knowledge, and so independent of the state θ. In this section we investigate to what extent our results are robust to the decision maker facing some limited uncertainty about the expert's preferences, including the possibility that preferences are correlated with the state θ.

[18] Krishna and Morgan (2004) and Aumann and Hart (2003) consider mixed strategies in multi-stage cheap talk games. Whereas the multi-stage aspect of communication is central to their results, in our game the communication may also be one-stage.

We model near transparency as follows. Suppose that the expert's preferences are given by a function $U(a,t)$, continuous in a, where $t \in \mathbf{T}$ is the type of the expert. The expert knows her own type t (in addition to θ), but the decision maker only has a prior Φ on t. Let $F(\theta|t)$ summarize the conditional distribution of θ given $t \in \mathbf{T}$. This approach allows t to be independent of θ or to be correlated with θ. It also covers the Crawford-Sobel model where $t = \theta$. More importantly, it allows us to conceptually separate uncertainty about the expert's motives t from that about the decision maker's ideal course of action θ. Notice that the expert's type t is fully specified by the pair $U(.,t)$ and $F(.|t)$.

We perform three distinct robustness tests in this section. The first concerns the case where the prior Φ is concentrated on a particular type $t^* \in \mathbf{T} = \{1, ..., T\}$. This captures a case where the decision maker is almost certain that the expert has preferences $U(a, t^*)$, although there are potentially many other unlikely possibilities. The second robustness test allows for arbitrary Φ but supposes that T is small. This captures a case where the decision maker attaches positive probability to only a few possible expert types, each of which might be quite probable. The third robustness test relaxes the equilibrium notion from cheap talk to epsilon cheap talk and focuses on the case where the expert's type space is rich, i.e., $t = \theta$, but her type has a limited effect on (differences in) her utility.

3.1 Almost Certain Motives

Suppose that the prior $\Phi = (\phi_1, ..., \phi_T)$ is close to the degenerate distribution Φ^* on type $t^* \in \mathbf{T}$ (i.e., $\phi^*_{t^*} = 1$). We use the implicit function theorem to look for equilibria in the neighborhood of the equilibria identified by Theorem 1 for the degenerate case, for general preferences and conditional distributions of θ. In such equilibria the decision maker anticipates that the low probability types will not be indifferent so they will always offer the same advice. For instance, the decision maker thinks that the expert is probably unbiased across dimensions, $U = a_1 + a_2$, but there is some chance that the expert has a relatively extreme slant, $U = 4a_1 + a_2$, in which case she will offer the same advice regardless of the state θ.[19]

To apply the implicit function theorem we assume that $U(a,t)$ is continuously differen-

[19] In contrast with related models of biased and unbiased types in one dimension (Sobel, 1985; Benabou and Laroque, 1992; Morris, 2001; Morgan and Stocken, 2003), the "unbiased" type here is unbiased across dimensions but still biased toward a higher action in each dimension.

tiable in a for each $t \in \mathbf{T}$ and, for simplicity, consider only the case $N = 2$.[20] We also suppose that the types $t \in \mathbf{T}$ have different preferences from each other in the following sense: for any two action profiles a, a' with $a \neq a'$, if $U(a, t^*) = U(a', t^*)$, then $U(a, t) \neq U(a', t)$ for all $t \neq t^*$, $t \in \mathbf{T}$. We call this condition (S). Notice that it will hold if, for instance, the indifference curves of the different types satisfy a single-crossing property in $\mathbb{R}^2$.

Proposition 1 *Suppose U satisfies (S) and $N = 2$. Generically in $U(., t^*)$, $F(.|t^*)$, there exists $\varepsilon > 0$ such that for each Φ with $||\Phi - \Phi^*|| < \varepsilon$ an influential cheap talk equilibrium exists.*

In the influential equilibrium of Proposition 1, type t^* discloses a two-message partition of $\boldsymbol{\Theta}$ similar to the equilibria of Theorem 1. By condition (S), no other type can be indifferent between two induced actions and so will send one of the two messages with probability one. Since Φ is close to Φ^*, the induced actions (and the equilibrium partition) is close to the equilibrium of the case where the expert's likely type t^* is common knowledge. The decision maker essentially ignores the implications of messages from unlikely types $t \neq t^*$ in determining his ideal course of action.[21]

3.2 Few Possible Motives

We now consider the case where no single type is most likely but the possible number of types T is small relative to N. For instance, a car magazine might be biased toward a particular car manufacturer. If the magazine's only information of interest to readers is the quality of each manufacturer, then credible communication is clearly a problem if the reader does not know which manufacturer the magazine favors. But if the magazine has information on multiple models for each manufacturer and/or on multiple attributes of each model, then the dimensionality of information increases and it would seem that comparative statements about different models or about the relative strengths of particular models should still be credible. Application of the Borsuk-Ulam theorem confirms this intuition quite generally.

[20] With some modifications, the arguments of this subsection extend to the case $N > 2$ via the arguments of the next subsection.

[21] Unlikely types always influence the decision maker in their preferred direction. This provides a multidimensional perspective on the finding by DellaVigna and Kaplan (2007) that, in an environment where few viewers expected a news network to have a conservative bias, Fox News had a large influence on voting behavior by viewers.

Proposition 2 *Suppose $T < N$. Then an influential cheap talk equilibrium exists for all* $\mathbf{T}$ *and* Φ.

When $T < N$, it is possible to find a two-message partition of $\boldsymbol{\Theta}$ that is an equilibrium for every type $t \in \mathbf{T}$. In other words, it is possible to find an informative communication strategy that induces actions in $\mathbb{R}^N$ that the T possible types of the expert agree on. This is not surprising for linear $U(.,t)$ since satisfying all the T experts simultaneously still leaves $N-1-T$ degrees of freedom. The result shows that this intuition extends to all continuous preferences, i.e., it suffices to simply count equations and unknowns.[22]

3.3 Epsilon Cheap Talk

For our third robustness test we consider preferences that are highly state-dependent, but that converge to state-independent preferences. As a notable example of such preferences, we consider Euclidean preferences where the expert's utility is based on the distance between the expert's ideal action and the decision maker's ideal action,

$$U(a;\theta) = -d(a, \theta + b) = -\left(\sum_{i=1}^{N}(a_i - (\theta_i + b_i))^2\right)^{1/2} \tag{2}$$

where $d(\cdot,\cdot)$ is the Euclidean distance function and $b = (b_1, ..., b_N) \in \mathbb{R}^N$ is the vector of known biases representing the distance between the decision maker's ideal action θ, and the expert's ideal action $\theta + b$. Distance preferences are used in the leading example from Crawford and Sobel and in a wide variety of applications.

As the expert's bias in each dimension increases, Euclidean preferences converge uniformly to state-independent linear preferences with known biases across dimensions equal to the ratios of these biases within dimensions. More precisely, if we write $b = \rho B$ for some vector $\rho \in \mathbb{R}^N$, $\rho \neq 0$, and real number $B \geq 0$, then for any θ, as B increases without bound the expert's ideal point $\theta + b$ becomes more and more distant from θ, and the circular indifference curves for Euclidean preferences become straighter, and converge to those of known linear preferences of the form $\rho \cdot a$ given by (1).[23]

[22] Proposition 2 also provides insight on how the result of Proposition 1 can be extended to the case where $N > 2$. In brief, one applies the implicit function theorem to the equilibrium of Proposition 2 constructed with $N-1$ types (including t^*) disclosing a partition and the remaining types sending only one message. Condition (S) then has to be suitably amended.

[23] All arguments go through unchanged if each b_i is of the form $b_i = \kappa_i + \rho_i B$, for constants κ_i. In such

To use this convergence, we modify the game so that the expert's payoff from any action a and message m given θ is $U(a, m; \theta) = -d(a, \theta + b)$ less an arbitrarily small cost $\varepsilon > 0$ of lying if the message m is not consistent with θ.[24] We study if influential equilibria exist in the modified game with distance preferences (2) and an arbitrarily small cost of lying. We say that a communication strategy is an *ε-cheap talk equilibrium for large biases* of the game with distance preferences if and only if for each $\varepsilon > 0$ there exists $\overline{B}$ such that for all $B > \overline{B}$ and any θ, the incentive to lie for an expert is at most ε. Our next result shows an equivalence between such equilibria and the cheap talk equilibria for linear preferences characterized by Theorem 3.

Proposition 3 *Suppose U is Euclidean. Then for all F and all k, a communication strategy is a k-message ε–cheap talk equilibrium for large biases if and only if it is a cheap talk equilibrium for the limiting linear U with slant ρ.*

We use this equivalence in our last application of the following section. As described in that application, Figure 2(d) shows Euclidean preferences that are close to linear preferences and shows how lack of full linearity creates a small incentive to deviate from a pure cheap talk equilibrium for the corresponding linear preferences.[25, 26]

---

cases, as B becomes large the expert becomes infinitely biased in dimensions where $\rho_i \neq 0$ but has finite (possibly, no) bias in dimensions where $\rho_i = 0$.

