Commitment and Randomization in Communication

Emir Kamenica and Xiao Lin*

February 2025

Abstract

When does Sender, in a Sender-Receiver game, strictly value commitment? In a setting with finitely many actions and states, we establish that, generically, commitment has no value if and only if a partitional experiment is optimal. Moreover, if Sender's preferred cheap-talk equilibrium necessarily involves randomization, then Sender values commitment. Our results imply that if a school values commitment to a grading policy, then the school necessarily prefers to grade unfairly. We also ask: how often (i.e., for what share of preference profiles) does commitment have no value? For any prior, any (independent, atomless) distribution of preferences, and any state space: if there are $|A|$ actions, the likelihood that commitment has no value is at least $\frac{1}{|A|^{|A|}}$. As the number of states grows large, this likelihood converges precisely to $\frac{1}{|A|^{|A|}}$.

Keywords: Bayesian persuasion; cheap talk
JEL Codes: D80, D83

*Kamenica: Booth School of Business, University of Chicago, 5807 S. Woodlawn Ave., Chicago, IL 60637 (emir@uchicago.edu); Lin: Department of Economics, University of Pennsylvania, 133 S. 36th St., Philadelphia PA 19104 (xiaolin7@sas.upenn.edu). We are grateful to Ben Brooks, Nima Haghpanah, Ilwoo Hwang, Yuhta Ishii, Jonathan Libgober, Elliot Lipnowski, Eric Mbakop, John Moore, Roger Myerson, Phil Reny, Daniel Rappoport, Ina Taneva, and Rakesh Vohra for helpful comments.

1 Introduction

Commitment is often valuable. In the context of communication, this fact is brought out by the contrast of Sender's payoff in Bayesian persuasion versus cheap talk. For any prior, and any profile of Sender and Receiver's preferences, Sender's payoff is always weakly higher under Bayesian persuasion than in any cheap-talk equilibrium.[1] In this paper, we ask: when does commitment make Sender *strictly* better off?

Answering this question would contribute to our understanding of circumstances that incentivize building strong institutions that are immune to influence (North 1993; Lipnowski, Ravid, and Shishkin 2022) or building a reputation for a degree of honesty (Best and Quigley 2024; Mathevet, Pearce, and Stacchetti 2024).

We focus exclusively on environments with finitely many states and actions. We show that, generically, Sender with commitment values that commitment if and only if he values randomization (Theorem 1). In other words, the Bayesian persuasion payoff is achievable in a cheap-talk equilibrium if and only if a partitional experiment is a solution to the Bayesian persuasion problem. Moreover, if Sender's preferred equilibrium in a cheap-talk game necessarily involves randomization, then Sender values commitment (Theorem 2).

For an application of these results, consider a school that assigns grades to students, each of whom is characterized by a vector of attributes. Some of the attributes are relevant, in the sense that an employer values those attributes or the school's value of placing a student depends on them. Other attributes are irrelevant. The school assigns a grade to each student based on her attributes. The school's grading policy is *fair* if it assigns the same grade to students with identical relevant attributes. Theorem 1 tells us that if the school values committing to a grading policy of any form (such as mandating a maximum GPA or mandating the exact distribution of grades), then the school prefers to grade unfairly. Conversely, if a fair grading scheme is optimal, there is no need for commitment: discretionary "cheap-talk" grades are as effective as those disciplined by a publicly declared grading policy.

We also address the question of "how often" (i.e., for what share of preferences), Sender finds

[1] In fact, Bayesian persuasion provides the upper bound on Sender's equilibrium payoff under any communication protocol, such as disclosure or signaling.

commitment (or, equivalently, randomization) valuable. Theorems 1 and 2 would be of substantially less interest if it were the case that (in the finite worlds we consider), commitment is almost always valuable, with only exceptions being knife-edge cases such as completely aligned or completely opposed preferences.[2] We show, however, that this is not the case. In fact, we uncover a potentially surprising connection between the likelihood that commitment has no value and the cardinality of the action set.

Let $|A|$ denote the cardinality of the action set. Suppose that for each action-state pair, we draw Sender's utility i.i.d. from some distribution F and we draw Receiver's utility i.i.d. from some distribution G. We assume that Sender's utility draw is independent of Receiver's. For any number of states and any atomless distributions of preferences (F and G), the likelihood that commitment has no value is bounded below by $\frac{1}{|A|^{|A|}}$; moreover, as the number of states grows large, the likelihood that commitment has no value converges precisely to $\frac{1}{|A|^{|A|}}$ (Theorem 3). So, if the action set is binary and there are many states, the share of preference profiles for which commitment has no value is approximately $\frac{1}{4}$.

Illustrative example

The workhorse example in the Bayesian-persuasion literature is a prosecutor (Sender) trying to convince a judge (Receiver) to convict a defendant who is guilty or innocent. The judge's preferences are such that she prefers to convict if the probability of guilt is weakly higher than the probability of innocence. The prosecutor has state-independent preferences and always prefers conviction. The prior probability of guilt is 0.3.

If the environment were cheap talk, the unique equilibrium outcome is that the judge ignores the prosecutor and always acquits the defendant. If the prosecutor can commit to an experiment about the state, however, he will conduct a stochastic experiment that indicates guilt whenever the defendant is guilty and indicates guilt with probability $\frac{3}{7}$ when the defendant is innocent (Kamenica and Gentzkow 2011). This experiment induces the judge to convict the defendant with 60% probability. The prosecutor is thus strictly better off than under cheap talk.

[2]Denoting Sender's utility by u_S and Receiver's utility by u_R, it is easy to see that when $u_S = u_R$, neither commitment nor randomization is valuable (because full information is optimal and achievable via a cheap-talk equilibrium). Similarly, when $u_S = -u_R$, neither commitment nor randomization are valuable (because no information is optimal and achievable via a cheap-talk equilibrium).

Our Theorem 1 tells us that the two facts, (i) the prosecutor's optimal experiment involves randomization and (ii) the prosecutor does better under commitment, imply each other.[3] Of course, the prosecutor-judge example was designed to be extremely simple, so in this particular example one can easily determine the optimal experiment and the value of commitment without our theorem. In more complicated environments, however, Theorem 1 can simplify the determination of whether commitment is valuable. Except in certain cases, such as uniform-quadratic (Crawford and Sobel 1982) or transparent preferences (Lipnowski and Ravid 2020), cheap-talk games can be difficult to solve. Theorem 1 can then be used to determine whether commitment is valuable without solving for cheap talk equilibria, simply by computing the Bayesian-persuasion optima and checking whether they include a partitional experiment.[4]

The prosecutor-judge example also illustrates the distinction between the if-and-only-if result in Theorem 1 and the unidirectional Theorem 2. Recall that Theorem 2 does not claim that the value of commitment is positive only if randomization is valuable in cheap talk. The prosecutor-judge example provides a counterexample to such a claim. In the cheap-talk game, the prosecutor has no value for randomization: with or without it, he never obtains any convictions. Yet, the prosecutor obviously values commitment.

Finally, the prosecutor-judge example also helps illustrate what Theorem 1 does *not* say. Prohibiting randomization would not mean commitment is not valuable. Suppose that the prosecutor is endowed with commitment, but is legally obliged to use only partitional experiments. In that case, the prosecutor will provide a fully informative experiment, obtaining a conviction with 30% probability. That is still better than his cheap-talk payoff of no convictions.

[3]Theorem 1 only states that (i) and (ii) imply one another for a generic – i.e., full Lebesgue measure – set of preferences. To apply the theorem here, we note that the preferences in the prosecutor-judge example belong to the generic set used in the proof of the theorem. (In particular, these preferences belong to what the proof terms partitional-unique-response and scant-indifferences environments). Moreover, in Section A.6 we show that Theorem 1 holds when Sender has state-independent preferences.

[4]Recent research provides a large toolbox for solving Bayesian-persuasion problems, including concavification (Kamenica and Gentzkow, 2011), price-theoretic approaches (Kolotilin, 2018; Dworczak and Martini, 2019), duality (Dworczak and Kolotilin, 2024), and optimal-transport theory (Kolotilin, Corrao, and Wolitzky, 2023). Bergemann and Morris (2016) show that persuasion problem can be formulated as linear programs; it is well known that linear programs can be computed in polynomial time. In contrast, Babichenko et al. (2023) establish that it is NP-Hard to approximate Sender's maximum payoff in cheap-talk, or even to determining if that payoff is strictly greater than in a babbling equilibrium. For a survey of computational approaches to Bayesian persuasion, see Dughmi (2017).

Related literature

Our paper connects the literatures on cheap talk (Crawford and Sobel 1982) and Bayesian persuasion (Kamenica and Gentzkow 2011). Min (2021) and Lipnowski, Ravid, and Shishkin (2022) examine environments with limited commitment that are a mixture of cheap talk and Bayesian persuasion. In contrast, we focus on the question of when cheap talk and Bayesian persuasion yield the same payoff to Sender.[5]

Glazer and Rubinstein (2006) and Sher (2011) consider disclosure games and derive conditions on preferences that imply that Receiver values neither commitment nor randomization.

Several papers examine value of commitment under the assumption that Sender has state-independent preferences. When the action space is finite, as in our framework, Lipnowski and Ravid (2020) show that (for almost every prior) Sender either: (i) obtains his ideal payoff in cheap talk, or (ii) values commitment; Best and Quigley (2024) show that (for almost every prior) Sender either: (i) obtains his ideal payoff under the prior, or (ii) values randomization. Corrao and Dai (2023) examine Sender's payoff under cheap talk, mediation, and Bayesian persuasion. They establish that Sender does not value commitment if and only if his payoffs are the same under mediation and Bayesian persuasion.

A number of papers examine when a monotone partition is optimal in Bayesian persuasion. For example, assuming posterior-mean preferences, Dworczak and Martini (2019) derive a condition on preferences that is equivalent to optimality of a monotone partition.

In the context of mechanism design, value of commitment and value of randomization have been studied separately. Mechanisms design with limited commitment has been studied by Akbarpour and Li (2020) and Doval and Skreta (2022), among others. Value of randomization in mechanism design has been widely recognized in single-agent multi-product monopolist settings (e.g., Manelli and Vincent 2006). In contrast, with two or more agents, Chen, He, Li, and Sun (2019) establish that if agents' types are atomless and independently distributed, randomization is never valuable.

[5] Perez-Richet (2014) and Koessler and Skreta (2023) examine the circumstances under which Sender attains his Bayesian persuasion payoff even if learns the state prior to selecting the experiment.

2 Set-up and definitions

Preference and beliefs

Receiver (she) has a utility function $u_R(a, \omega)$ that depends on her action $a \in A$ and the state of the world $\omega \in \Omega$. Both A and Ω are finite; our results rely heavily on this assumption.[6] For any finite set X, we denote its cardinality by $|X|$. Sender (he) has a utility function $u_S(a, \omega)$ that depends on Receiver's action and the state. The players share an interior common prior μ_0 on Ω. For each player i, we say action a^* is i's *ideal action in* ω if $a^* \in \arg\max_{a \in A} u_i(a, \omega)$.

Genericity

Our theorems will hold "generically"; we now formalize that notion. We refer to the pair (u_S, u_R) as the (preference) *environment*. The set of all environments is $\mathbb{R}^{2|A||\Omega|}$. A set of environments is *generic* if its complement has zero Lebesgue measure in $\mathbb{R}^{2|A||\Omega|}$. When we say that a claim holds *generically*, we mean that it holds for a generic set of environments.[7]

Cheap talk, Bayesian persuasion, and value of commitment

Let M be a finite message space with $|M| > \max\{|\Omega|, |A|\}$.[8] Sender chooses a messaging strategy $\sigma : \Omega \to \Delta M$. Receiver chooses an action strategy $\rho : M \to \Delta A$.

A profile of strategies (σ, ρ) induces expected payoffs

$$U_i(\sigma, \rho) = \sum_{\omega, m, a} \mu_0(\omega)\, \sigma(m|\omega)\, \rho(a|m)\, u_i(a, \omega) \quad \text{for } i = S, R.$$

A profile (σ^*, ρ^*) is *S-BR* if $\sigma^* \in \arg\max_\sigma U(\sigma, \rho^*)$. A profile (σ^*, ρ^*) is *R-BR* if $\rho^* \in \arg\max_\rho U_R(\sigma^*, \rho)$.

[6] At the risk of being excessively philosophical, we consider environments with finite A and Ω to be more realistic. The use of infinite sets often provides tractability, but rarely improves realism.

[7] Lipnowski (2020), who focuses on finite action and state spaces as we do, establishes that commitment has no value when Sender's value function over Receiver's beliefs is continuous. Such continuity holds for a zero-measure set of environments. In contrast, we focus on results that hold generically.

[8] Our results concern Sender's payoffs under cheap talk and Bayesian persuasion, and under a restriction to partitional strategies in those models. To derive Sender's maximal payoff, it is without loss of generality to set $|M| \geq |\Omega|$ under cheap talk (Matthews 1990), $|M| \geq \min\{|\Omega|, |A|\}$ under Bayesian persuasion (Kamenica and Gentzkow 2011), and $|M| \geq |\Omega|$ under partitional strategies (trivially). Therefore, assuming $|M| \geq |\Omega|$ would suffice for our results. However, further assuming $|M| \geq |A| + 1$ simplifies the proof of Lemma 5.

