Language Barriers*

Andreas Blume and Oliver Board
Department of Economics
University of Pittsburgh
Pittsburgh, PA 15260

October 5, 2009

Abstract

Private information about language competence drives a wedge between the indicative meanings of messages, the sets of states indicated by those messages, and their imperative meanings, the actions induced by those messages. When sender and receiver have common interests, optimal use of an imperfectly shared language subverts both the indicative and imperative meanings of utterances: Messages convey both directly payoff relevant information and instrumental information about the sender's language competence. Furthermore the actions induced by messages depend on the receiver's uncertain ability to decode them. With conflict of interest, an imperfectly shared language can substitute for mediated communication.

*We are grateful for comments received during seminar presentations at the Trimester Program on Mechanism Design at the Hausdorff Institute for Research in Mathematics (Universität Bonn) and the University of Pittsburgh.

1 Introduction

Language competence varies. There are instances where individuals from different countries do not understand each other; academics trying to communicate across disciplines have to cope with different terminologies; and, we frequently struggle to find the right words or to make sense of others' utterances. The following quote from a white paper on electric power transmission in the US illustrates the problem:

> *One of the many difficulties with discussing who should pay for transmission expansion is the surprising lack of a common language for conveying the critical underlying concepts. Important words such as "benefits" and "beneficiaries," and phrases such as "economic upgrades" and "participant funding" are too often used in radically different ways by different parties. At best, the meanings intended by some speakers are not transparent, and different meanings are inferred by different listeners. At worst, the same words have opposite meanings to different people.* (Baldick *et al.* [1], page 12.)

We propose a model of an imperfectly-shared language in which listeners and speakers are uncertain about each other's language competence. This reduces the precision of message meanings, even when there is common interest and, importantly, drives a wedge between the sender's and the receiver's perceived meaning of a message. In the leading examples of Lewis' book "Convention" [11], where he introduces simple communication games, a message can equally well be viewed as meaning that a particular state holds, a "signal-that," or as meaning to take a particular action, a "signal-to." Similarly, in the canonical sender-receiver model of Crawford and Sobel [4] each equilibrium message refers to a precise set of states and induces precisely one action. Messages in the CS model may be vague, when they indicate sets of states rather than individual states, but there is never an issue of words being "used in radically different ways by different parties" or "different meanings [being] inferred by different listeners," as in the above quote. Each messages has a definite "indicative meaning," the set of types who send that message, and a definite "imperative meaning," the action that is induced by the message.[1]

Our primary objective is to show that an imperfectly-shared language, when used optimally, subverts both the indicative and imperative meanings of utterances. Messages no longer have definite indicative meanings, because the payoff-relevant information that is conveyed with an utterance varies with the language competence of the speaker: The receiver would generally want to deviate from his equilibrium response to a message if the sender's

[1] The way Lewis puts it, in these instances the meaning of a message can be given both ways. In other instances, a message is properly viewed as indicative signal, when it allows receiver discretion, or as an imperative signal, when it allows sender discretion.

language competence were revealed to him. Similarly, messages no longer have definite imperative meanings because the action that the receiver takes in response to a message depends in part of his language competence: The sender would generally want to deviate from his equilibrium strategy if the receiver's language competence were revealed to her.[2]

A secondary goal of this paper is to show that when there is conflict of interest, an imperfectly-shared language can substitute for mediated communication. Specifically, with quadratic loss functions and a uniform distribution over payoff relevant information, there exist distributions over the language-competence types of either senders or receiver and an equilibrium that achieve the efficiency bound from mediated communication.

We give a game-theoretic account of private information about language competence within the standard model of strategic information transmission. In our model limited language competence is captured by limited availability of messages (on the sender side) and limited ability to discriminate among messages (on the receiver side). Coarse languages, in which senders have access to a limited set of messages, and their optimality properties, have recently been studied by Crémer et al [2] and Jäger et al [9]. Our principal formal innovation is to make these coarse languages private information.

We believe that this adds to our understanding of how the precise tool of game-theoretic equilibrium analysis can be fruitfully employed to think about imprecision of meaning of utterances. Our primary focus will be on common interest games. In these games mixing is not optimal, and therefore no rationale for imprecision. Pooling across states may occur in optimal equilibria, e.g. necessitated by limited access to messages, and is imprecise in the sense that it conceals the exact state. At the same time it is precise in the sense there is an exact boundary that separates the set of states that is indicated by a message from the remaining set of states. This kind of precision is unavoidable in an equilibrium approach to meaning: By definition, sender and receiver know each other's strategies. Nevertheless, we will show that with private information about language competence in equilibrium precise knowledge of the set of states that is conveyed through a message can coexist with imprecise knowledge about the directly payoff-relevant aspects of the states (as opposed to those aspects of the state that are only of instrumental interest to players). Similarly, precise knowledge of the receiver's strategy is compatible with uncertainty about the action induced by any given message.

In this setup the communication channel itself is free of noise and yet there are communicative pure-strategy equilibria in which the associations between payoff-relevant informa-

[2]The question we are asking is analogous to whether equilibria are *ex post* equilibria. Of course in a sender-receiver game generally equilibria will not be *ex post* unless they are fully revealing. In essence we are applying a notion of *partial ex-post equilibrium.* We show that when language competence is an issue equilibria typically do not pass the weaker test of being *ex post* with respect to information that is not directly payoff relevant.

tion and messages and those between messages and actions are stochastic. These equilibria can be interpreted as describing situations in which agents ascribe different meanings to words (sentences etc) in a language, misunderstandings may occur, and yet there is not complete communication failure. The language is imperfect but serviceable.

2 Language Competence of the Sender

In this section we focus on the the language competence of the sender, assuming for now that the language competence of the sender is not an issue.

2.1 The Model

A privately informed sender, S, communicates with a receiver, R by sending one of a finite number (greater than two) of messages $m \in M$. The payoffs, $U^S(a,t)$ and $U^R(a,t)$ of the sender and the receiver depend on the receiver's action, $a \in A = \mathbb{R}^\ell$, and the sender's payoff-relevant information $t \in T$, her *payoff type*; we assume that T is a convex and compact subset of $\mathbb{R}^\ell$ that has a nonempty interior. It is common knowledge that the sender's payoff type is drawn from a distribution F with density f on T. The function U^S is differentiable and strictly concave in a for every $t \in T$. Assume that the receiver has a unique best reply $\hat{\rho}(\mu)$ to any belief μ about t, where μ is a distribution over T, and for any measurable set $\Theta \subset T$, sightly abusing notation, denote by $\hat{\rho}(\Theta)$ his optimal response to his prior belief concentrated on Θ. Assume that for all $t' \neq t$, $\hat{\rho}(t') \neq \hat{\rho}(t)$. Note that for any set $\Theta \subset T$ that has positive probability and any set Θ^0 that has zero probability,

$$\hat{\rho}(\Theta) = \hat{\rho}(\Theta \setminus \Theta^0).$$

For any $\Theta \subset T$ and any two actions $a_1 \in A$ and $a_2 \in A$ define

$$\Theta_{a_1 \succsim a_2} := \{t \in \Theta | U^S(t, a_1) \geq U^S(t, a_2)\},$$

the set of types in Θ who prefer action a_1 to action a_2, and similarly define $\Theta_{a_1 \succ a_2}$ for strict preference, and $\Theta_{a_1 \sim a_2}$ for indifference. Note that for any measurable set $\Theta \subset T$ and for any pair $a_1, a_2 \in A$ with $a_1 \neq a_2$, the continuity of the sender's payoff function implies that the sets $\Theta_{a_1 \succ a_2}, \Theta_{a_1 \succ a_2}$ and $\Theta_{a_1 \sim a_2}$ are measurable. Assume that for any two $a_1, a_2 \in A$ with $a_1 \neq a_2$, $\text{Prob}(T_{a_1 \sim a_2}) = 0$. This implies that $\text{Prob}(\Theta) = \text{Prob}(\Theta_{a_1 \succ a_2} \cup \Theta_{a_2 \succ a_1})$. For any finite set of K actions $\{a_1, \ldots, a_K\}$ with $2 \leq K \leq M$ define $\Theta_{a_1 \succsim a_2, \ldots, a_K} := \cap_{n=2}^K \Theta_{a_1 \succsim a_n}$, the set of sender types who prefer action a_1 over actions $a_2, \ldots, a_K$, and use Ω to denote the collection of all such sets.

Assumption 1 *(A) For any* $\Theta \in \Omega$ *and any pair of actions* $a_1,\ a_2 \subset A$ *such that* $\Theta_{a_1 \succ a_2}$ *and* $\Theta_{a_2 \succ a_1}$ *both have positive probability,* $\hat{\rho}(\Theta_{a_1 \succ a_2}) \neq \hat{\rho}(\Theta)$*. (B) For any belief* μ*, there exists a type* $t(\mu)$ *such that* $\hat{\rho}(\mu) = \hat{\rho}(t(\mu))$*.*

Part (A) of Assumption 1 formalizes the idea that the optimal receiver response is sufficiently sensitive to beliefs. This is the key assumption that ensures that the receiver responds differently to a message, depending on knowing whether or not the sender has alternative attractive messages available. Part (B) requires that any best response to some belief is also the receiver's ideal point for some state of the world. Essentially it says that there are no gaps in the type space.

We will assume that not every message $m \in M$ may be available to the sender. Instead the sender privately learns a set $\lambda \subset M$ of available messages, her *availability type*.[3] One message, $m_0 \in M$ is assumed to be always available and could, for example be, interpreted as silence. Thus the sender's availability type λ is drawn, independently from her payoff type t, from a commonly known distribution π on $\Lambda = \{\lambda \in 2^M | m_0 \in \lambda\}$, the set of all subsets of M that contain the message m_0.

A sender strategy is a mapping $\sigma : T \times \Lambda \to \Delta(M)$ that satisfies the condition $\sigma(t, \lambda) \in \Delta(\lambda)$. A receiver strategy is a mapping $\rho : M \to A$.[4] We study perfect Bayesian Nash equilibria (σ, ρ, β) where β is a belief system that is derived from the sender's strategy σ by Bayes' rule whenever possible, the sender's strategy σ is a best reply to the receiver's strategy ρ, and ρ is a best reply after every message, given the belief system β.