[24] Since the meaning of a message is derived from a (candidate) equilibrium communication strategy, the notion of what constitutes a "lie" is endogenous here. Therefore this equilibrium notion is distinct from that of an "almost cheap talk" equilibrium (Kartik, 2008) or a "costly talk" equilibrium (Kartik, Ottaviani, and Squintanni, 2007) in which the sender's reports have an exogenous meaning corresponding to the true value of the state and any deviation from this true value is costly.

[25] In an extension of their analysis of lexicographic preferences, Levy and Razin (2007) show for Euclidean preferences that even slight asymmetries in distributions can preclude pure cheap talk for sufficiently large but finite B. Our result shows that for sufficiently large B any incentive to deceive the decision maker is arbitrarily close to 0.

[26] Epsilon equilibria are not invariant to monotonic transformations of the underlying preferences. For instance, with a quadratic variant of the Euclidean specification which drops the square root term in (2), the difference in the sender's utilities from actions a and a' is unbounded in B, implying that our equivalence result obtains only if the cost of lying also increases in the unit of payoffs B, e.g., if it is equal to εB for any $\varepsilon > 0$. Note however that indifference curves corresponding to such quadratic preferences are also linear *at* the limit of infinite biases.

4 Applications

In this section we present four situations where the transparency of an expert's preferences allows for credible and beneficial cheap talk. The first considers the effect of voting rules on the gains from information disclosure by a partisan expert. The second concerns credible and revenue enhancing disclosure policies for the seller in a private value auction. The third application shows that there is a role for "cheap talk advertising" about product attributes. The last application considers the interaction between bias, informativeness and the allocation of control rights inside organizations.

4.1 Influencing Voters

Following Condorcet's early analysis of jury voting, a central question in the voting literature is how voting rules affect the aggregation of information (e.g., Feddersen and Pesendorfer, 1998; Martinelli, 2002). Recently Coughlan (2000) and Austen-Smith and Feddersen (2006) extend this literature to consider cheap talk between voters before voting, and show that the unanimity rule for jury convictions discourages information flows among informed jurors with different preferences. In this example we consider the distinct but related question of information flows from an interested third party with transparent preferences to voters. In contrast with Farrell and Gibbons (1989), we allow this party to have multi-dimensional information.

Consider a jury trial with heterogeneous jurors that are influenced by two different aspects of a case. There are $N = 2$ groups of jurors where group $i = 1, 2$ prefers to vote for conviction if and only if $\theta_i > \tau_i$, where $\tau_i \in [0, 1]$ is a privately known threshold that is distributed uniformly and independently of θ and τ_{-i}. If conviction requires a unanimous vote, then the expected probability of conviction is $P(a) = a_1 a_2$, a strictly quasiconcave function in the interior of $[0, 1]^2$, so the probability of acquittal $1 - P(a)$ is strictly quasiconvex.[27] It follows from Theorem 1 that a defense lawyer who knows θ and is looking to maximize the probability of acquittal can engage in credible and influential cheap talk. And it follows from Theorem 2 that in any influential cheap talk equilibrium the defense benefits relative to no communication regardless of the merits of the case θ, i.e., the probability of a "wrongful" conviction is decreased but so is the probability of a

[27] More generally, the probability of acquittal is strictly quasiconvex for any N so the result extend to $N > 2$ different voters interested in different aspects of the case.

"correct" conviction.

Result 1 *The defense can strictly lower the probability of conviction via cheap talk under the unanimity rule.*

The unanimity rule encourages the defense to provide information that weakens the defendant's case in one dimension but strengthens it in another dimension in the hope that at least one type of voter will be persuaded to vote for acquittal. By the same token if the prosecution knew θ it would be worse off from any communication via cheap talk. Figure 2(a) shows a two-message equilibrium for i.i.d. uniform θ with the defense's quasiconvex preferences $U = 1 - a_1 a_2$ where utility is decreasing away from the origin as the probability of conviction increases. Note also that by Theorem 4 the defense has an arbitrarily "long" informative speech involving mixed messages where the probability of ultimate acquittal is strictly increasing in the length of the speech.

If only a majority rather than unanimity is required for conviction, the defense will want to appeal to the larger group, say group 1, so that $U = 1 - a_1$. Since preferences are effectively one-dimensional and monotonic, cheap talk cannot affect the outcome of the vote. However, this conclusion is sensitive to the assumption that all voters can be influenced. Suppose there is a third group of voters that is not influenced by additional information and which forms a majority if and only if it combines with either of the other groups. If the group will always vote for acquittal then the defense only needs to influence either of the other two groups to vote for acquittal. In such a case we again have $U = 1 - a_1 a_2$ and the defense can benefit from cheap talk. But if the third group will always vote for conviction then the defense needs to convince both of the other types to vote for acquittal, implying $U = 1 - (a_1 + a_2 - a_1 a_2)$, which is a quasiconcave function. The defense is then better off without communication while the prosecution is better off with it.

In general, influential cheap talk by a third party must have a divisive effect in equilibrium by attracting one group of voters at the expense of alienating the other. In a political context, this implies that a candidate seeking bipartisan support from both groups of influenceable voters is better off from babbling about the future. In contrast, another candidate who can get by with support from at most one group has an incentive to be informative about intended policies. In turn, the number of groups a candidate seeks support from may be affected by the relative strength or weakness of her position within the electorate.

A strong candidate with a relatively large group of dedicated followers can further benefit from being informative whereas a weak candidate with a smaller such group cannot.

4.2 Persuading Bidders

Should a seller reveal product information to buyers in a private value auction? Communication increases the chance that the buyer who values the product the most bids the most for it, but it can also soften competition by increasing the spread between buyer valuations. With a large enough number of bidders, the gains from allocating the good more efficiently dominate *on average*, so that the seller has higher expected revenues from committing to a disclosure policy (Ganuza, 2004; Board, 2006).[28] However, depending on the exact information that the seller has, revealing the information may increase or decrease revenues. In this application we show that with multidimensional information the seller can credibly communicate information through cheap talk without any need for commitment to a disclosure policy. Moreover, with sufficiently many bidders, the seller is *always* better off from such communication for any realization of her information.

Consider a seller with information on $N = 2$ attributes of an object who engages in cheap talk with $n \geq 2$ potential buyers prior to holding a second-price auction. Buyers have correlated private values, $v_j = \lambda_j \theta_1 + (1 - \lambda_j)\theta_2$, where the seller knows $\theta \in [0,1]^2$ and $\lambda_j \in [0,1]$ is the private information of bidder j. For instance we may think of θ_1 (θ_2) as the short-run (long-run) value of an asset, with λ_j capturing the time preference of buyer $j = 1, ..., n$. We assume that the λ_j's are stochastically independent of θ, but allow correlation between the λ_j's as long as ties are zero probability events.[29]

In a second price auction the buyers bid their expected values $E[v_j|m] = \lambda_j a_1 + (1-\lambda_j)a_2$ given the seller's message m (and the associated estimates $a = E[\theta|m]$), so that the seller's

[28] The ex ante gains in allocation efficiency due to seller communication are also addressed in Ganuza and Penalva (2006) and Bergemann and Pesendorfer (2007). Chakraborty, Gupta, and Harbaugh (2006) consider the gains from credible seller communication about multiple goods based on the linkage principle (Milgrom and Weber, 1982). Campbell (1998) and Miralles (2008) consider the gains to buyers from credible communication among them about which goods to bid on.

[29] This formulation of buyer valuations ensures that the identity of the winning bidder varies with (buyer estimates) of the seller's information θ. As long as this is guaranteed, the results remain qualitatively unchanged. See Board (2006) for similar conditions for the case where the seller can commit to a disclosure policy.

expected revenue is the expected second-highest bid

$$U(a) = E[2^{nd} \max_{\lambda_j}\{\lambda_j a_1 + (1-\lambda_j)a_2\}] = \begin{cases} E[\lambda_{2:n}]a_1 + (1 - E[\lambda_{2:n}])a_2 & \text{if } a_1 \leq a_2 \\ E[\lambda_{n-1:n}]a_1 + (1 - E[\lambda_{n-1:n}])a_2 & \text{if } a_1 > a_2 \end{cases}$$

where, following standard order statistic notation, $\lambda_{j:n}$ is the j-th lowest of the buyer signals. Since $U(a)$ is continuous in a, by Theorem 1 the seller can credibly disclose comparative information through pure cheap talk, for all n and any priors on θ and λ.

Notice from the expression for $U(a)$ that the (expected) price equals either the (expected) bid of the buyer with the second-lowest private signal $\lambda_{2:n}$ or that of the buyer with the second-highest one $\lambda_{n-1:n}$. The effect of informative cheap talk on expected revenues therefore depends on the relative magnitudes of these two variables. With $n = 2$, $\lambda_{2:n} > \lambda_{n-1:n}$ almost surely so that

$$U(a) = \min\{E[\lambda_{2:n}]a_1 + (1 - E[\lambda_{2:n}])a_2, E[\lambda_{n-1:n}]a_1 + (1 - E[\lambda_{n-1:n}])a_2\},$$

a concave function of a, since $\min\{.\}$ is a concave function. Since concavity implies quasiconcavity, from Theorem 2 we conclude that informative cheap talk cannot raise expected revenue when $n = 2$. Similarly, when $n = 3$, $\lambda_{2:n} = \lambda_{n-1:n}$ and $U(a)$ is linear in a, so that from Theorem 2 all cheap talk equilibria must yield the same expected revenue.