Sender's *ideal payoff* is the maximum U_S induced by any profile.

A *cheap-talk equilibrium* is a profile that satisfies S-BR and R-BR.[9] We define (Sender's) *cheap-talk payoff* as the maximum U_S induced by a cheap-talk equilibrium.[10]

A *persuasion profile* is a profile that satisfies R-BR. The (Bayesian) *persuasion payoff* is the maximum U_S induced by a persuasion profile.[11] We refer to a persuasion profile that yields the persuasion payoff as *optimal*.

We say that *commitment is valuable* if the persuasion payoff is strictly higher than the cheap-talk payoff. Otherwise, we say *commitment has no value*.

Partitional strategies and value of randomization

A messaging strategy σ is *partitional* if for every ω, there is a message m such that $\sigma(m|\omega) = 1$. A profile (σ, ρ) is a *partitional profile* if σ is partitional.[12] The *persuasion partitional payoff* is the maximum U_S induced by a partitional persuasion profile. The *cheap-talk partitional payoff* is the maximum U_S induced by a partitional cheap-talk equilibrium.[13]

We say that *committed Sender values randomization* if the persuasion payoff is strictly higher than the persuasion partitional payoff. We say that *cheap-talk Sender values randomization* if the cheap-talk payoff is strictly higher than the cheap-talk partitional payoff.

[9]This definition may seem unconventional since it uses Nash equilibrium, rather than perfect Bayesian equilibrium, as the solution concept. In cheap-talk games, however, the set of equilibrium outcomes (joint distributions of states, messages, and actions) is exactly the same whether we apply Nash or perfect Bayesian as the equilibrium concept. The formulation in terms of Nash equilibria streamlines the proofs.

[10]Throughout, we examine the value of commitment to Sender; hence the focus on Sender's payoff. The set of equilibrium payoffs is compact so a maximum exists. We are interested in whether Sender can attain his commitment payoff in *some* equilibrium, so we focus on Sender-preferred equilibria. Except when no information is the commitment optimum, it cannot be that *every* cheap-talk equilibrium yields the commitment payoff since every cheap-talk game admits a babbling equilibrium.

[11]Lipnowski, Ravid, and Shishkin (2024) establish that, with finite A and Ω, Sender's equilibrium payoff in a persuasion game is generically unique.

[12]Our focus is on the connection between Sender's value of commitment and Sender's randomization. Consequently, the definition of a partitional profile only concerns Sender's strategy. That said, along the way we will establish a result about Receiver playing pure strategies (see Lemma 5).

[13]A partitional cheap-talk equilibrium always exists because the babbling equilibrium outcome can be supported by Sender always sending the same message. Consequently, the cheap-talk partitional payoff is well-defined.

3 Value of commitment: willingness-to-accept

In this section, we consider a Sender with commitment power, who can choose his messaging strategy prior to being informed of the state. We ask whether this commitment power makes Sender strictly better off. We link the value of commitment to Sender's behavior under commitment, in particular to whether Sender has a strict preference for randomization.

Theorem 1. *Generically, commitment is valuable if and only if committed Sender values randomization.*

For an intuition about the only-if direction, suppose that there is a partitional optimal persuasion profile (σ, ρ). Let M_σ be the set of messages that are sent in equilibrium. For each $m \in M_\sigma$, let Ω_m be the set of states that lead to message m, and let μ_m be the belief induced by m. For a generic set of environments, Receiver's optimal action given belief μ_m (call it a_m) is unique. Since A is finite, a_m must also be the uniquely optimal action in a neighborhood of beliefs around μ_m. Subtly, this implies that every action a_m taken in equilibrium must be Sender's preferred action, among the actions taken in equilibrium, in all states where action a_m is taken. That is a mouthful, so in other words: let $A^* = \{a_m | m \in M_\sigma\}$; for each $a_m \in A^*$, we must have $u_S(a_m, \omega) \geq u_S(a_{m'}, \omega)$ for all $a_{m'} \in A^*$ and all $\omega \in \Omega_m$. Why does this hold? If it were not the case, Sender could attain a higher payoff with an alternative strategy. Suppose $u_S(a_m, \omega) < u_S(a_{m'}, \omega)$ for some $a_{m'} \in A^*$, $\omega \in \Omega_m$. Sender could send m' in ω with a small probability and still keep a_m optimal given m. Finally, the fact that $u_S(a_m, \omega) \geq u_S(a_{m'}, \omega)$ for all $a_{m'} \in A^*$ and all $\omega \in \Omega_m$ implies that (σ, ρ) also constitutes a cheap-talk equilibrium.[14] Hence, commitment is not valuable.

We postpone the sketch of the proof of the converse direction until the next section, as the intuition is related to the intuition for Theorem 2. Formal proofs are in the Appendix.

Theorem 1 can also easily be extended to establish a threefold equivalence. Generically, the following imply each other: (i) commitment is valuable, (ii) committed Sender values randomization, and (iii) any optimal persuasion profile induces a belief under which Receiver has multiple optimal actions (see Theorem 1$'$ in the Appendix).

[14] Deviating to an on-path message $m \in M_\sigma$ cannot be profitable by the inequality $u_S(a_m, \omega) \geq u_S(a_{m'}, \omega)$; for any off-path message $m \neq M_\sigma$, we can just set $\rho(\cdot|m) = \rho(\cdot|m^*)$ for some $m^* \in M_\sigma$, thus ensuring that this deviation is also not profitable.

Finally, Theorem 1 only tells us that, generically, commitment has *zero* value if and only if randomization has *zero* value. A natural question is whether, generically, small value of commitment implies or is implied by small value of randomization. The answer is no. In the Appendix, we construct a positive measure of environments where the value of commitment is arbitrarily large but the value of randomization is arbitrarily small (Section B.2.1), and a positive measure of environments where the value of randomization is arbitrarily large but the value of commitment is arbitrarily small (Section B.2.2).

4 Value of commitment: willingness-to-pay

In this section, we consider a Sender without commitment power who engages in a cheap-talk game. We ask whether he would be strictly better off if he had commitment power. We link the value of such commitment to Sender's behavior in Sender-preferred cheap-talk equilibria, in particular to whether Sender necessarily randomizes in such equilibria.

Theorem 2. *Generically, commitment is valuable if cheap-talk Sender values randomization.*

Theorem 2 and the if-direction of Theorem 1 both derive from the following result. Generically, if a cheap-talk equilibrium yields the persuasion payoff, then there is a partitional σ and a (pure-strategy) ρ such that (σ, ρ) is a cheap-talk equilibrium and yields the persuasion payoff. We build this result (Lemma 4) in two steps.

The first step (Lemma 5) shows that, generically, if (σ, ρ) is R-BR and yields the persuasion payoff, then ρ must be pure on-path. Consider toward contradiction that there is an m sent with positive probability under σ, and there are two distinct actions, say a and a', in the support of $\rho(\cdot|m)$. It must be that both Sender and Receiver are indifferent between a and a' under belief μ_m: Receiver has to be indifferent because (σ, ρ) is R-BR; Sender has to be indifferent because (σ, ρ) yields the persuasion payoff, which maximizes U_S over all persuasion profiles.[15] The result then follows from establishing that such a coincidence of indifferences generically cannot arise when Sender is optimizing. For some intuition for why this is the case, consider Figure 1 which illustrates this result when there are three states. Suppose a_1 and a_2 are in the support of $\rho(\cdot|m)$. Region

[15]If Sender strictly prefers one action over the other, say a over a', at μ_m, then Sender would obtain a higher payoff if Receiver always takes a following m (which would remain R-BR given Receiver's indifference).

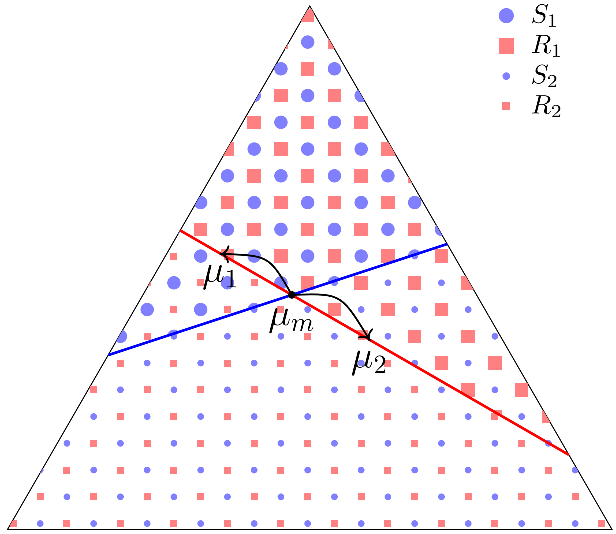

Figure 1: Indifference incompatible with optimality

R_i denotes beliefs where Receiver prefers a_i. Region S_i denotes beliefs where Sender prefers a_i. Generically, the border between R_1 and R_2 is distinct from the border between S_1 and S_2 and thus the two borders have at most one intersection, μ_m. Moreover, generically μ_m (if it exists) is an interior belief. But now, Sender could deviate to an alternate strategy that induces beliefs μ_1 and μ_2 instead of μ_m, with Receiver still indifferent between a_1 and a_2 at both μ_1 and μ_2. Suppose that Receiver takes action a_i following belief μ_i. This strategy is still R-BR for Receiver and gives Sender a strictly higher payoff. Thus, we have reached a contradiction. With more than three states and more than two actions, the proof that the coincidence of indifferences generically cannot arise is conceptually similar but notationally more involved. It is presented in the Appendix as Lemma 2.

The second step (Lemma 6) shows that, generically, if (σ, ρ) is a cheap-talk equilibrium that yields the persuasion payoff, and ρ is a pure strategy on-path, then there is a partitional cheap-talk equilibrium that yields the persuasion payoff. This is easy to see. Generically, for any ω and any $a \neq a'$, we have $u_S(a, \omega) \neq u_S(a', \omega)$. Now, consider some cheap-talk equilibrium (σ, ρ) that yields the persuasion payoff with ρ is a pure strategy on-path. If σ is partitional, our result is immediate. Suppose to the contrary that in some ω, both m and m' are sent with positive probability. Then, m and m' must induce the same action: if m induces some a and m' induces a distinct a', the fact that $u_S(a, \omega) \neq u_S(a', \omega)$ would mean that σ cannot be S-BR. Given that any two messages sent in ω induce the same action, we can define $\rho(\sigma(\omega))$ as *the* action that Receiver takes in state ω given (σ, ρ).

Now, we can consider an alternative, partitional profile $(\hat{\sigma}, \hat{\rho})$. Let f be any injective function from A to M. Let $\hat{\sigma}(\omega) = f(\rho(\sigma(\omega)))$ and $\rho(f(a)) = a$. It is immediate that $(\hat{\sigma}, \hat{\rho})$ is also a cheap-talk equilibrium and yields the persuasion payoff.

It is perhaps worth noting that 1 and 2 jointly imply the following:

Corollary 1. *Generically, if cheap-talk Sender values randomization, then committed Sender values randomization.*

5 Application to grading

For an application of our results, we consider their implications to grading policies. The application also clarifies a sense in which "randomization" in the statement of our results need not be interpreted literally.

Suppose Sender is a school that assigns grades to its students. We interpret M as the set of potential grades. Each student is characterized by a vector of attributes. We say an attribute is *relevant* if an employer values it or the school's value of placing the student with an employer depends on it. We interpret Ω as the set of all possible configurations of the relevant attributes. We maintain the assumption that Ω is finite.

Students also have irrelevant attributes. We denote by X as the set of all possible configurations of the irrelevant attributes. We assume that the distribution over X is atomless. The school utilizes a deterministic *grading scheme* $g : \Omega \times X \to M$. We say a grading scheme g is *fair* if $g(\omega, x) = g(\omega, x')$ for every ω, x, x'. Otherwise, the scheme is *unfair*.

For this application, instead of envisioning a single Receiver, we assume that each student applies to a distinct employer. Each employer observes the grade $m \in M$ of its applicant and chooses one of finitely many actions $a \in A$ (e.g., whether to hire the student and if so for what position). All employers have the same utility function $u_R(a, \omega)$ that depends on the employer's action and the relevant attributes of the applicant. (If there were a single employer who observed the grades of all of the applicants, this would effectively provide Sender with some commitment power because the distribution of messages would be directly observable to Receiver.) The school's utility is additive across its students; for each student, the school's payoff $u_S(a, \omega)$ depends on that student's outcome and that student's relevant attributes.

Under *discretionary grading*, the school freely chooses which grade to assign to each student, i.e., the school selects any grading scheme it wishes. The employer only observes the grade but not the grading scheme that was used.

Alternatively, the school could implement a (publicly observable) *grading policy* that restricts the set of schemes that it can use.

A grading policy could be a restriction to one specific grading scheme. This would make the situation equivalent to Bayesian persuasion. This is the case even though the grading scheme is deterministic because by conditioning the grade on the irrelevant attributes, the school can implement any distribution of grades conditional on each ω.[16]

A more commonly observed type of grading policy is one where school commits to a given distribution of grades (Lin and Liu, 2024). We refer to such a policy as a *mandated curve*. For example, the University of Chicago Law School mandates a pre-specified share of students that will receive any given grade.