2.2 Examples

The following two examples illustrate how both indicative and imperative meanings of messages may be compromised when there is a private information about language competence. In both examples, the focus is on the language competence of the sender.

In the first example, the receiver is uncertain about the indicative meaning of equilibrium messages because he is uncertain about the sender's language competence. There will be a message for which he is unable to determine whether the sender sent this message because no other message was available or because she preferred to send it in lieu of another available

[3] Our distinction between payoff types and availability types is a convenient terminological device. Of course, one could follow Harsanyi [7] and express the inability of the sender to send a particular message by assigning an arbitrarily large negative payoff to doing so. From this viewpoint, all types would be payoff types. This would not affect our results but would, in our view, obscure the fact that ultimately both parties are interested in communicating information about t. Any information transmission about language availability is merely instrumental. Finally, note that we will leave the analysis of a still more general model in which different messages are available at different privately known costs for later work.

[4] The restriction to pure strategies for the receiver is without loss of generality because of our assumption that the receiver has a unique best reply given any belief.

message. Her equilibrium response to that message will be a compromise that averages over the possibility that the sender pooled over all payoff types and the alternative that the sender used the message to indicate a strict subset of the set of payoff types.

In the second example, not only is the receiver (potentially) uncertain about the sender's language competence but the sender is herself is uncertain about the receiver's belief. The receiver is with some probability informed about the sender's language competence. As the result the sender can longer be certain about which action will be induced by his message. If the receiver is uninformed his action will be determined as the optimal response to a belief that averages over different sender competencies; if instead he is informed, there is no need to average and he will take a different action from when he is uninformed.

Example 1 *Assume that the sender's payoff type is drawn from a uniform distribution on the interval [0,1]. Sender and receiver have common interests and receive identical payoffs* $-(t-a)^2$ *when the sender's payoff type is* t *and the receiver takes action* a*. The message space is* $M = \{m_0, m_1\}$ *and the availability distribution* π *assigns positive probability to two availability types,* $\lambda_0 = \{m_0\}$ *and* $\lambda_1 = \{m_0, m_1\}$*, where* $\pi(\lambda_1) = p$ *and* $\pi(\lambda_0) = 1-p$*. Consider an equilibrium in which the sender adopts a strategy of the following form:*

- *if the sender's availability type is* λ_0*: send message* m_0 *for all* $t \in [0,1]$
- *if the sender's availability type is* λ_1*: send message* m_0 *for* $t \in [0, \theta_1)$ *and message* m_1 *for* $t \in [\theta_1, 1]$

The receiver's best response to this strategy is to choose action a_0 *if he received message* m_0 *and* a_1 *if* m_1*, where* a_0 *and* a_1 *are given by*

$$a_0 = \frac{(1-p)\frac{1}{2} + p\theta_1\frac{\theta_1}{2}}{(1-p) + p\theta_1}, \text{ and}$$
$$a_1 = \frac{\theta_1 + 1}{2}.$$

(Note that these actions are equal to the receiver's expectation of t *conditional on the message received.) We have an equilibrium if the sender of type* θ_1 *is indifferent between* a_0 *and* a_1*, i.e.*

$$\theta_1 = \frac{a_0 + a_1}{2}$$
$$\Rightarrow \quad \theta_1 = \frac{4p + \sqrt{9-8p} - 3}{4p}$$

Figure 1 plots the equilibrium actions a_0 *and* a_1 *chosen by the receiver (dashed red), and the cutoff type* θ_1 *for the sender (solid blue) as functions of* p*, the probability that the second*

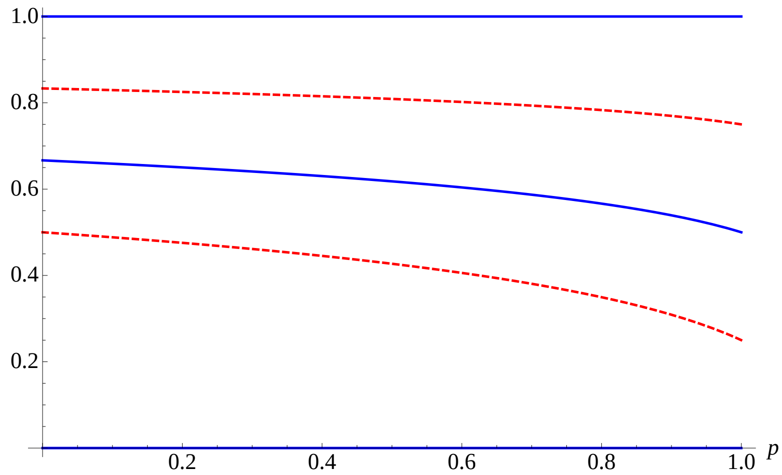

Figure 1: a_0, θ_1, and a_1

message available. Notice that for low values of p, there is considerable distortion *in the choice of a_0 compared with what it would be if the receiver knew that both messages were available ($\theta_1/2$); similarly, for high values of p, there is significant distortion compared with what the receiver would choose if he knew that only one message were available ($1/2$). There is no such distortion in the choice of a_1, because if m_1 is observed, the receiver knows that both message were available.*

Example 2 *Consider the a game with the same preferences and the same distribution over payoff types as in Example 1. Assume that the set of messages is $M = \{m_0, m_1, m_2\}$. Two availability types of the sender have positive probability according to the availability-type distribution π. These are $\lambda_0 = \{m_0, m_1\}$ and $\lambda_1 = \{m_0, m_1, m_2\}$, where $\pi(\lambda_1) = p$ and $\pi(\lambda_0) = 1 - p$. The main departure from the previous example is that the receiver learns the availability type of the sender with probability q, in which case we say that he has an information type τ_ℓ. If instead he does not learn the sender's availability type, his information type will be τ_n. Denote the sender's set of availability types by $\Lambda = \{\lambda_0, \lambda_1\}$ and the receiver's set of information types by $\mathcal{T} = \{\tau_\ell, \tau_n\}$. A (behavior) strategy for the sender is a function $\sigma : T \times \Lambda \to \Delta(M)$ that satisfies supp$(\sigma(t, \lambda_j)) \subset \lambda_j$ for $j = 0, 1$. A strategy for the receiver is a function $\rho : M \times \mathcal{T} \times \Lambda \to \mathbb{R}$ that satisfies $\rho(m, \tau_n, \lambda_0) = \rho(m, \tau_n, \lambda_1)$, which reflects the fact that when the receiver is uninformed about the sender's availability type, he cannot make his action contingent on the this information.*

Define:

$$\begin{aligned}
a_{01} &= \rho(m_0, \tau_\ell, \lambda_1)\\
a_{11} &= \rho(m_1, \tau_\ell, \lambda_1)\\
a_{02} &= \rho(m_0, \tau_\ell, \lambda_2)\\
a_{12} &= \rho(m_1, \tau_\ell, \lambda_2)\\
a_{22} &= \rho(m_2, \tau_\ell, \lambda_2)\\
a_{0n} &= \rho(m_0, \tau_n, \lambda_i)\\
a_{1n} &= \rho(m_1, \tau_n, \lambda_i)\\
a_{2n} &= \rho(m_2, \tau_n, \lambda_i)
\end{aligned}$$

and note that we must have $a_{2n} = a_{22}$ since receiving message m_2 reveals the sender's availability type to the receiver, regardless of the receiver's information type.

Consider an equilibrium in which

1. *λ_1 senders send m_0 for $t \in [0, \frac{1}{2}]$ and m_1 otherwise;*

2. *λ_2 senders send m_0 for $t \in [0, \theta]$, m_1 for $t \in [1-\theta, 1]$, and m_2 otherwise.*

There is such an equilibrium if θ satisfies the condition:

$$-\left(\frac{1}{2} - \theta\right)^2 = -\left(q\left(\theta - \frac{\theta}{2}\right)^2 + (1-q)\left(\theta - \left(\frac{\theta p}{\theta p + \frac{1}{2}(1-p)}\frac{\theta}{2} + \frac{\frac{1}{2}(1-p)}{\theta p + \frac{1}{2}(1-p)}\frac{1}{4}\right)\right)^2\right).$$

Figure 2 shows the solution to this equation as a function of the receiver's probability of being informed of the sender's availability type, q, assuming that the probability of availability type λ_1 is $p = 0.5$. Note that the more likely it is that the receiver knows the sender's availability type the closer the critical type will be to $\frac{1}{3}$, which corresponds to the optimal division of the payoff type space into three equal-length intervals in the even that three messages are available to the sender. Away from this limit, even if the sender has three available message, the sender has to make allowance for the possibility that this is not the case therefore will tend to take a somewhat higher action after m_0, which shifts up the critical type, who needs to be indifferent between sending m_0 and sending m_2.

