When Can Communication Be Informative?*

Yuichiro Kamada† and Anand Siththaranjan‡

March 28, 2025

Work in Progress

This paper is frequently updated. The latest version can be found at:

`https://anands29.github.io/`

Abstract

When do cheap talk games admit informative equilibria? In a binary-action model, we show that the alignment between Sender and Receiver preferences determines the existence of informative equilibria, or lack thereof. We provide a characterization of the equilibrium payoff set. When moving beyond the binary-action setting, we find that the preference alignment no longer captures the existence of an informative equilibrium in general. We demonstrate that even if the preferences are perfectly misaligned, there may exist an informative equilibrium. With a restriction to binary states, we find that such preferences do not allow the Receiver to improve their payoffs over a babbling equilibrium, yet this intuitive prediction does not hold beyond the binary-state assumption. We identify a different set of alignment conditions between preferences such that Bayesian persuasion realizes payoffs the same as, or different to, some equilibrium in cheap talk.

*We thank Attila Ambrus, Nemanja Antic, and Shintaro Miura for their helpful comments.

†University of California, Berkeley, Haas School of Business, Berkeley, 2220 Piedmont Avenue, CA 94720, USA and University of Tokyo, Faculty of Economics, 7-3-1 Hongo, Bunkyoku, Tokyo, 113-0033, Japan. E-mail: `y.cam.24@gmail.com`

‡University of California, Berkeley, Department of Electrical Engineering and Computer Science, Berkeley, 2594 Hearst Ave, Berkeley, CA 94720, USA. E-mail: `anandsranjan@berkeley.edu`

1 Introduction

Communication is essential in economic, business, and social activities. Sellers may want to convey information about their new product to buyers. An entrepreneur may want to explain the nature of their business to investors. A politician may want to explain her political view to the public. When can such communication be informative?

The seminal work by Crawford and Sobel (1982) studies a setting in which one agent, the Sender, knows the state of the world and talks to the other agent, the Receiver, who then takes an action that determines the welfare of the two agents. They demonstrate that informative communication can be an equilibrium outcome if and only if the bias, which measures the two agents' preference misalignment, is sufficiently small. This insight was obtained under a rather specific setting, where the optimal action by the two agents move in the same direction as the state of the world changes. Such a model would apply to various scenarios, but there are real-world settings that the model may not be able to capture.

To fix ideas, consider the following seller-buyer relationship. The seller has multiple products in her store. The buyer comes to the store, and he is particularly interested in two products: an old product that he knows well, and a new product that he does not know as much. Depending on the new product's quality or its match to his taste, the buyer may like to buy the new product or the old product. Likewise, depending on the supply cost of the new product, the seller may like to sell the new product or the old product. One could "order" the states in terms of the buyer's willingness to pay for the new product, where the buyer would be more eager to buy at higher states. However, being different from the models as in Crawford and Sobel (1982), this does not necessarily guarantee that the seller would also like to sell the new product when the states are high: the match to the buyer's taste may not be positively related to the cost, and when a "horizontal" aspect like the taste is absent, the seller would be more eager to sell when the states are lower in a natural case where the willingness to pay and the cost are postitively related.

To analyze this type of situations, this paper considers a general model of cheap talk and asks when informative communication can be an equilibrium outcome. We first examine the case of binary-actions—the cases like the seller-buyer relationship—and obtain a result that is (at least superficially) similar to the one in Crawford and Sobel (1982): communication can be informative if and only if preferences are aligned, where the meaning of "aligned" is suitably defined to draw this conclusion. We then turn to the case with more actions and show that communication can be informative even under an extreme form of preference misalignment, which we call perfect preference misalignment: this is a setting under which our result for the binary-action case implies that there would not be any informative equilibrium. We will also characterize the welfare to the agents and how the possibility of Bayesian persuasion, i.e., the situation where the Sender commits to a messaging strategy before observing the state of the world, changes the scope for information transmission.

For some of our examples, randomization by the Sender is necessary for there to exist an informative equilibrium. The insight that randomization helps communication goes back to Myerson (1991)'s example where the Sender has an option to send a pigeon that gets to its destination with some probability less than one. An analogous idea appears in Blume et al. (2007) that allows for noisy information transmission, and Green and Stokey (2007) where the Sender receives a noisy signal. The randomization is exogenous in those papers, but in our examples it is a best response for the Sender to randomize. In this sense, our examples are similar in the spirit of Ambrus et al. (2013) and Ivanov (2010) who consider models that introduce intermediaries and show that the intermediaries' incentive-compatible randomization may induce informative communication.

Miura (2014) considers information transmission with binary-actions and provided a characterization of the equilibria. Despite some similarity, the model is quite different: he considers the situation where the Sender's message is restricted to be a subset of the state space including the true state, while we consider cheap talk. Moreover, Miura (2014)'s characterization is for pure-strategy equilibria and he shows that the characterization continues to hold with mixing under certain assumptions, while mixing plays an important role in our model.

Like us, Kamenica and Gentzkow (2011) study how the degree of alignment between Sender and Receiver preferences affects the Sender's ex-ante payoff in Bayesian persuasion. Their model differs from ours since they allow the Sender to commit to a messaging strategy prior to observing the state. Moreover, the definitions of preference alignment are different from each other as well.

2 Model

We consider two agents, a Sender S and Receiver R. Let $I = \{S, R\}$. There is a finite state space Θ, a finite message space M with $|M| > 2$, and a finite action space A. The game proceeds as follows. First, a state of the world $\theta \in \Theta$ is realized according to a prior distribution $\mu_0 \in \Delta(\Theta)$. The Sender observes the state and sends a message $m \in M$ to the

Receiver who does not observe θ. After observing the message m, the Receiver takes an action $a \in A$. The Sender's strategy is a mapping $\sigma_S : \Theta \to \Delta(M)$, where we denote by $\sigma_S(m|\theta)$ the probability of a message m given state θ. The Receiver's strategy is a mapping $\sigma_R : M \to \Delta(A)$, where we denote by $\sigma_R(a|m)$ the probability of an action a given message m. An agent $i \in \{S, R\}$ has a utility function over the Receiver's action and the state of the world: $u_i(a, \theta)$. We assume that for all $i \in I$ and $\theta \in \Theta$, $u_i(a, \theta) \neq u_i(a', \theta)$ if $a \neq a'$. Define a_i^θ to be the unique optimal action for agent i in state $\theta \in \Theta$. The dynamic game specified here can be characterized by (Θ, M, A, μ_0, u).

The solution concept of interest is a *perfect Bayesian equilibrium* (PBE). Given a game (Θ, M, A, μ_0, u), we say that a strategy profile $\sigma = (\sigma_S, \sigma_R)$ is a PBE if the following hold:

1. (Sender's optimality) For every $\theta \in \Theta$ and for all $m \in \text{supp}(\sigma_S(\theta))$,
$$m \in \arg\max_{m' \in M} \mathbb{E}_{a \sim \sigma_R(m')} \left[u_S(a, \theta) \right].$$
2. (Receiver's optimality) There exists $\mu : M \to \Delta(\Theta)$ such that, for all $m \in M$ and $a \in \text{supp}(\sigma_R(m))$,
$$a \in \arg\max_{a' \in A} \mathbb{E}_{\theta \sim \mu} \left[u_R(a', \theta) | m \right],$$
and for any $m \in \cup_{\theta \in \Theta} \text{supp}(\sigma_S(\theta))$ and $\theta \in \Theta$,
$$\mu(\theta|m) = \frac{\sigma_S(m|\theta)\mu_0(\theta)}{\sum_{\theta' \in \Theta} \sigma_S(m|\theta')\mu_0(\theta')}. \qquad \text{(Bayes rule)}$$

The following is trivially a PBE. The Sender sends an arbitrary message, and the Receiver chooses an action that maximizes the expected payoff given the prior. We say that a PBE is babbling if the Receiver's action is independent of any message sent by the Sender. Otherwise, we say that the PBE is informative.

To compare preferences over strategy profiles σ, σ' ex-ante, write $\sigma \succeq_i \sigma'$ if i's (ex-ante) expected payoff from σ is weakly greater than the one form σ'.

3 Cheap Talk with a Binary-Action Set

In this section, we consider the case with binary-actions, e.g., when the Receiver's actions are either to buy or not. Let $A = \{B, N\}$.

Prior		B	N
$\frac{1}{3}$	θ_1	1, 1	0, 0
$\frac{1}{3}$	θ_2	0, 0	1, 1
$\frac{1}{3}$	θ_3	ϵ, 0	0, ϵ

(a) Game

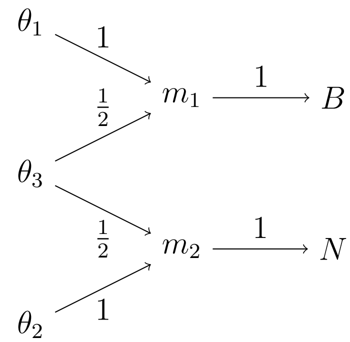

(b) PBE for $\epsilon = 0$

Figure 1: A game where the profile of utility functions is consistent with a preference profile that is not strongly aligned, but is aligned. Let $\epsilon \geq 0$.

In Section 3.1, we characterize the relationship between the agents' preferences and the existence of an informative PBE. Then, Section 3.2 considers the effect of communication on the payoffs.

3.1 Preference Alignment and Informative Equilibria

Our characterization depends on the ordinal rankings over actions, so we first introduce a notion that captures such ordinal rankings and relate it to the utility functions.

For each $\theta \in \Theta$ and $i \in I$, let $\succ_i^\theta$ be a strict linear order over A for player i at state θ. Let $\succ = (\succ_i^\theta)_{i \in I, \theta \in \Theta}$ be the preference profile. We restrict attention to preference profiles such that the Receiver does not have a strictly dominant action., i.e., there is no $a \in A$ such that $a \succ_R^\theta a'$ for all $a' \neq a$ and $\theta \in \Theta$.

Definition 1. A profile of utility functions u is **consistent** with $\succ$ if for all $i \in I, \theta \in \Theta, a, a' \in A$, we have $u_i(a, \theta) > u_i(a', \theta)$ if and only if $a \succ_i^\theta a'$.

The following definition provides a characterization of alignment between preference profiles:

Definition 2. A preference profile $\succ$ is **aligned** if there exist $\theta, \theta' \in \Theta$ with $\theta \neq \theta'$ such that $B \succ_i^\theta N$ and $N \succ_i^{\theta'} B$ for all $i \in I$.

To understand the condition, we divide the possible preference profiles into cases that satisfy certain properties - BB, BN, NB, and NN - where for $a, a' \in A$ we say a preference profile satisfies aa' if there exists $\theta \in \Theta$ such that $a_S^\theta = a$ and $a_R^\theta = a'$. See Figure 2. By definition, preferences are aligned if they satisfy both BB and NN.

We now provide some general results on the existence of informative PBE. Say that a prior μ_0 has full support if $\mu_0(\theta) > 0$ for every $\theta \in \Theta$.

Theorem 1. *Fix* (Θ, M, A) *such that* $|A| = 2$*, and a preference profile* $\succ$*.*

1. *If* $\succ$ *is aligned, then for all full-support priors* μ_0*, there exists* u *consistent with* $\succ$ *such that the game* (Θ, M, A, μ_0, u) *has an informative PBE.*

2. *If* $\succ$ *is not aligned, then for all full-support priors* μ_0 *there does not exist* u *consistent with* $\succ$ *such that the game* (Θ, M, A, μ_0, u) *has an informative PBE.*

We note that the proof of part 1 of Theorem 1 in fact shows that the profile of utility functions satisfying the stated property is not a knife-edge case: the result holds for some open set of profiles of utility functions, i.e., there exists $u = (u_S, u_R)$ and $\epsilon > 0$ such that for all u' in the ϵ-neighborhood of u, u' is consistent with $\succ$ and has an informative PBE.

The first part says that as long as the preference profile is aligned, there is a profile of utility functions that rationalizes the preferences and induces an informative PBE. This part is rather straightforward. Indeed, it is an intuitive statement when at *every* state of the world, the Sender and Receiver share an optimal action, such as in Figure 1 for $\epsilon = 0$. In this case, there is an informative PBE where the Sender truthfully reveals the state and the receiver chooses the mutually optimal action. The result in part 1 can be obtained by a continuity argument: we can observe that for $\epsilon > 0$ sufficiently small, the strict incentives are preserved, and hence we still obtain the existence of an informative PBE.

In contrast, part 2 of Theorem 1 is less straightforward because such continuity argument cannot be used. To prove the result, we first use the binary-action assumption to show that, in any informative PBE, there are exactly two action distributions that can be played by the Receiver on the path of play. To see this, fix an informative PBE and let $\mathcal{A}$ be the set of action distributions that can be played on the path of play. By the definition of informative PBE, $\mathcal{A}$ must constitute of multiple action distributions. The Sender has strict preferences at each state, and so either strictly prefers action B or strictly prefers action N. In the former types of states, the Sender must be sending a message that induces the action distribution assigning the highest probability to B in $\mathcal{A}$ (denote such an action distribution by α^B).[1] Similarly, in the latter types of states, the Sender must be sending a message that induces the action distribution assigning the highest probability to N in $\mathcal{A}$ (denote such an action distribution by α^N). This shows that no message can induce

[1]Such a distribution must exist, as otherwise the fixed strategy profile would not be an equilibrium.

an action distribution that is not α^B or α^N. Having shown that there are two possible action distributions, we now use the non-alignment assumption to show that there cannot exist an informative PBE. For example, if the preference profile does not satisfy BB, then whenever the Sender sends a message that induces α^B (such states must exist because α^B must be played with positive probability in equilibrium), the state must be such that the Receiver strictly prefers N, and hence the Receiver must assign probability 1 to action N. This implies that $\alpha^B(B) = 0$, which is a contradiction because $\alpha^B(B) > \alpha^N(B)$ must hold.

Theorem 1 identifies the condition under which it is *possible* to have the existence of an informative PBE is a possibility. We next take one step forward to ask when we can *guarantee* the existence of an informative PBE. To this end, we consider a stronger condition on the alignment of preferences.

Definition 3. $\succ$ is **strongly aligned** if it is aligned and for all $\theta \in \Theta$, either $B \succ_i^\theta N$ for all $i \in I$ or $N \succ_i^\theta B$ for all $i \in I$.

In terms of Figure 2, preferences are strongly aligned if they satisfy BB and NN, but nothing else.

Theorem 2. *Fix* (Θ, M, A) *such that* $|A| = 2$*, and a preference profile* $\succ$*.*

1. *If* $\succ$ *is strongly aligned, then for all full support priors* μ_0 *and for all* u *consistent with* $\succ$*, the game* (Θ, M, A, μ_0, u) *has an informative PBE.*

2. *If* $\succ$ *is not strongly aligned, then for all full support priors* μ_0*, there exists* u *consistent with* $\succ$ *such that the game* (Θ, M, A, μ_0, u) *does not have an informative PBE.*

As with Theorem 1, the proof of Theorem 2 shows that both parts hold for some open set of profiles of utility functions. In the proof of part 1, we find that a strongly aligned preference profile admit a PBE where both agents receive the maximum possible payoff. On the other hand, when considering preferences that are not strongly aligned, since preferences could be aligned, it is possible to find utility functions that admit informative PBE. However, part 2 of Theorem 2 states that we can always find a profile of utility functions. consistent with the preference profile such that no informative PBE exists. The idea is to construct payoffs such that the payoffs in the states with some misalignment are significant, which jeopardizes the possibility of communication.

To understand the first part of Theorem 2, observe that when agents share an optimal action at every state of the world, the following is a PBE: the Sender sends m_1 at every state where she prefers B (and hence, the Receiver prefers B as well by strong alignment),

	B	N
B	BB	BN
N	NB	NN

Figure 2: Properties that can hold for a preference profile. The rows represent the Sender's preferred action, and the columns represent the Receiver's preferred action.

and m_2 at every state where she prefers N (and hence, the Receiver prefers B as well by strong alignment). Both types of states exist because the preferences are strongly aligned. The Receiver then chooses B given m_1 and N given m_2. This is an equilibrium because at each state, the agents receive the maximum possible payoff given the state.

The second part of Theorem 2 states that a preference profile that is not strongly aligned can rationalize a profile of utility functions that admit no informative equilibria. To see why this is always possible, note that since preferences are not strongly aligned, there is a state where they disagree. Assume without loss that there is θ_{NB} such that the Sender prefers N while the Receiver prefers B. For contradiction, assume that there is an informative equilibrium σ^u for any consistent utility function u. Recall that there are two action distributions induced by the Sender in equilibrium, α^B and α^N with $\alpha^B(B) > \alpha^N(B)$, and the Sender at θ_{NB} assigns positive probability to message m^u that induces α_N. By taking $u_R(B, \theta_{NB})$ high enough, B becomes a unique best response for the Receiver after observing m^u. This means that $\alpha^N(B) = 1$, which is a contradiction because $\alpha^B(B) \leq 1$ must hold.

3.2 The Payoff Consequences of Cheap Talk

The argument so far characterizes the behavior of agents under different preferences. What is the effect of such behavior on the welfare? Define A_i^0 for $i \in I$ as the set of ex-ante optimal actions for i: $A_i^0 = \arg\max_{a \in A} \mathbb{E}_{\theta \sim \mu_0} [u_i(a, \theta)]$. Let $\delta : X \to \Delta(X)$ output a Dirac distribution centered at its input. When clear, we let X be inferred from context.

We first start with the following two observations, which should be straightforward:

Observation 1. *Fix (Θ, M, A, μ_0, u). There is $u_R^0 \in \mathbb{R}$ such that any babbling PBE gives the Receiver the payoff of u_R^0.*

Hence, there is a unique babbling payoff for the Receiver. This follows because, otherwise, the Receiver would have no incentive to play the action distribution that induces the lower babbling payoff.

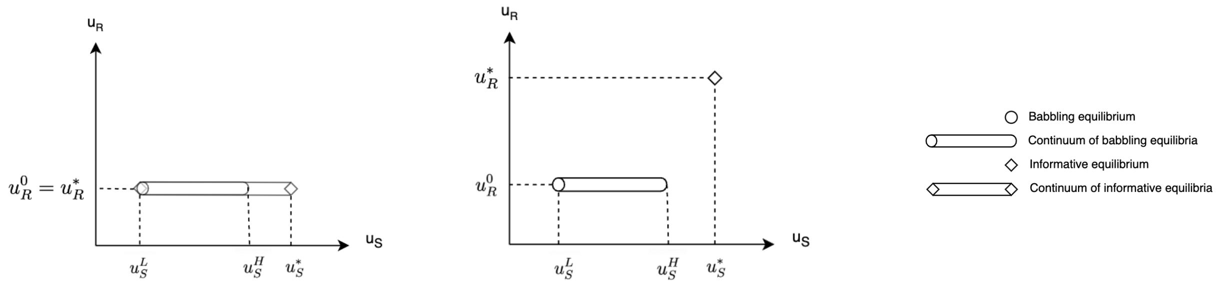

(a) $u_R = u_R^0$ (b) $u_R > u_R^0$ (c) Legend

Figure 3: Possible PBE payoffs when there exists an informative PBE σ with payoff profile (u_S, u_R). Example games that generate such payoffs can be found in Figure 20 in the Appendix. Note that Figure 3a and 3b cannot happen at the same time, which is shown in Theorem 7 in the Appendix.

Observation 2. *Fix* (Θ, M, A, μ_0, u). *Under any PBE, the Receiver receives an ex-ante payoff that is weakly greater than* u_R^0.

This follows because the Receiver always has a choice to ignore the message by the Sender and take an ex-ante optimal action. Let u_S^L and u_S^H denote the worst and best payoff, respectively, for the Sender under a babbling PBE. Let $u^* = (u_R^*, u_S^*)$ be such that $u_i^* = \mathbb{E}_{\theta\sim\mu_0}\left[u_i(a_S^\theta, \theta)\right]$, i.e., the utility profile when the Sender's optimal action is played in every state. Define $\rho_\sigma : \Theta \to \Delta(A)$ as a mapping from the state to the distribution of actions induced by σ at that state, i.e. $\rho_\sigma(a|\theta) = \mathbb{E}_{m\sim\sigma_S(\theta)}\left[\sigma_R(a|m)\right]$. We say that two strategy profiles σ and σ' are equivalent if $\rho_\sigma = \rho_{\sigma'}$.