However, for $n \geq 4$, $\lambda_{2:n} < \lambda_{n-1:n}$ almost surely and so

$$U(a) = \max\{E[\lambda_{2:n}]a_1 + (1 - E[\lambda_{2:n}])a_2, E[\lambda_{n-1:n}]a_1 + (1 - E[\lambda_{n-1:n}])a_2\},$$

a convex function of a since $\max\{.\}$ is a convex function. Since convexity implies quasiconvexity, from Theorem 2 we conclude that informative cheap talk cannot lower expected revenue when $n \geq 4$. Since $U(a)$ is not strictly quasiconvex however we cannot directly use Theorem 2 to conclude that revenues from communication are strictly higher in this case. But we can still ensure this by choosing the reference point c to be at the kink of the seller's indifference curve through $E[\theta]$, i.e., with $U(c) = U(E[\theta])$. The induced actions in an informative equilibrium must then lie on different linear segments of the same indifference curve, implying that for each θ (and associated equilibrium message), the seller's expected revenue must be strictly higher than the expected revenues from silence, $U(E[\theta])$. The next result follows from this analysis.

Result 2 *The seller can strictly increase auction revenues via cheap talk if and only if there are at least four bidders.*

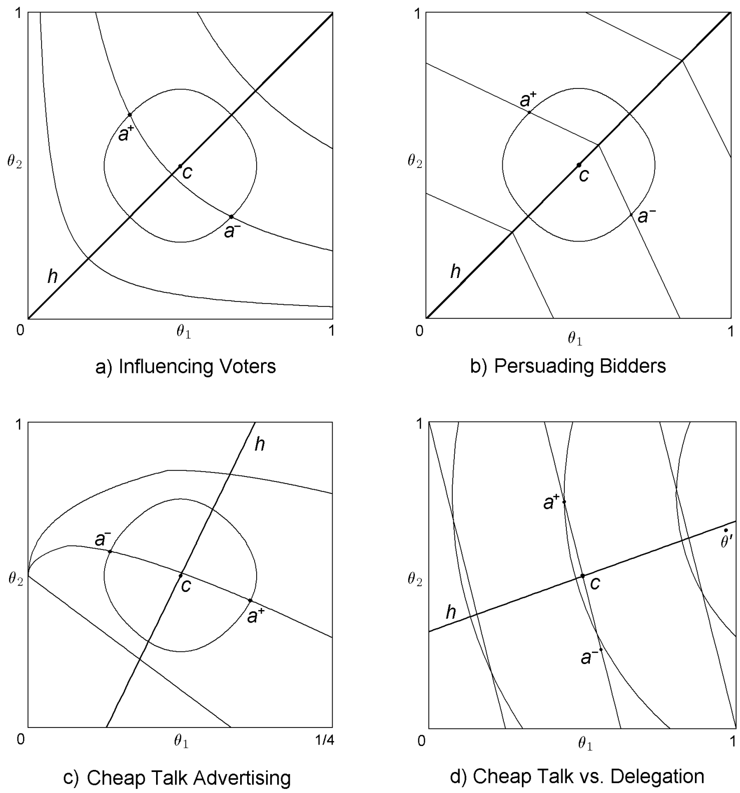

Figure 2: Applications

As an example consider the perfectly symmetric case where both the θ_i and the λ_j are i.i.d. uniformly distributed on $[0, 1]$. If there is no communication then $a_1 = a_2 = 1/2$ so seller revenues are $1/2$ for any n. If the seller discloses whether or not $\theta_1 \geq \theta_2$, then $a_1 = 2/3$ and $a_2 = 1/3$ or vice versa. Since $E[\lambda_{j:n}] = \frac{j}{n+1}$, the expected second highest bid in either case is $\frac{2}{n+1}\frac{1}{3} + \frac{n-1}{n+1}\frac{2}{3} = \frac{2}{3}\frac{n}{n+1}$ which is increasing in n and exceeds $1/2$ for $n \geq 4$. So, regardless of the actual realization of θ, the seller increases expected revenues through cheap talk if and only if $n \geq 4$.[30] Figure 2(b) shows the seller's quasiconvex preferences for the case of $n = 5$ so that the indifference curves have slope $-1/2$ above the diagonal and slope -2 below the diagonal.

In the case where the λ's are i.i.d., if instead of a second price auction we had a first price auction, then by the revenue equivalence theorem, for each a the expected revenue of

[30] If we exploit the stronger properties of concavity and convexity derived above, we obtain an additional but similar result, where revenues are computed in expectation over θ, for almost every realized value of λ.

the seller would still be given by the expression for $U(a)$ derived above. It follows that the set of informative cheap talk equilibria and their revenue implications are identical across all auction formats for which revenue equivalence obtains.[31]

4.3 Cheap Talk Advertising

The previous example indicates that a seller can benefit from cheap talk about product attributes. We now develop a model of cheap talk advertising with competition between sellers. Due to the conflict of interest between an advertiser and buyers, the literature on persuasive advertising has not examined whether the content of an advertisement can be credible even when it cannot be verified. As an example of how multidimensional information allows for cheap talk to have a role in advertising, we consider a price competition model where consumers are unsure of the extent of vertical and horizontal product differentiation between two firms.

Suppose that a unit mass of consumers is distributed uniformly on the unit line and the incumbent, firm 0, is located at 0 while the entrant, firm 1, is located at 1. A consumer located at $x \in [0,1]$ values the product of firm $i = 0,1$ at $v_i - t\,|x - i|$ less the price p_i. Any difference in the values v_0 and v_1 captures vertical differentiation and the parameter t measures horizontal differentiation. Marginal costs are set to zero. Consistent with the cheap talk assumption, any costs of advertising are fixed and do not vary with the content m. In what follows we assume that such fixed costs are small.

To capture uncertainty over the extent of vertical and horizontal differentiation of the entrant's product, assume the incumbent's value v_0 is common knowledge but t and v_1 are private information of the entrant. For simplicity we assume that the firms are ex-ante symmetric, i.e., $E[v_1] = v_0$. Prior to setting prices simultaneously with the incumbent, the entrant can engage in cheap talk advertising about t and v_1. If we interpret t and v_1 as the estimates of these objects commonly held by consumers and the incumbent given a message m from the entrant, then standard derivations imply that in equilibrium firm profits (gross

[31] For simplicity we assume there is no reserve price. This is optimal when the seller's reservation value is lower than the lowest possible virtual valuation of the buyers (e.g., Myerson, 1981).

of any advertising costs) given a message are[32]

$$\Pi_i = \frac{t}{2} + \frac{v_i - v_j}{3} + \frac{(v_i - v_j)^2}{18t}.$$

Since Π_1 is continuous, by Theorem 1 there exists an equilibrium of this game where cheap talk is informative.

As seen in the two-message equilibrium of Figure 2(c) for the case where t $(= \theta_1)$ is distributed uniformly on $[0, 1/4]$ and v_1 $(= \theta_2)$ is independently distributed uniformly on $[0, 1]$, the entrant can credibly emphasize either vertical or horizontal differentiation, but not both. Emphasizing that the product is better than the incumbent's positions the product to do well in direct competition for consumers but both firms price aggressively. Emphasizing that the product is well-differentiated from the incumbent's weakens the entrant's ability to compete directly for consumers but allows both firms to have more captive consumers. Checking the Hessian of the profit function confirms it is convex in (t, v_1) and therefore quasiconvex, so entrant profits increase due to communication regardless of the state. In fact, since all players' payoffs are convex, including the surplus of consumer $x \in [0, 1]$, $\max_i\{v_i - t\,|x - i| - p_i\}$, both firms and some consumers are strictly better off, while no consumer is worse off, in ex-ante expected terms from a better understanding of the entrant's product.

Result 3 *Informative cheap talk advertising is Pareto improving.*

Unlike in a signaling model of advertising (e.g., Milgrom and Roberts, 1986), the entrant's costs and benefits of advertising are not state-dependent so consumers do not learn about the product from the expense of advertising, but rather learn from the content of the advertising. And unlike in a disclosure/persuasion game model of advertising (e.g., Anderson and Renault, 2006) there is no restriction that the content must be verifiable. Although the content is pure cheap talk, it affects prices, profits and market shares.

The incumbent firm in the model also learns about the entrant from the advertising so this result is also relevant to the large literature on how firms communicate information to each other. This literature typically assumes a pre-commitment to reveal information (e.g., Gal-Or, 1986), an unraveling incentive to reveal verifiable information (e.g., Okuno-Fujiwara, Postlewaite, and Suzumura, 1990), or a signaling incentive to show off favorable

[32] This assumes that the parameters (and associated expectations) are such that prices, $p_i = t + (v_i - v_j)/3$, and market shares, $s_i = 1/2 + (v_i - v_j)/6$, are positive and that all consumers consume.

information (e.g., Ziv, 1993).[33] With multidimensional information, cheap talk provides another avenue by which firms can communicate information profitably.