A sender with payoff type θ and availability type λ_1 is indifferent between sending the message m_2, which reveals her availability type and induces the receiver action $a_{22} = a_{2n} = \frac{1}{2}$, and sending the message m_0, which induces a lottery over the actions $a_{02} = \frac{\theta}{2}$, which an informed receiver will take, and the action $a_{0n} = \left(\frac{\theta p}{\theta p + \frac{1}{2}(1-p)}\frac{\theta}{2} + \frac{\frac{1}{2}(1-p)}{\theta p + \frac{1}{2}(1-p)}\frac{1}{4}\right)$, which an

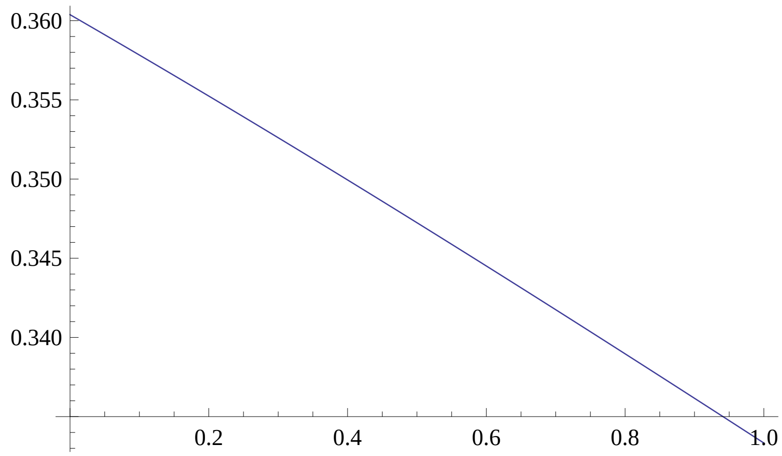

Figure 2: The critical type θ as a function of q for $p = 0.5$

uninformed receiver will take whose posterior probability of the sender having availability type λ_1 *equals* $\frac{\theta p}{\theta p + \frac{1}{2}(1-p)}$. *Note that the same condition that makes type* θ *indifferent between sending messages* m_0 *and* m_2*, makes type* $1 - \theta$ *indifferent between sending messages* m_1 *and* m_2. *There is distortion of indicative meaning because the uninformed receiver who takes action* a_{0n}*, which is an average of* $\frac{\theta}{2}$ *and* $\frac{1}{4}$*, would want to change his action if he could learn the senders availability type either to* $\frac{\theta}{2}$*, if the sender's availability type is* λ_1*, or to* $\frac{1}{4}$*, if the sender's availability type is* λ_0. *There is distortion of imperative meaning because the sender is uncertain about which action message* m_0 *induces. If her availability is* λ_1 *types close to the critical type* θ *would benefit from discovering the receiver's information type.*

2.3 Results

Inspired by Lewis [11], we refer to the *indicative meaning* of a message as the (payoff-relevant) information about the sender that is conveyed by the message. Distortions of indicative meaning arise when the receiver's strategy fails to be optimal given the sender's language competence.

Definition 1 *There is* distortion of indicative meaning *in equilibrium* (σ, ρ, β) *if there exists an availability type* λ *and* $m \in \lambda$ *that is used with positive probability by* λ *such that* $\rho(m)$ *is not optimal for the receiver conditional on the availability type* λ *being revealed.*

Distortions of indicative meaning need not arise if only a few actions are induced in equilibrium and the sender is never constrained by his language ability so that for every action that can be induced she always has a message that induces that actions. This is, trivially, the case in pooling equilibria. Intuitively, however, the more information is transmitted and the more actions are induced in equilibrium the more likely it is that there will be distortions of indicative meaning. Those availability types of the sender who have access to fewer message will sometimes find themselves language constrained and forced to send messages that they would prefer not to send if they had access to a larger set of messages. Thus different availability types will pool on the same message for different sets of payoff types. When receiving such messages the receiver best responds by averaging over these sets of payoff types and will generally take an action that differs from the action he would take if he knew the sender's availability type and therefore did not have to average. The following proposition formalizes this observation.

Proposition 1 *There will be distortion of indicative meaning in any equilibrium* (σ, ρ, β) *for which there is a message* $m^* \in M$ *and a pair of availability types* $\lambda^* \neq \tilde{\lambda}$ *such that* $\lambda^* = \tilde{\lambda} \cup \{m^*\}$, $\pi(\tilde{\lambda}) \neq 0$, $\pi(\lambda^*) \neq 0$, λ^* *uses all of her available messages with positive probability and all those messages induce distinct actions.*

Proof: Since m_0 is always available, the set $\tilde{\lambda}$ is not empty. The fact that λ^* uses all of her messages with positive probability and all of those messages induce distinct actions, and using the fact that for any two $a_1, a_2 \in A$ with $a_1 \neq a_2$ we have $\text{Prob}(T_{a_1 \sim a_2}) = 0$, implies that availability type $\tilde{\lambda}$ also uses all her messages with positive probability. Hence, there must be a set of payoff types, that has positive probability, who use m^* when their availability type is λ^* and use a message $\tilde{m} \neq m^*$ when their availability type is $\tilde{\lambda}$. Use a^* to denote the action that is induced by m^* and $\tilde{a}$ the action that is induced by $\tilde{m}$. Let $\tilde{\Theta}$ denote the set of payoff types who use message $\tilde{m}$ when their availability type is $\tilde{\lambda}$. Since $\tilde{\lambda}$ uses all of its messages with positive probability, the set $\tilde{\Theta}$ has positive probability. Similarly, since λ^* uses all of its message with positive probability the set $\tilde{\Theta}_{a^* \succ \tilde{a}}$ of types who switch to message m^* and the set $\tilde{\Theta}_{\tilde{a} \succ a^*}$ of types who continue to send $\tilde{m}$ both have positive probability. The set $\tilde{\Theta}_{\tilde{a} \succ a^*}$ differs at most by a set that has probability zero from the set of payoff types who send message $\tilde{m}$ when their availability type is λ^*. Hence, if there is no distortion in the equilibrium (σ, ρ, μ), then $\rho(\tilde{m}) = \hat{\rho}(\tilde{\Theta}_{\tilde{a} \succ a^*})$. Also, in the equilibrium (σ, ρ, μ) by assumption $\tilde{\Theta}$ is the set of payoff types who send message $\tilde{m}$ when their availability type is $\tilde{\lambda}$. Therefore, if there is no distortion, then $\rho(\tilde{m}) = \hat{\rho}(\tilde{\Theta})$. By Assumption 1 however,

$$\hat{\rho}(\tilde{\Theta}_{\tilde{a} \succ a^*}) \neq \hat{\rho}(\tilde{\Theta}),$$

which is inconsistent with having no distortion. □

We observe next that Propostion 1 holds in the setup of Crawford and Sobel [4] (CS). Recall that in the CS model the sender's payoff type t is drawn from a differentiable distribution F on $[0, 1]$ with a density f that is everywhere positive on $[0, 1]$. The receiver takes an action $a \in \mathbb{R}$. It is assumed that the functions U^S and U^R are twice continuously differentiable and, using subscripts to denote partial derivatives, the remaining assumptions are that for each realization of t there exist an action a_t^* such that $U_1^S(a_t^*, t) = 0$; for each t there exists an action a_t' such that $U_1^R(a_t', t) = 0$; $U_{11}^S(a, t) < 0 < U_{12}^S(a, t)$ for all a, t; and, $U_{11}^R(a, t) < 0 < U_{12}^R(a, t)$ for all a and t.

Corollary 1 *Proposition 1 holds for the CS model.*

Proof: CS preferences satisfy all the conditions we have imposed on sender and receiver utilities. Specifically, Assumption 1 is satisfied because sender and receiver preferences satisfy the single-crossing condition, $U_{12}^S, U_{12}^R > 0$: Single-crossing for the sender implies that for any positive-probability set $\Theta \subset T$ the set $\Theta_{a_1 \succ a_2}$ is of the form $\Theta \cap T_{a_1 \succ a_2}$ where $T_{a_1 \succ a_2}$ is an interval that is either of the form $(-\infty, t)$ or of the form (t, ∞). Hence, the distribution that is the prior probability concentrated on $\Theta \cap T_{a_1 \succ a_2}$ either stochastically dominates or is stochastically dominated by the distribution that is the prior probability concentrated on Θ. Therefore the single-crossing condition for the receiver implies that $\hat{\rho}(\Theta_{a_1 \succ a_2}) \neq \hat{\rho}(\Theta)$. $\square$

Another environment in which Proposition 1 holds is one where payoffs can be expressed in terms of convex loss functions and the sender's payoff type space T is permitted to be multi-dimensional. Suppose the sender's and receiver's payoffs are given by $U^S(a, t) = \nu_S(||t + b - a||)$ and $U^R(a, t) = \nu_R(||t - a||)$ respectively, where $|| \; ||$ is the Euclidean norm and $-\nu_S$ and $-\nu_R$ are strictly increasing convex functions.[5]

Corollary 2 *Proposition 1 holds when sender and receiver have convex loss functions.*

Proof: With convex loss functions every set Θ in Ω will be convex. For any pair of distinct actions a_1 and a_2, the set $T_{a_1 \succsim a_2}$ is a halfspace and thus if $\Theta_{a_1 \succsim a_2} = \Theta \cap T_{a_1 \succsim a_2}$ and $\Theta_{a_2 \succsim a_1} =$

[5] Jäger et al [9] have examined the optimal equilibria of this environment for the common-interest case, where $b = 0$. There are well-defined indicative meanings ("categories" in their terminology). In any optimal equilibrium categories are shown to be convex giving rise to a Voronoi tessalation of the type space, and all messages are used with positive probability and induce distinct actions. In the present paper the indicative meanings of messages become more fluid: While it is still the case that in equilibrium each availability type partitions the set of payoff types into convex sets, at the same time for a given message these sets will generally differ for different availability types and it is no longer the case that the set of payoff types is partitioned into categories with fixed boundaries. The receiver's posterior distributions after different messages will generally have overlapping supports. For an extreme example, if instead of always permitting silence, we required the availability distribution to have full support on the power set of M, then trivially in any equilibrium the receiver's posterior would have full support on T after every message. We will show below that in our setting with common interests it remains true that all messages (that are available to some availability type) will be used and that therefore by Proposition 1 there will be distortion of indicative meaning in optimal equilibria.