Even more common is commitment to a *GPA cap*. For example, the University of Chicago Booth School of Business mandates that the average grade assigned in a given course must not exceed B+.

We say that *the school values commitment* if it strictly prefers to implement *any* grading policy (full commitment, mandated curve, GPA cap, etc.) over discretionary grading. We know that any policy must yield a payoff that is weakly lower than full commitment and weakly higher than discretionary grading. Consequently, if any grading policy yields a strictly higher payoff than discretionary grades, we know that the persuasion payoff (full commitment) exceeds the cheap talk payoff (discretionary grades).

We say that *the school prefers to grade unfairly* if its ideal grading scheme is unfair. In other words, if the school were able to commit to a particular grading scheme, it would select an unfair

[16] The formulation of experiments as deterministic functions of an expanded state space was introduced by Green and Stokey (1978, 2022) and Gentzkow and Kamenica (2017). It has been further studied in Brooks et al. (2022, 2024).

one.

Theorem 1 tells us that, generically, the school values commitment if and only if it prefers to grade unfairly. Thus, whenever we observe a school mandating a curve or a GPA cap, we know that the school's ideal policy treats students unfairly.[17]

Note, however, that even if we observe a school mandating a curve, Theorem 1 does not imply that the school will implement an unfair scheme if it can only commit to a mandated curve (i.e., is unable to fully commit to a particular scheme). Consequently, in Appendix B.1, we analyze whether partial commitment being valuable (i.e., mandating a curve yields a strictly higher payoff than discretionary grades) implies that randomization under partial commitment is valuable (i.e., among the schemes that yield the mandated curve, every scheme that is optimal is unfair). Under the assumption that the school's preferences are supermodular, we establish that this is indeed the case (Theorem 6). Whether the conclusion of the theorem holds when preferences are not supermodular remains an open question.

6 How often is commitment valuable?

Theorems 1 and 2 would not be particularly interesting if it turned out that both commitment and randomization are almost always valuable.

When $u_S = u_R$ or $u_S = -u_R$, it is easy to see that neither commitment nor randomization are valuable, but those are knife-edge cases and it is important to show that commitment has no value in a broader class of environments. We do so in this section.

To formalize our result, we generate random environments by drawing Sender's utility for each action-state i.i.d. from some atomless distribution F and Receiver's utility for each action-state i.i.d from some atomless distribution G. We further assume that for each (a, ω), the random variables $u_S(a, \omega)$ and $u_R(a, \omega)$ are independent from one another.

We should note that this structure does not preclude any particular configuration of preferences.

[17] Our analysis views the school (that cares about student placements) and the professor (who is assigning grades) as a single agent. A distinct motivation for a grading policy such as a GAP cap, outside of our Sender-Receiver framework, would be an agency conflict between the school and the professor. For example, the professor may wish to give uniformly high grades in order to avoid student complaints so the school might impose a GPA cap to mitigate that temptation (Frankel, 2014). Moreover, a grading policy could have a distinct benefit of aiding equilibrium coordination about the meaning of grades; our focus on Sender-preferred equilibria assumes miscoordination away.

For any F and G, with some probability the environment will be such that Sender's and Receiver's preferences are perfectly aligned, with some probability they will be completely opposed, with some probability they will be aligned in some states but not others, etc.

Fixing A and Ω, we thus generate stochastic environments and can ask: what is the probability that commitment (or equivalently randomization) has no value. Our main theorem in this section establishes results about Pr (commitment has no value) that turn out to be independent of F and G.

Theorem 3. *For any atomless F and G:*

- $\Pr(\textit{commitment has no value}) \geq \frac{1}{|A|^{|A|}}$.
- *as* $|\Omega| \to \infty$, $\Pr(\textit{commitment has no value}) \to \frac{1}{|A|^{|A|}}$.

Denote the action space by $A = \{a_1, a_2, ..., a_{|A|}\}$ and denote some $|A|$ elements of M by m_1 through $m_{|A|}$. Let Ω_i be the set of states where a_i is Sender's ideal action. The *requesting* messaging strategy sets $\sigma(\omega) = m_i$ for $\omega \in \Omega_i$.[18] A *compliant* action strategy sets $\rho(m_i) = a_i$. A profile that consists of the requesting and a compliant strategy yields Sender's ideal payoff.

Say that an environment is *felicitous* if for each Ω_i and each $a_j \in A$, we have

$$\sum_{\omega \in \Omega_i} \mu_0(\omega)(u_R(a_i, \omega) - u_R(a_j, \omega)) \geq 0. \tag{1}$$

If the environment is felicitous, a profile that consists of the requesting and an compliant strategy clearly constitutes a cheap-talk equilibrium. Since such a profile yields Sender's ideal payoff, commitment clearly has no value if the environment is felicitous.[19]

Now, for any Ω_i that is not empty, the probability that inequality (1) is satisfied is $\frac{1}{|A|}$, since for each a_j, $u_R(a_i, \omega)$ and $u_R(a_j, \omega)$ are i.i.d. with an atomless distribution. Moreover, given two non-empty Ω_i and Ω_j, the probability that inequality (1) is satisfied for Ω_j is independent of the

[18] Generically, Ω_i and Ω_j do not intersect.

[19] The felicity condition also appears in Antic, Chakraborty, and Harbaugh (2022) and Aybas and Callander (2024). In Antic, Chakraborty, and Harbaugh (2022), it is a necessary condition for the possibility of subversive conversations: without it, a third-party (Receiver) with veto power would prevent a committee (Sender) from implementing a project solely based on the information that the committee wants to do so. Aybas and Callander (2024) consider preferences of the form $u_R(a, \omega(\cdot)) = \omega(a)^2$ and $u_S(a, \omega(\cdot)) = (\omega(a) - b)^2$ for some $b > 0$ where $\omega : A \to \mathbb{R}$ is the realized path of a Brownian motion. They identify features of b and A that make the environment felicitous.

probability that it is satisfied for Ω_i. Thus if all of Ω_i's are non-empty, the probability that the environment is felicitous is $\left(\frac{1}{|A|}\right)^{|A|}$, or $\frac{1}{|A|^{|A|}}$.

If an Ω_i is empty, inequality (1) is satisfied vacuously for that Ω_i. Thus, for any atomless F and G, the overall probability that the environment is felicitous must be weakly greater than $\frac{1}{|A|^{|A|}}$. Since commitment has no value in felicitous environments, we conclude that $\Pr\left(\text{commitment has no value}\right) \geq \frac{1}{|A|^{|A|}}$.

We establish the second part of the theorem by showing that as $|\Omega|$ grows large: (i) the likelihood that an Ω_i is empty converges to zero so $\Pr\left(\text{felicity}\right)$ converges to $\frac{1}{|A|^{|A|}}$, and (ii) $\Pr\left(\text{commitment has no value}\right)$ converges to $\Pr\left(\text{felicity}\right)$.

Part (i) is easy to see. For any $a \in A$, as Ω grows large, the chance that there is *no* state where a is Sender's ideal action vanishes.

To establish part (ii), say that an environment is *jointly-inclusive* if for every action a, there is some state ω such that a is the ideal action for both Sender and Receiver in ω. Analogously to part (i), it is easy to see that as Ω grows large, the probability that the environment is jointly-inclusive converges to 1. To complete the proof, we argue that, generically, if the environment is jointly-inclusive and commitment has no value, then the environment must be felicitous. First, we know from Theorem 2, that there is a partitional profile (σ, ρ) that is a cheap-talk equilibrium and yields the persuasion payoff. Next, we note that every action $a \in A$ must be induced by (σ, ρ): there is a state ω where a is both Sender's and Receiver's ideal action, so if a were never taken, the committed Sender could profitably deviate by sometimes[20] revealing ω and thus inducing a, thus contradicting the fact that (σ, ρ) yields the persuasion payoff. This in turn implies that, for every ω, the action induced in ω, $\rho\left(\sigma\left(\omega\right)\right)$, must be Sender's ideal action in ω. If Sender strictly preferred some other a' in ω, (σ, ρ) could not be S-BR as the cheap-talk Sender would profitably deviate and set $\sigma(\omega)$ to be whatever message induces a'; since all actions are induced by (σ, ρ), there must be such a message. Taking stock, we have established that (σ, ρ) is a partitional profile that is R-BR (since it is a cheap-talk equilibrium) and induces Receiver to take Sender's ideal action in every state. But this means that every message sent under σ fully reveals what action is ideal for Sender, and Receiver complies and takes that action. Hence, the environment is felicitous.

[20] Sender could reveal ω with some probability ϵ; Receiver's response to all other messages would remain unchanged if ϵ is sufficiently small.

We conclude this section with a few comments.

First, whether an environment (u_S, u_R) is felicitous and whether commitment has value in this environment depends on the prior μ_0. Yet, we fix an arbitrary prior μ_0 at the outset. Theorem 3 holds not only for an arbitrary F and G, but also for an arbitrary μ_0.

Second, as the sketch of the proof makes clear, when the state space is large, Sender does not value commitment only if he can obtain his ideal payoff in a cheap-talk equilibrium.[21] With a smaller state space, however, cheap-talk and persuasion payoffs can coincide even if they are substantially lower than the ideal payoff.

Third, the assumption that Sender's and Receiver's utility are drawn from distributions that are i.i.d. across action-state pairs is more palatable if we think of A and Ω as being not merely finite but also "unstructured," without a natural metric. For example, if A includes actions such as "buy one apple" and "buy two apples", or Ω includes states such as "temperature will be 88 Fahrenheit" and "temperature will be 89 Fahrenheit," then assuming that $u_S(a, \omega)$ is independent of $u_S(a', \omega')$ as soon as $a \neq a'$ or $\omega \neq \omega'$ would be unreasonable.

Fourth, in the second part of Theorem 3, a reader might be concerned that, by keeping F and G fixed as Ω grows, we are "squishing" utilities together and making the difference in payoffs become vanishingly small. We could let F and G depend on $|A|$ and $|\Omega|$ in an arbitrary way, however, and the Theorem would still hold. We formulate the Theorem with a fixed F and G solely for ease of exposition.

Finally, the felicity condition seems to have some flavor of alignment of Sender and Receiver's preferences. While that may the case, the felicity condition does not preclude the possibility that Receiver is much worse off than she would be if Sender and Receiver's preferences were fully aligned. For instance, consider the prosecutor-judge example and suppose that the prior is 0.7 rather than 0.3; then, the environment is felicitous but Receiver obtains no information.

[21] Formally, as $|\Omega|$ goes to infinity, the probability of an environment such that Sender does not value commitment but does not obtain his ideal payoff converges to zero.

References

Mohammad Akbarpour and Shengwu Li. Credible auctions: A trilemma. *Econometrica*, 88(2): 425–467, 2020.

Nemanja Antic, Archishman Chakraborty, and Rick Harbaugh. Subversive conversations. Technical report, 2022.

Yunus C Aybas and Steven Callander. Cheap talk in complex environments. 2024.

Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi. Algorithmic cheap talk. *arXiv preprint arXiv:2311.09011*, 2023.

Dirk Bergemann and Stephen Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.

James Best and Daniel Quigley. Persuasion for the long run. *Journal of Political Economy*, 132(5): 1740–1791, 2024.

Benjamin Brooks, Alexander Frankel, and Emir Kamenica. Information hierarchies. *Econometrica*, 90(5):2187–2214, 2022.

Benjamin Brooks, Alexander Frankel, and Emir Kamenica. Representing type spaces as signal allocations. *Economic Theory Bulletin*, pages 1–7, 2024.

Yi-Chun Chen, Wei He, Jiangtao Li, and Yeneng Sun. Equivalence of stochastic and deterministic mechanisms. *Econometrica*, 87(4):1367–1390, 2019.

Roberto Corrao and Yifan Dai. The bounds of mediated communication. *arXiv preprint arXiv:2303.06244*, 2023.

Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

Laura Doval and Vasiliki Skreta. Mechanism design with limited commitment. *Econometrica*, 90 (4):1463–1500, 2022.

Shaddin Dughmi. Algorithmic information structure design: a survey. *ACM SIGecom Exchanges*, 15(2):2–24, 2017.

Piotr Dworczak and Anton Kolotilin. The persuasion duality. *arXiv preprint arXiv:1910.11392*, 2024.

Piotr Dworczak and Giorgio Martini. The simple economics of optimal persuasion. *Journal of Political Economy*, 127(5):1993–2048, 2019.

Alexander Frankel. Aligned delegation. *American Economic Review*, 104(1):66–83, 2014.

Matthew Gentzkow and Emir Kamenica. Bayesian persuasion with multiple senders and rich signal spaces. *Games and Economic Behavior*, 104:411–429, 2017.

Jacob Glazer and Ariel Rubinstein. A study in the pragmatics of persuasion: a game theoretical approach. *Theoretical Economics*, 1:395–410, 2006.

Jerry R Green and Nancy L Stokey. Two representations of information structures and their comparisons. 1978.