4.4 Delegation versus Cheap Talk

Should a decision maker rely on a biased expert for cheap talk advice, or instead just delegate the decision to the expert? Dessein (2002) finds that the flexibility gains from delegation often exceed the information gains from cheap talk, and that this is always true in the standard one-dimensional uniform-quadratic example over the range of biases such that informative cheap talk is an equilibrium.[34] This latter result does not extend to the corresponding multi-dimensional model if the environment is sufficiently symmetric (Chakraborty and Harbaugh, 2007). Since full delegation is worse than the no communication outcome for sufficiently large biases in each dimension, and since comparative cheap talk that beats no communication is an equilibrium for any biases if the environment is sufficiently symmetric, comparative cheap talk beats full delegation for sufficiently large biases.

To see how this ordering extends to arbitrary asymmetries, consider the standard Euclidean preferences in (2). Following Proposition 3, as the biases within each dimension increase, these preferences converge uniformly to linear preferences with biases across dimensions asymptotically proportional to the ratios of biases within dimensions. It therefore follows that there exists an epsilon cheap talk equilibrium for sufficiently large biases which has the properties described by Theorem 3, i.e., that reveals detailed comparative information. For the decision maker this information always beats no communication and it beats full delegation if the biases are sufficiently large.

Result 4 *There exists an epsilon cheap talk equilibrium that offers the decision maker strictly higher payoffs than either full delegation or no communication for sufficiently large biases.*

For example, suppose that a board, acting in the interest of shareholders, is deciding

[33] An exception is Baliga and Morris (2001) who consider cheap talk about the cost side under investment complementarities.

[34] An alternative to full delegation is partial delegation in which the expert is given some range of discretionary choice or is otherwise constrained (e.g., Holmstrom, 1984; Dessein, 2002; and Alonso and Matouschek, 2008).

whether or not to give full control over investment decisions on two projects to a manager who knows the optimal level of investments, θ_1 and θ_2 distributed uniformly on $[0,1]^2$, and has a preference for inefficiently high investments, particularly on the first project, as represented by Euclidean preferences with biases $b_1 = 1$ and $b_2 = 1/4$.[35] Figure 2(d) shows such preferences for a manager with information θ' along with the equilibrium hyperplane and indifference curves for linear preferences for $\rho_1/\rho_2 = 4$ as in Figure 1(a). Since θ' is slightly closer to a^+ than a^-, the manager receives higher utility from action a^+ even though θ' is in the region for action a^-, so the hyperplane cannot be part of a pure cheap talk equilibrium.[36] However, any such incentive to lie goes to zero as as b_1 and b_2 increase in the ratio $b_1/b_2 = 4$ and the resulting indifference curves more and more closely approximate linear indifference curves for $\rho_1/\rho_2 = 4$, so the partition is an epsilon cheap talk equilibrium.

By Theorem 3, such epsilon cheap talk can reveal detailed information that slices up the unit box so that the equilibrium actions "fill" the indifference line through $E[\theta] = (1/2, 1/2)$. Essentially, the board recognizes that the manager prefers more investment on the first project and in the epsilon cheap talk equilibrium the manager recommends a pair of actions that satisfy the "budget constraint" $\rho \cdot a = \rho \cdot E[\theta]$. For any small cost of lying it is in the interest of a biased manager to engage in such informative communication and this is beneficial for shareholders. In contrast, giving full control of the investment decision to the informed manager would result in decisions $a_1 = a_2 = 1$ regardless of her information θ, while no communication would result in decisions $a_1 = a_2 = 1/2$, both of which are worse for shareholders.

5 Conclusion

We show that transparency of an expert's motives ensures the existence of influential cheap talk equilibria in multidimensional environments. When the expert's motives are known, our results provide an intuitive solution where the decision maker treats comparative statements with enough skepticism for communication to be credible. The decision maker always

[35] These preferences could reflect private benefits and/or empire-building motives. Harris and Raviv (2005) consider a one-dimensional version of this model in which the board also has its own information.

[36] The biases are large enough that pure cheap talk cannot be an equilibrium for each project considered separately, but the environment is sufficiently symmetric that there is a pure comparative cheap talk equilibrium in which h has intercept .16 and slope .11. This equilibrium is less informative than the epsilon cheap talk equilibria we now consider.

benefits from such communication, and the expert benefits in equilibrium if her preferences are quasiconvex. We show that such preferences arise naturally in many standard economic situations, including auctions, voting, and product differentiation.

These results add to the previous literature which has shown how transparency affects communication when there is uncertainty over an expert's bias in one dimension. In many environments the key issue is not an expert's bias within dimensions, e.g., whether a salesperson wants a consumer to buy a product or not, but an expert's bias or slant across dimensions, e.g., whether she gains more from pushing one product over another. Our results provide a theoretical foundation for regulations and social conventions that promote increased transparency in such environments.

6 Appendix

Proof of Theorem 1: We look for an influential cheap talk equilibrium involving a single hyperplane $h_{s,c}$ of orientation $s \in \mathbb{S}^{N-1}$ passing through $c \in int(\mathbf{\Theta})$ that partitions $\mathbf{\Theta}$ into two non-empty sets $\mathbf{R}^+(h_{s,c})$ and its complement $\mathbf{R}^-(h_{s,c})$, with corresponding receiver actions $a^+(h_{s,c})$ and $a^-(h_{s,c})$. Let $\mathbf{R}^+(h_{s,c})$ be the region that contains the point $s+c$.

Notice first that under the assumed conditions on priors, $a^+(h_{s,c}) \in int(\mathbf{R}^+(h_{s,c}))$ and $a^- \in int(\mathbf{R}^-(h_{s,c}))$, implying in particular that $a^+(h_{s,c}) \neq a^-(h_{s,c})$ so that any such equilibrium, if it exists, is influential. Furthermore, $a^+(h_{s,c})$ and $a^-(h_{s,c})$ are continuous functions of s (with the subspace topology for $\mathbb{S}^{N-1}$) for any fixed $c \in int(\mathbf{\Theta})$. Notice next that for any two antipodal orientations $-s, s \in \mathbb{S}^{N-1}$, we must have $\mathbf{R}^+(h_{s,c}) = \mathbf{R}^-(h_{-s,c})$ and so $\mathbf{R}^-(h_{s,c}) = \mathbf{R}^+(h_{-s,c})$. It follows that $a^+(h_{s,c}) = a^-(h_{-s,c})$ implying in particular that the map $G(\cdot; c) : \mathbb{S}^{N-1} \to \mathbb{R}$ defined by

$$G(s,c) = U(a^-(h_{s,c})) - U(a^+(h_{s,c})) \tag{3}$$

is a continuous odd function of s. By the Borsuk-Ulam theorem, there exists $s^* \in \mathbb{S}^{N-1}$, such that $G(s^*, c) = 0$. The hyperplane through c with orientation s^* generates a two-message convex partitional equilibrium. ■

Proof of Theorem 2: Consider any k message equilibrium with induced actions $a^1, ..., a^k$ such that the actions satisfy $a^j = E[\theta|m^k]$ where m^k is a message that induces action a^k. By the law of iterated expectations, $E[\theta] = \sum_{j=1}^k p^j a^j$, where $p^j > 0$ is the probability that action a^j is induced in equilibrium, $\sum_{j=1}^k p^j = 1$. Since all the induced actions belong to the same level set $\{a|U(a) = U(a^1)\}$ in equilibrium, and since the prior expectation $E[\theta]$ is a convex combination of these actions, $E[\theta]$ must belong to the convex lower contour set $\{a|U(a) \leq U(a^1)\}$ whenever U is quasiconvex. In such cases, $U(a^1) = ... = U(a^k) \geq U(E[\theta])$ and communication makes the expert weakly better off relative to no information. Indeed, the last inequality must be strict whenever U is strictly quasiconvex since in such cases $E[\theta]$ must be in the interior of $\{a|U(a) \leq U(a^1)\}$. Symmetric remarks apply to the case where the U is quasiconcave with convex upper contour sets. ■

Proof of Theorem 3. Formalizing the discussion in the text, note that for the first hyperplane h obtained via Theorem **1** the region $\mathbf{R}^+(h)$ is a convex set since it is the intersection of a convex set $\boldsymbol{\Theta}$ and one of the half-spaces associated with h. Further, it has a non-empty interior since $a^+(h) \in int(\mathbf{R}^+(h))$. Picking $c \in int(\mathbf{R}^+(h))$ and applying Theorem **1** again, we are guaranteed the existence of a second hyperplane h^+ and associated convex regions $\mathbf{R}^+_+(h^+)$ and $\mathbf{R}^+_-(h^+) = \mathbf{R}^+(h) \backslash \mathbf{R}^+_+(h^+)$, and corresponding actions $a^+_+(h^+) \in int(\mathbf{R}^+_+(h^+)), a^+_-(h^+) \in int(\mathbf{R}^+_-(h^+))$, such that $U(a^+_+(h^+)) = U(a^+_-(h^+))$. Since the actions are all conditional expectations, by the law of iterated expectations,

$$a^+(h) = qa^+_+(h^+) + (1-q)a^+_-(h^+). \tag{4}$$

for $q = \Pr[\theta \in \mathbf{R}^+_+(h^+)|\theta \in \mathbf{R}^+(h)] \in (0,1)$. Since U is linear, we conclude that

$$U(a^+_+(h^+)) = U(a^+_-(h^+)) = U(a^+(h)) = U(a^-(h)) \tag{5}$$

so that a three-message influential equilibrium exists with induced actions $a^+_+(h^+)$, $a^+_-(h^+)$ and $a^-(h)$. The k-message case uses the argument above as an inductive step. This shows that a k-message equilibrium exists for all $k \geq 2$.