Theorem 3. *Fix* (Θ, M, A, μ_0, u) *such that* $|A| = 2$. *The following holds:*

1. *The set of PBE is Pareto-ranked, i.e., there does not exist PBE* σ *and* σ' *such that* $\sigma \succ_S \sigma'$ *and* $\sigma' \succ_R \sigma$.

2. *If there exists an informative PBE* σ, *then there exists an informative PBE* σ' *that induces a payoff profile* u^*.

3. *All PBE payoff profiles are weakly greater than* (u_S^L, u_R^0).

4. *All PBE are equivalent to a PBE that uses at most two messages with positive probability, i.e., for all PBE* σ *there exists PBE* σ' *such that* $\rho_\sigma = \rho_{\sigma'}$ *and* $|\cup_{\theta\in\Theta}$ *supp*$(\sigma'_S(\theta))| \leq 2$.

Part 1 shows that PBE profiles can be ranked by a linear order such that no agent is worse off when choosing between PBE according to this order. Part 2 shows that if there exists an informative PBE within a given game, we can guarantee the existence of one that gives the Sender her ex-post optimal payoff. Part 3 implies that since all payoff profiles from some PBE are weakly greater than that of a babbling PBE, there does not exist informative PBE where the Sender's payoff is worse than all of her possible payoffs from babbling PBE. This means that allowing the Sender to send messages cannot introduce a PBE where she is worse off than in a PBE in a setting where she cannot send messages, which gives her a payoff the same as those in some babbling PBE. For a more precise theoretical characterization as well as illustrations of the set of PBE payoff profiles, that subsumes these results, see Theorem 7 and Figure 17 in Appendix C.

To prove part 1 of Theorem 3, we assume for contradiction that there exist PBE σ and σ' such that $\sigma \succ_S \sigma'$ and $\sigma' \succ_R \sigma$. Observe that σ' must be an informative PBE. This is because, if it were a babbling PBE, then $\sigma^0 \sim_R \sigma' \succ_R \sigma$ would hold where σ^0 is a babbling PBE, but $\sigma^0 \succ_R \sigma$ contradicts Observation 2. Recall that for an informative PBE, there are only two action distributions induced by the Sender in σ': α^B and α^N, such that $\alpha^B(B) > \alpha^N(B)$, where the Sender at each $\theta \in \Theta$ induces $\alpha^{a_S^\theta}$. The proof in the Appendix first shows that $\alpha^a = \delta(a)$ for each $a \in \{B, N\}$ must hold. That is, the Sender's preferred action in every state of the world is taken with probability 1 in any equilibrium. To see why, note that if this is not the case, then there is an action $a \in A$ such that for all messages m sent with positive probability, $\sigma'_R(a|m) > 0$. However, if after any message the Receiver is indifferent between playing a and playing $\sigma'_R(a|m)$, then it must be that the utility the Receiver achieves under σ' is the same as the one in the strategy profile in which the Receiver plays a given any message. This strategy profile can be shown to be a PBE, and it contradicts our previous observation that $\sigma' \succ_R \sigma^0$, where σ^0 gives payoff u_R^0, hence it must be that $\alpha^a = \delta(a)$ for all $a \in A$. As the Sender gets her best payoff in every state of the world when in the PBE σ' is played, it cannot be the case that $\sigma \succ_S \sigma'$. Hence we have a contradiction, and so we can conclude that the set of PBE is Pareto-ranked.

Note that this argument shows a result that is stronger than the claim of part 1. This is because the argument shows that if there is a PBE σ where the Receiver's payoff is strictly greater than u_R^0, then the payoff profile induced must be u^*. This also provides a proof of part 2 when the Receiver's payoff in σ is strictly greater than u_R^0. To complete the proof of part 2, we thus only need to consider σ such that the payoff to the Receiver is u_R^0. Without loss, suppose B is played with positive probability in some babbling PBE. Given σ, consider α^B and α^N as previously described. If both distributions have full support over actions, then the Receiver is indifferent between the two actions after any message, and it must be that any action distributions $(\tilde{\alpha}^B, \tilde{\alpha}^N)$ that satisfy $\tilde{\alpha}^B(B) > \tilde{\alpha}^N(B)$ can be induced by an informative PBE. This is because the action distributions preserve the Receiver optimality due to his indifference between actions, and at any state, no deviation by the Sender would induce an action distribution that plays her preferred action with a strictly greater probability. This in particular implies that the action distributions that play the Sender's optimal action with probability one at each state (i.e., $\tilde{\alpha}^B = 1$ and $\tilde{\alpha}^N = 0$) is also part of a PBE. Hence, the Sender's ex-post optimal payoff is a PBE payoff. In fact, any payoff for the Sender in $(u_S^L, u_S^*]$ is an equilibrium payoff. Similarly, if $\alpha^B(B) = 1$ and α^N has full support, then the Receiver is indifferent between all actions when $m^N \in M$ that induces α^N is sent. Thus any $\tilde{\alpha}$ such that $\tilde{\alpha}^B(B) = 1$ and $\tilde{\alpha}^N(B) < 1$ gives the Sender payoffs in the set $(u_S^L, u_S^*]$. Finally, if $\alpha^B(B) = \alpha^N(N) = 1$, then since the Receiver's ex-ante payoff is u_R^0, when he has observed m^N, he must be indifferent between playing B with positive probability and playing α^N. Given this action distribution, the Sender has her preferred action played in every state, so the Sender receives payoff u_S^* in σ.

If the Receiver gets a strictly higher payoff than u_R^0 in an informative PBE, then we can recall by the explanation of part 1 in Theorem 3 that u^* is the payoff profile induced. These results are illustrated in Figure 3. Hence, the payoff profile in any informative PBE is weakly greater than some babbling PBE's payoff profile. Since this is also true of any babbling PBE, and all PBE are either babbling or informative PBE, we can conclude that the payoff profile of any PBE is weakly greater than the payoff profile of any babbling PBE that induces the payoff profile (u_S^L, u_R^0). This proves the claim of part 3.

Finally, part 4 in Theorem 3 follows from the following observations. First, from any PBE σ we can construct another PBE σ' where all messages sent induce different action distributions by combining messages that induce the same action distribution. This doesn't change the incentives of the Sender, and if it was optimal for the Receiver to play the same distribution given different messages, it is also optimal when these messages are bundled. As there are only two actions and the Sender has strict preference, if there were three messages that induced different distributions each then one of them must be suboptimal for some preference as action distributions can be ranked by the Sender. Thus, in σ' there must be at most two messages sent with positive probability.

Prior		L	C	R
$\frac{1}{4}$	θ_1	0, 4	1, 1	4, 0
$\frac{1}{2}$	θ_2	4, 0	2, 2	0, 4
$\frac{1}{4}$	θ_3	0, 4	3, 3	4, 0

(a) Game

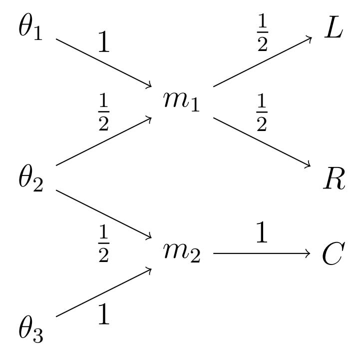

(b) PBE

Figure 4: A game with a perfectly misaligned preference profile but an informative PBE.

4 Beyond Binary-Action Sets

In this section we explore further results when we relax the binary-action assumption. We begin with counterexamples that contrast our results on the payoffs of equilibria and the existence of informative equilibria in the binary-action setting. We then provide a characterization of equilibrium strategies and payoffs when the preference profile is perfectly misaligned.

4.1 Counterexamples

Part 2 of Theorem 1 shows, in the binary-action setup, that there does not exist an informative PBE when preferences are not aligned. Even beyond the binary-action case, non-existence of an informative PBE seems intuitive in settings with *perfectly misaligned* preference profiles.

Definition 4. A preference profile $\succ$ is **perfectly misaligned** if for all $\theta \in \Theta$ and actions $a, a' \in A$, $a \succ_S^\theta a'$ if and only if $a' \succ_R^\theta a$.

The following example shows that an informative PBE can exist even when preferences are perfectly misaligned.

Example 1 (Informative PBE under perfectly misaligned preferences)**.** Consider the game in Figure 4a. It has a profile of utility functions consistent with a perfectly misaligned preference profile. However, one can show that the following strategy profile σ is an

Prior		L	C	R
$\frac{1}{3}$	θ_1	2, 3	1, 2	0, 0
$\frac{1}{3}$	θ_2	-99, 0	1, 1	-101, ϵ
$\frac{1}{3}$	θ_3	0, 0	1, 2	2, 3

(a) Game

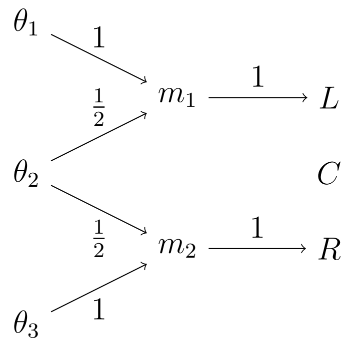

(b) PBE

Figure 5: A game where the results of Theorem 3 do not extend. Let $\epsilon > 0$ be sufficiently small.

informative PBE:

$$\sigma_S(\theta_1) = \delta(m_1), \quad \sigma_S(\theta_2) = \begin{cases} m_1 & \text{with probability } \frac{1}{2} \\ m_2 & \text{with probability } \frac{1}{2} \end{cases}, \quad \sigma_S(\theta_3) = \delta(m_2);$$

$$\sigma_R(m_1) = \begin{cases} L & \text{with probability } \frac{1}{2} \\ R & \text{with probability } \frac{1}{2} \end{cases}, \quad \sigma_R(m_2) = \delta(C).$$

Figure 4b provides a graphical representation of this strategy profile. The utility profile induced by σ is $(\frac{11}{4}, \frac{11}{4})$, whereas the unique babbling PBE induces $(2, 2)$. △

In the above example, we find that, even under perfectly misaligned preferences, there can exist an informative PBE where the Sender and Receiver are both strictly better off than in the babbling PBE. We later show in Theorem 4 that this is not possible in games with binary states.

We can also observe that the set of PBE does not satisfy all the properties found in Theorem 3. Consider the following example.

Example 2 (Sender is worse off in informative PBE)**.** Consider the game in Figure 5a.

Prior		L	C	R
$\frac{1}{2}$	θ_1	4, 4	0, 0	6, 1
$\frac{1}{2}$	θ_2	0, 0	4, 4	6, 1

(a) Game

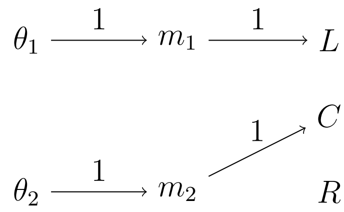

(b) PBE

Figure 6: A game where the result of part 2 in Theorem 3 do not extend, i.e. the Sender does not achieve her optimal payoff.

The following is an informative PBE:

$$\sigma_S(\theta_1) = \delta(m_1), \quad \sigma_S(\theta_2) = \begin{cases} m_1 & \text{with probability } \frac{1}{2} \\ m_2 & \text{with probability } \frac{1}{2} \end{cases}, \quad \sigma_S(\theta_3) = \delta(m_2)$$

$$\sigma_R(m_i) = \begin{cases} \delta(L) & \text{if } m_i = m_1 \\ \delta(R) & \text{if } m_i = m_2 \end{cases}.$$

Figure 5b provides a graphical representation of this strategy profile. We can observe that the ex-ante payoff profile for this informative PBE is $(-32, 2 + \frac{\epsilon}{6})$. In the unique babbling PBE, where C is played with probability 1, the payoff profile is $(1, \frac{5}{3})$. Hence the Sender is strictly worse off in the informative PBE than in the babbling PBE, whereas the Receiver is strictly better off. This implies that the conclusion of part 1 of Theorem 3, i.e., that the set of PBE is Pareto-ranked, does not generally hold when there are more than two actions. Moreover, notice that 1 is the lowest babbling PBE payoff because the babbling PBE is unique. Hence, the conclusion of part 3 of Theorem 3, that all PBE payoff profiles are weakly above (u_S^L, u_R^0), does not hold, either.[2] △

Furthermore, the following example shows that part 2 of Theorem 3 does not hold in general either:

Example 3 (Sender does not achieve her optimal payoff)**.** Consider the game in Figure

[2]We note that a counterexample for parts 1 and 3 of Theorem 3 does not require that there are at least 3 states. A more contrived example with two states can be found in Appendix F.

6a. Note that $u^* = (6, 1)$ holds. The following is an informative PBE:

$$\sigma_S(\theta_1) = \delta(m_1), \quad \sigma_S(\theta_2) = \delta(m_2)$$
$$\sigma_R(m_1) = \delta(L), \quad \sigma_S(m_2) = \delta(C)$$

Figure 6b provides a graphical representation of this strategy profile. The payoff profile induced by this informative PBE is $(4, 4)$.Although an informative PBE exists, there is no PBE that induces the payoff of u^*. This is because no posterior distribution over states would induce R as an optimal action by the Receiver. In fact, the best payoff that the Sender can achieve in this game is 4, which can be induced by σ. To see why 4 is the best payoff, observe that no posterior distribution over the states would make R a best response for the Receiver. Hence, there cannot be a PBE that plays R with positive probability. This implies that L and C are the only possible actions that may be played in equilibrium, and so it is clear that 4 is the best possible PBE payoff for the Sender. △

To see why a counterexample to part 2 of Theorem 3 exists when there are more than 2 actions, first note that the Sender at each state induces the best action distribution among those the Receiver would play in the given equilibrium. In a binary-action setting where the Receiver's payoff in σ is greater than his babbling payoff, the Sender gets her highest feasible payoff in each state. In Example 3, however, each of the Receiver's action distributions places probability one on different actions, while zero probability is placed on the Sender's preferred action. In the proof of the binary-action case, we relied on the property that in an informative equilibrium, in every state of the world the Sender's preferred action is played with positive probability because there are only two actions. From this, we can construct an equilibrium where the Receiver plays with probability one the Sender's preferred action in every state. In contrast, the case with more than two actions does not generally have this property because there does not need to be any probability placed on the Sender's preferred action in each state.

4.2 Perfectly Misaligned Preference Profiles

When preferences are perfectly misaligned, we can show that the state is never fully revealed by the Sender in arbitrary finite action spaces. We say the Sender **pools in all states** for a given strategy profile σ if for all $m \in \cup_{\theta \in \Theta} \text{supp}(\sigma_S(\theta))$, $|\{\theta \in \Theta | m \in \text{supp}\,(\sigma_S(\theta))\}| > 1$. That is, for every message sent with positive probability, there are multiple states that induce such a message. We say the Receiver **mixes on path** for a given strategy profile σ

if there exists $\theta \in \Theta$ and $m \in M$ such that $\sigma_S(m|\theta) > 0$ and $|\text{supp}(\sigma_R(m))| > 1$. That is, there is a state where the receiver plays a mixed action.

Proposition 1. *Fix (Θ, M, A), and consider a preference profile $\succ$. If $\succ$ is perfectly misaligned, then for all full support priors μ_0, u consistent with $\succ$, and any informative PBE σ in the game (Θ, M, A, μ_0, u), the Sender pools in all states and the Receiver mixes on path.*

In general, when there are more than two actions, the results of Section 3 no longer hold, as the previous section illustrated. As such, we restrict our attention in the following theorem to studying perfectly misaligned preference profiles when there are only two states:

Theorem 4. *Fix (Θ, M, A) such that $|\Theta| = 2$, and consider a preference profile $\succ$. If $\succ$ is perfectly misaligned, then for all full support priors μ_0 and all u consistent with $\succ$, the Receiver is indifferent between all PBE in the game (Θ, M, A, μ_0, u).*

When preferences are perfectly misaligned and there are more than two states, Example 1 shows that there is an informative PBE where the Receiver is strictly better off than in a babbling PBE.

The idea behind the proof is to first note that if there are binary states θ_1 and θ_2, then sending a message that induces a posterior different from the prior necessarily places higher probability on the state which sent it, and lower probability on the other state. If there exists such a message, then there must also exist another message that places correspondingly higher probability on the other state and lower probability on the given state. If the Receiver's optimal action is different under the posterior induced by each message, then the action distribution induced by θ_1 must be strictly worse for the Sender at θ_1 than the distribution at θ_2. This follows from the Receiver having perfectly misaligned preferences to the Sender. In such a case, the Sender would have an incentive to deviate, hence it must be that the Receiver's optimal action distribution does not change. Thus we can conclude that he his indifferent amongst all PBE.

5 Alternative Models of Strategic Communication

In this section, we compare alternative models of strategic communication to cheap talk. In particular, we study how equilibrium payoffs vary when considering different modes of communication such as Bayesian persuasion, arbitration, mediation and negotiation. We will see that the result that the preference alignment enables informative communication is robust, while the details of the effect may be different across models.

5.1 Bayesian Persuasion

In Bayesian persuasion (Kamenica and Gentzkow, 2011), the Sender commits to her messaging strategy σ_S prior to observing the state of the world. This is in contrast to cheap talk, where the Sender cannot commit and therefore chooses her message to maximize her utility given the realized state, rather than choosing a messaging strategy that maximizes her utility in expectation of the state. Following Kamenica and Gentzkow (2011), we say a strategy profile $\sigma = (\sigma_S, \sigma_R)$ is a **Bayesian Persuasion Equilibrium** (BPE) if σ_R satisfies Receiver optimality, as in a PBE, and σ_S satisfies *Sender's ex-ante optimality*:

$$\sigma_S \in \arg \max_{\sigma'_S:\Theta\to\Delta(M)} \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{m'\sim\sigma'_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m')}\left[u_S(a,\theta)\right].$$

Furthermore, we assume that when the Receiver is indifferent between actions, he takes an action that is optimal for the Sender.We say σ is a **Receiver-preferred BPE** if there is no other BPE that the Receiver strictly prefers. We can note the following observation comparing the Sender's welfare according to a Sender-preferred BPE and a PBE in the same game:

Observation 3. *Fix* (Θ, M, A, μ_0, u)*. Any BPE is ex-ante weakly better for the Sender than any PBE in the game* (Θ, M, A, μ_0, u)*.*

This follows from the fact that in Bayesian persuasion, the Sender can always commit to a messaging strategy used in any PBE, and hence her optimal strategy must be weakly better. In part 2 of Theorem 3, we found that when there exists an informative PBE, there is a PBE where the Sender gets her ex-post best utility in every state. Hence, when there is an informative PBE, the Sender-preferred BPE gives the same ex-ante utility as some PBE for both the Sender and Receiver. Can we characterize how equilibria and payoff profiles may differ between cheap talk and Bayesian persuasion in general? To do so, we define a new alignment condition on a preference profile.

Definition 5. A preference profile $\succ$ is **weakly misaligned** if there exists $\theta \in \Theta$ such that $a_R^\theta \neq a_S^\theta$.