Jerry R Green and Nancy L Stokey. Two representations of information structures and their comparisons. *Decisions in Economics and Finance*, 45(2):541–547, 2022.

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

Frédéric Koessler and Vasiliki Skreta. Informed information design. *Journal of Political Economy*, 131(11):3186–3232, 2023.

Anton Kolotilin. Optimal information disclosure: A linear programming approach. *Theoretical Economics*, 13(2):607–635, 2018.

Anton Kolotilin, Roberto Corrao, and Alexander Wolitzky. Persuasion and matching: Optimal productive transport. *arXiv preprint arXiv:2311.02889*, 2023.

Xiao Lin and Ce Liu. Credible persuasion. *Journal of Political Economy*, 132(7):2228–2273, 2024.

Elliot Lipnowski. Equivalence of cheap talk and bayesian persuasion in a finite continuous model, 2020.

Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. *Econometrica*, 88(4): 1631–1660, 2020.

Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Persuasion via weak institutions. *Journal of Political Economy*, 130(10):000–000, 2022.

Elliot Lipnowski, Doron Ravid, and Denis Shishkin. Perfect bayesian persuasion. *arXiv preprint arXiv:2402.06765*, 2024.

Alejandro M Manelli and Daniel R Vincent. Bundling as an optimal selling mechanism for a multiple-good monopolist. *Journal of Economic Theory*, 127(1):1–35, 2006.

Laurent Mathevet, David Pearce, and Ennio Stacchetti. Reputation and information design. *Working Paper*, 2024.

Steven A. Matthews. The number of messages required in sender-receiver games. Technical report, Northwestern University, 1990.

Daehong Min. Bayesian persuasion under partial commitment. *Economic Theory*, 72:743–764, 2021.

Douglass C North. Institutions and credible commitment. *Journal of Institutional and Theoretical Economics (JITE)/Zeitschrift für die gesamte Staatswissenschaft*, pages 11–23, 1993.

Eduardo Perez-Richet. Interim bayesian persuasion: First steps. *American Economic Review*, 104 (5):469–474, 2014.

Itai Sher. Credibility and determinism in a game of persuasion. *Games and Economic Behavior*, 71(2):409–419, 2011.

A Appendix

A.1 Notation and terminology

Let $A = \{a_1, ..., a_{|A|}\}$. Let $\Omega = \{\omega_1, ..., \omega_{|\Omega|}\}$.

Given a messaging strategy σ, let $M_\sigma = \{m \in M | \sigma(m|\omega) > 0 \text{ for some } \omega\}$ be the set of messages that are sent with positive probability under σ. For any ω, if $\sigma(\cdot|\omega)$ is degenerate (i.e., there exists a message m such that $\sigma(m|\omega) = 1$), we abuse notation and let $\sigma(\omega)$ denote the message that is sent in state ω. Similarly, if $\rho(\cdot|m)$ is degenerate, we let $\rho(m)$ denote the action taken following message m.

We say that ρ is *pure* if $\rho(\cdot|m)$ is degenerate for all $m \in M$. Given a profile (σ, ρ), we say ρ is *pure-on-path* if $\rho(\cdot|m)$ is degenerate for all $m \in M_\sigma$.

We denote a vector all of whose elements are equal to r by $\boldsymbol{r}$. We sometimes use μ and sometimes use $[\mu]$ for an element of $\Delta\Omega$.

A.2 Generic environments for the proofs

We now introduce two generic sets of environments that will play important roles in the proofs.

A.2.1 Partitional-unique-response environments

An environment (u_S, u_R) satisfies *partitional-unique-response* if for every non-empty $\hat{\Omega} \subseteq \Omega$,

$$\arg\max_{a \in A} \sum_{\omega \in \hat{\Omega}} \mu_0(\omega) u_R(a, \omega)$$

is a singleton.

Note that whether an environment satisfies partitional-unique-response does not depend on Sender's preferences. The partitional-unique-response property requires that, at the finitely many beliefs induced by partitional experiments, Receiver has a unique best response at those beliefs.

Lemma 1. *The set of partitional-unique-response environments is generic.*

Proof. Given a triplet $\left(\hat{\Omega}, a_i, a_j\right)$ such that $\hat{\Omega} \subseteq \Omega$, $a_i, a_j \in A$, and $a_i \neq a_j$, let $Q(\hat{\Omega}, a_i, a_j)$ denote the set of u_R such that

$$\sum_{\omega\in\hat{\Omega}} \mu_0(\omega)u_R(a_i,\omega) = \sum_{\omega\in\hat{\Omega}} \mu_0(\omega)u_R(a_j,\omega). \tag{2}$$

An environment (u_S, u_R) does not satisfy partitional-unique-response only if $u_R \in \cup_{a_i\neq a_j,\hat{\Omega}\subseteq\Omega} Q(\hat{\Omega}, a_i, a_j)$.

We wish to show that $\cup_{a_i\neq a_j,\hat{\Omega}\subseteq\Omega} Q(\hat{\Omega}, a_i, a_j)$ has measure zero in $\mathbb{R}^{|\Omega|\times|A|}$, which implies that the set of partitional-unique-response environments is generic.

Fix any $a_i \neq a_j$ and $\hat{\Omega} \subseteq \Omega$. The set of u_R that satisfy (2) can be written as:

$$\sum_{\omega,a} u_R(a,\omega)\eta(a,\omega) = 0 \tag{3}$$

where

$$\eta(a,\omega) = \begin{cases} \mu_0(\omega) & \text{if } a = a_i, \omega \in \hat{\Omega} \\ -\mu_0(\omega) & \text{if } a = a_j, \omega \in \hat{\Omega} \\ 0 & \text{otherwise.} \end{cases}$$

So (3) defines a hyperplane of $\mathbb{R}^{|\Omega|\times|A|}$, and thus has Lebesgue measure zero in $\mathbb{R}^{|\Omega|\times|A|}$. □

A.2.2 Scant-indifferences environments

For each $a_i \in A$, let $\mathbf{u}_S(a_i) = u_S(a_i,\cdot) \in \mathbb{R}^{|\Omega|}$ and $\mathbf{u}_R(a_i) = u_R(a_i,\cdot) \in \mathbb{R}^{|\Omega|}$ denote the payoff vectors across the states.

For each a_i, define the *expanded-indifference matrix* T^i as follows. Let T^i_S be the matrix with $|A|-1$ rows and $|\Omega|$ columns, with each row associated with $j \neq i$ and equal to $\mathbf{u}_S(a_j) - \mathbf{u}_S(a_i)$. Let T^i_R be the matrix with $|A|-1$ rows and $|\Omega|$ columns, with each row associated with $j \neq i$ and equal to $\mathbf{u}_R(a_j) - \mathbf{u}_R(a_i)$. Let I be the identity matrix of size $|\Omega|$. Then, let

$$T^i = \begin{bmatrix} T^i_S \\ T^i_R \\ I \end{bmatrix}.$$

Given any matrix T, a *row-submatrix* of T is a matrix formed by removing some of the rows of T.

Finally, we say that an environment satisfies *scant-indifferences* if every row-submatrix of every expanded-indifference matrix T^i is full rank.

We anticipate that the reader might find this definition mysterious, so we now try to provide some intuition by connecting this definition to the proof sketch we gave in the body of the paper for Theorem 2 in the case with two actions and three states.

Recall, that in Figure 1, the argument behind Lemma 5 relied on two facts that must hold generically. First, the border between R_1 and R_2 is distinct from the border between S_1 and S_2 and thus the two borders have at most one intersection, μ_m. Second, generically μ_m (if it exists) is an interior belief. Moreover, the argument behind Lemma 6 relied on the fact that, generically, for any ω and $a_i \neq a_j$, $u_S(a_i, \omega) \neq u_S(a_j, \omega)$.

We now illustrate why these three facts hold in any scant-indifferences environment. With only two actions, we can look at T^1 only, since the argument for T^2 is identical. We have

$$T^1 = \begin{bmatrix} \overset{\triangle}{u}_S(\omega_1) & \overset{\triangle}{u}_S(\omega_2) & \overset{\triangle}{u}_S(\omega_3) \\ \overset{\triangle}{u}_R(\omega_1) & \overset{\triangle}{u}_R(\omega_2) & \overset{\triangle}{u}_R(\omega_3) \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

where $\overset{\triangle}{u}_S(\omega_1) = u_S(a_2, \omega_1) - u_S(a_1, \omega_1)$ and analogously for other states and $\overset{\triangle}{u}_R$.

First, consider the row-submatrix

$$\overset{\triangle}{T} = \begin{bmatrix} \overset{\triangle}{u}_S(\omega_1) & \overset{\triangle}{u}_S(\omega_2) & \overset{\triangle}{u}_S(\omega_3) \\ \overset{\triangle}{u}_R(\omega_1) & \overset{\triangle}{u}_R(\omega_2) & \overset{\triangle}{u}_R(\omega_3) \end{bmatrix}.$$

Note that both Sender and Receiver are indifferent between the two actions at a belief μ if and only if $\overset{\triangle}{T}\mu = 0$. Thus, requiring that $\overset{\triangle}{T}$ be full-rank is equivalent to requiring that the border between R_1 and R_2 not be parallel to the border between S_1 and S_2. A fortiori, the environment satisfying scant-indifferences implies that the two borders do not coincide.

Second, consider the row-submatrix

$$\begin{bmatrix} \overset{\Delta}{u}_S(\omega_1) & \overset{\Delta}{u}_S(\omega_2) & \overset{\Delta}{u}_S(\omega_3) \\ \overset{\Delta}{u}_R(\omega_1) & \overset{\Delta}{u}_R(\omega_2) & \overset{\Delta}{u}_R(\omega_3) \\ 1 & 0 & 0 \end{bmatrix}.$$

Requiring that this matrix be full-rank yields that μ_m puts strictly positive probability on ω_1. Considering the row-submatrices that alternatively include the other two rows of the identity matrix yields that μ_m puts strictly positive probability on ω_2 and ω_3.

Finally, suppose that in, say state ω_1, $\overset{\Delta}{u}_S(\omega_1) = 0$. Consider the row-submatrix

$$\begin{bmatrix} 0 & \overset{\Delta}{u}_S(\omega_2) & \overset{\Delta}{u}_S(\omega_3) \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

Clearly, this matrix is not full-rank, so scant-indifferences rules out the possibility that $u_S(a_1, \omega_1) = u_S(a_2, \omega_2)$.

Having motivated the definition of scant-indifferences environments (and given some intuition for why our results hold in such environments), we now establish that the set of such environments is generic.

Lemma 2. *The set of scant-indifferences environments is generic.*

Proof. We seek to show that the set of (u_S, u_R) such that every row-submatrix of every expanded-indifference matrix is full-rank has full Lebesgue measure on $\mathbb{R}^{|\Omega|\times|A|\times 2}$.

First, observe that given any expanded-indifference matrix T^i, if every square row-submatrix of T^i is full-rank, than every row-submatrix of T^i is full-rank. To see why, suppose every square row-submatrix of T^i is full-rank. Now, consider an arbitrary row-submatrix $\hat{T}$ of T^i. If $\hat{T}$ square, it obviously has full-rank. Suppose that $\hat{T}$ has more than $|\Omega|$ rows. In that case, every square row-submatrix of $\hat{T}$ is also a square row-submatrix of T^i. This row-submatrix has rank $|\Omega|$. Therefore, $\hat{T}$ has rank $|\Omega|$ and is thus full-rank. Finally, suppose hat $\hat{T}$ has fewer than $|\Omega|$ rows. We know that $\hat{T}$ is a row-submatrix of some square row-submatrix $\tilde{T}$ of T^i. We know $\tilde{T}$ has full-rank so all of its rows are linearly independent. Consequently, the subset of its rows that constitute $\hat{T}$ is also linearly independent.

Now that we can consider only square row-submatrices of T^i, we recall that a square matrix is full-rank if and only if its determinant is non-zero. Hence, it will suffice to show that for a full Lebesgue measure set of (u_S, u_R), the determinant of every square row-submatrix of every expanded-indifference matrix is non-zero. Given (u_S, u_R), consider some square row-submatrix $\hat{T}$ of some expanded-indifference matrix. The determinant of $\hat{T}$ is a non-zero polynomial function of $(u_S, u_R) \in \mathbb{R}^{|\Omega| \times |A| \times 2}$. The zero set of any non-zero polynomial function has Lebesgue measure zero, so the set of (u_S, u_R) for which $\hat{T}$ does not have full rank has Lebesgue measure zero. Since there are only finitely many square row-submatrices of expanded-indifference matrices, the fact that any one of them is generically full-rank implies that all of them are generically full-rank (a union of finitely many sets of Lebesgue measure zero has Lebesgue measure zero). □

As we noted above (for the three state, two action case), in scant-indifferences environments, there is no state in which Sender is indifferent between two distinct actions.