$\Theta \cap T_{a_2 \succsim a_1}$ have positive probability, they are convex and have a nonempty interior. If we denote the interior of a set X by $\text{int}(X)$ then $\hat{\rho}(\Theta_{a_1 \succsim a_2}) \in \text{int}(\Theta_{a_1 \succsim a_2})$ and $\hat{\rho}(\Theta_{a_2 \succsim a_1}) \in \text{int}(\Theta_{a_2 \succsim a_1})$. To see this, let

$$V(a, K) = \int_K \nu_R(||t - a||) f(t) dt$$

for a convex set K and consider a point $\overline{a}$ on the boundary of K. By the supporting hyperplane theorem, there exists a vector $c \neq 0$ with $c \cdot t \geq c \cdot \overline{a}\ \forall t \in K$. Furthermore, $c \cdot t > c \cdot \overline{a}\ \forall t \in \text{int}(K)$. The derivative of $V(\cdot, K)$ at $\overline{a}$ in the direction c satisfies

$$\nabla V(\overline{a}, K) \cdot \frac{c}{||c||} = \int_K \nu_R'(||t - \overline{a}||) \frac{1}{2} ||t - \overline{a}||^{-\frac{1}{2}} (\overline{a} - t) \cdot \frac{c}{||c||} f(t) dt > 0$$

because ν_R is increasing and $(\overline{a} - t) \cdot \frac{c}{||c||} > 0$ for almost all $t \in K$. Use a_{12} to denote $\hat{\rho}(\Theta_{a_1 \succsim a_2})$ and a_{21} to denote $\hat{\rho}(\Theta_{a_2 \succsim a_1})$. Since $a_{12} \notin \Theta_{a_2 \succsim a_1}$, there exists a vector $d \neq 0$ with $d \cdot t \geq d \cdot a_{12}\ \forall t \in \Theta_{a_2 \succsim a_1}$ (and $>$ for all $t \in \text{int}(\Theta_{a_2 \succsim a_1})$). Consider the derivative of $V(\cdot, \Theta)$ at a_{12} in the direction d:

$$\begin{aligned}
\nabla V(a_{12}, \Theta) \cdot \frac{d}{||d||} &= \nabla V(a_{12}, \Theta_{a_1 \succ a_2}) \cdot \frac{d}{||d||} + \nabla V(a_{12}, \Theta_{a_2 \succ a_1}) \cdot \frac{d}{||d||} \\
&= \nabla V(a_{12}, \Theta_{a_2 \succ a_1}) \cdot \frac{d}{||d||} \\
&= \int_{\Theta_{a_2 \succ a_1}} \nu_R'(||t - a_{12}||) \frac{1}{2} ||t - a_{12}||^{-\frac{1}{2}} (a_{12} - t) \cdot \frac{d}{||d||} f(t) dt > 0,
\end{aligned}$$

which shows that $\hat{\rho}(\Theta_{a_1 \succsim a_2}) \neq \hat{\rho}(\Theta)$. □

2.3.1 Common Interest

In this section we consider the case where sender and receiver have identical preferences, $U^S \equiv U^R \equiv U$. We show that an optimal equilibrium exists. Furthermore, in any optimal equilibrium all availability types use all their messages with positive probability and all available messages induce distinct actions. It is interesting that this holds despite the fact that, as we showed above, different availability types using all their messages may lead to distortion of indicative meaning.

First-order intuition for why every availability type uses all of her messages is simple: unused messages can be introduced to refine the information that the sender transmits. A complication arises because other availability types may already use that message and may see their payoffs reduced as the action induced by that message changes. We will show, however, that the magnitude of such losses is of second order in comparison to the gains of the availability type who begins using that message.

We proceed by first establishing existence of an optimal strategy profile. Here we argue in terms of the receiver's strategy ρ which, as we will see, can be viewed as a point in the compact set T^M.[6] We construct a function that assigns to each strategy of the receiver the payoff that results from the sender using a best response to that strategy. Under our assumptions this function is continuous. Hence, we face the problem of maximizing a continuous function over a compact set, which has a solution. Therefore an optimal strategy profile exists and since we have a common interest game, this profile must be part of an equilibrium profile.

For each availability type and any optimal receiver strategy we can partition the set of payoff types into subsets for whom the same message is optimal. If there is an availability type who does not use all of her messages, we can take a pair of messages that induce the same action a, one of which is used by the availability type under consideration and one of which is not. Split the subset of payoff types who induce action a into two positive-probability subsets and have one of these subsets continue to use the message they used before and while the other subset switches to the formerly unused message m. Other availability types may already have been using message m, but note that since we are considering an optimal strategy profile the receiver's response to message m was itself optimal. Therefore an infinitesimal change in the response to m results in a first-order common loss that is zero when the expectation is taken over the types who used message m to begin with. At the same time there is a positive first-order gain for the availability type who starts using message m because she transmits useful information to the receiver. The following results formalize this intuition.

Lemma 1 *With common interests, there exists an optimal strategy profile.*

Proof: Without loss of generality we can confine attention to receiver strategies for which each action is a best response to some belief. Then, by Assumption 1 each receiver strategy prescribes only actions that are optimal for some type. Thus receiver strategies can be thought of as associating with each message m the type for whom the action $\rho(m)$ is optimal, i.e. it suffices to think of receiver strategies as elements of T^M. Suppose that for any given strategy ρ of the receiver, the sender uses a best reply; that best reply exists because given the receiver's strategy each sender type maximizes his payoff over a finite set of alternatives. Then the resulting payoff for type (t, λ) equals

$$\max_{m \in \lambda} \{U(\rho(m), t)\}.$$

[6] This result generalizes the corresponding one of Jäger et al [9] to environments with private information about language competence.

Given this behavior of the sender, we can assign the following expected payoff to the receiver's strategy ρ:

$$Q(\rho) = \sum_{\lambda \in \Lambda} \pi(\lambda) \int_T \max_{m \in \lambda} \{U(\rho(m), t)\} f(t) dt.$$

Since U and the max operator are continuous functions, the integrand is continuous and therefore by the Lebesgue dominated convergence theorem, Q is continuous. Therefore, by Weierstrass's theorem, Q achieves a maximum on the compact set T^M. □

Note that in an optimal profile the receiver's response after unsent messages is entirely arbitrary and therefore it is without loss of generality to assume that it is the same as after one of the sent messages; if it were not arbitrary, then for some specification the sender would have a profitable deviation which would contradict optimality.

Lemma 2 *In an optimal profile, each availability type induces every action a' for which she has a message m' with $\rho(m') = a'$ on a set of payoff types that contains an open set and therefore has positive probability.*

Proof: By Assumption 1 and common interest, $\rho(m)$ is some type's ideal point for all $m \in M$. Hence, a' is the ideal action of some type t'. Strict concavity implies that type t' strictly prefers a' to any of the finitely many other actions she can induce. By continuity this remains true for an open set of types $\mathcal{O}(t')$ containing t' and since f is everywhere positive the set $\mathcal{O}(t')$ has positive probability. □

For CS preferences, the single-crossing condition implies that the set of actions that are optimal for some type is of the form $[\underline{a}, \overline{a}]$ and that with common interest for any belief μ of the sender we have $\hat{\rho}(\mu) \in [\underline{a}, \overline{a}]$. Therefore $\rho(m) \in [\underline{a}, \overline{a}]$ for all $m \in M$, as required by Assumption 1. The assumption also holds for convex loss functions.

Lemma 3 *In an optimal profile all messages of an availability type λ with $\pi(\lambda) > 0$ induce distinct actions.*

Proof: In order to derive a contradiction, suppose not, i.e. there is an availability type λ^* with $\pi(\lambda^*) > 0$ with two or more messages that induce the same action. It is without loss of generality to consider an optimal strategy profile in which the sender of any given availability type uses only one out of any set of available messages that induce identical actions. Thus, suppose that $m^0, m^1 \in \lambda^*$, $\rho(m^1) = \rho(m^0)$, and λ^* uses m^0, but not m^1. The common *ex ante* payoff from the optimal strategy profile (σ, ρ) equals

$$\sum_{m \in M} \sum_{\lambda \in \Lambda} \pi(\lambda) \int_T U(\rho(m), t) \sigma(m|t, \lambda) f(t) dt.$$

Since all messages that type λ^* uses induce distinct actions, Lemma 2 implies that each of those messages is sent by an open set of types that has positive probability. Let Θ_0 be the set of payoff types for which availability type λ^* sends message m^0. Recall that different types have different best replies. Therefore we can find a type t_1 that is an element of an open subset of Θ_0 and that satisfies $\hat{\rho}(t_1) \neq \rho(m^1)$. By continuity, for a sufficiently small open ball Θ_1 containing t_1 and satisfying $\Theta_1 \subset \Theta_0$, we have $\hat{\rho}(\Theta_1) \neq \rho(m^1)$. Now alter (only) type λ^*'s behavior by having her split the set Θ_0 on which she sends m^0 into two subsets so that she sends m^1 on Θ_1 and continues to send m^0 on $\Theta_0 \setminus \Theta_1$. Denote the resulting sender strategy by $\tilde{\sigma}$ to distinguish it from the original strategy σ. Note that as long as we do not also modify the receiver strategy, this change in the sender strategy has no effect on the common *ex ante* payoff. If we use a^1 to denote the action that is induced by message m^1, we can define the contribution to the expected payoff from message m^1 as

$$\begin{aligned} W(m^1, a^1) &:= \sum_{\lambda \in \Lambda} \pi(\lambda) \int_T U(a^1, t)\tilde{\sigma}(m^1|t, \lambda) f(t) dt \\ &= \pi(\lambda^*) \int_T U(a^1, t)\tilde{\sigma}(m^1|t, \lambda^*) f(t) dt \\ &\quad + \sum_{\lambda \in \Lambda \setminus \lambda^*} \pi(\lambda) \int_T U(a^1, t)\sigma(m^1|t, \lambda) f(t) dt \\ &= \pi(\lambda^*) \int_T U(a^1, t)\tilde{\sigma}(m^1|t, \lambda^*) f(t) dt \\ &\quad + \sum_{\lambda \in \Lambda} \pi(\lambda) \int_T U(a^1, t)\sigma(m^1|t, \lambda) f(t) dt \end{aligned}$$

.

Observe that when we change a^1 we affect the contribution to the *ex ante* payoff from message m^1 only. Also, since a^1 was optimal for m^1 given the original sender strategy, we have

$$\nabla_a W(m^1, a^1) = \pi(\lambda^*) \int_T \nabla_a U(a^1, t)\tilde{\sigma}(m^1|t, \lambda^*) f(t) dt.$$

It follows from our choice of Θ_1 that $\nabla_a W(m^1, a^1) \neq 0$. This implies that the original profile (σ, ρ) was not optimal. □

The following result summarizes our findings and connects them to distortion of indicative meaning.