Note now that for any k-message equilibrium, all induced actions belong to the $N-1$ dimensional compact set $\mathbf{A}^* = \{a|\rho \cdot a = \rho \cdot E[\theta]\}$. We wish to demonstrate that for every $\varepsilon > 0$, there exists a k-message equilibrium with k large enough, with an induced action $a^* \in \mathbf{A}^*$ that is within ε distance of a, for all $a \in \mathbf{A}^*$. In this sense, the set $\mathbf{A}^*$ is asymptotically (in k) dense in the induced actions and we say that the expert reveals all information in the $N-1$ dimensions corresponding to $\mathbf{A}^*$.

First consider the case where $N = 2$ and fix $\varepsilon > 0$. Consider the ball $B_\varepsilon(a)$, open in $\mathbb{R}^2$, of radius ε and centered around $a \in \mathbf{A}^*$. Notice next that for k large enough, there exists an element $\mathbf{P}$ of the equilibrium partition such that $\mathbf{P} \cap \mathbf{A}^* \subset B_\varepsilon(a)$. This follows from the centerpoint theorem (see, e.g., Grunbaum, 1960) that each element has probability mass at most $(1-1/(N+1))^{\log_2 k} < \min_{a\in\mathbf{A}^*} \Pr\{\theta \in B_\varepsilon(a)\}$, for k large enough, provided that each stage of creating successively finer partition in the equilibrium construction above we choose the interior point c through which the corresponding hyperplane h passes as the centerpoint of the corresponding convex compact partition element. But then the equilibrium action $a^* \in \mathbf{P} \cap \mathbf{A}^*$, corresponding to the element $\mathbf{P}$, must lie within ε of a.

Next consider the case where $N > 2$, and construct a k-message equilibrium with the desired property as follows. Introduce $N-2$ fictitious linear preferences $U^{j_1}(a) = \rho^{j_1} \cdot a$ for $j_1 = 1, ..., N-2$. We use the index $j_1 = 0$ to denote the actual expert preferences, i.e., $\rho^0 = \rho$, and assume that all ρ^{j_1} are linearly independent. Using the Borsuk-Ulam theorem and arguments similar to the first part of the proof, we can construct a k-element partition of $\mathbf{\Theta}$ with the property that the resulting k action profiles lie on a unique line $\mathbf{A}_1^* = \{a|\rho^{j_1} \cdot a = \rho^{j_1} \cdot E[\theta], j_1 = 0, ..., N-2\} \subset \mathbf{A}^*$. We can then choose a second distinct set of $N-2$ linearly independent fictitious types and construct a k-element partition, for each element $i = 1, ..., k$ of the k-element partition obtained in the previous step, treating that element as the entire state-space, following the same procedure as above. The resulting k action profiles lie on a unique line $\mathbf{A}_{2,i}^* = \{a|\rho^{j_{2,i}} \cdot a = \rho^{j_{2,i}} \cdot E[\theta], j_{2,i} = 0, ..., N-2\} \subset \mathbf{A}^*$ that can be chosen to be orthogonal to $\mathbf{A}_1^*$ for each $i = 1, ..., k$, for a total of k^2 actions. Repeat this procedure $N-1$ times, at each step using a new set of $N-2$ distinct fictitious preferences to obtain successively the lines $\mathbf{A}_1^*, \{\mathbf{A}_{2,i}^*\}_{i=1}^k, ..., \{\mathbf{A}_{N-1,i}^*\}_{i=1}^{k^{N-2}}$. Notice that the resultant actions and associated partition of $\mathbf{\Theta}$ is a k^{N-1}-message equilibrium with each induced action on some line $\mathbf{A}_{N-1,i}^*$, $i = 1, ..., k^{N-2}$. For suitable choices of c, by the centerpoint theorem it follows that for any $\varepsilon > 0$ and k large enough, the element $\mathbf{P}$ of the partition with $a \in \mathbf{P}$, any $a \in \mathbf{A}^*$, must satisfy $\mathbf{P} \cap \mathbf{A}^* \subset B_\varepsilon(a)$, implying the existence of an equilibrium action $a^* \in \mathbf{P} \cap \mathbf{A}^*$ within ε of the given point a. ■

Proof of Theorem 4: Suppose, as part of the inductive hypothesis, that we have a 2^k-message equilibrium associated with a 2^k-element partition of $\mathbf{\Theta}$ created by 2^{k-1} hyperplanes, $k \geq 1$. Identify the j-th element of the partition μ^j (a compact convex subset of θ with non-empty interior) by the message m^j and the corresponding induced action by a^j. We suppose that message m^j is sent by all $\theta \in int\{\mu^j\}$ with probability $p^j > 0$ that does not depend on θ. Message m^j may also be sent by other types $\theta \notin \mu^j$ with positive probability. Let $z_{in}^j = E[\theta|m^j, \theta \in \mu^j]$ and $z_{out}^j = E[\theta|m^j, \theta \notin \mu^j]$, whenever defined. By the law of iterated expectations a^j is a probability weighted average of z_{in}^j and z_{out}^j, so that the line joining the latter two points must pass through a^j. Since we have an equilibrium, $U(a^1) = ... = U(a^{2^k})$. We proceed by induction on k by first creating 2^{k+1}-element partition of $\mathbf{\Theta}$ from the given 2^k-element partition. Next, we adjust the induced actions via mixed strategies in order obtain an equilibrium with 2^{k+1} messages.

For each m^j, consider a hyperplane through $z^j_{in} \in int(\mu^j)$ of orientation $s \in \mathbb{S}^{N-1}$ that splits μ^j into two regions $\mu^{j+}(s)$ and $\mu^{j-}(s)$ and expected values $z^{j+}_{in}(s)$ and $z^{j-}_{in}(s)$. Relative to the original 2^k-element partition, we think of this new 2^{k+1}-element partition where each corresponding message m^j is split into two messages m^{j+} and m^{j-} such that (i) each $\theta \in \mu^j$ sends message m^{j+} (resp. m^{j-}) with the same probability $p^j > 0$ as the original message m^j if $\theta \in \mu^{j+}$ (resp., μ^{j-}) and does not send the other message m^{j-} (resp., m^{j+}); and, (ii) each $\theta \notin \mu^j$ who sent m^j with positive probability, now splits that probability equally between the messages m^{j+} and m^{j-}. Accordingly, the corresponding actions can be written as

$$\begin{aligned} a^{j+}(s) &= E[\theta|m^{j+}] = \pi^{j+}(s)z^{j+}_{in}(s) + (1-\pi^{j+}(s))z^j_{out} \\ a^{j-}(s) &= E[\theta|m^{j-}] = \pi^{j-}(s)z^{j-}_{in}(s) + (1-\pi^{j-}(s))z^j_{out} \end{aligned}$$

where the conditional probabilities are

$$\begin{aligned} \pi^{j+}(s) &= \Pr[\theta \in \mu^{j+}(s)|m^{j+}] = \frac{p^j \Pr[\theta \in \mu^{j+}(s)]}{p^j \Pr[\theta \in \mu^{j+}(s)] + \frac{1}{2}P^j} = 1 - \Pr[\theta \notin \mu^j|m^{j+}] \\ \pi^{j-}(s) &= \Pr[\theta \in \mu^{j-}|m^{j-}] = \frac{p^j \Pr[\theta \in \mu^{j-}(s)]}{p^j \Pr[\theta \in \mu^{j-}(s)] + \frac{1}{2}P^j} = 1 - \Pr[\theta \notin \mu^j|m^{j-}] \end{aligned}$$

and $P^j = \Pr[m^j, \theta \notin \mu^j]$.