Lemma 3. *In any scant-indifferences environment, for any ω and $a_i \neq a_j$, $u_S(a_i, \omega) \neq u_S(a_j, \omega)$.*

Proof. Suppose, toward a contradiction, that there exist some ω, a_i, and a_j such that $u_S(a_i, \omega) = u_S(a_j, \omega)$. Without loss, suppose this holds for ω_1. Then, the vector $\mathbf{u}_S(a_i) - \mathbf{u}_S(a_j)$ has zero as its first element. Now consider the $|\Omega| \times |\Omega|$ row sub-matrix of T^j

$$\begin{bmatrix} \mathbf{u}_S(a_i) - \mathbf{u}_S(a_j) \\ e_2 \\ \dots \\ e_{|\Omega|} \end{bmatrix}.$$

This matrix is not full-rank because the first row can be expressed as a linear combination of the other rows. □

A.3 Key Lemma

In this section we establish a key lemma.

Lemma 4. *In a scant-indifferences environment, if commitment has no value, then there is a partitional* $\hat{\sigma}$ *and a pure strategy* $\hat{\rho}$ *such that* $(\hat{\sigma}, \hat{\rho})$ *is a cheap-talk equilibrium and yields the persuasion payoff (and* $|M_{\hat{\sigma}}| \leq |A|$*).*

Lemma 4 will be useful for proofs of Theorems 1, 2, and 3. The parenthetical remark that $|M_{\hat{\sigma}}| \leq |A|$ will be useful in the proof of Theorem 3.

To establish the Lemma, we first show that if a cheap-talk equilibrium yields the persuasion payoff, then Receiver must not randomize on path in that equilibrium. Second, we show that if Receiver does not randomize on path, Sender also need not randomize.

Lemma 5. *In a scant-indifferences environment, if* (σ, ρ) *is R-BR and yields the persuasion payoff, then* ρ *must pure-on-path.*

Proof. Suppose by contradiction that the environment satisfies scant-indifferences, profile (σ, ρ) is R-BR and yields the persuasion payoff, yet there exists a message $m \in M_\sigma$ such that $|supp(\rho(\cdot|m))| = k > 1$.

We first note that both Sender and Receiver must be indifferent among all the actions in $supp(\rho(\cdot|m))$ given μ_m, the belief induced by message m. In other words, for all $a_i, a_j \in supp(\rho(\cdot|m))$,

$$\sum_\omega \mu_m(\omega) u_R(a_i, \omega) = \sum_\omega \mu_m(\omega) u_R(a_j, \omega), \tag{4}$$

$$\sum_\omega \mu_m(\omega) u_S(a_i, \omega) = \sum_\omega \mu_m(\omega) u_S(a_j, \omega). \tag{5}$$

Equation (4) follows immediately from R-BR. Equation (5) follows from the fact that (σ, ρ) yields the persuasion payoff: if say $\sum_\omega \mu_m(\omega) u_S(a_i, \omega) > \sum_\omega \mu_m(\omega) u_S(a_j, \omega)$, an alternative strategy profile where Receiver breaks ties in favor of Sender would still satisfy R-BR while strictly improving Sender's payoff.

For each belief $\mu \in \Delta\Omega$, let $A_R^*(\mu)$ denote the set of Receiver's optimal actions under belief μ; that is, $A_R^*(\mu) = \arg\max_{a \in A} \mathbf{u}_R(a) \cdot \mu$. Clearly, $supp(\rho(\cdot|m)) \subseteq A_R^*(\mu_m)$, meaning that $A_R^*(\mu_m)$ contains at least the k actions in the support of $\rho(\cdot|m)$, but may also contain additional optimal actions that are not played following m. Without loss of generality, let $supp(\rho(\cdot|m)) = \{a_1, ..., a_k\}$ and $A_R^*(\mu) = \{a_1, ..., a_k, a_{k+1}, ..., a_{k+r}\}$ for some $r \geq 0$. Note that for any $i = 2, ..., k + r$, $\mathbf{u}_R(a_1) \cdot$

$\mu_m = \mathbf{u}_R(a_i) \cdot \mu_m$.

In addition, Equation (5) implies that for any $i = 2, ..., k$, $\mathbf{u}_S(a_1) \cdot \mu_m = \mathbf{u}_S(a_i) \cdot \mu_m$. Combining both Sender's and Receiver's indifference conditions, we have

$$\begin{bmatrix} \mathbf{u}_S(a_2) - \mathbf{u}_S(a_1) \\ ... \\ \mathbf{u}_S(a_k) - \mathbf{u}_S(a_1) \\ \\ \mathbf{u}_R(a_2) - \mathbf{u}_R(a_1) \\ ... \\ \mathbf{u}_R(a_{k+r}) - \mathbf{u}_R(a_1) \end{bmatrix} \mu_m = \mathbf{0}. \tag{6}$$

Let $\hat{\Omega} = \{\omega | \mu_m(\omega) = 0\}$, the (potentially empty) set of states that are not in the support of μ_m. Without loss, suppose that $\hat{\Omega} = \{\omega_1, ...\omega_l\}$ where $\ell \geq 0$. If $\ell > 0$ (i.e., $\hat{\Omega} \neq \emptyset$), then we have

$$\begin{bmatrix} e_1 \\ ... \\ e_\ell \end{bmatrix} \mu_m = \mathbf{0}. \tag{7}$$

Let $\hat{T}_S = \begin{bmatrix} \mathbf{u}_S(a_2) - \mathbf{u}_S(a_1) \\ ... \\ \mathbf{u}_S(a_k) - \mathbf{u}_S(a_1) \end{bmatrix}$, $\hat{T}_R = \begin{bmatrix} \mathbf{u}_R(a_2) - \mathbf{u}_R(a_1) \\ ... \\ \mathbf{u}_R(a_{k+r}) - \mathbf{u}_R(a_1) \end{bmatrix}$, $\hat{E} = \begin{bmatrix} e_1 \\ ... \\ e_\ell \end{bmatrix}$, and $\hat{T} = \begin{bmatrix} \hat{T}_S \\ \hat{T}_R \\ \hat{E} \end{bmatrix}$. Note that $\hat{T}$ is a row-submatrix of the expanded-indifference matrix T^1.

Combining (6) and (7), we know $\hat{T}\mu_m = \mathbf{0}$. Moreover, since $\mu_m \in \Delta\Omega$, we know $\mathbf{1}\mu_m = 1$.

Next we make two observations: (i) $rank(\hat{T}) < |\Omega|$, otherwise the unique solution to $\hat{T}\mu = \mathbf{0}$ is $\mu = \mathbf{0}$. Since we are in a scant-indifferences environment, this means that $\hat{T}$ has full row rank; (ii) vector $\mathbf{1}$ can not be represented as a linear combination of rows of $\hat{T}$. To see why, assume toward contradiction that there exists a row vector $\lambda \in \mathbb{R}^{2k+r+\ell-2}$ such that $\lambda\hat{T} = \mathbf{1}$. This would lead to a contradiction that $1 = \mathbf{1}\mu_m = \lambda\hat{T}\mu_m = \lambda\mathbf{0} = 0$.

Observations (i) and (ii) together imply that the matrix $\begin{bmatrix} \hat{T} \\ \mathbf{1} \end{bmatrix}$ has full row rank. Consequently, we know $rank\left(\begin{bmatrix} \hat{T} \\ \mathbf{1} \end{bmatrix}\right) > rank\left(\begin{bmatrix} \hat{T}_R \\ \hat{E} \\ \mathbf{1} \end{bmatrix}\right)$.

Now, we claim that there exists $x \in \mathbb{R}^n$ such that

$$\begin{bmatrix} \hat{T}_R \\ \hat{E} \\ \mathbf{1} \end{bmatrix} x = 0 \tag{8}$$

and

$$\hat{T}_S\, x \neq 0. \tag{9}$$

To see this, suppose by contradiction that for any x that solves (8), we have $\hat{T}_S\, x = 0$. This would imply that the set of solutions to (8) and the set of solutions to

$$\begin{bmatrix} \hat{T} \\ \mathbf{1} \end{bmatrix} x = 0 \tag{10}$$

coincide. By the Rank-Nullity Theorem, however, the subspace defined by (10) has dimension $|\Omega| - rank\left(\begin{bmatrix} \hat{T} \\ \mathbf{1} \end{bmatrix}\right)$, while the subspace defined by (8) has a higher dimension $|\Omega| - rank\left(\begin{bmatrix} \hat{T}_R \\ \hat{E} \\ \mathbf{1} \end{bmatrix}\right)$.

Consider two vectors, $[\mu_m + \varepsilon x]$ and $[\mu_m - \varepsilon x]$, where $\varepsilon \in \mathbb{R}_{>0}$. First we verify that for sufficiently small ε, $[\mu_m \pm \varepsilon x] \in \Delta\Omega$. Since $\mathbf{1}x = 0$, it follows that $\mathbf{1}\,[\mu_m \pm \varepsilon x] = \mathbf{1}\,[\mu_m] = 1$. For $\omega_j \notin \hat{\Omega}$, we have $[\mu_m]_j > 0$, so for small enough ε, $[\mu_m \pm \varepsilon x]_j \geq 0$. For $\omega_j \in \hat{\Omega}$, we know e_j is a row of $\hat{E}$, so $e_j x = 0$. Consequently, $[\mu_m \pm \varepsilon x]_j = e_j\,[\mu_m \pm \varepsilon x] = [\mu_m]_j = 0$. Thus, $[\mu_m \pm \varepsilon x] \in \Delta\Omega$.

Observe that $A^*_R(\mu_m) = A^*_R(\mu_m \pm \varepsilon x)$. First, for any $a \notin A^*_R(\mu_m)$, if ε is sufficiently small, $a \notin A^*_R(\mu_m \pm \varepsilon x)$. Therefore, $A^*_R(\mu_m \pm \varepsilon x) \subseteq A^*_R(\mu_m)$. But, $\hat{T}_R\, x = 0$ implies that $[\mu_m \pm \varepsilon x] \cdot \mathbf{u}_R(a)$ is constant across $a \in A^*_R(\mu_m)$, so $A^*_R(\mu_m \pm \varepsilon x) = A^*_R(\mu_m)$.

Consider an alternative messaging strategy $\hat{\sigma}$ that is identical to σ, except that the message

m is split into two new messages, m^+ and m^-, which induce the beliefs $\mu_m + \varepsilon x$ and $\mu_m - \varepsilon x$, respectively.[22] We consider $\hat{\rho}$ that agrees with ρ on messages other than $\{m, m^+, m^-\}$ and leads Receiver to break indifferences in Sender's favor following m^+ and m^-. We will show that $(\hat{\sigma}, \hat{\rho})$ yields a strictly higher payoff to Sender, thus contradicting the assumption that (σ, ρ) yields the persuasion payoff.

Since $\hat{T}_S\, x \neq 0$, we know there is an $a_i \in \{a_2, ..., a_k\}$ such that $x \cdot [\mathbf{u}_S(a_i) - \mathbf{u}_S(a_1)] \neq 0$.

Because $a_1 \in A^*_R(\mu_m \pm \varepsilon x) = A^*_R(\mu_m)$, we have

$$\max_{a \in A^*(\mu_m)} [\mu_m + \varepsilon x] \cdot [\mathbf{u}_S(a) - \mathbf{u}_S(a_1)] \geq 0$$

and

$$\max_{a \in A^*(\mu_m)} [\mu_m - \varepsilon x] \cdot [\mathbf{u}_S(a) - \mathbf{u}_S(a_1)] \geq 0.$$

We now establish that at least one of these inequalities has to be strict. Suppose toward contradiction that both hold with equality. The first equality implies $[\mu_m + \varepsilon x] \cdot [\mathbf{u}_S(a_i) - \mathbf{u}_S(a_1)] \leq 0$, which combined with the fact that $\mu_m \cdot \mathbf{u}_S(a_i) = \mu_m \cdot \mathbf{u}_S(a_1)$ implies that $x \cdot [\mathbf{u}_S(a_i) - \mathbf{u}_S(a_1)] \leq 0$. Similarly, the second equality implies that $-x \cdot [\mathbf{u}_S(a_i) - \mathbf{u}_S(a_1)] \leq 0$. Together, this yields that $x \cdot [\mathbf{u}_S(a_i) - \mathbf{u}_S(a_1)] = 0$, a contradiction. Hence, one of the inequalities has to be strict.

Consequently, Sender's interim payoff under $\hat{\sigma}$ (in the event that m is sent under σ) is

$$\begin{aligned} &\frac{1}{2} \max_{a \in A^*(\mu_m)} [\mu_m + \varepsilon x] \cdot \mathbf{u}_S(a) + \frac{1}{2} \max_{a \in A^*(\mu_m)} [\mu_m - \varepsilon x] \cdot \mathbf{u}_S(a) \\ >&\frac{1}{2} [\mu_m + \varepsilon x] \cdot \mathbf{u}_S(a_1) + \frac{1}{2} [\mu_m - \varepsilon x] \cdot \mathbf{u}_S(a_1) \\ =&\mu_m \cdot \mathbf{u}_S(a_1) \end{aligned}$$

Thus, $(\hat{\sigma}, \hat{\rho})$ yields a strictly higher payoff to Sender, contradicting the assumption that (σ, ρ) yields the persuasion payoff. □

Lemma 6. *In a scant-indifferences environment, if a cheap-talk equilibrium (σ, ρ) yields the persuasion payoff and ρ is pure-on-path, then there exists a partitional $\hat{\sigma}$ and a pure strategy $\hat{\rho}$ such that*

[22] It is possible for $M_\sigma = M$, but we can consider an alternative strategy that induces the same outcome as σ and uses only $|A|$ messages. We can also let m play the role of m^+ or m^-, so our assumption that $|M| \geq |A| + 1$ suffices.