This condition of a preference profile holds when the Sender and Receiver disagree on

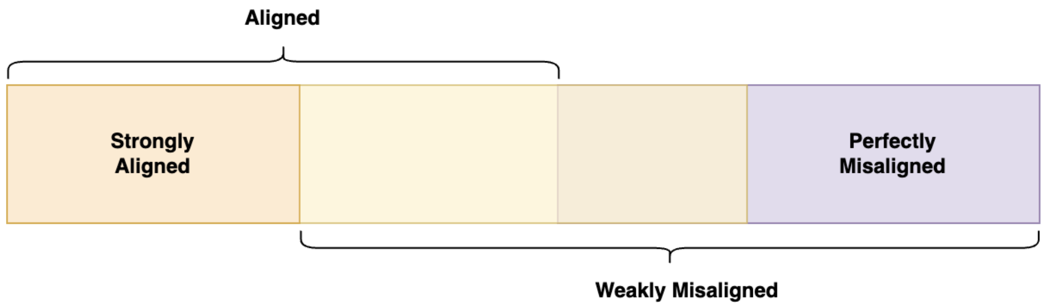

Figure 7: This diagram depicts the space of preference profiles considered in this work for the binary-action setting, and how the different conditions on profiles relate to one another. Note that we assumed that the Receiver does not have a dominant action when constructing this diagram.

Prior		B	N
$\frac{1}{2}$	θ_1	4, 0	0, $\frac{3}{2}$
$\frac{1}{2}$	θ_2	0, 1	1, 0

(a) Part 1 (perfectly misaligned)

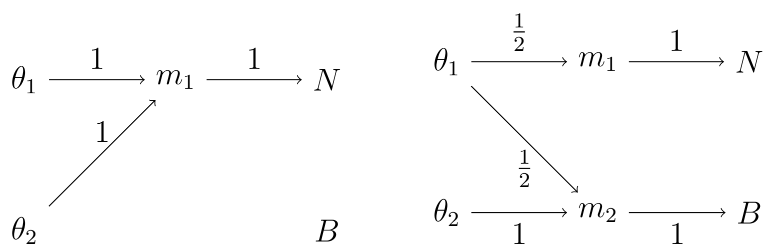

(b) PBE for 8a

(c) Sender prefers this to 8b

Figure 8: Example for Part 1 of Theorem 5.

their preferred action in some state of the world. To see how this condition relates to the other alignment definitions in the binary-action setting, see Figure 7.

Theorem 5. *Fix* (Θ, M, A) *such that* $|A| = 2$*, and a preference profile* $\succ$.

1. *If* $\succ$ *is weakly misaligned, then for all full-support priors* μ_0*, there exists* u *consistent with* $\succ$ *such that for the game* (Θ, M, A, μ, u)*, the Receiver-preferred BPE has an ex-ante payoff profile different from all PBE payoff profiles.*

2. *If* $\succ$ *is weakly misaligned, then for all full-support priors* μ_0 *there exists* u *consistent with* $\succ$ *such that for the game* (Θ, M, A, μ, u)*, the Receiver-preferred BPE payoff profile is the same as some as PBE payoff profile.*

3. *If* $\succ$ *is not weakly misaligned, then for all full-support priors* μ_0 *and for all* u *consistent with* $\succ$*, the Receiver-preferred BPE payoff profile is the same as some PBE payoff profile in the game* (Θ, M, A, μ, u).

Prior		B	N
$\frac{1}{3}$	θ_1	1, 1	0, 0
$\frac{1}{3}$	θ_2	0, 0	1, 3
$\frac{1}{3}$	θ_3	0, 1	1, 0

(a) Part 2 (aligned)

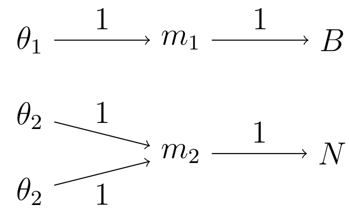

(b) PBE and BPE for 9a

Prior		B	N
$\frac{1}{3}$	θ_1	2, 3	0, 0
$\frac{1}{3}$	θ_2	2, 0	0, 1
$\frac{1}{3}$	θ_3	0, 1	1, 0

(c) Part 2 (not aligned)

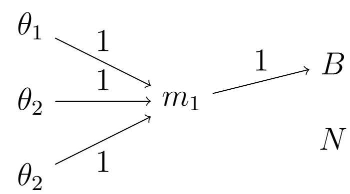

(d) PBE and BPE for 9c

Figure 9: Examples for Part 2 of Theorem 5.

Parts 1 and 2 of Theorem 5 show that under weakly misaligned preferences, we can always find consistent profiles of utility functions where Bayesian persuasion gives a different payoff profile to cheap talk, as well as profiles where the payoff profile is the same. To see why the former is possible, observe that we can always construct profiles such that we have no informative PBE (as in Figure 8b). In particular, we can consider two cases: either preferences are perfectly misaligned or there is a state where both agents agree on their preferred action. In the latter case, we can choose the payoff functions where the Sender and Receiver agrees on the optimal action to be sufficiently high such that full information disclosure - the Sender at each state sends a unique message - is ex-ante more profitable than the babbling PBE.

When the preference profile is instead perfectly misaligned, we can leverage the idea that noisy communication can improve outcomes. For example, say the Receiver's babbling strategy is N with probability one, as in Figure 8a. We can construct the Receiver's utility function such that there is a BPE where the Receiver's strategy has positive probability on both actions. Thus we can arbitrarily increase the payoff in the state where the Sender prefers B so that her ex-ante payoff in this BPE is greater than her payoff in the babbling PBE. Consider the following example:

Example 4. With respect to the game in Figure 8a, there is a unique PBE given by Figure 8b, which gives the babbling payoff profile where the Sender's ex-ante payoff is $\frac{1}{2}$. Consider the following strategy profile in Figure 8c:

$$\sigma_S(\theta_1) = \begin{cases} m_1 & \text{with probability } \frac{1}{2} \\ m_2 & \text{with probability } \frac{1}{2} \end{cases}, \quad \sigma_S(\theta_2) = \delta(m_2)$$

$$\sigma_R(m_i) = \begin{cases} \delta(N) & \text{if } m_i = m_1 \\ \delta(B) & \text{if } m_i = m_2 \end{cases}$$

where $m_1, m_2 \in M$ are distinct. The ex-ante payoff for the Sender under this strategy profile is 1. As σ_R is a best response to σ_S, the Sender can guarantee a payoff of 1 in a BPE. Hence in a BPE, the Sender can achieve an ex-ante payoff higher than in any PBE. △

To see part 2, we can consider two possible cases: either the preference profile is aligned or not. In the latter case, by choosing the Receiver's payoff to be sufficiently high in states of agreement between both agents, we can ensure that there is a PBE where the Sender achieves her ex-post optimal payoff. This is as in Figure 9, and since the outcome is ex-post optimal for the Sender, a BPE cannot do better ex-ante.

In the case when the preference profile is not aligned, there are only babbling PBE as per part 2 of Theorem 1. Furthermore, when preference are not perfectly misaligned, we can always ensure that the babbling action is unique and identical to the preferred action in the state where both agents agree. Without loss, let the babbling action be B, and let θ_3 be the state where they disagree with the Sender preferring N, as in Figure 9c. Note that in Θ_N, there is only ever disagreement, and that we cannot improve ex-ante payoffs by improving payoff in states in Θ_B. Hence, improving Sender payoffs in states $\theta \in \Theta_N$ would require pooling of messages from states in Θ_N with states in Θ_B. Yet we can choose the payoff in Θ_N to be sufficiently small such that for any message sent that pools between states, the reduction in payoff for states in Θ_B is more than any gain, if it exists, in payoff for states in Θ_N. As such, the Sender cannot improve her ex-ante payoff over the babbling PBE. A similar idea can be leveraged when the preference profile is perfectly misaligned.

Finally, when preferences are not weakly misaligned, the arguments for part 3 of Theorem 5 are the same as that for the proof of part 2 in Theorem 1 and the proof of part 1 in Theorem 2. In particular, when the preference profile is not weakly misaligned, since we assume that the Receiver does not have a dominant action then it must be the case that the Sender agrees with the Receiver in every state of the world. This implies that the preference profile is strongly aligned, and thus the Sender achieves her ex-post optimal payoff in some PBE. Hence, commitment cannot improve on this payoff.

5.2 Arbitration, Mediation and Negotiation

Other models of strategic communication are found in the literature on dispute resolution. In this section, we study arbitration, mediation and negotiation as described in Goltsman et al. (2009). These models share a common feature whereby there is a third-party T that facilitates the interaction between agents.

Arbitration. In arbitration, T acts in the Receiver's stead and commits to a strategy that maximizes his ex-ante expected utility while ensuring that the Sender has an incentive to truthfully reveal the state of the world. This can also be viewed as the Receiver committing to a state-contingent strategy that incentivizes truth telling. Consider a strategy $\sigma_T : \Theta \to \Delta(A)$, where $\sigma_T(\theta)$ denotes the distribution of actions played in state $\theta \in \Theta$. σ_T is an **Optimal Arbitration Rule** (OAR) if it is a solution to the following optimization problem:

$$\max_{p:\Theta\to\Delta(A)} \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{a\sim p(\theta)}\left[u_R(a,\theta)\right]$$

subject to

$$\theta \in \arg\max_{\hat{\theta}\in\Theta} \mathbb{E}_{a\sim p(\hat{\theta})}\left[u_S(a,\theta)\right] \tag{Sender-IC}$$

When comparing between OAR and PBE, we compare their ex-ante payoff profiles. Note the following observation:

Observation 4. *Fix an OAR σ_T and a babbling PBE σ^0. We have $\sigma_T \succeq_R \sigma^0$.*

We say that σ_T is babbling if for all $\theta, \theta' \in \Theta$, $\sigma_T(\theta) = \sigma_T(\theta')$. Otherwise it is called informative. Note that a babbling OAR induces the same ex-ante payoff profile as some babbling PBE. Furthermore, if a babbling OAR exists, then for each babbling PBE, there is a babbling OAR with the same payoff profile.

Mediation. On the other hand, mediation considers a third-party that can only give non-binding recommendations to the Receiver about their action. An equilibrium in this setting must thus consider the Receiver's inference about the state from the recommendation.

This is equivalent to having an obedience condition, which ensures that the Receiver has an no incentive to deviate from the third-party's recommended action. σ_T is an **Optimal Mediation Rule** (OMR) if it solves the following:

$$\max_{p:\Theta\to\Delta(A)} \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{a\sim p(\theta)}\left[u_R(a,\theta)\right]$$

subject to

$$\theta \in \arg\max_{\hat{\theta}\in\Theta} \mathbb{E}_{a\sim p(\hat{\theta})}\left[u_S(a,\theta)\right] \tag{Sender-IC}$$

and for all $a \in A$ such that $\mathbb{E}_{\theta\sim\mu_0}\left[p(a|\theta)\right] > 0$,

$$a \in \arg\max_{a'\in A} \mathbb{E}_{\theta\sim\mu_0}\left[u_R(a',\theta)|p(\theta)=a\right] \tag{Obedience}$$

A babbling OMR is defined similarly as a babbling OAR.

Negotiation. Negotiation differs from arbitration and mediation in that the third-party does not act on behalf of the Receiver, but rather structures how communication between the agents occur. In particular, T chooses a protocol by which both agents send messages to each other simultaneously for a possibly infinite period of time. Once this period of time has ended, the Receiver takes an action. To formalize this, let M_S and M_R be the message spaces for the Sender and Receiver respectively. We say the protocol of a game is $P = (\tau, M_S, M_R)$, where $\tau \in \mathbb{N} \cup \{\infty\}$ is the time horizon. Let $\Sigma(P)$ denote the set of PBE of the extensive-form game defined by a given protocol P, and for a strategy profile $\sigma \in \Sigma(P)$, let $\rho_\sigma : \Theta \to \mathcal{A}$ denote the distribution of actions played by the Receiver at a given state induced by strategy profile σ . $N = (P, \sigma)$ is an **Optimal Negotiation Rule** (ONR) if it solves the following optimization problem:

$$\max_{P,\sigma} \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{a\sim\rho_\sigma(\theta)}\left[u_R(a,\theta)\right]$$

subject to $\sigma \in \Sigma(P)$. That is, T chooses a protocol that induces a PBE which maximizes the Receiver's ex-ante payoff.

We find that in a binary-action environment, the Sender-optimal solution for each mode of communication, apart from BPE, gives the same payoff for both agents:

Theorem 6. *Fix the game* (Θ, M, A, μ_0, u) *such that* $|A| = 2$*. If* σ^{PBE}*,* σ^{OAR}*,* σ^{OMR}*,* N

are Sender-preferred PBE, OAR, OMR, and ONR solutions respectively, then they induce the same payoff profile. In particular, this payoff profile is either (u_0^R, u_S^H) *if there is no informative PBE, or* u^* *otherwise.*

The proof follows by first showing that the claim holds for just PBE and OAR, and then showing that that ordered set of payoff profiles for PBE, ONR, OMR, OAR are actually written in (weakly) increasing order. We show the former claim by first determining the OAR payoff profile is either u^* or u^0. This is achieved by considering the different alignment cases for preference profiles, and showing that the obedience constraints either bind or the optimal action in each state can be chosen to be optimal for the Sender. This allows us to order payoffs between PBE and OAR. To see that the the remaining payoff profiles can be ordered as above, observe that the set of feasible strategies for PBE, ONR, OMR, and OAR are, in a sense, (weakly) decreasing in the subset order. For example, any feasible mediation strategy is a feasible arbitration strategy. In the case of ONR, what this means is that from any feasible negotiation strategy, we can construct a feasible arbitration or mediation strategy that preserves the ex-ante payoff profile of the negotiation strategy. Given that the OAR and PBE payoff profiles of interest are identical, we can use this to show that the payoff profiles of ONR and OMR, which are ranked between the OAR and PBE payoff profiles, are also identical.

We can use Theorem 6 in conjunction with our alignment results in Theorem 1 to provide results for these other communication models:

Corollary 1. *Fix* (Θ, M, A) *such that* $|A| = 2$*, and a preference profile* $\succ$*.*

1. *If* $\succ$ *is aligned, then for all full support priors* μ_0 *there exists* u *consistent with* $\succ$ *such that the game* (Θ, M, A, μ_0, u) *has an informative OAR, OMR, and ONR.*

2. *If* $\succ$ *is not aligned, then for all full support priors* μ_0 *and* u *consistent with* $\succ$*, the game* (Θ, M, A, μ_0, u) *does not have an informative OAR, OMR, or ONR. [Can we use the same uility function for all of these?]*

Beyond the binary-action setting, the payoff equivalence between the Sender-preferred solutions no longer holds. The following examples study games with three actions to show this.

Example 5 (PBE and OAR have different utility profiles)**.** Consider the game in Figure 10a. First observe that in any PBE, there are no two messages sent with positive probability such that the belief of the Receiver is different, i.e. there does not exist

Prior		L	M	R
$\frac{1}{2}$	θ_1	1, 2	0, 0	−10, 3
$\frac{1}{2}$	θ_2	0, −1	1, 1	−2, −2

(a) Game

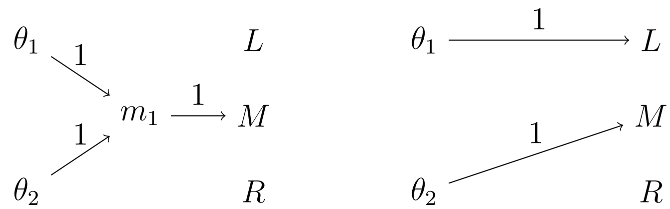

(b) PBE (c) σ_T

Figure 10: A game with three actions where the results of Theorem 6 does not hold for an OAR.

$m_1, m_2 \in \cup_{\theta\in\Theta}\text{supp}(\sigma_S(\theta))$ such that for some θ, $\mu(\theta_1|m_1) \neq \mu(\theta_1|m_2)$. If this were true, then we can assume that $\mu(\theta|m_1) > \frac{1}{2} > \mu(\theta|m_2)$ by Bayes' plausibility[3]. Hence, Receiver optimality would require that $\sigma_R(m_1) = \Delta(R)$ and $\sigma_R(R|m_2) = 0$. However the Sender at state θ_1 would have a profitable deviation by using the strategy $\sigma_S(m_1|\theta_1) = 1$. This contradicts σ being a PBE, hence $\mu(\theta|m)$ is independent of m for all $m \in \text{supp}(\sigma_S)$. By Bayes' plausibility, $\mu(\theta|m) = \mu_0(\theta)$ for all $m \in \text{supp}(\sigma_S)$. Hence given any message, the Receiver is indifferent over all actions, and we must have that $\sigma_S(\theta_1) = \sigma_S(\theta_2)$. If the latter did not hold, there would be a message where the belief changes from the prior. In a Sender-preferred PBE, the Receiver's strategy would not place positive probability on R. Finally, we show that in this PBE, the Sender is indifferent to a babbling PBE σ^* where $\sigma^*_R(M|m) = 1$ for all $m \in \text{supp}(\sigma_S)$. Since the Receiver is indifferent over playing L and M in a PBE, the only means by which the Sender's utility may improve is for there to be $m \in \text{supp}(\sigma_S(\theta_1))$ such that $\sigma_R(L|m_1) > 0$. If $\sigma_S(m_1|\theta_1) < 1$, then there would be a profitable deviation for the Sender at θ_1 to $\sigma_S(m_1|\theta_1) = 1$. However if $\sigma_S(m_1|\theta_1) = 1$, then it must be that $\sigma_S(m_1|\theta) = 1$ for all θ since the Receiver's belief cannot change. As such, we have that the only Receiver action distribution induced by any message is $\sigma_R(m_1)$, hence the ex-ante utility of the Sender is always $\frac{1}{2}$. Thus the Sender is indifferent between σ and σ^*, and the Sender-preferred PBE payoff profile is $(\frac{1}{2}, \frac{1}{2})$.

Now we show that Sender-preferred OAR has a different utility profile by showing that there is a feasible arbitration rule that has higher utility that the Sender-preferred PBE. First note that the following σ_T satisfies the Sender-IC constraint for an arbitration rule:

$$\sigma_T(\theta) = \begin{cases} L & \text{if } \theta = \theta_1 \\ M & \text{if } \theta = \theta_2 \end{cases}$$

[3]That is, $\mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{m\sim\sigma_S(\theta)}[\mu(\theta|m)] = \mu_0(\theta)$ for all $\theta \in \Theta$.

Prior		L	M	R
$\frac{1}{2}$	θ_1	$\frac{1}{2}$, 10	1, 9.9	0, 1
$\frac{1}{2}$	θ_2	0, 0	1, $\frac{1}{5}$	$\frac{1}{2}$, 1

(a) Game

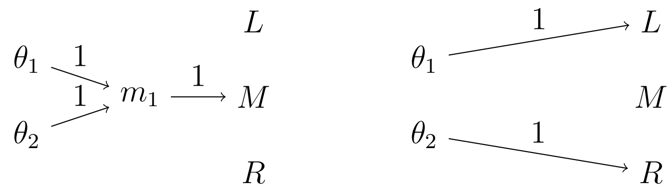

(b) PBE (c) σ_T

Figure 11: A game with three actions where the results of Theorem 6 does not hold for an OMR.

Furthermore, the utility profile for this strategy is $(1, \frac{3}{2})$. Hence the Sender-preferred OAR is distinct from the Sender-preferred PBE. △

Example 6 (PBE and OMR have different utility profiles)**.** Consider the game in Figure 11a. Observe that the following mediation rule is the Sender-preferred OMR σ_T:

$$\sigma_T(\theta) = \begin{cases} L & \text{if } \theta = \theta_1 \\ R & \text{if } \theta = \theta_2 \end{cases}$$

Note that it gives the Receiver their ex-post optimal payoff $\frac{11}{2}$, and the Sender's payoff is $\frac{1}{2}$. Furthermore, this is the unique Sender-preferred OMR as no other feasible mediation strategy gives the Receiver this payoff or better. Now note that the following strategy profile σ^* is a Sender-preferred PBE:

$$\sigma_S(\theta) = \delta(m_1)\ ,\ \sigma_R(m) = \delta(M)$$

The Sender receives their ex-post optimal payoff 1, and the Receiver's payoff is 10.1. Since the Sender must receive 1 in every Sender-preferred PBE, and the unique Sender payoff in the OMR is $\frac{1}{2}$, then there is no Sender-preferred PBE and OMR that induce the same payoff profile. △

To show PBE and ONR can have different utility profiles, we refer to Example 2.6 in Aumann and Hart (2003). In particular, they construct an example that shows how two stages of conversation can give an equilibrium strategy with higher payoff for the Receiver than just one stage of conversation. The former provides a lower bound for the ONR payoff, and latter corresponds to the payoffs achieved in some PBE.