Proposition 2 *In any common interest game, there exists an optimal equilibrium; in any such equilibrium all messages of an availability type that has positive probability induce distinct actions; all such availability types use each of their messages with positive probability; and, if the availability type distribution π has full support on Λ, there will be distortion of* *indicative meaning.*

Proof: The first three parts of the proposition summarize Lemmas 1-3. This sets the stage for invoking Proposition 1, which proves the fourth part of the proposition: If the availability distribution π has full support on Λ, there will be pairs of availability types both of which have positive probability and which differ only by one available message and by Lemmas 1-3 all of these messages are used by both availability types and induce distinct actions. $\square$

Proposition 2 is our key result. It demonstrates the ubiquity of distortion of indicative meaning that results from a combination of private information about language competence and close incentive alignment. With congruence of incentives, optimality requires that a large variety of messages will be used; private information about language competence then implies that the receiver cannot always be sure whether a message was sent out of necessity, because more preferable message were not available, or out of a desire to communicate payoff-relevant information.

It should be clear that while the common-interest case is emblematic for what can go wrong with private information about language competence, the insight that there will be distortion of indicative meaning generally also holds when there is conflict of interest, as long as there is not so much conflict as to rule out all communication in equilibrium. There are, however, other more subtle interactions between conflicts of interest and private information about language competence. These we turn to next.

2.3.2 Conflict of Interest

It is well-known that when there is conflict of interest, access to a noisy channel or, more generally, a nonstrategic mediator can improve communication outcomes in sender-receiver games. In this section we show that private information about message availability can substitute for communication through a nonstrategic mediator. Specifically, in the leading example of the CS model, with a uniform payoff-type distribution and quadratic payoff functions, the efficiency gains from mediated communication can be fully replicated through direct communication when there is private information about message availability.

Myerson [14] gives an example in which there is no communicative equilibrium when the communication technology is perfect, but there is one when agents have to rely on sending a carrier pigeon that gets lost with positive probability. Blume, Board and Kawamura [3] (henceforth BBK) consider communication through a noisy channel that lets the sender's message pass through with probability ϵ and otherwise transmits a random draw from a distribution G on the interval $[0,1]$. They show that with quadratic preferences, i.e.

$$\begin{aligned} U^S(a,t,b) &= -(t+b-a)^2, \\ U^R(a,t) &= -(t-a)^2, \end{aligned}$$

and a uniform type distribution on the interval $[0,1]$ (the "uniform quadratic model") for almost all values of the sender's bias $b \in (0, \frac{1}{2})$ there exists a value of the error probability ϵ and an equilibrium with higher *ex ante* payoffs than from the most efficient equilibrium in the model without noise. Goltsman, Hörner, Pavlov and Squintani (GHPS) [6], also in the uniform quadratic model, investigate the limits from mediated communication; that is, they permit agents to send messages to a correlation device and to receive instructions from the device. This amounts to finding the payoffs from optimal *communication equilibria*, as defined by Forges [5] and Myerson [13]. Using the revelation principle (Myerson [12]) one can characterize the set of communication equilibria in the CS model as corresponding to a family of conditional distributions on $\mathbb{R}$, $\{p(\cdot|t)\}_{t\in T}$, that satisfies:

$$\begin{aligned} &t = \arg\max_{t' \in T} \left[-\int_{\mathbb{R}} (t+b-a)^2 dp(a|t') \right], \ \forall t \in T \\ &a = \mathrm{E}_t[t|a] \ \forall a \in A. \end{aligned}$$

Goltsman, Hörner, Pavlov and Squintani (GHPS) [6] use this characterization to show that the receiver's *ex ante* payoff in any communication equilibrium of the CS model is bounded above by $-\frac{1}{3}b(1-b)$. Since the *ex ante* payoffs of the receiver, V_R and the sender, V_S, are related through $V_R = V_S + b^2$, this is also the efficiency bound for communication equilibria in the CS model.

BBK provide a mechanism that attains this efficiency bound.[7] For any b there exists a noise level $\epsilon(b)$ and an equilibrium of the corresponding $\epsilon(b)$-noise game, $\Gamma(\epsilon(b))$, that achieves the GHPS bound. We will show that this bound can also be attained with private information about message availability. In that case, all communication between the players is direct and misunderstandings arise exclusively because of receiver uncertainty about the sender's repertoire of messages: When receiving a message, the receiver does not know to what degree the sender was forced to use that message rather than some other message that she would have preferred had it been in his repertoire. Our proof strategy is to show that for any so-called "front-loading equilibrium" of BBK that achieves the efficiency bound there exists an outcome-equivalent equilibrium in the model with private information about message availability.

As background it is useful briefly to recall the key elements of the construction of the

[7] Ivanov [8] has recently demonstrated how to attain this bound through a strategic mediator.

front-loading equilibria in BBK. In such an equilibrium the type set, $[0, 1]$, is partitioned into a finite number K of intervals Θ_k (with left endpoint θ_{k-1} and right endpoint θ_k) that are indexed from left to right; for any partition element Θ_k with $k > 1$ there is a single message m_k that is sent by types in that partition element; and, types in the leftmost partition element, Θ_1, uniformly randomize over all the remaining messages. As a result, when the receiver observes one of the messages m_k he believes with probability one that there was no transmission error, that the sender's type belongs to the interval Θ_k and takes action

$$a_k = \frac{\theta_{k-1} + \theta_k}{2}.$$

When the receiver observes any of other messages, his posterior probability of an error having occurred is

$$\frac{\epsilon}{\epsilon + \theta_1(1-\epsilon)}$$

and he takes action

$$a_1 = \frac{\theta_1(1-\epsilon)\frac{\theta_1}{2} + \epsilon\frac{1}{2}}{\epsilon + \theta_1(1-\epsilon)},$$

which is the average of the actions he would have taken with and without error weighted by the posterior probabilities of error and no error respectively.

Proposition 3 below is proven by translating this BBK front-loading construction into the present environment through substituting private information for transmission errors. For example, in the BBK equilibrium, when the receiver observes a message that is voluntarily sent by the lowest interval of payoff types, he must average over the two possibilities that the message was sent in error and that it was sent intentionally. In the present environment, analogously, we have the receiver be uncertain between the possibility that a type from the lowest interval deliberately sent the message that is always available and the possibility that another payoff type sent the message because no other message was available to her.

Proposition 3 *With a uniform type distribution, quadratic preferences and sender bias* $b > 0$*, there exists a message space* M*, an availability distribution* π *on* $\Lambda = \{\lambda \in 2^M | m_0 \in \lambda\}$ *and an equilibrium in the corresponding game that attains the efficiency bound for communication equilibria.*

Proof: Suppose that the optimal BBK-front-loading equilibrium $\mathcal{E}(b)$ has K steps, $\Theta_1, \ldots,$ Θ_K. Pick any message space that satisfies $\#(M) \geq K$. Let there be an availability type $\tilde{\lambda} \subset M$ with $\#(\tilde{\lambda}) \geq K$ and choose an availability distribution π that satisfies the conditions $\pi(\tilde{\lambda}) = 1 - \epsilon(b)$ and $\pi(\lambda) > 0 \Rightarrow \lambda \cap \tilde{\lambda} = \{m_0\}\ \forall \lambda \neq \tilde{\lambda}$. Then we can induce the outcome of the the optimal BBK-front-loading equilibrium $\mathcal{E}(b)$ in our environment by prescribing the following sender strategy. Whenever the realized availability type is $\tilde{\lambda}$, payoff types in the interval Θ_1 pool on the message m_0 (which is always available) and for each interval Θ_k, $k = 2, \ldots, K$, there is a message in $\tilde{\lambda}$ that is sent by payoff types in that interval, and only by those payoff types. All payoff types send the message m_0 whenever their availability type λ is not equal to $\tilde{\lambda}$. The receiver chooses his best response given this sender strategy for any of the messages that are sent with positive probability. Following any of the messages that are sent with probability zero by the sender the receiver's posterior is assumed to be the same as following m_0, and he takes the corresponding optimal action. □

2.4 A Universal Availability Structure

To establish our last result we chose Λ and π as a function of the sender's bias b. Tying Λ and π to the sender's bias is not necessary, if we allow infinite message spaces. Specifically, it is possible to find Λ and π that are universal in the sense that for any $b > 0$ there is an equilibrium that achieves the efficiency bound for communication equilibria.

For our next result, we will assume M to be infinite. We consider an *availability structure*, which is a 4-tuple $(M; (\Lambda, \mathcal{F}, \pi))$ that consists of a set of potential messages, M, a set of availability types $\Lambda \subset 2^M$, a sigma-algebra $\mathcal{F}$ of subsets of Λ, and a probability measure π on $(\Lambda, \mathcal{F})$. For our next result we are interested in a class of availability structures where availability types λ can be ordered in such a way that all messages available to a given type are also available to all lower types and the probability distribution μ can be described in terms of this ordering; in this availability structure there is a natural sense of the degree to which the sender knows the language.