$|M_{\hat{\sigma}}| \leq |A|$ and $(\hat{\sigma}, \hat{\rho})$ is a cheap-talk equilibrium and yields the persuasion payoff.

Proof. Suppose a cheap-talk equilibrium (σ, ρ) yields the persuasion payoff and ρ is pure-on-path.

First, we show that for any ω and any m, m' such that $\sigma(m|\omega), \sigma(m'|\omega) > 0$, $\rho(m) = \rho(m')$. The fact that both m and m' are sent in ω implies that $u_S(\rho(m), \omega) = u_S(\rho(m'), \omega)$. Moreover, by Lemma 3, there exist no distinct a and a' such that $u_S(a, \omega) = u_S(a', \omega)$, so it must be that $\rho(m) = \rho(m')$.

Let $A^* = \{a \in A | a = \rho(m) \text{ for some } m \in M_\sigma\}$ be the set of actions that are taken on-path. Without loss, let $A^* = \{a_1, ..., a_k\}$. For each a_i, let $M_i = \{m \in M_\sigma | \rho(m) = a_i\}$ be the set of on-path messages that induce action a_i, and $\Omega_i = \{\omega \in \Omega | supp(\sigma(\cdot|\omega)) \subseteq M_i\}$ be the set of states that induce action a_i. Note that $\{M_i\}_{i=1}^k$ is a partition of M_σ. Moreover, it is easy to see that $\{\Omega_i\}_{i=1}^k$ is a partition of Ω. First, Ω_i cannot be empty because every $a_i \in A^*$ is taken on-path. Second, every $\omega \in \Omega$ belongs to some Ω_i as only actions in A^* are taken on-path; hence, $\cup_i \Omega_i = \Omega$. Finally, the fact that for any ω and any m, m' such that $\sigma(m|\omega), \sigma(m'|\omega) > 0$ we have $\rho(m) = \rho(m')$ implies that if $i \neq j$, Ω_i and Ω_j are disjoint. To see why, suppose toward contradiction that some $\omega \in \Omega_i \cap \Omega_j$. The fact that $\omega \in \Omega_i$ implies there is a message $m \in M_i$ such that $\sigma(m|\omega) > 0$. The fact that $\omega \in \Omega_j$ implies there is a message $m' \in M_j$ such that $\sigma(m'|\omega) > 0$. But this cannot be since $\rho(m) = a_i \neq a_j = \rho(m')$.

Now select one message in each M_i, and label it as m_i.

Next, consider the following alternative strategy profile $(\hat{\sigma}, \hat{\rho})$:

- $\hat{\sigma}(m_i|\omega) = 1$ if $\omega \in \Omega_i$.
- $\hat{\rho}(m_i) = a_i$.
- $\hat{\rho}(m) = a_1$ if $m \in M \backslash \{m_1, ...m_k\}$.

Note that $\hat{\sigma}$ is well defined because $\{\Omega_i\}_{i=1}^k$ is a partition of Ω. By construction, $\hat{\sigma}$ is partitional, $|M_{\hat{\sigma}}| \leq |A|$, and $\hat{\rho}$ is a pure strategy. Moreover, under both (σ, ρ) and $(\hat{\sigma}, \hat{\rho})$, every state in Ω_i induces action a_i with probability 1. Thus, the two strategy profiles induce the same distribution over states and actions, so $(\hat{\sigma}, \hat{\rho})$ also yields the persuasion payoff. It remains to show that $(\hat{\sigma}, \hat{\rho})$ is a cheap-talk equilibrium.

Note that S-BR of (σ, ρ) implies that for any ω and $m \in supp(\sigma(\cdot|\omega))$, we have

$$u_S(\rho(m), \omega) \geq u_S(\rho(m'), \omega) \text{ for all } m' \in M_\sigma.$$

Therefore, for any $\omega \in \Omega_i$, $u_S(a_i, \omega) \geq u_S(a_j, \omega)$ for all $a_j \in A^*$. This implies that $u_S(\hat{\rho}(\hat{\sigma}(\omega)), \omega) \geq u_S(\hat{\rho}(m'), \omega)$ for all $m' \in M$. Hence, $(\hat{\sigma}, \hat{\rho})$ satisfies S-BR.

Fact (σ, ρ) is R-BR requires that for all $m \in M_\sigma$,

$$\sum_{\omega \in \Omega} \mu_0(\omega)\sigma(m|\omega)u_R(\rho(m), \omega) \geq \sum_{\omega \in \Omega} \mu_0(\omega)\sigma(m|\omega)u_R(a', \omega) \quad \text{for all } a' \in A.$$

For any $i \in \{1, ..., k\}$, we sum the inequality above over all $m \in M_i$. Since for $m \in M_i$ we have $\rho(m) = a_i$, this yields

$$\sum_{\omega \in \Omega} \mu_0(\omega) \sum_{m \in M_i} \sigma(m|\omega)u_R(a_i, \omega) \geq \sum_{\omega \in \Omega} \mu_0(\omega) \sum_{m \in M_i} \sigma(m|\omega)u_R(a', \omega) \quad \text{for all } a' \in A.$$

Since for any $m \in M_i$ and $\omega \in \Omega_i$, we have $\sigma(m|\omega) = 0$, the inequality above implies

$$\sum_{\omega \in \Omega_i} \mu_0(\omega) \sum_{m \in M_i} \sigma(m|\omega)u_R(a_i, \omega) \geq \sum_{\omega \in \Omega_i} \mu_0(\omega) \sum_{m \in M_i} \sigma(m|\omega)u_R(a', \omega) \quad \text{for all } a' \in A.$$

Since $\sum_{m \in M_i} \sigma(m|\omega) = 1$ if $\omega \in \Omega_i$, we have

$$\sum_{\omega \in \Omega_i} \mu_0(\omega)u_R(a_i, \omega) \geq \sum_{\omega \in \Omega_i} \mu_0(\omega)u_R(a', \omega) \quad \text{for all } a' \in A. \tag{11}$$

To establish $(\hat{\sigma}, \hat{\rho})$ is R-BR, we need to show that for any $m_i \in M_{\hat{\sigma}}$, we have

$$\sum_{\omega \in \Omega} \mu_0(\omega)\hat{\sigma}(m_i|\omega) \sum_{a \in A} \hat{\rho}(a|m_i) u_R(a, \omega) \geq \sum_{\omega \in \Omega} \mu_0(\omega)\hat{\sigma}(m_i|\omega) u_R(a', \omega) \quad \text{for all } a' \in A.$$

But, by definition of $(\hat{\sigma}, \hat{\rho})$, we know that $\hat{\sigma}(m_i|\omega) = 0$ for $\omega \notin \Omega_i$ and that $\hat{\rho}(a_i|m_i) = 1$. Hence, the inequality above is equivalent to Equation (11). □

A.4 Proof of Theorem 1

We present and prove a result that generalizes Theorem 1 into a threefold equivalence.

Theorem 1′. *Generically, the following statements are equivalent:*

1. *Commitment is valuable.*
2. *Committed Sender values randomization.*
3. *For any optimal persuasion profile (σ, ρ), there exists $m \in M_\sigma$ such that*

$$|\arg\max_{a \in A} \sum_{\omega} \mu_m(\omega) u_R(a, \omega)| \geq 2,$$

where μ_m is defined as $\mu_m(\omega) = \frac{\mu_0(\omega)\sigma(m|\omega)}{\sum_\omega \mu_0(\omega)\sigma(m|\omega)}$.

Proof. We establish the equivalence for any environment that satisfies both partitional-unique-response and scant-indifferences. Since the set of partitional-unique-response environments is generic (Lemma 1) and the set of scant-indifferences environments is generic (Lemma 2), the set of environments that satisfy both properties is also generic.

We will establish that (2) implies (1), then that (1) implies (3), and finally that (3) implies (2).

Since we are in a scant-indifferences environment, (2) implies (1) by Lemma 4.

Next we wish to show that (1) implies (3). We do so by establishing the contrapositive. Suppose that there exists an optimal persuasion profile (σ, ρ) such that for every $m \in M_\sigma$, $\arg\max_{a \in A} \sum_\omega \mu_m(\omega) u_R(a, \omega)$ is unique. This implies that ρ must be pure-on-path. We will construct an optimal persuasion profile $(\sigma, \hat{\rho})$ that it is a cheap-talk equilibrium. Consider the following $\hat{\rho}$: for all $m \in M_\sigma$, let $\hat{\rho}(m) = \rho(m)$; for $m \notin M_\sigma$, let $\hat{\rho}(m) = \rho(m_0)$ for some $m_0 \in M_\sigma$. Since $\hat{\rho}$ and ρ coincide on path, (σ, ρ) and $(\sigma, \hat{\rho})$ yield the same payoffs to both Sender and Receiver. Therefore, $(\sigma, \hat{\rho})$ satisfies R-BR and yields the persuasion payoff. It remains to show that $(\sigma, \hat{\rho})$ is S-BR, which is equivalent to Sender's interim optimality: for each ω,

$$\sum_m \sigma(m|\omega) u_S(\hat{\rho}(m), \omega) \geq u_S(\hat{\rho}(m'), \omega) \tag{12}$$

for all $m' \in M$. First, note that it suffices to show that Equation (12) holds for $m' \in M_\sigma$. Once we establish that, we know $\sum_m \sigma(m|\omega)u_S(\hat{\rho}(m),\omega) \geq u_S(\hat{\rho}(m_0),\omega)$ since $m_0 \in M_\sigma$. Therefore, since $\hat{\rho}(m') = \rho(m_0) = \hat{\rho}(m_0)$ for $m' \notin M_\sigma$, Equation (12) holds for $m' \notin M_\sigma$.

Now, suppose toward contradiction that there exist $\hat{\omega}$ and $\hat{m} \in M_\sigma$ such that $\sum_m \sigma(m|\hat{\omega})u_S(\hat{\rho}(m),\hat{\omega}) < u_S(\hat{\rho}(\hat{m}),\hat{\omega})$. Consider an alternative messaging strategy $\hat{\sigma}$: $\hat{\sigma}(\omega) = \sigma(\omega)$ for $\omega \neq \hat{\omega}$ while $\hat{\sigma}(\hat{\omega})$ sends the same distribution of messages as $\sigma(\hat{\omega})$ with probability $1-\varepsilon$ and otherwise sends message $\hat{m}$. Formally, $\hat{\sigma}(m|\hat{\omega}) = \begin{cases} (1-\varepsilon)\,\sigma(m|\hat{\omega}) & \text{if } m \neq \hat{m} \\ (1-\varepsilon)\,\sigma(m|\hat{\omega}) + \varepsilon & \text{if } m = \hat{m} \end{cases}$.

Fix any $m \in M_\sigma$. Since A is finite, the fact that $\hat{\rho}(m) = \rho(m)$ is the unique $\arg\max_{a\in A} \sum_\omega \mu_m(\omega)u_R(a,\omega)$ implies that $\hat{\rho}(m)$ remains the best response for a neighborhood of beliefs around μ_m. Therefore, for sufficiently small ε, $(\hat{\sigma},\hat{\rho})$ is R-BR. Hence, $(\hat{\sigma},\hat{\rho})$ is a persuasion profile and yields the payoff

$$\begin{aligned} U_S(\hat{\sigma},\hat{\rho}) &= U_S(\sigma,\hat{\rho}) + \varepsilon[u_S(\hat{\rho}(\hat{m}),\hat{\omega}) - \sum_m \sigma(m|\hat{\omega})u_S(\hat{\rho}(m),\hat{\omega})] \\ &> U_S(\sigma,\hat{\rho}). \end{aligned}$$

This contradicts the fact that $(\sigma,\hat{\rho})$ yields the persuasion payoff.

Finally, since we are considering a partitional-unique-response environment, the fact that (3) implies (2) is immediate. □

A.5 Proof of Theorem 2

Lemmas 2 and 4 jointly imply Theorem 2.

A.6 State-independent preferences

As we mentioned in the discussion of related literature, several papers examine value of commitment under the assumption that Sender has state-independent preferences. To connect to that literature, it is worth asking whether our results hold under that assumption. When $|A| \geq 3$, state-independent preferences by Sender mean that we are not in a scant-indifferences environment, so the proofs above do not apply. Nonetheless, Theorems 1 and 2 indeed remain true.

To formalize this, say that environment (u_S, u_R) is transparent if there exists some function

$v : A \to \mathbb{R}$ such that $u_S(a,\omega) = u_S(a,\omega') \equiv v(a)$ for any a, ω, ω'. The set of all transparent environments is $\mathbb{R}^{|A|\,(|\Omega|+1)}$. A set of transparent environments is *transparently-generic* if it has full Lebesgue measure in $\mathbb{R}^{|A|\,(|\Omega|+1)}$. We say that a claim holds *generically in transparent environments*, if it holds for a transparently-generic set of environments.