References

Ambrus, Attila, Eduardo M Azevedo, and Yuichiro Kamada (2013), "Hierarchical cheap talk." *Theoretical Economics*, 8(1), 233–261.

Aumann, Robert J and Sergiu Hart (2003), "Long cheap talk." *Econometrica*, 71(6), 1619–1660.

Blume, Andreas, Oliver J Board, and Kohei Kawamura (2007), "Noisy talk." *Theoretical Economics*, 2(4), 395–440.

Crawford, Vincent P. and Joel Sobel (1982), "Strategic information transmission." *Econometrica*, 50(6), 1431–1451.

Goltsman, Maria, Johannes Hörner, Gregory Pavlov, and Francesco Squintani (2009), "Mediation, arbitration and negotiation." *Journal of Economic Theory*, 144(4), 1397–1420.

Green, Jerry R. and Nancy L. Stokey (2007), "A two-person game of information transmission." *Journal of Economic Theory*, 135(1), 90–104.

Ivanov, Maxim (2010), "Communication via a strategic mediator." *Journal of Economic Theory*, 145(2), 869–884.

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101(6), 2590–2615.

Miura, Shintaro (2014), "A characterization of equilibrium set of persuasion games with binary actions." *Journal of Mathematical Economics*, 51, 55–68.

$\succ$	**Case(s)**	**Result**
Aligned	Does not contain BB and NN	Obs 6
Aligned	(BB, NN)	Prop 7
Aligned	(BB, BN, NB, NN)	Prop 2
Aligned	$(BB, NB, NN), (BB, BN, NN)$	Prop 3
Misaligned	Contains BB and NN	Obs 6
Misaligned	$(BB, BN), (NB, NN)$	Prop 4
Misaligned	$BB, NN, BN, NB, (BB, NB), (BN, NN)$	Prop 5
Misaligned	$(BN, NB), (BB, BN, NB), (BN, NB, NN)$	Corollary 2
Strongly Aligned	All	Prop 7
Weak Misaligned	$(BB, NB, NN), (BB, BN, NN)$	Prop 8
Weak Misaligned	(BB, NB, BN, NN)	Prop 9
Weak Misaligned	Not (BB, NB, NN), (BB, BN, NN) or (BB, NB, BN, NN)	Obs 7

Figure 12: Table of results and cases that prove Theorem 1 and 2. Cases refer to some subset of $\{BB, BN, NB, NN\}$. For example, if BN holds, then this refers to there existing $\theta \in \Theta$ such that $B \succ_S^\theta N$ and $B \succ_R^\theta N$.

A Proof of Theorem 1

The proofs of Theorem 1 and 2 are shown in various cases, as seen in Figure 12.

Observation 5. *All preference profiles satisfy at least one of the following, and at most all of them: BB, BN, NB, NN.*

Given this observation, we can prove our theorem by considering these various cases. As our results differ based on whether the preference profile is aligned or not, we can narrow down which cases are necessary to consider depending on this property:

Observation 6. *$\succ$ is aligned if and only if BB and NN simultaneously hold.*

Thus for aligned profiles, we can just consider the cases where BB and NN hold simultaneously, and for misaligned profiles we only need to consider all other cases. For $c = aa'$, where $a, a' \in A$, let $\Theta_c = \{\theta \in \Theta : a_S^\theta = a, a_R^\theta = a'\}$. We first prove the following useful lemma:

Lemma 1. *Consider some preference profile $\succ$ for a given Θ. Let $\emptyset \neq C \subseteq \{BB, BN, NB, NN\}$ be the set of cases satisfied by $\succ$. Let $|\tilde{\Theta}| = |C|$. The following hold:*

1. *If for all support priors $\tilde{\mu}_0 \in \Delta(\tilde{\Theta})$ there exists $\tilde{u}$ satisfying C that has an informative PBE, then for all support priors $\mu_0 \in \Delta(\Theta)$ there exists u satisfying C that has an informative PBE.*

2. *If for all support priors $\tilde{\mu}_0 \in \Delta(\tilde{\Theta})$ there exists $\tilde{u}$ satisfying C that has no informative PBE, then for all support priors $\mu_0 \in \Delta(\Theta)$ there exists u satisfying C that has no informative PBE.*

Proof. Consider the first part, and assume the antecedent is true. Let $\mu_0 \in \Delta(\tilde{\Theta})$ be a full-support prior. Let $\tilde{\Theta} = \{\theta_c\}_{c\in C}$, and define $\tilde{\mu}_0(\theta_c) = \mu_0(\Theta_c)$. Then there exists $\tilde{u} : \tilde{\Theta} \to \mathbb{R}$ that satisfies C with an informative PBE $\tilde{\sigma}$ by assumption. Define $u : \Theta \to \mathbb{R}$ as follows: if $\theta \in \Theta$ is such that $c \in C$ holds, then $u_i(\cdot, \theta) = \tilde{u}_i(\cdot, \theta_c)$. Define the Sender's strategy as follows: if $\theta \in \Theta$ is such that $c \in C$ holds, then $\sigma_S(\theta) = \tilde{\sigma}_S(\cdot, \theta_c)$. Define the Receiver's strategy as $\sigma_R = \tilde{\sigma}_R$. We now verify that σ is a PBE for the full state space Θ. Observe that there is no profitable deviation for σ_S at each θ since $\tilde{\sigma}_S$ is a PBE. As for σ_R, note for $m \in \text{supp}(\sigma_S(\theta))$ for $\theta \in \Theta_c$, we have that

$$\begin{aligned}
\mathbb{E}_{\theta\sim\mu_0}\left[u_R(a,\theta)|m\right] &= \sum_{c\in C}\sum_{\theta\in\Theta_c}\mu(\theta|m)u_R(a,\theta) \\
&= \sum_{c\in C}\tilde{u}_R(a,\theta_c)\sum_{\theta\in\Theta_c}\mu(\theta|m) \\
&= \sum_{c\in C}\tilde{u}_R(a,\theta_c)\sum_{\theta\in\Theta_c}\frac{\sigma_S(m|\theta)\mu(\theta)}{\sum_{\theta'\in\Theta}\sigma_S(m|\theta')\mu(\theta')} \\
&= \sum_{c\in C}\tilde{u}_R(a,\theta_c)\frac{\tilde{\sigma}_S(m|\theta_c)}{\sum_{\theta'\in\Theta}\sigma_S(m|\theta')\mu(\theta')}\sum_{\theta\in\Theta_c}\mu(\theta) \\
&= \sum_{c\in C}\tilde{u}_R(a,\theta_c)\frac{\tilde{\sigma}_S(m|\theta_c)\mu(\Theta_C)}{\sum_{\theta'\in\Theta}\sigma_S(m|\theta')\mu(\theta')} \\
&= \sum_{c\in C}\tilde{u}_R(a,\theta_c)\tilde{\mu}(\theta_c|m) \\
&= \mathbb{E}_{\theta\sim\tilde{\mu}_0}\left[\tilde{u}_R(a,\theta)|m\right].
\end{aligned}$$

Observe that since $\tilde{\sigma}_R$ is an optimal solution to the above since it is a PBE, hence σ_R is also optimal. Thus σ is a PBE. To see that this is an informative PBE, first note that the ex-ante utility of R is the same in both u and $\tilde{u}$:

$$\begin{aligned}
\mathbb{E}_\mu\left[u_R(a,\theta)\right] &= \sum_{\theta\in\Theta} u_R(a,\theta)\mu(\theta)\\
&= \sum_{c\in C}\sum_{\theta\in\Theta_c} u_R(a,\theta)\mu(\theta)\\
&= \sum_{c\in C}\sum_{\theta\in\Theta_c} \tilde{u}_R(a,\theta_c)\mu(\theta)\\
&= \sum_{c\in C} \tilde{u}_R(a,\theta_c)\sum_{\theta\in\Theta_c}\mu(\theta)\\
&= \sum_{c\in C} \tilde{u}_R(a,\theta_c)\mu(\Theta_c)\\
&= \sum_{c\in C} \tilde{u}_R(a,\theta_c)\tilde{\mu}(\theta_c)\\
&= \mathbb{E}_{\tilde{\mu}}\left[\tilde{u}_R(a,\theta_c)\right].
\end{aligned}$$

where we observe that $\Theta = \oplus_{c\in C}\Theta_c$. Furthermore, the ex-ante utility correspond to σ and $\tilde{\sigma}$ are the same:

$$\begin{aligned}
\mathbb{E}_\mu\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m)}\left[u_R(a,\theta)\right] &= \sum_{\theta\in\Theta}\mu(\theta)\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m)}\left[u_R(a,\theta)\right]\\
&= \sum_{c}\sum_{\theta\in\Theta_c}\mu(\theta)\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m)}\left[u_R(a,\theta)\right]\\
&= \sum_{c}\sum_{\theta\in\Theta_c}\mu(\theta)\mathbb{E}_{m\sim\tilde{\sigma}_S(\theta)}\mathbb{E}_{a\sim\tilde{\sigma}_R(m)}\left[\tilde{u}_R(a,\theta_c)\right]\\
&= \sum_{c}\mathbb{E}_{m\sim\tilde{\sigma}_S(\theta)}\mathbb{E}_{a\sim\tilde{\sigma}_R(m)}\left[\tilde{u}_R(a,\theta_c)\right]\sum_{\theta\in\Theta_c}\mu(\theta)\\
&= \sum_{c}\mathbb{E}_{m\sim\tilde{\sigma}_S(\theta)}\mathbb{E}_{a\sim\tilde{\sigma}_R(m)}\left[\tilde{u}_R(a,\theta_c)\right]\tilde{\mu}(\theta_c)\\
&= \mathbb{E}_{\tilde{\mu}}\mathbb{E}_{m\sim\tilde{\sigma}_S(\theta_c)}\mathbb{E}_{a\sim\tilde{\sigma}_R(m)}\left[\tilde{u}_R(a,\theta_c)\right].
\end{aligned}$$

Since σ is distinct from a babbling PBE since it is an informative PBE, the same can be said of $\tilde{\sigma}$. Observe that if σ_R does not take the same pure action after every message sent in the support of σ_S, then this is the case of $\tilde{\sigma}$. Hence $\tilde{\sigma}$ is an informative PBE.

For the second claim, we proceed by contraposition. Assume there exists μ_0 such that for all u satisfying C there is an informative PBE. For $\tilde{\theta}_c \in \tilde{\Theta}$, define $\tilde{\mu}_0(\tilde{\theta}_c) = \mu_0(\Theta_c)$. Consider some $\tilde{u}$ satisfy C. Define $u : \Theta \to \mathbb{R}$ as follows: if $\theta \in \Theta_c$ for $c \in C$, then

$u_i(\cdot,\theta) = \tilde{u}_i(\cdot,\tilde{\theta}_c)$. Observe that u satisfies C, hence there exists σ an informative PBE. Define the Sender's strategy as follows: for $\tilde{\theta}_C \in \tilde{\Theta}$, $\tilde{\sigma}_S(m|\tilde{\theta}_c) = \frac{1}{\mu(\Theta_c)} \sum_{\theta\in\Theta_c} \mu_0(\theta)\sigma_S(m|\theta)$. Define the Receiver's strategy as $\tilde{\sigma}_R = \sigma_R$. Observe that $\tilde{\sigma}_S$ is optimal since the utility of each agent in Θ_c is the same, hence if there were a deviation for S at some state in $\tilde{\theta}_c$, then there would also be a deviation for $\sigma(\theta)$ for $\theta \in \Theta_c$. We now verify that $\tilde{\sigma}_R$ is optimal. Consider some $m \in M$ sent with positive probability. We can see that optimization problem given m for the Receiver in $\tilde{\Theta}$ is the same as that in Θ:

$$
\begin{aligned}
\mathbb{E}_\mu \left[u(a,\theta)|m\right] &= \sum_{\theta\in\Theta} u_R(a,\theta)\mu(\theta|m) \\
&= \sum_{c\in C}\sum_{\theta\in\Theta_c} \tilde{u}_R(a,\theta_c)\frac{\mu(\theta)\sigma_S(m|\theta)}{\sum_{\theta'\in\Theta}\mu(\theta')\sigma_S(m|\theta')} \\
&= \sum_{c\in C} \tilde{u}_R(a,\theta_c)\frac{1}{\sum_{\theta'\in\Theta}\mu(\theta')\sigma_S(m|\theta')}\sum_{\theta\in\Theta_c}\mu(\theta)\sigma_S(m|\theta) \\
&= \sum_{c\in C} \tilde{u}_R(a,\theta_c)\frac{\mu(\Theta_c)\tilde{\sigma}_S(m|\theta_c)}{\sum_{c'\in C}\sum_{\theta'\in\Theta_{c'}}\mu(\theta')\sigma_S(m|\theta')} \\
&= \sum_{c\in C} \tilde{u}_R(a,\theta_c)\frac{\tilde{\mu}(\theta_c)\sigma_S(m|\theta_c)}{\sum_{c'\in C}\tilde{\mu}(\theta_{c'})\tilde{\sigma}_S(m|\theta_{c'})} \\
&= \sum_{c\in C} \tilde{u}_R(a,\theta_c)\tilde{\mu}(\theta_c|m) \\
&= \mathbb{E}_{\tilde{\mu}}\left[\tilde{u}_R(a,\theta_c)|m\right].
\end{aligned}
$$

As such, we can see that $\tilde{\sigma}_R$ is optimal as σ_R is a PBE. Thus $\tilde{\sigma}$ is a PBE. We now verify that $\tilde{\sigma}$ is an informative PBE. Note that the babbling PBE between $\tilde{u}$ and u give the same ex-ante utility to both agents by appealing to the same argument as before. The following shows that the ex-ante utility of R is the same between $\tilde{\sigma}$ and σ, and the same argument applies for S:

$$\begin{aligned}
\mathbb{E}_{\mu}\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m)}\left[u_R(a,\theta)\right] &= \sum_{\theta\in\Theta}\mu(\theta)\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m)}\left[u_R(a,\theta)\right] \\
&= \sum_{c\in C}\sum_{\theta\in\Theta_c}\mu(\theta)\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m)}\left[\tilde{u}_R(a,\theta_c)\right] \\
&= \sum_{c\in C}\sum_{\theta\in\Theta_c}\mu(\theta)\sum_{m\in M}\sigma_S(m|\theta)\mathbb{E}_{a\sim\sigma_R(m)}\left[\tilde{u}_R(a,\theta_c)\right] \\
&= \sum_{c\in C}\sum_{m\in M}\sum_{\theta\in\Theta_c}\mu(\theta)\sigma_S(m|\theta)\mathbb{E}_{a\sim\sigma_R(m)}\left[\tilde{u}_R(a,\theta_c)\right] \\
&= \sum_{c\in C}\sum_{m\in M}\mathbb{E}_{a\sim\sigma_R(m)}\left[\tilde{u}_R(a,\theta_c)\right]\sum_{\theta\in\Theta_c}\mu(\theta)\sigma_S(m|\theta) \\
&= \sum_{c\in C}\sum_{m\in M}\mathbb{E}_{a\sim\sigma_R(m)}\left[\tilde{u}_R(a,\theta_c)\right]\mu(\Theta_c)\tilde{\sigma}_S(m|\theta_c) \\
&= \sum_{c\in C}\tilde{\mu}(\theta_c)\sum_{m\in M}\tilde{\sigma}_S(m|\theta_c)\mathbb{E}_{a\sim\sigma_R(m)}\left[\tilde{u}_R(a,\theta_c)\right] \\
&= \mathbb{E}_{\tilde{\mu}}\mathbb{E}_{m\sim\tilde{\sigma}_S(\theta_c)}\mathbb{E}_{a\sim\sigma_R(m)}\left[\tilde{u}_R(a,\theta_c)\right].
\end{aligned}$$

Observe that again, if σ_R does not take the same pure action after every message sent in the support of σ_S, then this is the case of $\tilde{\sigma}$. Hence $\tilde{\sigma}$ is an informative PBE. Thus we can conclude that $\tilde{\sigma}$ is an informative PBE since σ is an informative PBE. □

This allows us to study a small finite state space instead of one of arbitrary size. Now we prove the first result for an aligned preference profile by studying the following cases.

Proposition 2. *If $\succ$ is aligned and satisifes (BB, NB, BN, NN), then for all full support priors μ_0 there exists u consistent with $\succ$ such that there is an informative PBE. Furthermore, this holds generically.*

Proof. As per Lemma 1, we can just consider a $\tilde{\Theta}$ that satisfies the same conditions. Fix μ_0. Let $\tilde{\Theta} = \{\theta_{BB}, \theta_{BN}, \theta_{NB}, \theta_{NN}\}$. Let u' be the profile of utility functions in Figure 13. Let u be defined as follows:

$$u_i(a,\theta) = \frac{u_i'(a,\theta)}{\mu(\theta)}.$$

Observe that u satisfies (BB, NB, BN, NN) since u' does and u is a scaling of it at each

Prior		B	N
$\frac{1}{4}$	θ_{BB}	1, 2	0, 0
$\frac{1}{4}$	θ_{BN}	1, 0	0, 1
$\frac{1}{4}$	θ_{NB}	0, 1	1, 0
$\frac{1}{4}$	θ_{NN}	0, 0	1, 2

Figure 13: A game with a profile of utility functions that is consistent with an aligned, but not strongly aligned, preference profile that admit an informative PBE. This example is utilized in the proof of Theorem 1 for the case where BB, BN, NB, and NN hold for a preference profile.

Prior		B	N
$\frac{1}{3}$	θ_{BB}	1, 1	0, 0
$\frac{1}{3}$	θ_{NB}	0, 1	1, 0
$\frac{1}{3}$	θ_{NN}	0, 0	1, 2

Figure 14: (BB,NB,NN)

state. Consider the following candidate PBE: for all $a, a' \in A$

$$\sigma_S(\theta_{aa'}) = \delta(m_a)$$
$$\sigma_R(m_a) = \delta(a)$$

where $\{m_a\}_{a \in A} \subseteq M$ distinct. Observe that the Sender gets her optimal action at each state, hence there is no profitable deviation for them. To verify the optimality of the Sender's strategy, observe that after observing $m_{a'}$, her expected utility is the following:

$$\bar{u}_R(a) = \frac{1}{\mu(\theta_{a'B}) + \mu(\theta_{a'N})} \begin{cases} 2 & a = a' \\ 1 & \text{otherwise} \end{cases} .$$

Hence σ_R is optimal, and σ is a PBE. It has the ex-ante utility for the Receiver of 4. On the other hand, the babbling utility is upper bounded by 3. Hence this is an informative PBE. Furthermore, we can note that any small perturbation of the utilities in u, and thus $\bar{u}$, still have the same properties as above. Thus this property holds generically. □

The next case of interest is the following:

Proposition 3. *If $\succ$ is aligned and satisifes (BB, NB, NN), then for all full support priors μ_0 there exists u consistent with $\succ$ such that there is an informative PBE. Furthermore,* *this holds generically.*

Proof. As per Lemma 1, we can just consider a $\tilde{\Theta}$ that satisfies the same conditions. Fix μ_0. Let $\tilde{\Theta} = \{\theta_{BB}, \theta_{NB}, \theta_{NN}\}$. Let u' be the profile of utility functions in Figure 14. Let u be defined as follows:

$$u_i(a, \theta) = \frac{u'_i(a, \theta)}{\mu(\theta)}.$$

Observe that u satisfies (BB, NB, NN). Consider the following candidate PBE: for all $a, a' \in A$

$$\sigma_S(\theta_{aa'}) = \delta(m_a)$$
$$\sigma_R(m_a) = \delta(a)$$

where $\{m_a\}_{a\in A} \subseteq M$ distinct. Observe that the Sender gets her optimal action at each state, hence there is no profitable deviation for them. To verify the optimality of the Receiver's strategy, observe that after observing m_B, his expected utility is

$$\bar{u}_R(a) = \begin{cases} 1 & a = B \\ 0 & \text{otherwise} \end{cases}$$

and for m_N it is

$$\bar{u}_R(a) = \frac{1}{\mu(\theta_{NB}) + \mu(\theta_{NN})} \cdot \begin{cases} 2 & a = N \\ 1 & a = B \end{cases}.$$

Thus σ_R is optimal, and σ is a PBE. Observe that the ex-ante utility is 3, whereas the babbling utility is upper bounded by 2. Hence this is an informative PBE. Observe that this example is invariant to the labelling of actions, hence we can reproduce a similar example for the case where (BB, BN, NN) holds. Furthermore, we can note that any small perturbation of the utilities in u, and thus $\bar{u}$, still have the same properties as above. Thus this property holds generically. □

Finally, we can consider the case of only (BB, NN) holding by observing that this satisfies the definition of *strongly aligned*, and thus we can appeal to Proposition 7 as used in the proof of Theorem 2.