Since $\mu^{j+}(-s) = \mu^{j-}(s)$ for all s, we have $\pi^{j+}(-s) = \pi^{j-}(s)$. Since in addition $z^{j+}_{in}(-s) = z^{j-}_{in}(s)$, it follows that $a^{j+}(-s) = a^{j-}(s)$ for all s, and symmetrically $a^{j-}(-s) = a^{j+}(s)$. But then the difference $U(a^{j+}(s)) - U(a^{j-}(s))$ is a continuous odd function of s, so that by the Borsuk-Ulam theorem there exists $s^{j*} \in \mathbb{S}^{N-1}$ such that $U(a^{j+}(s^{j*})) - U(a^{j-}(s^{j*})) = 0$, for each $j = 1, ..., 2^k$. Furthermore, by the law of iterated expectations, there exists $\delta^j \in (0,1)$ such that $a^j = \delta^j a^{j+}(s^{j*}) + (1-\delta^j)a^{j-}(s^{j^*})$ for each j. Since the orientation s^{j*} will be fixed for all j for the remainder of the proof we suppress it in what follows. Figure 3(a) depicts the typical situation with respect to the new actions and expectations obtained for the j-th element of the original 2^k-element partition and it will be useful for the reader to consult it for the rest of the proof.[37]

[37] The figure depicts the typical situation where the expectations z^{j+}_{in} and z^{j-}_{in} are not co-linear with the

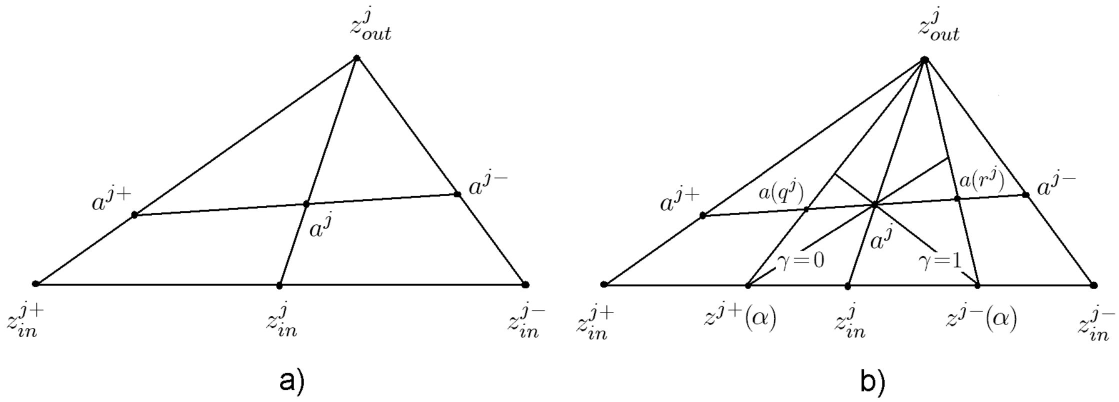

Figure 3: Mixed message construction

If $U(a^{j+}) = U(a^{j-})$ does not vary with j, we have created a 2^{k+1}-message equilibrium. If not, suppose without loss of generality that for $j = 1$, $U(a^{j+}) = U(a^{j-})$ is the lowest. Since for the original 2^k-message equilibrium $U(a^j)$ did not depend on j, exploiting the strict quasi-convexity of U we have for all $j > 1$,

$$U(a^{j+}) = U(a^{j-}) \geq U(a^{1+}) = U(a^{1-}) > U(a^1) = U(a^j). \tag{6}$$

For each j for which the first inequality in (6) holds with equality we do not alter the probabilities with which the messages m^{j+} and m^{j-} are sent. In contrast, for j for which the first inequality in (6) is strict, we adjust the induced actions after messages m^{j+} and m^{j-} respectively by suitably altering the probabilities by which these messages are sent, as follows.

First, since preferences are continuous, from (6) there must exist $q^j, r^j \in (0,1)$ and actions $a(q^j) = q^j a^{j+} + (1 - q^j)a^j$ and $a(r^j) = r^j a^{j+} + (1 - r^j)a^{j-}$ such that

$$U(a(q^j)) = U(a(r^j)) = U(a^{1+}) = U(a^{1-}).$$

Indeed, we must have $1 > q^j > \delta^j > r^j > 0$, i.e., $a(q^j)$ and $a(r^j)$ both lie on the line joining a^{j+} and a^{j-} that passes through a^j, on either side of a^j. We wish to adjust the induced actions to $a(q^j)$ and $a(r^j)$, after messages m^{j+} and m^{j-} respectively, by suitably altering

original expecations z^j_{in}, z^j_{out} and a^j. The arguments go through in the co-linear case, except for non-generic situations where, in addition, $a^{j+} = a^{j-} = a^j$. For $N > 2$, this can be ruled out generally not only generically.

the probabilities by which these messages are sent. To this end, let $\alpha^{j+}p^j$ (resp., $\alpha^{j-}p^j$) be the probability with which any $\theta \in \mu^{j+}$ (resp., μ^{j-}) sends message m^{j+} (resp., m^{j-}), with the remaining probability $(1-\alpha^{j+})p^j$ (resp., $(1-\alpha_{j-})p_j$) on the other message m^{j-} (resp., m^{j+}). Similarly, let $\theta \notin \mu^j$, divide the probability with which they sent the original message m^j into the messages m^{j+} and m^{j-} in the ratio $\frac{\gamma}{1-\gamma}$, $\gamma \in (0,1)$. We wish to find α^{j+}, α^{j-} and γ such that $a(q^j) = E[\theta|m^{j+}]$ and $a(r^j) = E[\theta|m^{j-}]$. As Figure 3(b) depicts, this is always possible, using properties of conditional expectations.

To understand Figure 3(b), notice first that the point $z^{j+}(\alpha)$ is obtained by producing the line joining z^j_{out} with $a(q^j)$ till it meets the line joining z^{j+}_{in} and z^{j-}_{in} that must pass through z^j_{in}; and similarly for the point $z^{j-}(\alpha)$. Indeed, there exists $q, q', q'' \in (0,1)$ such that

$$\begin{aligned} z^j_{in} &= q z^{j+}_{in} + (1-q) z^{j-}_{in} \\ z^{j+}(\alpha) &= q' z^{j+}_{in} + (1-q') z^{j-}_{in} \\ z^{j-}(\alpha) &= q'' z^{j+}_{in} + (1-q'') z^{j-}_{in} \end{aligned}$$

with $q' > q > q''$. Using Bayes' Rule we wish to choose α^{j+}, α^{j-} such that $E[\theta|m^{j+}, \theta \in \mu^j] = z^{j+}(\alpha)$ and $E[\theta|m^{j-}, \theta \in \mu^j] = z^{j-}(\alpha)$. This yields

$$\begin{aligned} \alpha^{j+} &= \frac{q'(q-q'')}{q(q'-q'')} \in (0,1) \\ \alpha^{j-} &= 1 - \frac{(1-q'')(q'-q)}{(1-q)(q'-q'')} \in (0,1) \end{aligned}$$

It remains to choose γ such that

$$\begin{aligned} a(q^j) &= E[\theta|m^{j+}] \equiv \Pr[\theta \in \mu^j | m^{j+}] z^{j+}(\alpha) + \Pr[\theta \notin \mu^j | m^{j+}] z^j_{out} \\ a(r^j) &= E[\theta|m^{j-}] \equiv \Pr[\theta \in \mu^j | m^{j-}] z^{j-}(\alpha) + \Pr[\theta \notin \mu^j | m^{j-}] z^j_{out} \end{aligned}$$

For any γ, the expectations $E[\theta|m^{j+}]$ and $E[\theta|m^{j-}]$ correspond to the endpoints of some line going through a^j, on the inner triangle formed by the points z^j_{out}, $z^{j+}(\alpha)$ and $z^{j-}(\alpha)$. For $\gamma = 0$ this is the line so labeled in Figure 3(b) going through $E[\theta|m^{j+}] = z^{j+}(\alpha)$, while for $\gamma = 1$ it is the line so labeled going through $E[\theta|m^{j-}] = z^{j+}(\alpha)$. As γ varies from 0 to 1, the line will slide around in a clockwise direction continuously, so that for some $\gamma \in (0,1)$ we must have $E[\theta|m^{j+}] = a(q^j)$ and $E[\theta|m^{j-}] = a(r^j)$.