A transparent environment (u_S, u_R) satisfies *no-duplicate-actions* if for any $a \neq a'$, $v(a) \neq v(a')$. Clearly, the set of no-duplicate-actions transparent environments is transparently-generic.

A strategy profile (σ, ρ) is a *simple babbling cheap-talk equilibrium* if it is a cheap-talk equilibrium in which $|M_\sigma| = 1$ and $\rho(m) = a_0$ for all $m \in M$ and for some $a_0 \in A$.

Lemma 7. *In a no-duplicate-actions transparent environment, if commitment has no value, then there exists a simple babbling cheap-talk equilibrium that yields the persuasion payoff.*

Proof. If commitment has no value, a cheap-talk equilibrium, denoted by (σ, ρ), yields the persuasion payoff. First, we claim that ρ must be pure on-path. Suppose by contradiction that at some on-path message m, $|supp(\rho(\cdot|m))| > 1$. R-BR then implies that Receiver must be indifferent among all actions in $supp(\rho(\cdot|m))$. Since the environment satisfies no-duplicate-actions, Sender must strictly prefers one of the actions in $supp(\rho(\cdot|m))$ over all others. Therefore, an alternative strategy profile where Receiver breaks ties in favor of Sender would still satisfy R-BR while strictly improving Sender's payoff.

Since (σ, ρ) is a cheap-talk equilibrium and the environment is transparent, S-BR implies that for any $m, m' \in M_\sigma$, we have $v(\rho(m)) = v(\rho(m'))$. Moreover, because the environment satisfies no-duplicate-actions, it follows that $\rho(m) = \rho(m')$ for any $m, m' \in M_\sigma$; that is, only a single action is induced in equilibrium.

We now construct a simple babbling cheap-talk equilibrium $(\hat{\sigma}, \hat{\rho})$ that yields the same payoff as (σ, ρ). Choose an arbitrary message $m_0 \in M_\sigma$. Let $\hat{\sigma}(m_0|\omega) = 1$ for all $\omega \in \Omega$, and $\hat{\rho}(m) = \rho(m_0)$ for all $m \in M$. The strategy profile $(\hat{\sigma}, \hat{\rho})$ trivially satisfies S-BR and yields the same payoff as (σ, ρ). Since (σ, ρ) satisfies R-BR, it follows that for each $m \in M_\sigma$, $\sum_\omega \mu_0(\omega)\sigma_S(m|\omega)u_R(\rho(m),\omega) \geq \sum_\omega \mu_0(\omega)\sigma_S(m|\omega)u_R(a,\omega)$ for all $a \in A$. Summing over all m, we obtain

$$\sum_{\omega\in\Omega} \mu_0(\omega) \sum_{m\in M} \sigma_S(m|\omega)u_R(\rho(m),\omega) \geq \sum_{\omega\in\Omega} \mu_0(\omega) \sum_{m\in M} \sigma_S(m|\omega)u_R(a,\omega) \quad \text{for all } a \in A.$$

Since $\rho(m) = \rho(m_0)$ for all $m \in M_\sigma$, we can rewrite the inequality as

$$\sum_{\omega \in \Omega} \mu_0(\omega) u_R(\rho(m_0), \omega) \geq \sum_{\omega \in \Omega} \mu_0(\omega) u_R(a, \omega) \quad \text{for all } a \in A,$$

which then implies $(\hat{\sigma}, \hat{\rho})$ satisfies R-BR. Therefore, $(\hat{\sigma}, \hat{\rho})$ is a simple babbling cheap-talk equilibrium that yields the persuasion payoff. □

Theorem 4. *Generically in transparent environments, commitment is valuable if and only if committed Sender values randomization.*

Proof. We establish the equivalence for any transparent environment that satisfies both partitional-unique-response and no-duplicate-actions. Recall that whether an environment satisfies partitional-unique-response does not depend on Sender's preferences, so the same argument as in Lemma 1 implies that the set of partitional-unique-response transparent environments are transparently-generic. Moreover, the set of no-duplicate-actions transparent environments is transparently-generic. Therefore, the set of transparent environments that satisfy both properties is also transparently-generic.

The same arguments that establish (1) implies (3) and (3) implies (2) in Theorem 1' apply directly to any partitional-unique-response transparent environment, thereby proving the "only if" direction. The "if" direction follows immediately from Lemma 7. □

Theorem 5. *Generically in transparent environments, commitment is valuable if cheap-talk Sender values randomization.*

Proof. The theorem follows immediately from Lemma 7 and the fact that the set of no-duplicate-actions transparent environments is transparently-generic. □

A.7 Proof of Theorem 3

Recall that we consider a setting where for each (a, ω), $u_S(a, \omega)$ is drawn from F and $u_R(a, \omega)$ is drawn from G. Both F and G are atomless, and all variables $\{u_S(a, \omega), u_R(a, \omega)\}_{(a,\omega) \in A \times \Omega}$ are mutually independent. Throughout this section, we fix some atomless F and G. When we say that the probability of some property is q, we mean that when $u_S \sim F$ and $u_R \sim G$, the likelihood that (u_S, u_R) satisfies that property is q. We use the word *event* to refer to a set of environments.

Given u_S, let $\Omega_i^{u_S} = \{\omega \in \Omega | a_i \in \arg\max_{a \in A} u_S(a, \omega)\}$ denote the set of states where a_i is an ideal action for Sender.[23] Note that each ω must belong to at least one $\Omega_i^{u_S}$, but the same ω may appear in multiple $\Omega_i^{u_S}$. Say that u_S is *regular* if $\Omega_i^{u_S} \cap \Omega_j^{u_S} = \emptyset$ for $i \neq j$. Lemmas 2 and 3 jointly imply that the set of u_S that are regular has full Lebesgue measure in $\mathbb{R}^{|A||\Omega|}$. Since F is atomless, this in turn implies that u_S is regular with probability one.

Recall that an environment is felicitous if for each non-empty $\Omega_i^{u_S}$,

$$a_i \in \arg\max_a \sum_{\omega \in \Omega_i^{u_S}} \mu_0(\omega) u_R(a, \omega). \tag{13}$$

A.7.1 Arbitrary state space

In this section, we establish that for any Ω, $\Pr(\text{commitment has no value}) \geq \frac{1}{|A|^{|A|}}$.

Lemma 8. *In any felicitous environment, commitment has no value.*

Proof. Select $|A|$ elements from M and denote them by m_1 through $m_{|A|}$. Consider a pure strategy profile (σ, ρ) such that

- $\sigma(\omega) = m_i$ implies $\omega \in \Omega_i^{u_S}$;[24]
- $\rho(m) = a_i$ for $m = m_i$;
- $\rho(m) = a_1$ for $m \notin \{m_1, ..., m_{|A|}\}$.

From (13), this strategy profile satisfies R-BR. In addition, in every state, Sender achieves his ideal payoff, so S-BR is satisfied and the profile yields the persuasion payoff. Therefore, (σ, ρ) is a cheap-talk equilibrium that yields the persuasion payoff. □

Lemma 9. $\Pr(\textit{felicity}) \geq \frac{1}{|A|^{|A|}}$.

Proof. Fix some regular u_S. Consider any non-empty $\Omega_i^{u_S}$. Given independence and the fact that each $u_R(a, \omega)$ is drawn from the atomless G, each $a \in A$ has an equal chance, $1/|A|$, to maximize

[23] In the body of the paper we denoted this set as simply Ω_i, but for the formal proofs, it is helpful to keep track of the fact that this set depends on the randomly drawn u_S.

[24] If u_S is not regular, it could be that ω belongs to $\Omega_i^{u_S}$ and $\Omega_j^{u_S}$ for distinct i and j. If so, it does not matter whether $\sigma(\omega)$ is m_i or m_j. The fact that $\cup_i \Omega_i^{u_S} = \Omega$, implies there exists a σ such that $\sigma(\omega) = m_i$ implies $\omega \in \Omega_i^{u_S}$.

$\sum_{\omega \in \Omega_i^{u_S}} \mu_0(\omega) u_R(a,\omega)$. In particular,

$$\Pr\left(a_i \in \arg\max_a \sum_{\omega \in \Omega_i^{u_S}} \mu_0(\omega) u_R(a,\omega) \big| u_S\right) = \frac{1}{|A|}.$$

Moreover, this probability is independent across i. Therefore,

$$\Pr(\text{felicity}|u_S) = \prod_{i:\Omega_i^{u_S} \text{ is non-empty}} (1/|A|) \geq \frac{1}{|A|^{|A|}}. \tag{14}$$

This is an inequality because some $\Omega_i^{u_S}$ could be empty. So, we have established that for any regular u_S, $\Pr(\text{felicity}|u_S) \geq \frac{1}{|A|^{|A|}}$. Since u_S is regular with probability one, this in turn implies $\Pr(\text{felicity}) \geq \frac{1}{|A|^{|A|}}$. □

Lemmas 8 and 9 jointly imply that $\Pr(\text{commitment has no value}) \geq \frac{1}{|A|^{|A|}}$.

A.7.2 Limit as $|\Omega| \to \infty$

In this section, we establish that as $|\Omega| \to \infty$, $\Pr(\text{commitment has no value}) \to \frac{1}{|A|^{|A|}}$.

We first give an outline of the proof. The proof is broken up into two major parts. First, recall that felicity implies that commitment has no value, but the converse does not hold in general. We first show that generically, if the environment is jointly-inclusive,[25] then commitment having no value implies felicity (Lemma 10). We then show, that as $|\Omega| \to \infty$, the probability of joint-inclusivity converges to one (Lemma 11). Combining these two results, we conclude that as $|\Omega| \to \infty$, $\Pr(\text{commitment has no value}) \to \Pr(\text{felicity})$.

Second, recall that $\Pr(\text{felicity}) \geq \frac{1}{|A|^{|A|}}$ and that the reason this is an inequality is the possibility that some $\Omega_i^{u_S}$ might be empty. When no $\Omega_i^{u_S}$ is empty, it is indeed the case that $\Pr(\text{felicity}) = \frac{1}{|A|^{|A|}}$ (Lemma 13). We then show, that as $|\Omega| \to \infty$, the probability that some $\Omega_i^{u_S}$ is empty converges to zero (Lemma 14). Combining these two results, we conclude that as $|\Omega| \to \infty$, $\Pr(\text{felicity}) \to \frac{1}{|A|^{|A|}}$.

Lemma 10. *If commitment has no value in a jointly-inclusive environment that satisfies partitional-unique-response and scant-indifferences, then this environment is felicitous.*

[25] Recall that an environment is jointly-inclusive if for every action a, there is some state ω such that a is the unique ideal action for both Sender and Receiver in ω.

Proof. Consider a jointly-inclusive environment that satisfies partitional-unique-response and scant-indifferences and suppose that commitment has no value. By Lemma 4, there is a partitional σ and a pure strategy ρ such that $|M_\sigma| \leq |A|$ and (σ, ρ) is a cheap-talk equilibrium and yields the persuasion payoff.

First note that every action is induced under (σ, ρ); that is, for any $a \in A$, there exists ω such that $a = \rho(\sigma(\omega))$. To see why, suppose toward contradiction that there is an $a^* \in A$ that is not induced. Since the environment is jointly-inclusive, there exists ω^* such that

$$u_S(a^*, \omega^*) > u_S(a, \omega^*) \text{ and } u_R(a^*, \omega^*) > u_R(a, \omega^*) \text{ for all } a \neq a^*. \tag{15}$$

Since $|M_\sigma| \leq |A| < |M|$, there is an unsent message, say m^*.

Consider the strategy profile $(\hat{\sigma}, \hat{\rho})$:

- $\hat{\sigma}(\omega) = \sigma(\omega)$ for $\omega \neq \omega^*$, and $\hat{\sigma}\left(m|\omega^*\right) = \begin{cases} (1-\varepsilon) & \text{if } m = \sigma(\omega^*) \\ \varepsilon & \text{if } m = m^* \\ 0 & \text{otherwise} \end{cases}$.
- $\hat{\rho}(m) = \rho(m)$ for $m \neq m^*$, and $\hat{\rho}(m^*) = a^*$.

We show that $(\hat{\sigma}, \hat{\rho})$ is R-BR for sufficiently small ε. For any $m \notin \{\sigma(\omega^*), m^*\}$, Receiver's belief upon observing m is unchanged, so $\hat{\rho}(m) = \rho(m)$ remains a best response. For $m = m^*$, (15) implies that $\hat{\rho}(m^*) = a^*$ is the best response. For $m = \sigma(\omega^*)$, the fact the environment satisfies partitional-unique-response implies that $\hat{\rho}\left(m\right) = \rho(m)$ is the unique best response to μ_m. Moreover, since A is finite, this further implies that $\hat{\rho}(m)$ remains the best response for a neighborhood of beliefs around μ_m. Therefore, for sufficiently small ε, $\hat{\rho}(m)$ remains a best response.

Now, note that $\rho(\sigma(\omega^*)) \neq a^*$ because a^* is not induced under (σ, ρ). By (15),

$$\begin{aligned} U_S(\hat{\sigma}, \hat{\rho}) &= U_S(\sigma, \rho) + \varepsilon[u_S(a^*, \omega^*) - u_S(\rho(\sigma(\omega^*)), \omega^*)] \\ &> U_S(\sigma, \rho). \end{aligned}$$

This contradicts the fact that (σ, ρ) yields the persuasion payoff. Hence, we have established that every action is induced under (σ, ρ).