Now we prove the following statement: if $\succ$ is misaligned, then for all full-support priors $\mu_0 \in \Delta(\Theta)$ there does not exist (u_S, u_R) consistent with $\succ$ such that there exists σ informative. Assume for contradiction that for that there exists a full-support prior $\mu_0 \in \Delta(\Theta)$ and there exists (u_S, u_R) consistent with $\succ$ such that there exists σ an informative PBE. Note that a misaligned preference profile implies the following: for all θ, θ' distinct, one of the following must hold: there exists $i \in I$ such that $B \succ_i^\theta N$, or there exists $i' \in I$ such that $N \succ_{i'}^{\theta'} B$. By our previous observation, we can see that any case can hold for a misaligned preference profile except for BB and NN simultaneously. We now prove these remaining cases. First we prove the following useful lemma:

Lemma 2. *In an informative PBE, there exists $\alpha^B, \alpha^N \in \Delta(A)$, $\theta^B, \theta^N \in \Theta$, and $m^B \in supp(\sigma_S(\theta^B)), m^N \in supp(\sigma_S(\theta^N))$ where for $\alpha^B = \sigma_R(m^B)$ and $\alpha^N = \sigma_R(m^N)$ we have that*

$$\left(\alpha^B \succ_S^{\theta^B} \alpha^N\right) \wedge \left(\alpha^N \succ_S^{\theta^N} \alpha^B\right).$$

Furthermore, we have that $\alpha^N \neq \alpha^B$, $B \succ_S^{\theta^B} N$, $N \succ_S^{\theta^N} B$, and

$$\left(B \succ_S^\theta N \implies m^N \notin supp\sigma_S(\theta)\right) \wedge \left(N \succ_S^\theta B \implies m^B \notin supp\sigma_S(\theta)\right).$$

Proof. Let σ be an informative PBE. First note that there exists $m, m' \in M' = \cup_{\theta \in \Theta} \text{supp}(\sigma_S(\theta))$ such that $\sigma_R(m) \neq \sigma_R(m')$. To see why, suppose not for contradiction. Thus we have that $\sigma_R(m) = \alpha \in \Delta(A)$ for all $m \in M'$. Hence the Receiver has the same utility as a babbling PBE, thus this would be an optimal response according to the prior distribution. This contradicts σ not being a babbling PBE by definition.

Since $\sigma_R(m) \neq \sigma_R(m')$ for some $m, m' \in M'$, it must be the case that each action is taken with positive probability. Let $\mathcal{A} = \{\alpha \in \Delta(A) : \exists m \in M' \ \alpha = \sigma_R(m)\}$ be the set of action distributions that are induced by some message sent with positive probability. Let $\alpha^B, \alpha^N \in \Delta(A)$ be such that for all $\alpha \in \mathcal{A}$

$$\alpha^B(B) \geq \alpha(B)$$
$$\alpha^N(N) \geq \alpha(N).$$

Note that $\alpha^B(B)$ and $\alpha^N(N)$ are positive by the previous observation. Furthermore, since all states induce the same distribution, there must exist $\theta^B, \theta^N \in \Theta$ and $m^B, m^N \in M'$ such that $\alpha^a = \sigma_R(m^a)$ for all $a \in A$. Since there exists multiple action distributions taken in an informative PBE, it must be that case $\alpha^B \neq \alpha^N$. Clearly it is the case that the Sender at θ^B and θ^N prefer B and N as the optimal action, since otherwise there would be a profitable deviation to send the other message. Since the action distributions are different, if the Sender at a state prefers B to N, then they would not send m^N since this induces a distribution over actions α^N that is not favourable compared to α^B. The same applies for the Sender at a state that prefers N to B, where they would not send m^B, hence this gives the final claim. □

Now consider the case where the Sender prefers the same action at every state:

Proposition 4. *If only one of the following holds, then there are no informative PBE:* $(BB, BN)(NB, NN)$.

Proof. First consider the case where (BB, BN) holds. Observe that this implies that for all $\theta \in \Theta$, $B \succ_S^\theta N$. Assume for contradiction that there is an informative PBE σ. Hence by Lemma 2 there must be more θ^B and θ^N that induce different action distributions α^B and α^N for some message m^B and m^N respectively such that:

$$B \succ_S^\theta N \implies \alpha^B \succ_S^\theta \alpha^N$$
$$N \succ_S^\theta B \implies \alpha^N \succ_S^\theta \alpha^B.$$

However by assumption, θ^N is such that $B \succ_S^{\theta^N} N$. Thus the Sender at θ^N has a profitable deviation from $\sigma_S(\theta_N)$ to $\delta(m^B)$, which contradicts the assumption of σ being an PBE. A similar argument holds for (NB, NN). □

Now consider the case where the Receiver has the same preference over actions for all states:

Proposition 5. *If only one of the following holds, then there are no informative PBE:* $BB, NN, BN, NB, (BB, NB), (BN, NN)$.

Proof. First assume only BB holds. Since preferences are strict this implies for all $\theta \in \Theta$, $B \succ_i^\theta N$ for all $i \in I$. Thus for all priors $\mu \in \Delta(\Theta)$,

$$\{B\} = \arg\max_{a \in A} \mathbb{E}_{\theta \sim \mu} \left[u_R(a, \theta)\right].$$

This holds for all posteriors induced by σ_S, hence $\{B\}$ is always optimal. Thus the optimal outcome is the same under the prior, hence it is equivalent to a babbling outcome. The

Prior		B	N
$\frac{1}{4}$	θ_{BB}	1, 2	0, 0
$\frac{1}{4}$	θ_{BN}	1, −1	0, 0
$\frac{1}{4}$	θ_{NB}	0, 2	1, 0
$\frac{1}{4}$	θ_{NN}	0, −1	1, 0

Figure 15: A game where the profile of utility functions is consistent with a preference profile that is not strongly aligned, but is aligned. Used in the proof of Theorem 2 for the case where BB, BN, NB, and NN hold for a preference profile.

argument follows similarly for the remaining cases since the Receiver's optimal action is always the same. □

Consider the following assumption.

Assumption 1. *Let* $\Theta_a^i = \{\theta \in \Theta : a \succ_i^\theta a'\}$ *for all* $i \in I$ *and for all* $a, a' \in A$ *such that* $a \neq a'$. *Then there exists* $a, a' \in A$ *such that for* $\Theta_{a'}^S \subseteq \Theta_a^R$.

This assumption effectively says that if there is some information uniquely sent by the Sender at a certain state, then the Receiver will be able to identify that they have the opposite preference over actions according to the actual state.

Proposition 6. *If Assumption 1 hold and* $|A| = 2$, *then there are no informative PBE.*

Proof. Assume for contradiction that there is an informative PBE σ. Thus there must exist α^B and α^N that satisfy the conditions in Lemma 2. Without loss of generality, assume that $\Theta_B^S \subseteq \Theta_N^R$. Thus we have that $\mu(\Theta_B^S | m^B) = 1$ since m^B is only sent by Senders with states in Θ_B^S. Thus $\mu(\Theta_N^R | m_B) = 1$, and the unique optimal action distribution is $\alpha = \delta(N)$. This contradicts property of $\alpha^B(B) > 0$ since $0 = \alpha(B) = \sigma_R(m^B)(B) = \alpha^B(B)$. Hence σ cannot be an informative PBE. A similar argument holds for $\Theta_N^S \subseteq \Theta_B^R$. □

Corollary 2. *Assumption 1 is satisified in the following cases:* (BN, NB), (BB, BN, NB), (BN, NB, NN).

Proof. By direct observation. □

B Proof of Theorem 2

First we consider the following claim: if $\succ$ is strongly aligned, then for all full-support priors $\mu_0 \in \Delta(\Theta)$ and for all (u_S, u_R) consistent with $\succ$, there exists σ informative. The following proposition proves this claim:

Prior		B	N
$\frac{1}{3}$	θ_{BB}	1, 1	0, 0
$\frac{1}{3}$	θ_{NB}	−1, 2	0, 0
$\frac{1}{3}$	θ_{NN}	−1, −1	0, 0

Figure 16: A game with utility functions consistent with a non-strongly aligned preference profile, and admits no informative equilibria.

Proposition 7. *If $\succ$ is strongly aligned if and only if (BB, NN) holds. Furthermore, truth telling is an informative PBE.*

Proof. The first statement follows directly from the definition of strongly aligned. For the second statement, recall that for $\delta(m_\theta) = \sigma_S(\theta)$, $\mu(\theta|m_\theta) = 1$. Hence $\sigma_R(m) = a_R^\theta$ is optimal. Furthermore, for all $\theta \in \Theta$ the Sender cannot find a profitable deviation by sending a different message since for all $m \in \text{supp}(\sigma_S(\theta)$, $\sigma_R(m) \succeq_S^\theta \alpha$ for all $\alpha \in \Delta(A)$. To see informativeness of the PBE, not that since different actions are optimal at different states, a babbling PBE is dis-preferred to an PBE that is ex-post optimal for the Receiver. □

Now we prove the second claim. The following observation suggests that all that remains to prove is the case when a preference profile is not strongly aligned but not misaligned:

Observation 7. *If a preference profile is misaligned, then by Theorem 1, any utility function consistent with the preference profile does not have an informative equilibrium. Hence we need only to consider the case where the preference profile is not strongly aligned but not misaligned. This corresponds to one of the following cases: (BB, NB, NN), (BB, BN, NN), and (BB, BN, NB, NN).*

Thus we have the following theorem for these remaining cases:

Proposition 8. *If $\succ$ satisfies (BB, NB, NN) or (BB, BN, NN), then for all full-support priors $\mu_0 \in \Delta(\Theta)$ there exists (u_S, u_R) consistent with $\succ$ such that there does not exist σ informative. Furthermore, this holds for an open set of profiles of utility functions.*

Proof. By Lemma 1, we can just consider the case where $|\Theta| = 3$. Assume without loss that (BB, NB, NN) holds, since the argument for (BB, BN, NN) is similar due to symmetry. Let $\tilde{u}$ be the same as that in Figure 16. Define $u : \Theta \to \mathbb{R}$ as follows:

$$u(a, \theta) = \frac{\tilde{u}(a, \theta)}{\tilde{\mu}_0(\theta)}.$$

Note that the babbling PBE σ_0 gives ex-ante utility of 2 to the Receiver, since B is uniquely taken in such an PBE. To see that there is no informative PBE, assume for contradiction that there is an informative PBE σ.

By Lemma 2, let m^N induce α^N. Note that m^N must be sent by the Sender at both θ_{NB} and θ_{NN} as S at θ_{NB} cannot send a message not sent by other states. To see why, if it does then θ_{NB} is revealed to the Receiver, who then takes the worst action for S at θ_{NB} (which is strictly less preferred than α^N, which takes action N with positive probability). Since θ_{NN} is the only other state to induce α^N since it is suboptimal for θ_{BB}, and θ_{NB} must share a message with another state, it must be that θ_{NB} sends some m that induces α^N. Let m^N be this message without loss of generality. Furthermore, θ_{NN} cannot send a message that is not sent by θ_{NB}. If they did, then the Receiver would take action N with probability 1 upon receiving such a message. For this not to be a profitable deviation for θ_{NB} to send the same message, for all messages sent by θ_{NB} or θ_{NN}, the Receiver must take action N. This gives a worse utility ex-ante than taking action B, a contradiction since in any PBE it must be that $\sigma \succeq_R \sigma_0$. Furthermore, they must send messages with the same probability. Assume they do not, if θ_{NB} sends a message with higher probability than θ_{NN}, then B is the optimal action by the Receiver. Thus deviating to sending m^N always would be a profitable deviation. If θ_{NN} sends a message with higher probability than θ_{NB}, then there must be a message it sends that is of lower probability than θ_{NB} since they share the same support, leading to the same issue. Thus $\sigma_S(\theta_{NB}) = \sigma_S(\theta_{NN}) = \beta_N$. We now observe the following:

$$\mu_0(\theta|m^N) = \begin{cases} 0 & \text{if } \theta = \theta_{BB} \\ \frac{\mu_0(\theta_{NB})}{\mu_0(\theta_{NB})+\mu_0(\theta_{NN})} & \text{if } \theta = \theta_{NB} \\ \frac{\mu_0(\theta_{NN})}{\mu_0(\theta_{NB})+\mu_0(\theta_{NN})} & \text{if } \theta = \theta_{NN} \end{cases}$$
$$\implies \{B\} = \arg\max \mathbb{E}_{\theta\sim\mu_0}\left[u_R(a,\theta)|m^N\right].$$

Hence $\alpha^N(N) = 0$, a contradiction. Thus it cannot be that σ was an informative PBE. Furthermore, we can note that any small perturbation of the utilities in $\tilde{u}$, and thus u, still have the same properties as above. Hence this property holds for an open set of profiles of utility functions.

□

We can provide a similar argument as Proposition 8 to show the following case:

Proposition 9. *If $\succ$ satisfies (BB, BN, NB, NN), then for all full-support priors $\mu_0 \in \Delta(\Theta)$ there exists (u_S, u_R) consistent with $\succ$ such that there does not exist σ informative. Furthermore, this holds for some open set of profiles of utility functions.*

Proof. By Lemma 1, we can just consider the case where $|\Theta| = 4$. Let $\tilde{u}$ be the same as that in Figure 15. Define $u : \Theta \to \mathbb{R}$ as follows:

$$u(a, \theta) = \frac{\tilde{u}(a, \theta)}{\tilde{\mu}_0(\theta)}.$$

Note that the babbling PBE σ_0 gives ex-ante utility of 2 to the Receiver, since B is taken in such an PBE. To see that there is no informative PBE, assume for contradiction that there is an informative PBE σ. By Lemma 2, let m^a induce α^a for $a \in A$. By the same argument in Proposition 8, we find that m^a is sent by θ_{aB} and θ_{aN}, and S at θ_{aB} and θ_{aN} share a message distribution. This give m^B or m^N, and R takes action B regardless. This contradicts $\alpha^N(N) > 0$, and thus there is no informative PBE. Furthermore, we can note that any small perturbation of the utilities in $\tilde{u}$, and thus u, still have the same properties as above. Hence this property holds for an open set of profiles of utility functions. □

C Characterization of Equilibrium Payoffs

Theorem 7. *Fix (Θ, M, A, μ_0, u) such that $|A| = 2$ and μ_0 has full support. The following hold:*

1. *(General PBE) Define*

$$\mathcal{U}^0(x) = \{(u_S, u_R^0) \in \mathbb{R}^2 | u_S \in [u_S^L, x]\}.$$

The set of PBE payoff profiles is either

   (a) $\mathcal{U}^0(u_S^H)$,

   (b) $\mathcal{U}^0(u_S^*)$, *or*

   (c) $\mathcal{U}^0(u_S^H) \cup \{u^*\}$ *where* $u_R^* > u_R^0$ *and* $u_S^H < u_S^*$.

2. *(Babbling PBE) The set of babbling PBE payoff profiles is* $\{(u_S, u_R^0) \in \mathbb{R}^2 | u_S \in [u_S^L, u_S^H]\}$.

3. *(Informative PBE) Suppose there is an informative PBE, σ, and let (u_S, u_R) be its payoff profile.*
   (a) i. *If $u_R = u_R^0$, then the set of informative PBE payoff profiles is $\{(\tilde{u}_S, u_R^0) \in \mathbb{R}^2 | \tilde{u}_S \in (u_S^L, u_S^*]\}$.*
   ii. *If $u_R > u_R^0$, then the set of informative PBE payoff profiles is $\{u^*\}$.*
   (b) *If the Receiver has a unique ex-ante optimal action, i.e., $|A_R^0| = 1$, then $u_S > u_S^H$ $(= u_S^L)$*
   (c) i. *If the Sender has a unique ex-ante optimal action, i.e., $|A_S^0| = 1$, then $u_S > u_S^L$.*
   ii. *If the Sender is ex-ante indifferent between the two actions, i.e., $|A_S^0| = 2$, then $u_S \geq u_S^H$ $(= u_S^L)$.*
4. *All PBE are equivalent to a PBE that use at most two messages with positive probability, i.e., for all PBE σ there exists PBE σ' such that $\rho_\sigma = \rho_{\sigma'}$ and $| \cup_{\theta \in \Theta} supp(\sigma'_S(\theta))| \leq 2$.*

Theorem 7 begins by providing a general characterization of the set of PBE payoff profiles. In particular, part 1 shows that this set can take three possible structures: firstly, it could just be the set of babbling PBE payoff profiles (Part 1a). Second, the set could be the payoff profiles such that the Receiver gets the babbling payoff u_R^0 and the Sender receives any payoffs in a range starting from her worst babbling payoff u_S^L to her best ex-post payoff u_S^* (Part 1b). The final case is that the set has two components - the babbling payoff profiles as well as the Sender optimal ex-post payoff u^* - but it must be that both agents' payoffs in u^* strictly improve over any babbling payoff (Part 1c). A notable consequence of Part 1 is that all equilibria are *Pareto-ranked*. That is, there does not exist PBE σ and σ' such that the Sender strictly prefers σ over σ', and the Receiver strictly prefers σ' over σ.

Parts 2 and 3 provide a specific description of the two main elements of the set of payoff profiles: the payoff profiles from the babbling and informative PBE. Part 2 concerns the former and provides a rather stratightforward characterization. In this case, the set of payoff profiles must be convex because if the set has multiple elements, then the Receiver is indifferent between the two actions, and this implies that the Receiver is indifferent across any mixture over the two actions.