This completes the construction of the communication strategies and induced actions. We now have a 2^{k+1}-message equilibrium with messages m^{j+}, m^{j-}, $j = 1, ..., 2^k$, where the induced actions satisfy (i) $a(m^{j+}) = a^{j+}$ and $a(m^{j-}) = a^{j-}$ if $U(a^{j+}) = U(a^{j-}) = U(a^{1+}) = U(a^{1-})$, (ii) $a(m^{j+}) = a(q^j)$ and $a(m^{j-}) = a(r^j)$ if $U(a^{j+}) = U(a^{j-}) > U(a^{1+}) = U(a^{1-})$. By construction, all induced actions yield the same payoff to the expert. Furthermore, each $\theta \in \mu^{j+}$ (resp., μ^{j-}) sends the corresponding message m^{j+} (resp., m^{j-}) with strictly positive probability $\alpha^{j+}p^j$ (resp., $\alpha^{j-}p^j$). Since, by Theorem 1 such an equilibrium exists for the case $k = 1$, this completes the induction. By the strict quasi-convexity of U, payoffs are strictly increasing in k. ■

Proof of Proposition 1: For any fixed $c \in int(\mathbf{\Theta})$, let s^* be the orientation of an equilibrium hyperplane through c with corresponding actions $a^+(h_{s^*,c})$ and $a^-(h_{s^*,c})$ when priors are degenerate on t^*, i.e., given by Φ^*. Such s^* exists by Theorem 1 and we have $U(a^+(h_{s^*,c}), t^*) = U(a^-(h_{s^*,c}), t^*)$. By condition (S), for all $t \neq t^*$, $U(a^+, t) \neq U(a^-, t)$. W.l.o.g., rename types so that $U(a^+, t) > U(a^-, t)$ for all $t > t^*$ and $U(a^+, t) < U(a^-, t)$ for all $t < t^*$ and consider the actual priors Φ. Pick a hyperplane of arbitrary orientation s through c (the same c as above) and, using usual notation, let the actions or expected values corresponding to each message be

$$\begin{aligned} \alpha^+(h_{s,c}; \Phi) &= \Pr[t^*|m^+]E[\theta|t^*, m^+] + \sum_{t>t^*} \Pr[t|m^+]E[\theta|t] \\ \alpha^-(h_{s,c}; \Phi) &= \Pr[t^*|m^-]E[\theta|t^*, m^-] + \sum_{t<t^*} \Pr[t|m^-]E[\theta|t] \end{aligned}$$

That is, we assume type t^* discloses the partition of $\mathbf{\Theta}$ associated with $h_{s,c}$ truthfully, while all types $t > t^*$ (resp., $t < t^*$) send only message m^+ (resp., m^-). Thus,

$$\Pr[t^*|m^+] = \frac{\Pr[\theta \in m^+|t^*]\phi_{t^*}}{\Pr[\theta \in m^+|t^*]\phi_{t^*} + \sum_{t>t^*} \phi_t}$$

and

$$\Pr[t|m^+] = \frac{\phi_t}{\Pr[\theta \in m^+|t^*]\phi_{t^*} + \sum_{t<t^*} \phi_t}$$

if $t > t^*$ and is 0 otherwise, and similarly for the message m^-.

Let

$$\Delta(s, \Phi, t) = U(\alpha^+(h_{s,c}; \Phi), t) - U(\alpha^-(h_{s,c}; \Phi), t),$$

a continuously differentiable function of s. We know that $\Delta(s^*, \Phi^*, t^*) = 0$. We wish to show via the implicit function theorem that there exists $\varepsilon > 0$, such that for $||\Phi - \Phi^*|| < \varepsilon$, there exists $s(\Phi)$ close to s^* for which $\Delta(s(\Phi); \Phi, t^*) = 0$. This is enough to show the result, since when $s(\phi)$ is close to s^* the corresponding actions $\alpha^+(s(\Phi);\phi)$ and $\alpha^-(s(\phi);\phi)$ are close to a^+ and a^- respectively, so that $\Delta(s^*,\Phi^*,t) < 0$ if $t < t^*$ and $\Delta(s^*,\Phi^*,t) > 0$ if $t > t^*$. Then type t^* has the right incentives to disclose the partition of $\mathbf{\Theta}$ associated with $h_{s(\Phi),c}$ truthfully, while all types $t > t^*$ (resp., $t < t^*$) send only message m^+ (resp., m^-).

We use the fact that the circle is locally like the line. That is, we set $s_1(z) = z \in [-1,1]$ and, when $s_2^* = \sqrt{1-s_1^{*2}}$ set $s_2(z) = \sqrt{1-z^2}$ (and, similarly, when $s_2^* = -\sqrt{1-s_1^{*2}}$ set $s_2(z) = -\sqrt{1-z^2}$). We then consider the function $\Delta(s(z),\Phi,t^*)$ as a function of z in a neighborhood of s^*. To apply the implicit function theorem we have to show that $\frac{\partial\Delta(s(z),\Phi^*,t^*)}{\partial z} \neq 0$. It is easy to see that this derivative consists of terms involving the derivative of $U(.,t^*)$ with respect to the actions (that do not depend on $F(.|t^*)$) and terms involving the expected actions α^+ and α^- (that depend on $F(.|t^*)$ but not on $U(.,t^*)$). Since we can vary $F(.|t^*)$ and $U(.,t^*)$ independently, it follows that $\frac{\partial\Delta(s(z),\Phi^*,t^*)}{\partial z}$ can vanish only in non-generic cases, establishing the result. ■

Proof of Proposition 2: The arguments are identical to that for Theorem 1. The only difference is that now we look for an s that simultaneously sets the $T \leq N-1$ maps $G(s,c,t) = U(a^-(h_{s,c}),t) - U(a^+(h_{s,c}),t) = 0$, one for each $t \in \mathbf{T}$. ■

Proof of Proposition 3: We show that for all F and all k, a communication strategy is a k-message ε–cheap talk equilibrium for Euclidean U with large B if and only if it is a cheap talk equilibrium for the limiting linear preferences $\rho \cdot a$. To do this first pick an arbitrary hyperplane $h_{s,c}$ of orientation $s \in \mathbb{S}^{N-1}$ passing through $c \in int(\mathbf{\Theta})$ and let L be the line joining the corresponding actions $a^+ = a^+(h_{s,c})$ and $a^- = a^-(h_{s,c})$. Pick any $\theta \in \mathbf{\Theta}$ and let $p(\theta)$ be the point where the perpendicular from $\theta + B\rho$ on to L meets L. Then

$$p(\theta) = q(\theta)a^+ + (1-q(\theta))a^- \tag{7}$$

where $q(\theta) \in \mathbb{R}$ is given by

$$q(\theta) = \frac{(\theta - a^-)\cdot(a^+ - a^-) + B\rho\cdot(a^+ - a^-)}{(a^+ - a^-)\cdot(a^+ - a^-)}. \tag{8}$$

Notice that this is well-defined since $a^+ \neq a^-$. Notice next that

$$\begin{aligned} d(a^+,\theta+b) - d(a^-,\theta+b) &= \frac{d^2(a^+,\theta+b) - d^2(a^-,\theta+b)}{d(a^+,\theta+b) + d(a^-,\theta+b)} \\ &= \frac{d^2(a^+,p(\theta)) - d^2(a^-,p(\theta))}{d(a^+,\theta+b) + d(a^-,\theta+b)}. \end{aligned} \tag{9}$$

For the if part consider first a two-message cheap talk equilibrium with induced actions a^+ and a^- when U is given by (1), so that $\rho \cdot (a^+ - a^-) = 0$. Then $q(\theta)$ and so $p(\theta)$ do not depend on B. Furthermore, using (9),

$$\begin{aligned} \left|U(a^+;\theta+b) - U(a^-;\theta+b)\right| &= \left|d(a^+,\theta+b) - d(a^-,\theta+b)\right| \\ &= \left|\frac{d^2(a^+,p(\theta)) - d^2(a^-,p(\theta))}{d(a^+,\theta+b)+d(a^-,\theta+b)}\right| \\ &\leq \max_{\theta\in\boldsymbol{\Theta}}\left|\frac{d^2(a^+,p(\theta)) - d^2(a^-,p(\theta))}{d(a^+,\theta+B\rho)+d(a^-,\theta+B\rho)}\right|. \end{aligned} \tag{10}$$

Let θ_B be the solution to the last maximization problem. As B rises, θ_B stays bounded in the compact set $\boldsymbol{\Theta}$, so that $p(\theta_B)$ stays bounded as well, implying that the numerator stays bounded. However the denominator becomes arbitrarily large. It follows that for any $\varepsilon > 0$, for B large enough, $|U(a^+,\theta+b) - U(a^-,\theta+b)| < \varepsilon$ for all θ. An analogous argument obtains for the k-message equilibria if we consider pairs of equilibrium actions that must all lie on the same line L and use the logic above.

For the only if part, suppose that two actions a^+ and a^- do not constitute a two-message cheap talk equilibrium when U is given by (1). W.l.o.g., suppose that $\rho \cdot a^+ < \rho \cdot a^-$. Consider type $\theta = a^+$ and observe via (9) that

$$\begin{aligned} \lim_{B\to\infty}\left[U(a^+;a^++b) - U(a^-;a^++b)\right] &= \lim_{B\to\infty}\left[\frac{d^2(a^+,p(a^+)) - d^2(a^-,p(a^+))}{d(a^+,a^++B\rho)+d(a^-,a^++B\rho)}\right] \\ &= \frac{\rho\cdot(a^- - a^+)}{\sqrt{\rho\cdot\rho}} \end{aligned} \tag{11}$$

where we have used (7) and (8) in the last line. It follows that when $\varepsilon < \rho\cdot(a^- - a^+)/\sqrt{\rho\cdot\rho}$, and B is large enough, type $\theta = a^+$ would gain by more than ε from lying (i.e., by inducing the decision maker to choose the action a^- instead of a^+), implying in turn that h is not an ε-cheap talk equilibrium for large B when U is given by (2). An identical argument obtains for the k-message equilibria of Theorem 3, $k \geq 2$ and finite, if we consider some pair of actions for which $\rho \cdot a^+ \neq \rho \cdot a^-$. ■

References

[1] Alonso, Ricardo and Niko Matouschek (2008), Optimal delegation, *Review of Economic Studies* 75, 259–293.