Definition 2 *An availability structure* $(M; (\Lambda, \mathcal{F}, \mu))$ *is* **nested** *if for each* $\alpha \in [0, 1]$ *there exists an infinite set* $M_\alpha \subset M$ *such that* $M_\alpha \cap M_{\alpha'} = \emptyset$, $\cup_{\alpha \in [0,1]} M_\alpha = M$, $\Lambda = \{\lambda_\alpha \subset M | \lambda_\alpha = \cup_{\alpha' \leq \alpha} M_{\alpha'}\}$, $\mathcal{F} = \{F \subset 2^\Lambda | F = \cup_{\alpha \in B} \lambda_\alpha \text{and} B \in \mathcal{B}\}$ *(where* $\mathcal{B}$ *denotes the set of Borel subsets of the interval* $[0, 1]$*) and there exists an atomless distribution* G *on* $[0, 1]$ *with density* g *such that* $g(\alpha) > 0$ *for all* $\alpha \in [0, 1]$ *and*

$$\mu(\{\lambda_{\alpha'} | \alpha' \leq \alpha\}) = G(\alpha) \ \forall \alpha \in [0, 1].$$

Example 3 *Let* M *be the unit square,* $M_\alpha = \{(x, y) \in M | x = \alpha\}$ *and* G *the uniform distribution. Then, given a draw* α *from* G*, the set of available messages is the rectangle* $[0, \alpha] \times [0, 1]$*. The probability that the messages in* $M_{\alpha*}$ *are available equals the probability that that* $\alpha \geq \alpha^*$*, i.e.* $1 - \alpha^*$*.*

Proposition 4 *With quadratic preferences, uniform type distribution and a nested availability structure there exists an equilibrium that achieves the efficiency bound for communication* *equilibria.*

Proof: For any $\epsilon \in [0,1]$, define $\alpha(\epsilon)$ as the (unique) solution of the equation $G(\alpha(\epsilon)) = \epsilon$. Thus, the probability that the messages in $M_{\alpha(\epsilon)}$ are not available is ϵ. Define $\epsilon(b)$ as the noise level for the BBK front-loading equilibrium that attains the GHPS efficiency bound when the sender's bias is b. Suppose that the optimal BBK-front-loading equilibrium has K steps, $\Theta_1, \ldots, \Theta_K$. Then we can replicate this outcome in our environment with a nested availability structure by prescribing the following sender strategy. Whenever the messages in $M_{\alpha(\epsilon(b))}$ are available, types in interval Θ_1 pool on one of the messages $m_0 \in M_0$ (which are always available) and for each interval Θ_k, $k = 2, \ldots, K$, there is a message in $M_{\alpha(\epsilon(b))}$ that is sent by types in that interval. All types send a message m_0 when the messages in $M_{\alpha(\epsilon(b))}$ are not available. The receiver chooses his best response given this sender strategy for any of the messages that are sent with positive probability. Following any of the messages that are sent with probability zero by the sender the receiver's posterior is assumed to be the same as following m_0, and he takes the corresponding optimal action. □

3 Language Competence of the Receiver

The recipient of a message is as likely limited by his language competence as the sender is by hers. In this section we propose a simple model in which the receiver's language competence is private information. We show that in general this gives rise to distortion of the imperative meanings of messages. When the receiver's language competence is his private information, then even if he uses a pure strategy and there is no randomness in the transmission channel, the sender can no longer be sure how her message will be interpreted; messages typically induce non-degenerate distributions over receiver actions; and, the sender's strategy is generally not optimal given the receiver's language competence.

For simplicity, in this section we focus exclusively on the receiver's language competence and assume that the sender's language competence is not an issue. We model the receiver's language competence as a partition $\mathcal{P}$ of the message set M, with the interpretation that the receiver cannot distinguish messages that belong to the same partition element $P \in \mathcal{P}$. Formally, we require the receiver's strategy to be measurable with respect to $\mathcal{P}$. The receiver's partition type $\mathcal{P}$ is private information and is drawn from a common-knowledge distribution π_R on the set $\mathbf{P}$ of partitions of M. We restrict attention to CS preferences. In this environment a sender strategy is a mapping $\sigma : T \to \Delta(M)$ and it is convenient to represent a receiver strategy as a mapping $\rho : 2^M \to A$. With CS preferences this is without loss of generality because the receiver has a unique best reply to any belief and therefore his best response to observing a partition element P, which we denote by $\rho(P)$, is the same regardless of the partition (type) to which the element P belongs.

For example, suppose sender and receiver have identical payoffs $-(t-a)^2$ from action a when the sender's payoff type is t, the sender's payoff types are uniformly distributed on the interval $[0,1]$, there are three messages m_1, m_2, and m_3 and the receiver has two possible partition types, the type $\{\{m_1\},\{m_2,m_3\}\}$ with probability p and the type $\{\{m_1\},\{m_2\},\{m_3\}\}$ with probability $1-p$. Then there is a three-step equilibrium in which the lowest interval $[0,\theta_1]$ uses message m_1 and the critical type θ_1 increases monotonically from $\frac{1}{3}$ to $\frac{1}{2}$ as p increases from 0 to 1.

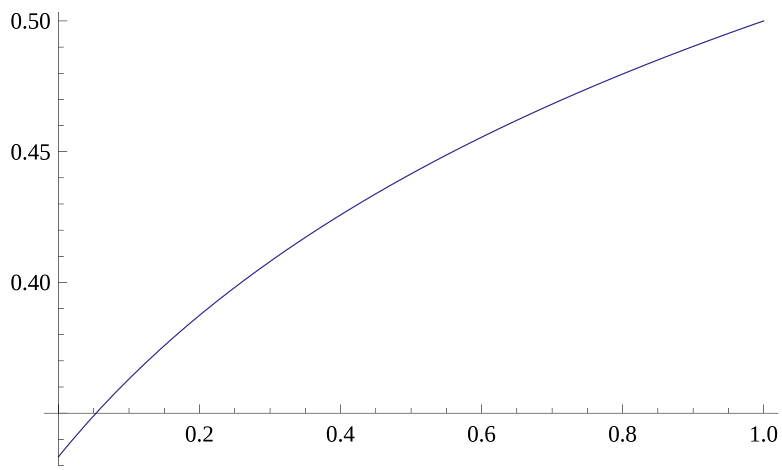

Figure 3: The critical type $\theta_1 = \frac{1}{3}\left(-1 + \frac{4p}{-1+p} - \frac{2\sqrt{1+3p}}{-1+p}\right)$

In this equilibrium there is distortion of imperative meaning: The sender would want to change her strategy conditional upon learning the receiver's availability type. For example, as p, the probability of the receiver having a limited ability to discriminate among messages m_2 and m_3 converges to one, the action a_2 that the receiver takes if he receives and identifies m_2, converges to $\frac{5}{8}$, whereas a_1, the action he takes in response to m_1, converges to $\frac{1}{4}$. Thus, in the limit the type who would be indifferent between sending messages m_1 and m_2 if he knew the receiver's partition type to be $\{\{m_1\},\{m_2\},\{m_3\}\}$ is $\frac{7}{16}$, while in equilibrium the critical type is $\frac{1}{2}$. Types in the interval $\left(\frac{7}{16},\frac{1}{2}\right)$ would want to switch from their equilibrium message m_1 to sending m_2 if they learned that the receiver can distinguish all messages.

There is also another equilibrium in which m_1 is used on the middle interval $(\frac{1}{3},\frac{2}{3})$. In this equilibrium there is no distortion of imperative meaning. Note, however, that in this equilibrium useful information is transmitted only if the receiver can distinguish all three messages.

For our next result we first formally define distortion of imperative meaning. Then we introduce the notion of a *varied receiver response* that lets us distinguish between the two equilibria in the above example that will allow us to give a sufficient condition for distortion of imperative meaning.

Definition 3 *There is* **distortion of imperative meaning** *in equilibrium* (σ, ρ, β) *if there exists a set of payoff types* $\Theta \subset T$ *that has positive probability, a message* $m \in M$ *with* $\sigma(m|t) > 0$ *for all* $t \in \Theta$ *and a partition type* $\mathcal{P}$ *of the receiver that has positive probability such that message* m *fails to be optimal for payoff types in* Θ *conditional on the receiver's partition type* $\mathcal{P}$.

For the case where the sender's language-competence is privately known we showed that it is sufficient for distortion of indicative meaning to occur that there is variety in the use of messages and in the support of the availability type distribution, i.e. when there are availability types that differ in just one message, who use all their messages and all of their messages induce distinct actions. In Definition 4 we introduce an analogous condition that requires the existence of multiple receiver types each of which responds differently to each of its partition elements and that suffices for distortion of imperative meaning when the receiver's language competence is the issue.

Definition 4 *There is a* **varied receiver response** *in equilibrium* $\mathcal{E} = (\sigma, \rho, \beta)$ *if there is a pair of partition types* $\mathcal{P}^* \neq \tilde{\mathcal{P}}$ *of the receiver with a common element* P_0 *such that* $\pi_R(\tilde{\mathcal{P}}) \neq 0$, $\pi_R(\mathcal{P}^*) \neq 0$ *and for every* $P \in \mathcal{P}^* \cup \tilde{\mathcal{P}}$ *the set* $\{t \in T | U^S(\rho(P), t) > U^S(\rho(P'), t), \forall P' \neq P, P' \in \mathcal{P}^* \cup \tilde{\mathcal{P}}\}$ *has positive probability.*

With a varied receiver response it becomes important for the sender to know exactly what the partition type of the sender is. The reason is that it guarantees that there will be at least one pair of receiver types for which a positive probability set of sender types would want to induce the action associated with a common partition element for one receiver type and another action for the other receiver type.

Proposition 5 *There will be distortion of imperative meaning in any equilibrium* $\mathcal{E} = (\sigma, \rho, \beta)$ *with a varied receiver response.*

Proof: Call two elements P_i and P_j of the set $\mathcal{P}^* \cup \tilde{\mathcal{P}}$ adjacent for equilibrium $\mathcal{E}$ if $\rho(P_i) < \rho(P_j)$ and there does not exist $P_k \in \mathcal{P}^* \cup \tilde{\mathcal{P}}$ with $\rho(P_k) \in (\rho(P_i), \rho(P_j))$. Since $\mathcal{P}^*$ and $\tilde{\mathcal{P}}$ have a common element and because $\mathcal{P}^* \neq \tilde{\mathcal{P}}$, there is (at least) one common element, P_C, that is adjacent to a non-common element, P_{NC}. With CS preferences, the sender's single-crossing condition implies that there is a unique type who is indifferent between the actions $\rho(P_C)$ and $\rho(P_{NC})$. Without loss of generality, let $\rho(P_C) < \rho(P_{NC})$ and $\rho(P_{NC}) \in \tilde{\mathcal{P}}$.