More nontrivial is part 3, which first shows in part 3(a)i that if there is an informative PBE where the Receiver gets his babbling payoff, then the set of informative PBE payoff

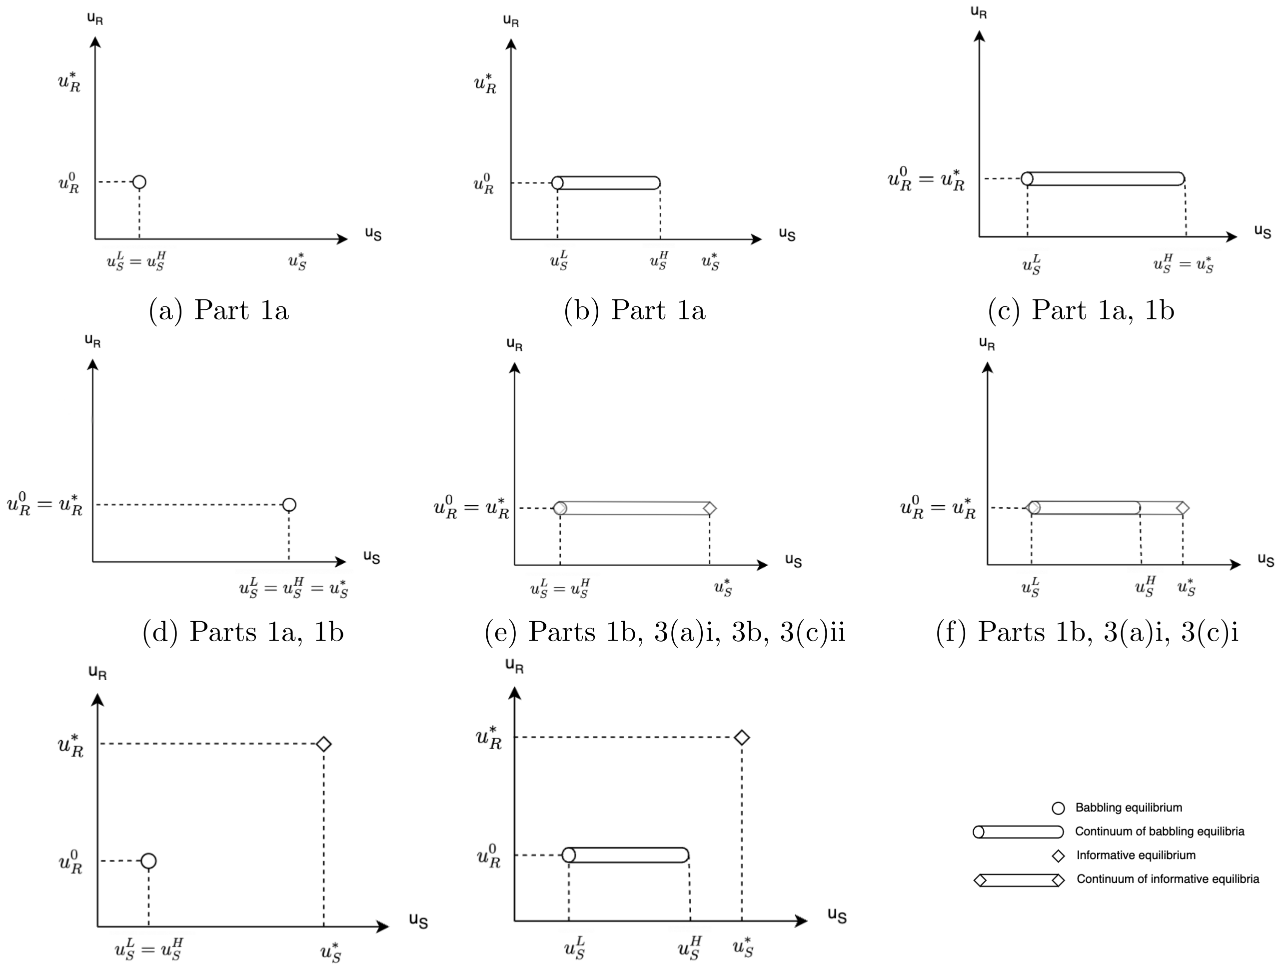

(a) Part 1a (b) Part 1a (c) Part 1a, 1b

(d) Parts 1a, 1b (e) Parts 1b, 3(a)i, 3b, 3(c)ii (f) Parts 1b, 3(a)i, 3(c)i

(g) Parts 1c, 3(a)ii, 3b, 3(c)ii (h) Parts 1c, 3(a)ii, 3(c)i (i) Legend

Figure 17: Possible PBE payoffs for binary-action cheap talk games, with the relevant parts of Theorem 7 listed. Note that part 2 applies to all examples. Example games that generate such payoffs can be found in Figure 20 in the Appendix.

Prior		B	N
$\frac{1}{4}$	θ_{11}	1, 1	0, 0
$\frac{1}{4}$	θ_{12}	1, 0	0, 1
$\frac{1}{4}$	θ_{21}	−1, 1	1, 0
$\frac{1}{4}$	θ_{22}	−1, 0	1, 1

(a) Game

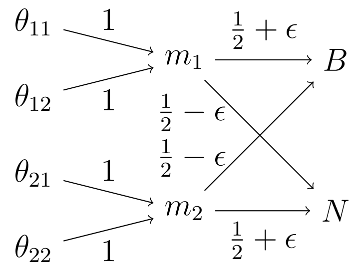

(b) PBE

Figure 18: A game with an informative PBE where the Sender is worse off than some babbling PBE.

profiles is exactly all payoff profiles where the Receiver attains his babbling payoff, and the Sender gets any payoff in $(u_S^L, u_S^*]$. On the other hand, part 3(a)ii states that if the Receiver's payoff strictly improves in some informative PBE over u_R^0, then all informative PBE give the same payoff u^*.

Based on the uniqueness of the ex-ante optimal actions for each agent, further structure can be imposed on the set of payoff profiles. Part 3b consider the case when the Receiver has a unique ex-ante optimal action. There is a single babbling PBE payoff profile in this case, and it is shown that in any informative PBE, the Sender must strictly improve upon her payoff from the babbling PBE. The condition of uniqueness is necessary for the conclusion of the claim to hold. To see this, consider the following example.

Example 7 (Sender prefers babbling to informative PBE)**.** Consider the game in Figure 18. The following is an informative PBE:

$$\sigma_S(\theta_{ij}) = \delta(m_i)$$

$$\sigma_R(m_1) = \begin{cases} B & \text{with probability } \frac{1}{2}+\epsilon \\ N & \text{with probability } \frac{1}{2}-\epsilon \end{cases}$$

$$\sigma_R(m_2) = \begin{cases} N & \text{with probability } \frac{1}{2}+\epsilon \\ B & \text{with probability } \frac{1}{2}-\epsilon \end{cases}$$

where $m_1, m_2 \in M$ are distinct and $\epsilon \in (0, \frac{1}{6})$. The ex-ante utility of the Sender is $\frac{1+6\epsilon}{4}$. The Sender gets ex-ante utility $\frac{1}{2}$ from the babbling PBE such that N is taken with probability 1, and this payoff is greater than the ex-ante utility of the informative PBE. △

This example shows that the Sender can receive a strictly lower payoff in an informative

PBE than in a babbling PBE. However, part 3c of Theorem 7 shows that the Sender is never strictly worse off in a informative equilibrium than in all babbling PBE. In particular, part 3(c)i states that if the Sender has a unique ex-ante optimal action, then she is always strictly better off than in the worst babbling PBE. Part 3(c)ii concerns the case where the Sender is ex-ante indifferent over actions and hence there is a unique babbling payoff for her (her best and the worst babbling payoffs are the same as each other). It shows that the Sender's payoff in an informative equilibrium is weakly greater than this unique babbling payoff. Furthermore, if the Sender is ex-ante indifferent over the two actions, then they are also indifferent to any distribution over actions. Hence, the Sender is also indifferent between all babbling PBE. Overall, part 3c implies that any informative PBE cannot give worse payoffs to the Sender than in all babbling PBE.

[Note for myself]

To prove part 1 of Theorem 7, we begin by noting that the set of PBE payoff profiles can be exactly as in part 1a. This follows because this is the set of babbling PBE payoff profiles, which always exist for any game. Furthermore, any other payoff that might exist must come from an informative PBE, hence in games where such PBE do not exist, the set of PBE payoff profiles are just the babbling payoff profiles (we can find examples of profiles of utility functions such that only babbling PBE exist). Now, consider the case where there are informative PBE. We consider two possibilities: there is an informative PBE that improves the Receiver's payoff over his babbling payoff, or there is no such PBE. According to parts 3(a)i and 3(a)ii, u^* is a PBE payoff profile under both possibilities. However, they cannot hold simultaneously as the former asserts that $u^*_R > u^0_R$, whereas the latter assumes that $u^*_R = u^0_R$. Thus, either the set of payoff profiles are as in 1b, or the set of payoff profiles are the babbling payoff profiles and u^*, where $u^*_R > u^0_R$ as in 1c.

To see why part 3a holds, first consider part 3(a)i. Assume there is an informative PBE where the Receiver gets u^0_R. Without loss, we let B be some action played with positive probability in a babbling PBE. Let α^B and α^N be the action distributions played in the informative PBE. We can recall that states such that the Sender strictly prefers $a \in A$ induce α^a. Note that it cannot be the case that $\alpha^B(B) = \alpha^N(B)$, since this is an informative equilbrium, hence it must be that $\alpha^B(B) > \alpha^N(B)$. If both distributions have full support over actions, then the Receiver is indifferent between actions after any message, and it must be that any action distribution $\tilde{\alpha}^a$ that satisfies $\tilde{\alpha}^B(B) > \tilde{\alpha}^N(B)$ can be induced by a informative PBE. Furthermore, this allows for the Sender's ex-post optimal payoff, hence any payoff for the Sender in $(u^L_S, u^*_S]$ is possible. Similarly, if $\alpha^B(B) = 1$ and

α^N has full support, then the Receiver is indifferent between all actions when m^N that induces α^N is sent. Thus any $\tilde{\alpha}$ such that $\tilde{\alpha}^B(B) = 1$ and $\tilde{\alpha}^N(B) < 1$ gives payoff profiles in the same set. Finally, if $\alpha^B(B) = \alpha^N(N) = 1$, then since the Receiver's payoff is u_R^0, they must be indifferent to playing B with positive probability after observing m^N since his payoff does not change. Thus we can use a similar construction of $\tilde{\alpha}^a$.

On the other hand, when the Receiver's payoff from a PBE is greater than his babbling payoff, as in part 3(a)ii, we find that u^* is the payoff profile induced. Note that given parts 1 and 3(a)i, this implies part 3(a)ii, however for exposition we prove 3a to show 1. First we can show that the distribution of actions induced in $\sigma_S(\theta_a)$ must be $\alpha^a = \delta(a)$ for $a_S^{\theta_a} = a$ and $a \in A$. To see why, note that if this is not the case then there is an action $a \in A$ such that for all messages m sent with positive probability, $\sigma_R(a|m) > 0$. However, if after any message the Receiver is indifferent between playing a or playing $\sigma_R(a|m)$, then it must be that the utility the Receiver achieves is the same as a babbling PBE. In particular, it must be that a is an action played with probability one in some babbling PBE, and thus the Receiver's payoff u_R is equal to u_R^0. This would contradict the assumption that $u_R > u_R^0$, hence it must be that $\alpha^a = \delta(a)$ for all $a \in A$. Furthermore, for this to be a PBE, it must be that if the Sender strictly prefers an action a to a' in state $\theta \in \Theta$, then all messages they send must induce α^a. Thus in all states the Sender gets her optimal payoff, and the ex-ante payoff profile is u^*.

Part 3b follows directly from 3a by first noting that if the Receiver has a unique ex-ante optimal action, then $u_S^L = u_S^H$. If there is an informative PBE, the Receiver's payoff in this PBE has either strictly improved over the babbling payoff or stayed the same. In the former case, by part 3(a)ii, the Sender gets her ex-post optimal payoff u_S^* this informative PBE. This must be different to u_S^H as the distribution of actions taken for each state in a PBE that induces u^* places probability one on the Sender's preferred action, and thus if $u_S^H = u_S^*$ then the PBE must take the same action in every message sent with positive probability. This contradicts the PBE being an informative PBE. In the latter case, where the Receiver's payoff is the same as a babbling PBE, part 3(a)i shows that every informative PBE must give the Sender payoff strictly greater than u_S^L, which is equal to u_S^H.

If the babbling PBE is unique, as in part 3b, then this also implies part 3c. Otherwise, we can note that in any informative PBE the Receiver's distribution over actions for all messages has full support, i.e., both actions are played with positive probability. Given that both actions induce a babbling PBE, if the Sender has a strict preference ex-ante over actions, then the informative PBE must be strictly better than one of the actions, and thus also strictly better than the babbling PBE it induces. If the Sender is indifferent, then the informative PBE is weakly better than one of the pure babbling PBE.

The explanation for the final claim, part 4, is as in the main text.

C.1 Proof of Theorem 7

Note that the arguments for parts 1 of Theorem 7 follows directly from the claims in parts 3 and 2, as explained in Section 3.2. Further observe that part 2 is direct from the definition of a babbling PBE. Consider part 3(a)ii, which is shown in the following proposition:

Proposition 10. *If $\sigma \succ_R \sigma_0$ for a PBE σ and babbling PBE σ_0, then σ is an informative PBE and σ induces the payoff profile u^*.*

Proof. Assume for contradiction that the Receiver is not indifferent to some babbling payoff, hence $\sigma \succ_R \sigma_0$ for σ_0 a babbling PBE. Note that this must be an informative PBE as the Receiver gets the same payoff as u_R^0 if it were a babbling PBE, which would contradict the assumption that $\sigma \succ_R \sigma_0$. Recall that by Lemma 2, there exists multiple action distributions take in PBE. Let α^N and α^B be said distributions, and note that for all $\theta \in \Theta$ and for all $m \in \text{supp}(\sigma_S(\theta))$

$$a \succ_S^\theta a' \implies \sigma_R(m) = \alpha^a$$

for $a, a' \in A$. As such, it is either the case that $\alpha^a = \delta(a)$ for all $a \in A$ or there exists $a \in A$ such that for all $m \in M' = \cup_{\theta\in\Theta}\text{supp}(\sigma_S(\theta))$ $\sigma_R(a|m) > 0$. The second case follows by leveraging the fact that $|A| = 2$. In the first case, we can note that σ induces the ex-post optimal utility for the Sender. Thus there cannot exist σ' a strategy profile strictly preferred by the Sender ex-ante. This is a contradiction, hence the other case must hold. However, if this holds then Lemma 3 implies that σ is a babbling PBE:

Lemma 3. *If there exists $a \in A$ such that for all $m \in M$ sent with positive probability $\sigma_R(a|m) > 0$, then $\sigma \sim_R \sigma_0$ for some σ_0 a babbling PBE and a is an action played with positive probability in some babbling PBE.*

Proof. We define a distribution of posteriors $\tau \in \Delta(\Delta(\Theta))$ to be Bayes plausible if

$$\sum_{\mu\in\text{Supp}(\tau)} \mu\tau(\mu) = \mathbb{E}_{\mu\sim\tau}\mu = \mu_0.$$

A Sender's strategy σ_S induces τ if $\text{Supp}(\tau) = \{\mu_m\}_{m \in M}$ where

$$\mu_m(\theta) = \frac{\sigma_S(m|\theta)\mu_0(\theta)}{\sum_{\theta' \in \Theta} \sigma_S(m|\theta')\mu_0(\theta')} \; \forall m \in M, \theta \in \Theta$$

and we have that

$$\tau(\mu) = \sum_{m:\mu_m = \mu} \sum_{\theta \in \Theta} \sigma_S(m|\theta)\mu_0(\theta) \; \forall \mu \in \Delta(\Theta)$$

Let $\tau \in \Delta(\Delta(\Theta))$ be the set of induced posterior distributions corresponding to σ. Observe that $\tau(\mu)$ is Bayes plausible. We have that

$$\begin{aligned}
\mathbb{E}_{\theta \sim \mu_0}\left[u_R(a, \theta)\right] &\leq \max_{a'} \mathbb{E}_{\theta \sim \mu_0}\left[u_R(a', \theta)\right] \\
&= \max_{a'} \mathbb{E}_{\mu \sim \tau}\mathbb{E}_{\theta \sim \mu}\left[u_R(a', \theta)\right] \\
&\leq \mathbb{E}_{\mu \sim \tau} \max_{a'} \mathbb{E}_{\theta \sim \mu}\left[u_R(a', \theta)\right] \\
&= \mathbb{E}_{\mu \sim \tau}\mathbb{E}_{\theta \sim \mu}\left[u_R(a, \theta)\right] \\
&= \mathbb{E}_{\theta \sim \mu_0}\left[u_R(a, \theta)\right]
\end{aligned}$$

where we observe that $\mu \sim \tau$ is some posterior distribution after observing $m \sim \sigma_S(\theta)$ for some θ, and that the following must hold since a occurs with positive probability for every message (and thus posterior distribution):

$$a \in \arg\max_{a'} \mathbb{E}_{\theta \sim \mu}\left[u_R(a', \theta)\right].$$

Thus we find that $\sigma \sim_R \sigma_0$. Furthermore, we can observe that a gives the same payoff as a babbling PBE to R, hence it is also a babbling PBE. □

This gives another contradiction, implying that our assumption is false. Thus the equilibrium action distribution is ex-post optimal for the Sender. As the Sender has strict preference is every state, any strategy profile that induces her ex-post optimal payoff for the Sender induces a unique payoff for the Receiver. Hence, σ induces the payoff profile u^*.

□

Note that though the statement of part 3(a)ii assumes that σ is an informative PBE, under the assumption that the Receiver's payoff strictly improves over his babbling payoff,

σ is necessarily an informative PBE.

Now we prove part 3(a)i of Theorem 7:

Proposition 11. *If $\sigma \sim_R \sigma_0$ for an informative PBE σ and babbling PBE σ_0, then the set of informative PBE payoff profiles is $\{(u_S, u_R^0) \in \mathbb{R}^2 | u_S \in \left(u_S^L, u_S^*\right]\}$*

Proof. Let σ be an informative PBE such that the Receiver gets a babbling payoff. Recall from Lemma 2 that there are two distinct action distributions α^B and α^N played in a PBE, such that states where the Sender prefers B induce α^B and those where they prefer N induce α^N. Without loss of generality let $\alpha^B(B) > \alpha^N(B)$. We proceed by considering various cases. Let m^B and m^N be messages sent by the Sender that induce α^B and α^N respectively.

First suppose $0 < \alpha^B(B) < 1$ and $0 < \alpha^N(B) < 1$. Given this, the Receiver is indifferent between B and N given either message, and any $\tilde{\sigma}$ such that the induced action distribution $\tilde{\alpha}^B, \tilde{\alpha}^N$ satisfy $\tilde{\alpha}^B(B) > \tilde{\alpha}^N(B)$ is a PBE. This is because it preserves Receiver optimality, and the Sender at each state has no profitable deviation. Furthermore, if the Receiver is indifferent after both messages, then this induces the ex-ante babbling payoff for the Receiver. Note that this allows for Sender's ex-post optimal payoff, since we can have $\tilde{\alpha}^B(B) = 1$ and $\tilde{\alpha}^N(B) = 0$. Note that there are some informative PBE with the same payoff as a babbling equilibirum, as per Lemma 3. Hence there are a continuum of informative equilibria, from but not including the worst babbling payoff for the Sender to the ex-post optimal payoff for the Sender, where the receiver is indifferent between them.

Now consider the case where $\alpha^B(B) = 1$ and $0 < \alpha^N(B) < 1$. R is indifferent between B and N given m^N. Thus for any $\tilde{\sigma}$ that induces $\tilde{\alpha}^B(B) = 1$ and any $\tilde{\alpha}^N(B) < 1$, we have an informative PBE that induces the babbling payoff. Furthermore, we can achieve the ex-post optimal sender payoff by having $\tilde{\alpha}^N(B) = 0$. Thus again we have a continuum of informative PBE payoff profiles that span from a babbling payoff for the Sender to ex-post optimal payoff. Observe that the case where $\alpha^N(B) = 0$ and $0 < \alpha^B(B) < 1$ is similar.

The final case to consider is where $\alpha^B(B) = 1$ and $\alpha^N(B) = 0$. Without loss of generality let B be one of the babbling action. If the Receiver gets the babbling payoff under σ, then if they play B with positive probability given m^N his payoff does not change. Hence we can apply the same construction of a continuum as in the case where $\alpha^B(B) = 1$ and $0 < \alpha^N(B) < 1$.

To see that there cannot be other informative PBE payoff profiles, that (u_R^0, u_S^L) cannot be an informative payoff as this is only induced by a babbling PBE. Furthermore, if $u^R > u_R^0$ is a payoff to the Receiver in some informative PBE, then by 3(a)ii such a PBE must induce a payoff profile u^*. However, since the Sender has strict preference in every state, if a PBE induces the payoff u_S^* to the Sender, then the Receiver must have payoff u_R^*. As shown above, u_S^* is an informative PBE payoff in the set described, and thus it must be that $u_R^* = u_R^0$. Hence this payoff is covered by the set described.