[2] Ambrus, Attila and Satoru Takahashi (2008), Multi-sender cheap talk with restricted state space, *Theoretical Economics*, 1–27.

[3] Anderson, Simon P. and Régis Renault (2006), Advertising content, *American Economic Review* 96, 93–113.

[4] Aumann, Robert J. and Sergio Hart (2003), Long cheap talk, *Econometrica* 71, 1619–1660.

[5] Austen-Smith, David (1993), Information acquisition and orthogonal argument, in *Political Economy: Institutions Competition and Representation*, proceedings of the seventh international symposium in economic theory and econometrics, eds. Barnet, W. A., Melvin. H., J. and N. Schofield, Cambridge University Press, 407–436.

[6] Austen-Smith, David and Timothy Feddersen (2006), Deliberation, preference uncertainty, and voting rules, *American Political Science Review* 100, 209–217.

[7] Battaglini, Marco (2002), Multiple referrals and multidimensional cheap talk, *Econometrica* 70, 1379–1401.

[8] Benabou, Roland and Guy Laroque (1992), Using privileged information to manipulate markets: insiders, gurus, and credibility, *Quarterly Journal of Economics* 107, 921–958.

[9] Bergemann, Dirk and Martin Pesendorfer (2007), Information structures in optimal auctions, *Journal of Economic Theory* 137, 580–609.

[10] Board, Simon (2006), On revealing information in auctions: the allocation effect, *Economic Theory*, forthcoming.

[11] Campbell, Colin M. (1998), Coordination in auctions with entry, *Journal of Economic Theory* 82, 425–450.

[12] Caplin, Andrew and Barry Nalebuff (1991), Aggregation and social choice: a mean voter theorem, *Econometrica* 59, 1–23.

[13] Chakraborty, Archishman, Nandini Gupta, and Rick Harbaugh (2006), Best foot forward or best for last in a sequential auction?, *RAND Journal of Economics* 37, 176–194.

[14] Chakraborty, Archishman and Rick Harbaugh (2007), Comparative cheap talk, *Journal of Economic Theory* 132, 70–94.

[15] Coase, Ronald H. (1979), Payola in radio and television broadcasting, *Journal of Law and Economics* 22, 269–328.

[16] Coughlan, Peter J. (2000), In defense of unanimous jury verdicts: mistrials, communication, and strategic voting, *American Political Science Review* 94, 375–393.

[17] Crawford, Vincent P. and Joel Sobel (1982), Strategic information transmission, *Econometrica* 60, 1431–1450.

[18] DellaVigna, Stefano and Ethan Kaplan (2007), The Fox News effect: media bias and voting, *Quarterly Journal of Economics* 122, 1187–1234.

[19] Dessein, Wouter (2002), Authority and communication in organizations, *Review of Economic Studies*, 69, 811–838.

[20] Dimitrakas, Vassilios and Yannis Sarafidis (2005), Advice from an expert with unknown motives, mimeo, INSEAD.

[21] Farrell, Joseph and Robert Gibbons (1989), Cheap talk with two audiences, *American Economic Review* 79, 1214–1223.

[22] Feddersen, Timothy and Wolfgang Pesendorfer (1998), Convicting the innocent: the inferiority of unanimous jury verdicts under strategic voting, *American Political Science Review* 92, 23–35.

[23] Gal-Or, Esther (1986), "Information transmission – Cournot and Bertrand equilibria," *Review of Economic Studies* 53, 85–92.

[24] Ganuza, Juan-José (2004), Ignorance promotes competition: an auction model with endogenous private valuations, *RAND Journal of Economics,* 35, 583–598.

[25] Ganuza, Juan-José and José Penalva (2006), On information and competition in private value auctions, mimeo, Universitat Pompeu Fabra.

[26] Gentzkow, Matthew and Jesse M. Shapiro (2006), Media bias and reputation, *Journal of Political Economy* 114, 280–316.

[27] Gick, Wolfgang (2006), Two experts are better than one: multi-sender cheap talk under simultaneous disclosure, mimeo, Dartmouth University.

[28] Gilligan, Thomas W. and Keith Krehbiel (1989), Asymmetric information and legislative rules with a heterogeneous committee, *American Journal of Political Science* 33, 459–490.

[29] Glazer, Jacob and Ariel Rubinstein (2004), On optimal rules of persuasion, *Econometrica* 72, 1715–1736.

[30] Gordon, Sid (2006), Informative cheap talk equilibria as fixed points, mimeo, University of Montreal.

[31] Grunbaum, B. (1960), Partitions of mass-distributions and of convex bodies by hyperplanes, *Pacific Journal of Mathematics* 10, 1257–1261.

[32] Harris, Milton and Artur Raviv (2005), Allocation of decision-making authority, *Review of Finance* 9, 353–383.

[33] Holmstrom, Bengt (1984), On the theory of delegation, in M. Boyer, and R. Kihlstrom (eds.) *Bayesian Models in Economic Theory*, New York: North-Holland, 115–141.

[34] Inderst, Roman and Marco Ottaviani (2007), (Mis)selling through agents, mimeo, Northwestern University and London School of Economics.

[35] Jain, Shailendra Pratap and Steven S. Posovac (2004), Valenced comparisons, *Journal of Marketing Research* 41, 46–58.

[36] Kartik, Navin (2008), Strategic communication with lying costs, mimeo, University of California San Diego.

[37] Kartik, Navin, Marco Ottaviani, and Francesco Squintani (2007), Credulity, lies, and costly talk, *Journal of Economic Theory* 134, 93–116.

[38] Krishna, Vijay and John Morgan (2001), A model of expertise, *Quarterly Journal of Economics* 116, 747–775.

[39] Krishna, Vijay and John Morgan (2004), The art of conversation: eliciting information from experts through multi-stage communication, *Journal of Economic Theory* 117, 147–179.

[40] Lau, Richard R., Lee Sigelman, Caroline Heldman, and Paul Babbitt (1999), The effects of negative political advertisements: a meta-analytic assessment, *American Political Science Review* 93, 851–875.

[41] Levy, Gilat and Ronny Razin (2007), On the limits of communication in multidimensional cheap talk: a comment, *Econometrica* 75, 885–893.

[42] Li, Ming and Kristof Madarasz (2008), When mandatory disclosure hurts: expert advice and conflicting interests, *Journal of Economic Theory* 139, 47–74.

[43] Lyon, Thomas and John W. Maxwell (2004), Astroturf lobbying and corporate strategy, *Journal of Economics and Management Strategy* 13, 561–597.

[44] Martinelli, Cesar (2002), Convergence results for unanimous voting, *Journal of Economic Theory* 105, 278–297.

[45] Matousek, Jiri (2003), *Using the Borsuk-Ulam Theorem*, Springer.

[46] Milgrom, Paul R. (1981), Good news and bad news: representation theorems and applications, *Bell Journal of Economics* 12, 380–391.

[47] Milgrom, Paul R. and John Roberts (1986), Price and advertising as signals of product quality, *Journal of Political Economy* 94, 796–821.

[48] Milgrom, Paul R. and Robert J. Weber (1982), A theory of auctions and competitive bidding, *Econometrica* 50, 1089–1122.

[49] Miralles, Antonio (2008), Self-enforced collusion through comparative cheap talk in simulaneous auctions with entry, mimeo, Boston University.

[50] Morgan, John and Phillip C. Stocken (2003), An analysis of stock recommendations, *RAND Journal of Economics* 34, 183–203.

[51] Morris, Stephen (2001), Political correctness, *Journal of Political Economy* 109, 231–265.

[52] Moscaroni, Giuseppe (2007), Competence implies credibility, *American Economic Review* 97, 37–63.

[53] Mullainathan, Sendhil and Andrei Shleifer (2006), The market for news, *American Economic Review* 95, 1031–1053.

[54] Myerson, Roger (1981), Optimal auction design, *Mathematics of Operations Research* 6, 58–73.

[55] Okuno-Fujiwara, Masahiro, Andrew Postlewaite, and Kotaro Suzumura (1990), Strategic information revelation, *Review of Economic Studies* 57, 25–47.

[56] Ottaviani, Marco and Peter Norman Sørensen (2006), Reputational cheap talk, *RAND Journal of Economics* 37, 155–175.

[57] Polborn, Mattias K. and David T. Yi (2006), Positive and negative campaigning, *Quarterly Journal of Political Science* 1, 351–371.

[58] Sobel, Joel (1985), A theory of credibility, *Review of Economic Studies* 52, 557–573.

[59] Spence, Michael (1973), Job market signaling, *Quarterly Journal of Economics* 87, 355-374.

[60] Stiglitz, Joseph E. (1975), The theory of "screening," education, and the distribution of income, *American Economic Review* 65, 283-300.

[61] Visser, Bauke and Otto H. Swank (2007), On committees of experts, *Quarterly Journal of Economics* 122, 337–372.

[62] Vives, Xavier (2001), *Oligopoly pricing: old ideas and new tools*, MIT Press.

[63] Ziv, Amir (1993), Information sharing in oligopoly: the truth-telling problem, *Rand Journal of Economics* 24, 455–465.