Define $P_+ := \arg\min\{\rho(P)|P \in \mathcal{P}^*$ and $\rho(P) \geq \rho(P_C)\}$. Suppose that $P_+ = P_C$. Since P_C is common to both partitions, we have $P_C \cap P_{NC} = \emptyset$. From the sender's single-crossing condition, it follows that those type who would want to induce $\rho(P_{NC})$ when learning $\tilde{\mathcal{P}}$, would want to induce $\rho(P_C)$ when learning $\mathcal{P}^*$. Since $P_C \cap P_{NC} = \emptyset$, they would want to send different message in both cases. Thus in one of the cases the message they would want to send differs from their equilibrium message, which establishes our claim. Now consider the case where $P_+ \neq P_C$. Since P_C and P_{NC} are adjacent, it must be the case that $\rho(P_+) > \rho(P_{NC})$. Since $\rho(P_C) < \rho(P_{NC}) < \rho(P_+)$, the sender's single crossing condition implies that there is a positive probability set of types (near the type who is indifferent between $\rho(P_C)$ and $\rho(P_+)$,) who would want to induce $\rho(P_{NC})$ when learning $\tilde{\mathcal{P}}$ and would want to induce $\rho(P_C)$ when learning $\mathcal{P}^*$. Thus, as before in one of these two cases the message these types would want to send differs from their equilibrium message, which establishes our claim. □

At this point one might be tempted to proceed as in the case where the language competence of the sender is the issue and to try to show that with common interests all messages will be used and that this in turn leads to having the varied-receiver-response condition satisfied in optimal equilibria. The following example, however, demonstrates that there is an interesting asymmetry in the effects of making the sender's language competence private information versus doing the same for the receiver. It shows that in the latter case optimality sometimes requires that there are messages that will never be used.

Example 4 *Suppose the sender's type is drawn from a uniform distribution on* $[0, 1]$ *and both players receive identical payoffs* $-(t-a)^2$ *when the receiver takes action* a *in state* t. *Let* $M = \{m_1, m_2, m_3, m_4\}$. *For any* $\epsilon \in [0, 1)$, *define a game* Γ^ϵ *by the property that each of the receiver types* $\{\{m_1, m_4\}, \{m_2\}, \{m_3\}\}$, $\{\{m_1\}, \{m_2, m_4\}, \{m_3\}\}$ *and* $\{\{m_1\}, \{m_2\}, \{m_3, m_4\}\}$ *has probability* $\frac{1-\epsilon}{3}$ *and the remaining receiver types are equally likely. Note that if* $\epsilon \in (0, 1)$, *the partition-type distribution* π_R *has full support.*

If $\epsilon = 0$, *then in any optimal equilibrium, the type space is partitioned into three equal-length intervals and the actions that are induced in equilibrium are* $\frac{1}{6}$, $\frac{1}{2}$ *and* $\frac{5}{6}$. *To see this, observe first that this holds if for the moment we make the receiver type common knowledge. This provides an upper bound. Then note that the same outcome that is optimal when the receiver type is common knowledge can be realized when the receiver type is private information. Denote the corresponding* ex ante *payoff by* $v^0_{\max}$.

With positive small ϵ, *the messages* m_1, m_2 *and* m_3 *must approximately induce the same set of actions in an optimal equilibrium as they do in an optimal equilibrium for* $\epsilon = 0$. *Otherwise, the* ex ante *payoff from optimal equilibria,* $v^\epsilon_{\max}$, *would remain bounded away from* v^0_{max}, *and we know that that cannot be the case because the strategy profile that results in* v^0_{max} *when* $\epsilon = 0$ *yields approximately* v^0_{max} *when* $\epsilon > 0$ *and since we have a common-interest game the optimal equilibrium strategy must do even better.*

For any ϵ, let $\mathcal{E}(\epsilon)$ be an optimal equilibrium for the game Γ^ϵ. We will argue that for sufficiently small $\epsilon > 0$ no type $t \in [0,1]$ of the sender sends message m_4 in the equilibrium $\mathcal{E}(\epsilon)$. For any $\delta > 0$, there exists $\epsilon(\delta) > 0$ such that for all $\epsilon \in (0, \epsilon(\delta))$ type t's payoff from sending message m_4 is bounded from above by

$$\overline{v}^\epsilon(t) = \left(\frac{1-\epsilon}{3}\right)\left(-\left(t-\frac{1}{6}\right)^2 - \left(t-\frac{1}{2}\right)^2 - \left(t-\frac{5}{6}\right)^2\right) + \epsilon \cdot 0 + \delta$$

while at the same time the payoff to t from sending the optimal message from the set $\{m_1, m_2, m_3\}$ is bounded from below by

$$\underline{v}^\epsilon(t) = (1-\epsilon)\left(-\left(\min\left\{\left(t-\frac{1}{6}\right), \left(t-\frac{1}{2}\right), \left(t-\frac{5}{6}\right)\right\}\right)^2\right) - \epsilon \cdot 1 - \delta.$$

For sufficiently small ϵ and δ, we have $\underline{v}^\epsilon(t) > \overline{v}^\epsilon(t)$ for all $t \in [0,1]$, which shows that there is no type of the sender who would be willing to send message m_4 in any optimal equilibrium of the game Γ^ϵ for sufficiently small $\epsilon \in (0,1)$.

The example shows that unlike in the case where only sender competence is the issue, when there is uncertainty about receiver competence, there may be instances when the sender may not want to use all messages in an optimal equilibrium. This will be the case when there are messages for which the probability is high that the receiver does not understand them. Therefore only a few of the receiver's partition types may be relevant. This undermines the varied-response condition from the previous proposition. On the other hand, in an optimal equilibrium of a common interest game, the sender will want to communicate some information. Thus, an optimal equilibrium will not be a pooling equilibrium and for the communicated information to have an impact, there will be receiver messages that induce distinct actions.

For the following result we adopt a slightly different perspective. Denote by $\mathcal{P}^f$ the finest partition of M, i.e. the type of the receiver who understands all messages. We will show that in any optimal equilibrium of a game that is near an optimal equilibrium of the game in which $\mathcal{P}^f$ has probability one but where π_R has full support there is distortion of imperative meaning.

Proposition 6 *With common interests, an optimal optimal equilibrium exists. For any class of games that differ only in the distributions π_R, if there are finitely many optimal equilibria in the game with $\pi_R(\mathcal{P}^f) = 1$ (e.g. if CS's condition M holds), then there exists an $\epsilon_0 > 0$ such that for all $\epsilon \in (0, \epsilon_0)$ and for every π_R that has full support and satisfies $\pi_R(\mathcal{P}^f) = 1 - \epsilon$, there will be distortion of imperative meaning in any optimal equilibrium.*

Proof: We begin by proving existence. Without loss of generality we can confine attention to receiver strategies for which each action is a best response to some belief. Then, by Assumption 1 each receiver strategy prescribes only actions that are optimal for some type of the sender. Thus receiver strategies can be thought of as associating with each receiver message P the type for whom the action $\rho(P)$ is optimal, i.e. it suffices to think of receiver strategies as elements of T^{2^M}, the set of functions from the powerset of M into the sender's type space. Suppose that for any given strategy ρ of the receiver, the sender uses a best reply; that best reply exists because given the receiver's strategy each sender type maximizes his payoff over a finite set of alternatives, the set of distributions over actions that are induced by each message. Then the resulting payoff for sender type t equals

$$\max_{m \in M} \sum_{\mathcal{P} \in \mathbf{P}} \sum_{P \in \mathcal{P}} U\left(\rho(P), t\right) 1_{\{m \in P\}}.$$

Given this behavior of the sender, we can assign the following expected payoff to the receiver's strategy ρ:

$$Q(\rho) = \int_T \max_{m \in M} \left\{ \sum_{\mathcal{P} \in \mathbf{P}} \sum_{P \in \mathcal{P}} U\left(\rho(P), t\right) 1_{\{m \in P\}} \right\} f(t) dt.$$

Since U and the max operator are continuous functions, the integrand is continuous and therefore by the Lebesgue dominated convergence theorem, Q is continuous. Therefore, by Weierstrass's theorem, Q achieves a maximum on the compact set T^{2^M}.

It remains to show that there is distortion of imperative meaning for sufficiently small positive ϵ. If the receiver's language competence is not an issue, which corresponds to $\epsilon = 0$, then any optimal equilibrium partitions T into M nonempty intervals I_m, $m \in M$, with types belonging to the same interval sending the same message and the receiver's optimal actions following any two messages $m \neq m'$ satisfy $a_m \neq a_{m'}$. For sufficiently small positive ϵ any optimal equilibrium $\mathcal{E}^\epsilon$ of a game in which π_R has full support must approximately induce the same set of actions in the event that messages are understood as in one of the optimal equilibria $\mathcal{E}^0$of the game where message are always understood. Without loss of generality, we can name the messages in ascending order of the actions they induce in $\mathcal{E}^0$. Now consider two receiver types, $\mathcal{P}^f$ and $\mathcal{P}^p$ who only differ in that the latter type cannot distinguish messages m_1 and m_2. With ϵ sufficiently small, the sets of type who send messages m_1 and m_2 respectively are approximately the same in $\mathcal{E}^0$ and $\mathcal{E}^\epsilon$ and the receiver responds in $\mathcal{E}^\epsilon$ to $\{m_1\}$, $\{m_1, m_2\}$ and $\{m_2\}$ with actions $a_1 < a_{12} < a_2$. Hence, the varied-response condition is satisfied. The result then follows from Proposition 5. □

We will conclude by showing that as in the case where the language competence of the sender is private information, private information about the language competence of the receiver can substitute for mediated communication. A particularly simple way of utilizing private information about the receiver's language competence replicates an equilibrium outcome from Krishna and Morgan's (KM) [10] study of multi-stage communication in the CS environment. This is the subject of the following observation.

Observation. *With a uniform type distribution, quadratic preferences and sender bias $b \in \left(0, \frac{1}{8}\right)$, there exists a finite message space M, an availability distribution π_R of the receiver on the set $\mathbf{P}$ of partitions of M that assigns positive probability to exactly two elements of $\mathbf{P}$ and an equilibrium in the corresponding game that attains the efficiency bound for communication equilibria.*

Proof: To verify the observation, first recall that KM showed that for $b \in \left(0, \frac{1}{8}\right)$ there is a class of equilibria that achieve an *ex ante* payoff of $-\frac{1}{3}b(1-b)$ for the receiver and that GHPS's showed that this is the efficiency bound for communication equilibria in this environment. It remains to show how to replicate KM's construction with private information about the receiver's language competence. For this we briefly summarize the key aspects of their construction: Communication proceeds in two stages. In the first stage the sender reveals whether her type t is less than some quantity x, or not. In the second stage, if $t < x$ then a partition equilibrium is played on the interval $[0, x]$; otherwise with probability p (that is generated by a jointly controlled lottery) the sender sends a message to indicate whether $t \in (x, z)$ or $t \in [z, 1]$, and with probability $1 - p$ no further message is sent.