□

For the parts 3b and 3c of Theorem 7, we assume without loss of generality that B is the action taken in some babbling PBE. This is possible as either R takes a single pure action in a babbling PBE, or mixes. If they mix, they must be indifferent and thus all pure actions are babbling. If they take a single pure action, we can refer to this as B without loss. Now we show both parts together.

Proposition 12. *If $|A_R^0| = 1$, then $\sigma \succ_S \sigma^0$ for all informative PBE σ and babbling PBE σ^0. If $|A_S^0| = 1$, then for all informative PBE σ there is some babbling PBE σ^0 such that $\sigma \succ_S \sigma^0$. If $|A_S^0| = 2$, then $\sigma \succeq_S \sigma^0$ and $\sigma^0 \sim_S \tilde{\sigma}^0$ for all informative PBE σ and babbling PBE σ^0 and $\tilde{\sigma}^0$.*

Proof. Assume there exists an informative PBE σ. First consider the following the case: suppose there exists a message $m \in M$ that is sent with positive probability, such that $\sigma_R(m) = \delta(B)$. Thus the Sender at each state must get at least her babbling payoff. Note that there must exist some $\theta \in \Theta$ such that S sends a message $m' \in M$ and $\sigma_R(N|m') > 0$. Otherwise every message sent induces the action B with probability 1, which is a contradiction since this would be the babbling PBE. At this θ, it must be the case that $N \succ_S^\theta B$. If the Sender at every state gets a weakly better payoff than babbling, and at some state gets a strictly better payoff, then the Sender gets higher utility ex-ante since all states occur with positive probability. This proves all claims if the assumption holds, that there is a message which induces B.

On the other hand, assume there does not exist a message that induces B with probability 1. It must be the case that after every message $m \in M$, $\sigma_R(N|m) > 0$. By Lemma 3, N is a babbling PBE as well, which is a contradiction if we assume that there is a unique, pure babbling PBE. This proves the first claim.

As for the second claim, if the Sender has a uniquely optimal ex-ante action, assume without loss that it is B. If B and N are babbling PBE, then any mixture is a babbling PBE. Hence the mixture α induced by σ, is a babbling PBE, and thus S is indifferent to it ex-ante. Note that α places positive probability on all actions, as it is a non-degenerate mixture of two action distributions, at least one of which is non-degenerate (as per Lemma

2 due to the informativeness of σ). For action distributions of the form $\alpha_\beta = \beta\delta(B) + (1-\beta)\delta(N)$, where $\beta \in [0,1]$, the Sender's ex-ante preference α_β must be increasing in β since its ex-ante optimal action is B uniquely. As α is strictly preferred to all α_β for $\beta < \bar{\beta}$ for some $\beta \in (0,1)$, if it is non-degenerate, then α is ex-ante strictly preferred by S to some babbling distribution since all distributions are babbling. If the Sender does not have a uniquely optimal ex-ante action, as in the third claim, then any action distribution is optimal. Thus, we can say that σ is weakly preferred ex-ante to all babbling PBE by the Sender, and they are indifferent to all babbling PBE.

We now consider the final claim in part 4. Let σ be some PBE. First note that if there are two messages that induce the same distribution over actions, that is there are distinct $m, m' \in \cup_\theta \text{supp}\sigma_S(\theta)$ such that $\sigma_R(m) = \sigma_R(m')$, we can construct a new equivalent PBE σ' that combines the messages into a single message. Formally, let $\sim_\sigma$ be an equivalence relation on $\cup_\theta \text{supp}\sigma_S(\theta)$ where $m \sim_\sigma m'$ if $\sigma_R(m) = \sigma_R(m')$. Let the equivalence classes be given by $[m] \in M_\sigma$, M_σ is the induced partition over messages sent with positive probability. We treat $[m]$ as a representative message when appropriate. Hence define σ as follows: for all $\forall [m] \in M_\sigma$,

$$\sigma'_S([m]|\theta) = \sum_{m' \in [m]} \sigma_S(m|\theta)$$
$$\sigma'_R([m]) = \sigma_R([m])$$

Observe that $\rho_\sigma = \rho_{\sigma'}$. To see σ' is a PBE, first note the Sender has no incentive to deviate given that the messages they send induce the same distribution over actions at each state as in σ. Furthermore, we can show that this strategy is still optimal for the Receiver.

When the Receiver observes $[m]$ maximizes of $a \in A$ the following

$$\begin{aligned}
\mathbb{E}_{\theta\sim\mu_0}\left[u_R(a,\theta)|[m]=\sigma'(\theta)\right] &= \sum_\theta \mathbb{P}(\theta|[m]=\sigma'(\theta))u_R(a,\theta)\\
&= \sum_\theta \frac{\sigma'([m]|\theta)\mu(\theta)}{\sum_{\theta'}\sigma'([m]|\theta')\mu(\theta')}u_R(a,\theta)\\
&= \frac{1}{\sum_{\theta'}\sigma'([m]|\theta')\mu(\theta')}\sum_\theta \sigma'([m]|\theta)\mu(\theta)u_R(a,\theta)\\
&= \frac{1}{\sum_{\theta'}\sigma'([m]|\theta')\mu(\theta')}\sum_\theta\sum_{m'\in[m]} \sigma_S(m'|\theta)\mu(\theta)u_R(a,\theta)\\
&= \frac{1}{\sum_{\theta'}\sigma'([m]|\theta')\mu(\theta')}\sum_{m'\in[m]}\sum_\theta \sigma_S(m'|\theta)\mu(\theta)u_R(a,\theta)
\end{aligned}$$

Observe that for each θ where $\sigma_S(m'|\theta) > 0$, because σ satisfied Receiver optimality, then for all $a \in \sigma_R(m')$ it is that $a \in \arg\max \sum_\theta \sigma_S(m'|\theta)\mu(\theta)u_R(a,\theta)$. Hence $\sigma'([m])$ satisfies Receiver optimality.

Hence we can assume with loss of generality that every message sent with positive probability in a PBE σ induces a different distribution over actions. Now assume for contradiction that there are at least three messages in the support of σ_S (across all θ). Let them be $\sigma_R(m_i)$. Let θ be some state such that Sender sends some message $\hat{m}$ with positive probability and $\sigma_R(B|\hat{m}), \sigma_R(N|\hat{m}) > 0$. This must exist as there only two actions and at least three distributions. Given that the distributions are distinct, let us rank the distributions as $\sigma_R(B|m_i) > \sigma_R(B|m_j)$ if $i < j$. Note that $\hat{m} = m_i$ for $i > 1$. Because the Sender has strict preferences, assume with loss that they prefer B to N at state θ. Furthermore, because there only two actions, the Sender at θ always prefers distributions over actions with higher probability on B than with lower probability. Hence the Sender strictly prefers $\sigma_R(m_1)$ over all $\sigma_R(m_i)$ for $i > 1$. As such, the Sender at θ has an incentive to deviate from sending message $\hat{m}$ with positive probability to sending message m_1 with probability 1. This is a contradiction with σ being a PBE. Thus it must be that σ uses at most two messages. □

D Proof of Theorem 4

If every signal distribution is the same, i.e., the $\sigma_S(\theta_1) = \sigma_S(\theta_2) = \hat{\sigma} \in \Delta(M)$, then this gives the babbling payoff since $\mu_0(\theta|m) = \mu_0(\theta)$ for all $\theta \in \Theta$ and messages $m \in M$ sent with positive probability. Let $M_i = \text{supp}\sigma_S(\theta_i)$. Now we consider the following cases: the support of each signal distribution has a non-overlapping point each, one support is a subset of the other, or the support is the same. Define $a^*_R(\theta_i) = \arg\max_{a\in A} u_R(a, \theta_i)$ and $\delta(x)$ refers to a Dirac distribution with probability 1 on x.

In the first case, we have that there exists $m_1 \in M_1 - M_2$ and $m_2 \in M_2 - M_1$. This leads to the posterior $\mu(\theta_i|m_i) = 1$, hence the Receiver takes the optimal action for θ_i: $\sigma_R(m_i) = \delta(a_R(\theta_i))$. There are two possible cases, either $a_R(\theta_i)$ are the same or not. If they are different, then for $j \neq i$ we have that $a_R(\theta_i) \succ_S^{\theta_j} a_R(\theta_j)$. Hence there is a profitable deviation for the Sender at θ_1 and θ_2. If they are the same then in every state of the world the ex-post Receiver optimal outcome is the same. Thus the babbling outcome is the same.

In the second case, we have that without loss $M_2 \subset M_1$. Consider $m_1 \in M_1 - M_2$. Thus $\mu(\theta_1|m_1) = 1$, and $\sigma_R(m_1) = \delta(a_R(\theta_1))$. Now we have that either for all $m_2 \in M_2$, $\sigma_R(m_2) = \delta(a_R(\theta_1))$ or not. In the former case, we have a contradiction by Lemma 2 since there is only one action distribution taken in PBE. If the latter case, then for some $m_2 \in M_2$, there exists $a \in \text{supp}\sigma_R(m_2)$ such that $a \neq a_R(\theta_1)$. Thus $\mathbb{E}_{a\sim\sigma_R(m_2)} u_S(a, \theta_1) > u_S(a_R(\theta_1), \theta_1)$ since the Sender's preferences are strict and $a_R(\theta_1)$ is the worst outcome for the Sender by anti-coordination. Thus there is a profitable deviation for the Sender at θ_1.

Now consider the case where $M_1 = M_2 = M$. Recall that we have already considered the case where the at each state the Sender has the same strategy. Thus we can consider the Sender to have different strategies at each state. Since there are only two states, there exists m_1, m_2 such that for $j \neq i$ $\mu_0(\theta_i|m_i) > \mu_0(\theta_i)$ and $\mu_0(\theta_j|m_i) < \mu_0(\theta_j)$.

As such, we have that $\sigma_R(m_i) \preceq_S^{\theta_i} \sigma_R(m_j)$ for all $i \neq j$. To see why, consider the following intuitive lemma:

Lemma 4. *If $a(\alpha) \in \arg\max_{x\in A} \alpha u(x) + (1-\alpha)v(x)$ for $\alpha \in (0, 1]$, then $\alpha \geq \alpha'$ implies that $u(a(\alpha)) \geq u(a(\alpha'))$.*

Proof. Let $\alpha \geq \alpha'$ such that $\alpha > 0$, and let $a = a(\alpha)$ and $a' = a(\alpha')$.Assume for contradiction that $u(a') > u(a)$. By optimality of a and a' for α and α' respectively, we have that

$$\alpha(u(a) - u(a')) + (1-\alpha)(v(a) - v(a')) \geq 0$$
$$\alpha'(u(a') - u(a)) + (1-\alpha')(v(a') - v(a)) \geq 0$$

Because $\alpha \geq \alpha'$, the last equation is such that $\alpha(u(a') - u(a)) + (1-\alpha')(v(a') - v(a)) \geq 0$. Adding this to the first equation gives $(v(a) - v(a'))(a' - a) \geq 0$. Hence we have that $v(a) \leq$

$v(a')$. This gives the following contradiction: $\alpha u(a')+(1-\alpha)v(a') > \alpha u(a)+(1-\alpha)v(a)$. $\square$

Thus we have that for all $a_i \in \text{supp}\sigma_R(m_i)$, we have that $u_R(a_i, \theta_i) \geq u_R(a_j, \theta_i)$. By the perfect misalignment structure on preferences, we have that $u_S(a_i, \theta_i) \leq u_S(a_j, \theta_i)$. Since this is true for all elements of the support, it must be that $\sigma_R(m_i) \preceq_S^{\theta_i} \sigma_R(m_j)$ for all $i \neq j$.

If for some a_i, a_j we have that $u_R(a_i, \theta_i) > u_R(a_j, \theta_i)$, then $u_S(a_i, \theta_i) < u_S(a_j, \theta_i)$. Thus we have a profitable deviation by the Sender at θ_i to m_j since $\sigma_R(m_i) \prec_S^{\theta_i} \sigma_R(m_j)$ for all $i \neq j$. Hence we can assume that $u_R(a_i, \theta_i) = u_R(a_j, \theta_i)$, meaning that $\sigma_R(m_1) \sim_R^{\theta_i} \sigma_R(m_2)$ for all i. Note that this has to be true for all $m_1, m_2 \in M$. As such we can pick some m_1 since the Receiver is indifferent between all messages at every state:

$$\begin{aligned}\max_{a\in\mathcal{A}} \mathbb{E}_{\theta\sim\mu_0} u_R(a,\theta) &\leq \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m)} u_R(a,\theta) \\ &= \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{m\sim\sigma_S(\theta)}\mathbb{E}_{a\sim\sigma_R(m_1)} u_R(a,\theta) \\ &= \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{a\sim\sigma_R(m_1)} u_R(a,\theta) \\ &\leq \max_{a\in\mathcal{A}} \mathbb{E}_{\theta\sim\mu_0} u_R(a,\theta).\end{aligned}$$

Hence the Receiver receives a babbling utility.

E Proof of Proposition 1

First we show that the Receiver mixes on path. Let σ be an informative PBE. Assume for contradiction that for all states $\theta \in \Theta$ and $m \in M$ such that $\sigma_S(m|\theta) > 0$, we have that $\sigma_R(m) = \delta(a)$ for $a \in A$. Since σ is an informative PBE, we have that there is some m^B and m^N that induce α^B and α^N as in Lemma 2. Note that $\alpha^B \neq \alpha^N$. By assumption that there is no mixing and that σ is informative, we have that $\alpha^a = \delta(a)$ for $a \in A$. Without loss of generality, let B be an action played in a babbling PBE σ^0. By perfect misalignment, the Receiver either takes the babbling action or some action that is optimal for the Sender, hence it is strictly worse for the Receiver. Thus $\sigma \prec_R \sigma^0$, which is a contradiction as all PBE must be ex-ante weakly better for the Receiver than babbling. Thus the assumption that the Receiver does not mix in a PBE is false.

Now we show that Sender pools in all states. Assume for contradiction that there is a state θ that does not pool in a PBE σ. Let $M_\theta = \text{supp}(\sigma_S(\theta))$. Given this, we have that $M_\theta \cap \cup_{\theta' \neq \theta} M_{\theta'} = \emptyset$. Thus $\forall m \in M_\theta$, $\mu(\theta|m) = 1$. By perfect misalignment of the preference profile, $\sigma_R(m) = \delta(a_\theta)$, where $a_\theta = \arg\min_a u_S(a, \theta)$. Since this is a non-babbling PBE, there must exist m' such that $\sigma_R(m') \neq \sigma_R(m)$. If a is the worst action for θ, and preferences are strict, then $\sigma_R(m') \succ_S^\theta \sigma_R(m)$. Furthermore $m' \notin M_\theta$ as $\forall m_1, m_2 \in M_\theta$, $\sigma_R(m_1) \sim_S^\theta \sigma_R(m_2)$. Thus there is a profitable deviation by the Sender at θ from $\sigma_S(\theta)$ to $\delta(m')$. This contradicts σ being a PBE.

F Additional Examples

Prior		L	M	R
$\frac{2}{3}$	θ_1	1, 3	2, 2	0, 0
$\frac{1}{3}$	θ_2	−6, −3	0, 0	−10, 2

Figure 19: Example of PBE not being Pareto ranked in a game with $|\Theta| = 2$.

Example 8. *Figure 19 provides a binary state example that, similar to Figure 5, illustrates how some of the claims in Theorem 7 no longer hold when the set of actions are no longer binary.* △

G Proof of Theorem 5

Note the following observation:

Observation 8. *A preference profile is weakly misaligned if and only if it satisfies one of the following:* (BB, BN), (NB, NN), (BB, BN, NB), (BB, BN, NN), (BN, NB, NN), (BB, NB, NN), (BB, BN, NB, NN).

First we prove part 1 and 2 of Theorem 5. First note that by a similar argument to Lemma 1, to show some condition $\emptyset \neq C \subseteq \{BB, BN, NB, NN\}$, we only need to show it true for the case where $|\Theta| = |C|$. Furthermore, we can scale utilities by $\mu_0(\theta)$ as needed, so we can assume uniform priors. See Figures 21 and 22 for the example profiles of utility functions that satisfy the various cases required by a preference being weakly misaligned, and note that the other cases follow by symmetry.

Now we prove part 3 of Theorem 5. Note that we only need to consider the following cases, as per the previous observation: BB, BN, NB, NN, (BB, NN), (BB, NB), (BN, NN). We can also see that the following cases are handled by Proposition 5, since for all posteriors the optimal action by the Receiver is identical: BB, BN, NB, NN,

Prior		L	M
$\frac{1}{2}$	θ_1	1, 0	0, 1
$\frac{1}{2}$	θ_2	0, 10	1, 0

(a) Single babbling.

Prior		L	M
$\frac{1}{2}$	θ_1	1, 0	0, 1
$\frac{1}{2}$	θ_2	0, 1	1, 0

(b) Continuum of babbling.

Prior		L	M
$\frac{1}{2}$	θ_1	1, 1	0, 0
$\frac{1}{2}$	θ_2	0, 0	1, 2

(c) One babbling, one ex-post.

Prior		L	M
$\frac{1}{4}$	θ_1	1, 2	0, 0
$\frac{1}{4}$	θ_2	1, 0	0, 1
$\frac{1}{4}$	θ_3	0, 1	2, 0
$\frac{1}{4}$	θ_4	0, 0	1, 2

(d) Babbling continuum + ex-post (improvement for R).

Prior		L	M
$\frac{1}{2}$	θ_1	1, 1	0, 0
$\frac{1}{2}$	θ_2	1, 1	0, 0

(e) Babbling is ex-post.

Prior		L	M
$\frac{1}{4}$	θ_1	1, 1	0, 0
$\frac{1}{4}$	θ_2	1, 0	0, 1

(f) Babbling continuum up to expost but same for receiver.

Prior		L	M
$\frac{1}{4}$	θ_1	1, 1	0, 0
$\frac{1}{4}$	θ_2	1, 0	0, 1
$\frac{1}{4}$	θ_3	0, 1	2, 0
$\frac{1}{4}$	θ_4	0, 0	1, 1

(g) Continuum babbling + continuum informative but same for receiver.

Prior		L	M
$\frac{1}{4}$	θ_1	1, 1	0, 0
$\frac{1}{4}$	θ_2	1, 0	0, 1
$\frac{1}{4}$	θ_3	0, 1	2, 0
$\frac{1}{4}$	θ_4	0, 0	1, 2

(h) Single babbling + continuum informative but same for receiver.

Figure 20: Games that achieve the PBE payoff profiles in Figure 17.

(BB, NB), (BN, NN), (BN, NB). Hence a babbling PBE would implement the same outcome as any BPE.

The remaining case is (BB, NN). By Proposition 7, we can observe that truth-telling by the Sender results in a unique response by the Receiver, where both agents get their preferred action in every state of the world. This would be the unique BPE, which is also a PBE.

H Dispute Resolution

This section studies the processes of arbitration, mediation and negotiation from Goltsman et al. (2009) within our model, giving proofs of results in the main text. In models with a third-party, we will refer to said third-party via the pronoun "they". Fix a game

	B	N
θ_1	10, 1	0, 0
θ_2	1, 0	0, 2

(a) (BB, BN)

	B	N
θ_1	100, 0	0, $\frac{3}{2}$
θ_2	0, 1	1, 0

(b) (BN, NB)

	B	N
θ_1	10, 1	0, 0
θ_2	1, 0	0, 10
θ_3	0, 1	1, 0

(c) (BB, BN, NB)

	B	N
θ_1	1, 1	0, 0
θ_2	0, 0	10, 1
θ_3	0, 2	1, 0

(d) (BB, NN, NB)

	B	N
θ_1	1, 2	0, 0
θ_2	1, 0	0, 1
θ_3	0, 0	20, 1
θ_4	0, 2	1, 0

(e) (BB, NN, NB, BN)

Figure 21: Examples for proof of Part 1 of Theorem 5.

	B	N
θ_1	1, 2	0, 0
θ_2	2, 0	0, 1

(a) (BB, BN)

	B	N
θ_1	1, 0	0, 1
θ_2	0, 1	1, 0

(b) (BN, NB)

	B	N
θ_1	2, 3	0, 0
θ_2	2, 0	0, 1
θ_3	0, 1	1, 0

(c) (BB, BN, NB)

	B	N
θ_1	1, 1	0, 0
θ_2	0, 0	1, 3
θ_3	0, 1	1, 0

(d) (BB, NN, NB)

	B	N
θ_1	1, 3	0, 0
θ_2	1, 0	0, 1
θ_3	0, 1	1, 0
θ_4	0, 0	1, 3

(e) (BB, NN, NB, BN)

Figure 22: Examples for proof of Part 2 of Theorem 5.

(Θ, M, A, μ_0, u). We denote the **neutral** third party via J for *judge*. A strategy by the Judge is $\sigma_J : \Theta \to \Delta(A)$.

H.1 Arbitration

σ_J is an optimal arbitration rule (OAR) if it solves the following:

$$\max_{p:\Theta\to\Delta(A)} \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{a\sim p(\theta)}\left[u_R(a,\theta)\right]$$

subject to

$$\theta \in \arg\max_{\hat{\theta}\in\Theta} \mathbb{E}_{a\sim p(\hat{\theta})}\left[u_S(a,\theta)\right] \quad \text{(Sender-IC)}$$

Remark 1. *This is equivalent to the definition of arbitration in Goltsman et al. (2009).*

An OAR can be interpreted as a solution whereby the Sender is truthful, and the Receiver commits to an action strategy before observing the message. When comparing between OAR and PBE, we compare their ex-ante payoff profiles. Note that following observation:

Observation 9. *Fix σ_J an OAR and σ^0 a babbling PBE. Then $\sigma_J \succeq_R \sigma^0$.*

We say that σ_J is a babbling if for all $\theta, \theta' \in \Theta$, $\sigma_J(\theta) = \sigma_J(\theta')$. Otherwise it is informative. Note that σ_J a babbling OAR induces the same ex-ante payoff profile as some babbling PBE. Furthermore, if such a babbling OAR exists, then any babbling PBE has a corresponding babbling OAR with the same payoff profile.

First we show the following result:

Theorem 8. *Fix an environment ... TODO: formal. Then the following are true:*

1. *σ_J is an informative OAR if there exists $\alpha^B, \alpha^N \in \Delta(A)$ such that $\alpha^B(B) > \alpha^N(B)$ and if $\theta \in \Theta$ such that $a \succ_S^\theta a'$ for $a, a' \in A$, then $\sigma_J(\theta) = \alpha^a$.*

2. *σ_J is an OAR iff it induces (α^B, α^M) that our solutions to the following optimization problem:*

$$\max_{\alpha^B,\alpha^N\in\Delta(A)} \sum_{a,a'\in A} \mathbb{E}_{\theta\sim\mu_0}\left[\mathbb{E}_{a\sim\alpha^a} u_R(a,\theta)|\theta \in \Theta_{aa'}\right]\mu_0(\Theta_{aa'})$$

   subject to $\alpha^B(B) \geq \alpha^N(B)$. Furthermore, it is an informative OAR iff solves the above optimization without the constraint binding, i.e. $\alpha^B(B) > \alpha^N(B)$.

3. *A Sender-preferred OAR σ_J gives payoff profile u^* or u^0, where u^* corresponds to the Sender-optimal payoff profile and u^0 is the of payoff profile of the Sender-preferred babbling equilibrium.*

4. *If the profile of utility functions u is consistent with a preference profile $\succ$ that is not aligned, then all σ_J are babbling OARs.*

5. *If there exists an informative PBE, then σ_J induces the utility profile u^*.*

6. *If the profile of utility functions u is consistent with a preference profile $\succ$ that is aligned, but there are no informative PBE, then σ_J is a babbling OAR.*

Proof. First, consider the claim that an informative OAR induces two distributions. Since σ_J is not babbling, then there must be $\theta, \theta' \in \Theta$, $\sigma_J(\theta) \neq \sigma_J(\theta')$. It can't be that $\theta, \theta' \in \Theta_S^a$ for some $a \in A$. If it were, then we can assume with loss that $\sigma_J(\theta) \succ_S^{\theta'} \sigma_J(\theta')$. Hence the Sender at θ' has a profitable deviation to report θ, which violates the incentive compatibility constraint. Now assume for contradiction that there is $\tilde{\theta} \in \Theta$ such that $\sigma_J(\tilde{\theta}) \neq \sigma_J(\theta), \sigma_J(\theta')$. Since preferences are strict, we can assume without loss that $\theta, \tilde{\theta} \in \Theta_S^a$ for some $a \in A$. By the same reasoning as before, there is a profitable deviation for either θ or $\tilde{\theta}$. Thus any action distribution induced by a state is either $\sigma_J(\theta)$ or $\sigma_J(\theta')$. Since these action distributions are distinct, for each $a \in A$, we can define $\alpha^a = \sigma_J(\theta)$ if $\theta \in \Theta_S^a$. As before, we require that there is no incentive for the Sender to misreport, hence it must be that $\alpha^B(B) \geq \alpha^N(B)$. Furthermore, they can't be the same since σ_J is informative, so this inequality holds strict. This completes the proof.

Second, consider the claim that an OAR can be characterized as the solution to an optimization problem with respect to (α^B, α^N). Let p be an OAR. First observe that if $\theta, \theta' \in \Theta_S^a$ for some a, then it must be that $p(\theta) = p(\theta')$. Otherwise, without loss of generality, we would have that $p(a|\theta) > p(a|\theta')$. This contradicts Sender-IC, as the Sender at θ' would report θ. Hence for all $a \in A$ there is α^a such that for all $\theta \in \Theta_s^a$, $p(\theta) = \alpha^a$. We can rewrite our objective by conditioning on Θ_S^a, as the distribution over actions is identical:

$$\begin{aligned}\mathbb{E}_{\theta \sim \mu_0} \mathbb{E}_{a \sim p(\theta)} \left[u_R(a, \theta)\right] &= \sum_{a \in A} \mathbb{E}_{\theta \sim \mu_0} \left[\mathbb{E}_{a \sim p(\theta)} u_R(a, \theta) | \theta \in \Theta_S^a\right] \mu(\Theta_S^a) \\ &= \sum_{a \in A} \mathbb{E}_{\theta \sim \mu_0} \left[\mathbb{E}_{a \sim \alpha^a} u_R(a, \theta) | \theta \in \Theta_S^a\right] \mu(\Theta_S^a)\end{aligned}$$

Furthermore, we can note that the Sender at $\theta \in \Theta_S^a$ pick θ that induces the highest $p(a|\theta)$. As there are only two strategies played, it will report θ iff $\alpha^a(a) \geq \alpha^{a'}(a)$ for $a' \neq a$. For a', the analogous condition is $\alpha^{a'}(a') \geq \alpha^a(a')$, but we can observe that this is if and only $\alpha^a(a) \geq \alpha^{a'}(a)$ because $\alpha^a(a) = 1 - \alpha^a(a')$ as there are only two actions. Hence we just need one of these conditions, giving us the optimization problem we wanted to show.

Third, consider the claim that a Sender-preferred OAR gives either payoff profile u^* or u^0. Let $v_a(a') = \mathbb{E}_{\theta \sim \mu_0} \left[u_R(a', \theta) | \theta \in \Theta_S^a\right] \mu(\Theta_S^a)$. We consider the following cases.

- $v_B(B) \geq v_B(N)$, $v_N(B) \geq v_N(N)$: We can (weakly) increase utility by having $\alpha^B(B) = 1$. If $v_N(B) > v_N(N)$ we can increase utility by having $\alpha^N(B) = 1$, inducing u^0. If $v_N(B) = v_N(N)$, we can induce u^* by having $\alpha^N(N) = 1$.

- $v_B(B) \geq v_B(N)$, $v_N(B) < v_N(N)$: We can (weakly) increase utility by having $\alpha^B(B) = \alpha^N(N) = 1$, thus inducing u^*.

- $v_B(B) < v_B(N)$: We can (weakly) increase utility by having $\alpha^B(B) = \alpha^N(B)$. Since the constraint binds, this corresponds to a babbling OAR that induces u^0.

Since this comprises of all possible cases, we've proven our claim.

Fourth, consider the claim that a profile of utility functions that are consistent with a preference profile that is not aligned induces only babbling OAR. Given that the preference profile is not aligned, we can find $a \in A$ such that for some $a' \in A$, $a \succ_S^\theta a' \iff a' \succ_R^\theta a$. As such, for all $\theta \in \Theta_S^a$, $u_R(a', \theta) > u_R(a, \theta)$. Thus it must be the case that $v_a(a') > v_a(a)$. Assume without loss of generality that $a' = N$ and $a = B$. Then it cannot be the case that $\alpha^B(B) > \alpha^N(B)$, because the objective of the optimization problem can be strictly improved by having $\alpha^B(B) = \alpha^N(B)$. This corresponds to a babbling OAR.

Fifth, consider the claim that if there exists an informative PBE, then the OAR σ_J induces u^*. Recall from part 2 of Theorem 3, this implies there exists a PBE that induces u^*. Observe that all PBE are feasible solutions to the optimization problem at hand. Hence the Receiver's ex-ante utility must be at least u_R^*. Given our previous result that only u^0 or u^* can be induced as payoff profiles from an OAR in general, it must be that u^* is the payoff profile induce in all OAR in this case.

Finally, consider the claim that a profile of utility functions that are consistent with a preference profile that is aligned but only admits babbling PBE must only induce babbling OAR. Assume for contradiction the OAR is not babbling, hence it must induce payoff profile u^*. Let (α^B, α^N) be the corresponding strategy distributions for the OAR. Consider the following strategy profile $\sigma = (\sigma_S, \sigma_R)$:

$$\forall a \in A \ , \ \theta \in \Theta_S^a \implies \sigma_S(\theta) = \delta(m_a)$$
$$\forall a \in A \ , \ \sigma_R(m_a) = \delta(a)$$

where $m^B \neq m^N$. This must not be a PBE, since if it were then there is an informative equilibrium and thus a contradiction. Since it satisfies Sender optimality, it must not satisfy Receiver optimality. Hence, without loss of generality, the following holds:

$$\begin{aligned}\mathbb{E}_{\theta\sim\mu_0}\left[u_R(B,\theta)|\theta\in\Theta_S^N\right] &> \mathbb{E}_{\theta\sim\mu_0}\left[u_R(N,\theta)|\theta\in\Theta_S^N\right]\\ \implies v_N(B) &> v_N(N)\end{aligned}$$

Note that $\alpha^N(B) < \alpha^B(B)$ since this OAR is not babbling, yet by the above the Receiver can improve their payoff by increasing $\alpha^N(B)$ while maintaining the Sender-obedience constraint $\alpha^N(B) \leq \alpha^B(B)$. This contradicts the optimality of (α^B, α^N) for the OAR, hence it can't have been that the OAR was babbling. □

Corollary 3. *Consider a Sender-preferred OAR σ_J and a payoff profile $\bar{u}$ that it induces. Then there is a PBE σ that induces $\bar{u}$.*

Proof. Let u be a profile of utility functions consistent with a preference profile $\succ$. If $\succ$ is not aligned, by the previous theorem we have that all OARs are babbling. Clearly there is some PBE that induces the same payoff $\bar{u}$. If $\succ$ is aligned but induces only babbling PBE, then all OAR are also babbling. Thus we get the same conclusion as before. If $\succ$ is aligned and induces an informative PBE, then there is a PBE and σ_J that induces u^*. The latter holds by the previous theorem, and the latter holds by part 2 of Theorem 3. Since this considers all possible cases, then all OAR induce a payoff profile that can also be induced by some PBE. □

H.2 Mediation

σ_J is an optimal mediation rule (OMR) if it solves the following:

$$\max_{p:\Theta\to\Delta(A)} \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{a\sim p(\theta)}\left[u_R(a,\theta)\right]$$

subject to

$$\theta \in \arg\max_{\hat{\theta}\in\Theta} \mathbb{E}_{a\sim p(\hat{\theta})}\left[u_S(a,\theta)\right] \tag{Sender-IC}$$

and for all $a \in A$ such that $\mathbb{E}_{\theta\sim\mu_0}\left[p(a|\theta)\right] > 0$,

$$a \in \arg\max_{a'\in A} \mathbb{E}_{\theta\sim\mu_0}\left[u_R(a',\theta)|p(\theta)=a\right] \tag{Obedience}$$

An OMR solves a similar optimization problem to an OAR except with an additional *obedience* constraint that ensures that the Receiver his willing to follow the prescribed strategy. The interpretation is that there is a mediator who offers a recommendation to the Receiver of what action to take, hence via a revelation principle argument, it is without loss of generality to consider strategies that ensure the Receiver is willing to follow the recommendation.

Observe the following relationship between PBE, OMR and OAR:

Proposition 13. *Let σ^{PBE}, σ^{OMR}, and σ^{OAR} be any PBE, OAR, and OMR solution respectively. Then $\sigma^{PBE} \preceq_R \sigma^{OMR} \preceq_R \sigma^{OAR}$.*

Proof. Observe that σ^{OMR} is a feasible solution to the optimization problem for a OAR, hence we must have that $\sigma^{\text{OMR}} \preceq_R \sigma^{\text{OAR}}$. We proceed by showing that $\sigma^{\text{PBE}} \preceq_R \sigma^{\text{OMR}}$. First consider the case that there are only babbling PBE in our environment. Observe that babbling OMR, defined similarly to babbling OAR, gives the Receiver the same utility as a babbling PBE. Hence, in such a case, $\sigma^{\text{PBE}} \preceq_R \sigma^{\text{OMR}}$. On the other hand, consider the case that there is an informative PBE. By part 2 of Theorem 3, we know there exists $\sigma^* = (\sigma_S^*, \sigma_R^*)$ that induces the payoff profile u^*. Note that σ^* can be written as follows: $\forall a \in A$,

$$\theta \in \Theta_S^a \implies \sigma_S^*(\theta) = \delta(m_a)$$
$$\sigma_R^*(m_a) = \delta(a)$$

where $m_B, m_N \in M$ and $m_B \neq m_N$. Now define $\sigma_J : \Theta \to \Delta(A)$ as follows:

$$\sigma_J(\theta) = \mathbb{E}_{m \sim \sigma_S^*(\theta)}\left[\sigma_R^*(m)\right]$$

Note that for $a \in A$ and $\theta \in \Theta_S^a$, $\sigma_J(\theta) = \delta(a)$, hence this clearly satisfies the Sender's incentive constraint. Furthermore, we have that

$$\mu_0(\theta | \sigma_J(\theta) = a) = \frac{\mu_0(\theta)}{\mu_0(\Theta_S^a)} \cdot 1\{\theta \in \Theta_S^a\} = \mu_0(\theta | \sigma_S^*(\theta) = m_a)$$

Hence the Receiver's optimality criterion in a PBE is the same as the obedience constraint in an OMR. Given that the former is satisfied by assumption, then so is the latter. Hence σ_J satisfies the constraints of the optimization problem for an ORM. Observe that σ_J also gives the Receiver the same ex-ante payoff as σ^*:

$$\begin{aligned}\mathbb{E}_{\theta\sim\mu_0}\left[\mathbb{E}_{a\sim\sigma_J(\theta)}u_R(a,\theta)\right] &= \mathbb{E}_{\theta\sim\mu_0}\left[\sum_{a\in A}u_R(a,\theta)\sigma_J(a|\theta)\right]\\ &= \mathbb{E}_{\theta\sim\mu_0}\left[\sum_{a\in A}u_R(a,\theta)\mathbb{E}_{m\sim\sigma_S^*(\theta)}\left[\sigma_R^*(a|m)\right]\right]\\ &= \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{m\sim\sigma_S^*(\theta)}\left[\sum_{a\in A}u_R(a,\theta)\sigma_R^*(a|m)\right]\\ &= \mathbb{E}_{\theta\sim\mu_0}\mathbb{E}_{m\sim\sigma_S^*(\theta)}\mathbb{E}_{a\sim\sigma_R^*(m)}\left[u_R(a,\theta)\right]\end{aligned}$$

Observe that the Sender also has the same ex-ante payoff in σ_J as in σ^*. Because an ORM maximizes the Receiver's ex-ante utility subject to certain constraints, and σ_J is a feasible solution, we have that $\sigma^{\text{OMR}} \succeq_R \sigma_J \sim_R \sigma^*$. Thus in either case, $\sigma^{\text{ORM}} \succeq_R \sigma^{\text{PBE}}$ given that the $\sigma^* \succeq_R \sigma$ for all σ PBE, which follows by Theorem 3. This concludes the proof. □

Corollary 4. *Consider a Sender-preferred OMR σ_J and a payoff profile $\bar{u}$ that it induces. Then there is a PBE strategy profile and OAR strategy that induces $\bar{u}$.*

Proof. Observe from the previous result that OAR give the same utility profile as some PBE, hence OAR, OMR and such a PBE must give the same Receiver utility by the previous proposition. When there is an informative PBE, we saw that u^* could be induced by a feasible mediation strategy. Since the best Receiver payoff for OMR is u_R^*, then the Sender-preferred OMR induces u^*. When there is no informative PBE, then the OMR gives the Receiver a babbling payoff. For contradiction, assume that the Sender receives a strictly higher payoff in the OMR then that in the Sender-preferred OAR. Given that feasible mediation rules are also feasible for arbitration, then this would be a contradiction to Sender-preferredness of the OAR rule as there is an arbitration rule that gives the same payoff to the Receiver but strictly higher payoff to the Sender. □

H.3 Negotiation

The model of negotiation in Goltsman et al. (2009) is the same as that in Aumann and Hart (2003). Let $T \in \mathbb{N} \cup \{\infty\}$ be the (possibly infinite) time horizon, and M_S and M_R the space of messages that the Sender and Receiver can send, respectively. Consider an extensive form game where the state of the world is revealed at time 0 to the Sender, and the Sender and Receiver simultaneously send messages from their respective message spaces to one another.

Proposition 14. *Let σ^{PBE}, σ^{OMR}, σ^{OAR}, and σ^{ONR} be the Sender-preferred PBE, OAR, OMR, and ONR solution respectively. They all induce the same payoff profile, which is u^* if there is an informative PBE, and the Sender-preferred babbling PBE otherwise.*

Proof. Observe that by the previous results, $\sigma^{\text{PBE}} \sim_i \sigma^{\text{OMR}} \sim_i \sigma^{\text{OAR}}$ for all $i \in \{S, R\}$. We proceed by showing that $\sigma^{\text{PBE}} \preceq_i \sigma^{\text{ONR}} \preceq_i \sigma^{\text{OAR}}$ for all $i \in \{S, R\}$, thus proving the desired result.

First we show that this holds for $i = R$. Note that by choosing $\tau = 1$ and $M_S = M_R = M$, then $P = (\tau, M_S, M_R)$ represents the same game as the original cheap talk model. Given that the third party is maximizing the payoff of the Receiver, and all σ^{PBE} are feasible strategies with respect to P, then $\sigma^{\text{Neg}} \succeq_R \sigma^{\text{PBE}}$. Furthermore, for any P, let $\sigma : \Theta \to A$ be the distribution over actions induced by σ^{Neg}. We can observe that this is a feasible arbitration strategy. If this did not hold, that is the Sender-IC constraint failed, that means that the Sender at state θ has a profitable deviation by reporting some $\theta' \neq \theta$. However this would mean that in negotiation, the Sender would also have a profitable deviation to pretend the state is θ'. This contradicts σ^{ONR} is an equilibrium. Thus it must be that $\sigma^{\text{OAR}} \succeq_R \sigma^{\text{ONR}}$. These same arguments apply for the Sender. We can conclude from this that $\sigma^{\text{PBE}} \sim_i \sigma^{\text{ONR}} \sim_i \sigma^{\text{OAR}}$ for all $i \in \{S, R\}$.

□

Note that we can also show that σ^{ONR} is also feasible mediation strategy, though this is not required for the proof.