The outcome of any such equilibrium can be induced with private information about the receiver's language as follows: If the partition equilibrium on the interval $[0, x]$ has $K - 2$ steps, let M contain K messages. The receiver's partition type is either the finest partition of M, denoted $\mathcal{P}^f$, or it is the finest partition that contains the element $\{m_{K-1}, m_K\}$, denoted $\mathcal{P}^c$. The receiver's type distribution π_R is given by $\pi_R(\mathcal{P}^f) = p$ and $\pi_R(\mathcal{P}^c) = 1 - p$. Sender types in the kth interval of the partition equilibrium on $[0, x]$ send message m_k, sender types in (x, z) send message m_{K-1}, and sender types in the interval $[z, 1]$ send message m_K. The key observation is that the distinction between the messages m_{K-1} and m_K is activated only if the receiver's type is $\mathcal{P}^f$, which happens with probability p. Otherwise, the receiver uses an action in response to these message that is optimal against prior beliefs concentrated on the interval $[z, 1]$. it is now easy to see that sender and receiver face the exact same incentives as in the KM construction. □

It is also possible, as in the case of private information about language competence of the sender, to translate the BBK construction into an outcome equivalent equilibrium when there is private information about language competence of the receiver. The construction works for all biases $b > 0$ and is universal in the sense that the message space and availability type distribution are independent of the bias, but comes at the cost of requiring an infinite message space. Before proving this result, we will show by way of examples how one can use the BBK construction in simple settings.

We begin by constructing the analog to a two-step front-loading equilibrium of BBK. Suppose that $M = \{m_0, m_1, m_2\}$ and that the receiver's partition types are either $\mathcal{P}_1 = \{\{m_1\}, \{M \setminus \{m_1\}\}\}$ or $\mathcal{P}_2 = \{\{m_2\}, \{M \setminus \{m_2\}\}\}$, with equal probability. Types in the low step always send m_0. Types in the high step randomize uniformly over m_1 and m_2. Thus when a $\mathcal{P}_i$ type observes $\{M \setminus \{m_i\}\}$ he does not know whether this is the result of a low-step sender having sent m_0 or a high-step sender's randomization having failed to result in m_i. This scheme is analogous to having an error probability of $\frac{1}{2}$ in BBK. Higher error probabilities can be simulated as follows: Let $M = \{m_0, m_1, m_2, m_3\}$. There are three equally likely receiver types $\mathcal{P}_i = \{\{m_i\}, \{M \setminus \{m_i\}\}\}$ $i = 1, 2, 3$. Types in the low step always send m_0. Types in the high step randomize uniformly over m_1, m_2 and m_3. This scheme is analogous to having an error probability of $\frac{2}{3}$ in BBK. Finally, lower error probabilities can be simulated as follows: Let $M = \{m_0, m_1, m_2, m_3\}$. There are three equally likely receiver types $\mathcal{P}_i = \{\{m_i, m_{i+1}\}, \{M \setminus \{m_i, m_{i+1}\}\}\}$ $i = 1, 2, 3$, where addition is mod3. If the high-step sender randomizes uniformly over over m_1, m_2 and m_3, this scheme is analogous to having an error probability of $\frac{1}{3}$ BBK. It should be clear now how all two step equilibria with rational error probabilities in BBK can be simulated with private information about the receiver's language ability.

The following result extends these ideas to all error probabilities and all BBK front-loading equilibria.

Proposition 7 *With a uniform type distribution and quadratic preferences, there exists a message space M and an availability distribution π_R on the set of partition types of the receiver such that for every $b > 0$ there is an equilibrium in the corresponding game that attains the efficiency bound for communication equilibria.*

Proof: We will first show that there exists a message space M such that for every b there is a π_R and and corresponding equilibrium with the desired property.

Suppose that the optimal BBK front-loading equilibrium has K steps, $\Theta_1, \ldots, \Theta_K$ and that the associated error rate is $\epsilon(b)$. Denote the finest partition of a set S by $\mathcal{F}(S)$. Let $\widetilde{M}$ be the unit square, and for any $\alpha \in [0, 1]$ define

$$\widetilde{M}_\alpha := \left\{(x, y) \in \widetilde{M} \middle| x \in \left(\left(\alpha - \frac{\epsilon(b)}{2}\right)(\text{mod}1), \left(\alpha + \frac{\epsilon(b)}{2}\right)(\text{mod}1)\right)\right\},$$

$M = \widetilde{M} \cup \{m_0\}$, $M_\alpha = \widetilde{M}_\alpha \cup \{m_0\}$, $\mathcal{P}_\alpha = \mathcal{F}(M \setminus M_\alpha) \cup \{M_\alpha\}$, and let α be drawn from a uniform distribution on $[0, 1]$. The realization of α determines the receiver's partition type $\mathcal{P}_\alpha$.

To replicate the outcome from the optimal BBK front-loading equilibrium, consider the following sender strategy: Select $K - 1$ distinct values $y_2, \ldots, y_k \in [0, 1]$ and fix $x \in [0, 1]$. Sender types in the lowest step, Θ_1, send message m_0 and types in step Θ_k with $k > 1$ send message (x, y_k). Since all α are equally likely, the sender cannot foresee or control which pairs of messages (x, y) the receiver can distinguish from m_0, because they belong to $\widetilde{M} \setminus \widetilde{M}_\alpha$, and which ones he cannot distinguish from m_0, because they belong to $\widetilde{M}_\alpha$. Observe that given this strategy of the sender regardless of the value α, a receiver with partition type $\mathcal{P}_\alpha$ will receive a message $\{M_\alpha\}$ with probability $\epsilon(b)$ when a message other than m_0 is sent. Messages sent by any step Θ_k with $k > 1$ are understood by the receiver as sent with probability $1 - \epsilon(b)$ and otherwise the receiver cannot distinguish these messages from m_0. Thus, exactly as in the optimal BBK front-loading equilibrium, messages sent by types in the lowest step induce the intended action with probability one, and for any $k > 1$ the message sent by types in step k is correctly identified as coming from that set of types with probability $1 - \epsilon(b)$ and otherwise pooled with the message sent by the step Θ_1. Therefore, for all messages sent and received in the candidate equilibrium, both sender and receiver face the exact same incentives as in the optimal BBK front-loading equilibrium. Finally, assume that the receiver believes that any other (off-equilibrium) message was sent by type $t = 0$. Then no type will want to send this message because in any equilibrium that implements the efficiency bound for communication equilibrium (including the BBK equilibrium) type $t = 0$ induces his ideal action. Therefore, we have an equilibrium that induces the same outcome as the optimal front-loading BBK equilibrium.

Finally, we can make both the messages space M and the receiver's availability distribution π_R independent of b by replicating the above construction for every b, thus adding a dimension to the message space, making it the union of the unit cube and the always available message m_0. □

4 Conclusion

What is the meaning of the expression "dark red"? A sensible answer might be that it corresponds to a particular RGB value (in the case of "dark red" (139,0,0)) or to a well defined set of such values when used in a natural language that has only finitely many color words. In this paper we point out that in general one may want to think of meaning in more general terms to allow for example "dark red" to mean "I am saying 'dark red,' but keep in mind that this may only be due to the fact that I am lacking the word for maroon."[8] This drives a wedge between the indicative and imperative meanings of messages. The same is true for the related phenomenon that the recipient of a message may be unable to discriminate this message from other messages, e.g. "dark red" from "maroon." If language is imperfectly shared in this sense, there may be different perceptions of which obligations are entailed by agreements or contracts.

[8] The RGB for maroon is (128,0,0).

References

[1] Baldick, Ross Ashley Brown, James Bushnell, Susan Tierney and Terry Winter [2007], "A National Perspective on Allocating the Costs of New Transmission Investment: Practice and Principles," White paper prepared for WIRES (working group for investment in reliable and economic electric systems).

[2] Crémer, Jacques, Luis Garicano and Andrea Prat [2007], "Language and the Theory of the Firm," *Quarterly Journal of Economics* **122**, 373-407.

[3] Blume, Andreas, Oliver J. Board and Kohei Kawamura [2007], "Noisy Talk," *Theoretical Economics* **2**, 395–440.

[4] Crawford, V.P.and J. Sobel [1982], " Strategic Information Transmission," *Econometrica* **50**, 1431–1451.

[5] Forges, F. [1986], "An Approach to Communication Equilibria," *Econometrica* **54**, 1375–1385.

[6] Goltsman, M., J. Hörner, G. Pavlov, and F. Squintani [2007], "Mediation, Arbitration and Negotiation," *Journal of Economic Theory*, forthcoming.

[7] Harsanyi, J.C. [1967-8], "Games of Incomplete Information Played by Bayesian Players, I, II, and III," *Management Science* **14**, 159-182, 320-334, 486-502.

[8] Ivanov, M. [2009], "Communication via a Strategic Mediator," *Journal of Economic Theory*, forthcoming.

[9] Jäger, Gerhard, Lars Koch-Metzger and Frank Riedel [2009], "Voronoi Languages," Working paper, University of Tübingen and University Bielefeld.

[10] Krishna, V. and J. Morgan [2004], "The Art of Conversation: Eliciting Information from Experts Through Multi-Stage Communication," *Journal of Economic Theory*, **117**, 147-179.

[11] Lewis, David [1969], *Convention: A Philosophical Study*, Harvard University Press, Cambridge, MA.

[12] Myerson, R.B. [1982], "Optimal Coordination Mechanisms in Generalized Principal-Agent Problems," *Journal of Mathematical Economics* **10**, 67–81.

[13] Myerson, R.B. [1986], "Multistage Games with Communication," *Econometrica* **54**, 323–358.

[14] Myerson, R.B. [1991], *Game Theory: Analysis of Conflict*, Harvard University Press, Cambridge, MA.