Next, we show that this fact, coupled with the maintained assumptions, implies that the environment is felicitous. Recall that (σ, ρ) is a cheap-talk equilibrium; hence for each ω,

$$u_S(\rho(\sigma(\omega)), \omega) \geq u_S(\rho(m), \omega) \text{ for all } m \in M.$$

Since every action is induced under (σ, ρ), the inequality above is equivalent to

$$u_S(\rho(\sigma(\omega)), \omega) \geq u_S(a, \omega) \text{ for all } a \in A.$$

Moreover, since the environment satisfies scant-indifferences, Lemma 3 implies that

$$u_S(\rho(\sigma(\omega)), \omega) > u_S(a, \omega) \text{ for all } a \neq \rho(\sigma(\omega)). \tag{16}$$

Hence, $\Omega_i^{u_S} = \{\omega \in \Omega | \rho(\sigma(\omega)) = a_i\}$ and $\Omega_i^{u_S} \cap \Omega_j^{u_S} = \emptyset$ for $i \neq j$. Let $M_i = \{m \in M_\sigma | \rho(m) = a_i\}$. For each i and each $m \in M_i$, R-BR of (σ, ρ) implies

$$\sum_{\omega \in \{\omega : \sigma(\omega) = m\}} \mu_0(\omega) u_R(a_i, \omega) \geq \sum_{\omega \in \{\omega : \sigma(\omega) = m\}} \mu_0(\omega) u_R(a', \omega) \text{ for all } a' \in A.$$

Summing over all $m \in M_i$, and noting that $\cup_{m \in M_i} \{\omega : \sigma(\omega) = m\} = \{\omega \in \Omega | \rho(\sigma(\omega)) = a_i\} = \Omega_i^{u_S}$, we have

$$\sum_{\omega \in \Omega_i^{u_S}} \mu_0(\omega) u_R(a_i, \omega) \geq \sum_{\omega \in \Omega_i^{u_S}} \mu_0(\omega) u_R(a', \omega) \text{ for all } a' \in A.$$

Thus, the environment is felicitous. □

Lemma 11. *As* $|\Omega| \to \infty$*,* Pr(*joint-inclusivity*) $\to 1$.

Proof. Let $E_{a,\omega}$ denote the event that a is the unique ideal action for both Sender and Receiver in state ω. Let $E_a = \cup_{\omega \in \Omega} E_{a,\omega}$ denote the event that action a is the unique ideal action for both Sender and Receiver in some state. Let $E = \cap_{a \in A} E_a$ denote joint-inclusivity: each action is uniquely ideal for for both Sender and Receiver in some state. Our goal is to show that $\Pr(E) \to 1$ as $|\Omega| \to \infty$.

Since F and G are atomless and payoffs are independent, in each state ω, the probability that any given action a is the unique ideal action for Sender is $1/|A|$, and the same holds for Receiver. Hence, $\Pr(E_{a,\omega}) = 1/|A|^2$ for any a and ω.

Moreover, for each a, the events $E_{a,\omega}$ are independent across ω. Therefore,

$$\begin{aligned}\Pr(E_a) &= \Pr(\cup_\omega E_{a,\omega}) \\ &= 1 - \Pr(\cap_\omega E^c_{a,\omega}) \\ &= 1 - \prod_{\omega\in\Omega} \Pr(E^c_{a,\omega}) \\ &= 1 - \left(1 - \frac{1}{|A|^2}\right)^{|\Omega|}.\end{aligned}$$

Finally,

$$\begin{aligned}\Pr(E) &= \Pr(\cap_{a\in A} E_a) \\ &= 1 - \Pr(\cup_{a\in A} E^c_a) \\ &\geq 1 - \sum_{a\in A} \Pr(E^c_a) \\ &= 1 - |A|\left(1 - \frac{1}{|A|^2}\right)^{|\Omega|} \\ &\to 1 \quad \text{as } |\Omega| \to \infty.\end{aligned}$$

□

Lemma 12. *As* $|\Omega| \to \infty$*,* Pr(*commitment has no value*) $\to$ Pr(*felicity*)*.*

Proof. Let *PPS* denote the event that the environment is jointly-inclusive and satisfies partitional-unique-response and scant-indifferences. We know that in any PPS environment, if commitment has no value, then the environment is felicitous (Lemma 10) Hence, Pr (felicity|PPS) $\geq$ Pr (commitment has no value|PPS). As $|\Omega| \to \infty$, Pr (PPS) $\to 1$ (Lemmas 1, 2, and 11). Hence, As $|\Omega| \to \infty$, Pr (felicity) $\geq$ Pr (commitment has no value). Moreover, in general Pr (commitment has no value) $\geq$ Pr (felicity). Thus, as $|\Omega| \to \infty$, Pr(commitment has no value) $\to$ Pr(felicity). □

Say an environment is *Sender-inclusive* if $\Omega_i^{u_S}$ is non-empty for all i.

Lemma 13. $\Pr(\textit{felicity}|\textit{Sender-inclusivity}) = \frac{1}{|A|^{|A|}}$

Proof. As noted earlier in Equation (14), $\Pr(\text{felicity}|\text{regular } u_S) = \prod_{i:\Omega_i^{u_S} \text{ is non-empty}} \frac{1}{|A|}$. If the environment is Sender-inclusive, no $\Omega_i^{u_S}$ is empty, so $\Pr(\text{felicity}|\text{Sender-inclusivity \& regular } u_S) = \frac{1}{|A|^{|A|}}$. Since u_S is regular with probability one, we have $\Pr(\text{felicity}|\text{Sender-inclusivity}) = \frac{1}{|A|^{|A|}}$. □

Lemma 14. *As* $|\Omega| \to \infty$, $\Pr(\textit{Sender-inclusivity}) \to 1$.

Proof. Obviously, any jointly-inclusive environment is Sender-inclusive. Thus, this Lemma is a corollary of Lemma 11. □

Lemma 15. *As* $|\Omega| \to \infty$, $\Pr(\textit{felicity}) \to \frac{1}{|A|^{|A|}}$.

Proof. This follows from Lemmas 13 and 14. □

Lemmas 12 and 15 jointly yield the fact that, a $|\Omega| \to \infty$, $\Pr(\text{commitment has no value}) \to \frac{1}{|A|^{|A|}}$.

B Online Appendix

B.1 Value of partial commitment

We consider a partial commitment setting where Sender can commit to a distribution of messages, as in Lin and Liu (2024). For any Sender's strategy σ, let $D(\sigma) := \{\sigma' : \Omega \to \Delta M | \sum_\omega \mu_0(\omega)\sigma'(m|\omega) = \sum_\omega \mu_0(\omega)\sigma(m|\omega), \forall m\}$ denote the set of messaging strategies that preserve the same distribution of messages.

We say a profile (σ^*, ρ^*) is a *curve equilibrium* if

$$\sigma^* \in \arg\max_{\sigma \in D(\sigma^*)} U_S(\sigma, \rho^*)$$

$$\rho^* \in \arg\max_{\rho} U_R(\sigma^*, \rho).$$

The first condition requires that Sender has no incentive to deviate to any other messaging strategy that preserves the same distribution of messages as σ^*, and the second condition requires Receiver to play a best response (i.e., R-BR).

The *curve payoff* is the maximum U_S induced by a curve equilibrium. The *curve partitional payoff* is the maximum U_S induced by a partitional curve equilibrium. We say Sender *values committing to a curve* if the curve payoff is strictly higher than the cheap-talk payoff, and that *curve-committed Sender values randomization* if the curve payoff is strictly higher than the curve partitional payoff.

A function $u_S : A \times \Omega \to \mathbb{R}$ is *strictly supermodular* if there exists a total order $>_A$ on A and a total order $>_\Omega$ on Ω such that for $a' >_A a$ and $\omega' >_\Omega \omega$,

$$u_S(a', \omega') - u_S(a, \omega') > u_S(a', \omega) - u_S(a, \omega).$$

To simplify notation, once we fix a strictly supermodular u_S, we use $>$ in place of $>_A$ and $>_\Omega$.

A set of Receiver's preferences is generic if it has full Lebesgue measure in $\mathbb{R}^{|A|\,|\Omega|}$.

We say that an *outcome distribution* $\pi : \Omega \to \Delta A$ is induced by a profile (σ, ρ) if $\pi(a|\omega) = \sum_{m \in M} \rho(a|m)\sigma(m|\omega)$ for all ω, a. The following lemma offers a characterization of outcome distributions that can be induced by a curve equilibrium.

Lemma 16. *Fix any strictly supermodular u_S. An outcome distribution $\pi : \Omega \to \Delta A$ is induced by some curve equilibrium (σ, ρ) where ρ is pure strategy on-path if and only if*

1. π is comonotone; that is, for any $a < a'$, if $\pi(a|\omega) > 0$ and $\pi(a'|\omega') > 0$, we must have $\omega \leq \omega'$;

2. π is u_R-obedient: for each $a, a' \in A$,

$$\sum_{\omega \in \Omega} \pi(a|\omega)\mu_0(\omega) \left[u_R(a, \omega) - u_R(a', \omega)\right] \geq 0.$$

Proof. The lemma restates Theorem 1 and Lemma 1 in Lin and Liu (2024), with the small caveat that in Lin and Liu (2024), ρ is restricted to be pure both on and off-path. However, note that off-path actions do not affect whether a profile is a curve equilibrium. Thus, the lemma follows. □

Theorem 6. *Fix any strictly supermodular u_S. For a generic set of Receiver's preferences, if Sender values committing to a curve, then a curve-committed Sender values randomization.*

Proof. We consider any Receiver preference that satisfies partitional-unique-response. By Lemma

1, this set of preferences is generic.

We prove the statement by contraposition. Suppose that a curve-committed Sender does not value randomization. This means there exists a partitional curve equilibrium, denoted by (σ, ρ), that yields the curve payoff. Since σ is partitional, partitional-unique-response and R-BR of (σ, ρ) imply that ρ is a pure strategy on-path. We will construct a strategy profile $(\sigma, \hat{\rho})$ that is a cheap-talk equilibrium and yields the curve payoff.

Consider the following $\hat{\rho}$: for all $m \in M_\sigma$, let $\hat{\rho}(m) = \rho(m)$; for $m \notin M_\sigma$, let $\hat{\rho}(m) = \rho(m_0)$ for some $m_0 \in M_\sigma$. Since $\hat{\rho}$ and ρ coincide on path, (σ, ρ) and $(\sigma, \hat{\rho})$ induce that same outcome distribution and yield the same payoffs to both Sender and Receiver. Therefore, $(\sigma, \hat{\rho})$ satisfies R-BR and yields the curve payoff. It remains to show that $(\sigma, \hat{\rho})$ is S-BR, which is equivalent to Sender's interim optimality: for each ω,

$$u_S(\hat{\rho}(\sigma(\omega)), \omega) \geq u_S(\hat{\rho}(m'), \omega) \tag{17}$$

for all $m' \in M$. Note that it suffices to show that Equation (17) holds for $m' \in M_\sigma$. Once we establish that, we know $\sum_m \sigma(m|\omega) u_S(\hat{\rho}(m), \omega) \geq u_S(\hat{\rho}(m_0), \omega)$ since $m_0 \in M_\sigma$. Therefore, since $\hat{\rho}(m') = \rho(m_0) = \hat{\rho}(m_0)$ for $m' \notin M_\sigma$, Equation (17) holds for $m' \notin M_\sigma$.

Since (σ, ρ) is a curve equilibrium and ρ is pure on-path, by Lemma 16, the induced outcome distribution, denoted by π, satisfies comonotonicity and u_R-obedience. In addition, since σ is partitional and ρ is pure on-path, the induced outcome distribution π is also a pure mapping. Moreover, π being comonotone can be strengthened to π being monotone partitional:

$$\forall a < a', \; \pi(a|\omega) > 0 \text{ and } \pi(a'|\omega') > 0 \text{ implies } \omega < \omega'. \tag{18}$$

Moreover, since the environment satisfies partitional-unique-response, u_R-obedience can be strengthened to strict u_R-obedience: for each $a \in A^* \equiv \cup_{\omega \in \Omega} \text{supp}(\pi(\cdot|\omega))$ and $a' \in A / \{a\}$,

$$\sum_{\omega \in \Omega} \pi(a|\omega) \mu_0(\omega) \left[u_R(a, \omega) - u_R(a', \omega) \right] > 0. \tag{19}$$

For each $a \in A^*$, let $\Omega_a = \{\omega | \hat{\rho}(\hat{\sigma}(\omega)) = a\}$ denote the set of states that induce action a. By (18),

$\{\Omega_a\}_{a\in A^*}$ forms a monotone partition of Ω: $\cup_{a\in A^*}\Omega_a = \Omega$, and for any $a < a'$, $\omega \in \Omega_a$, $\omega' \in \Omega_{a'}$, we have $\omega < \omega'$.

To establish Equation (17), it suffices to show that for each $\omega \in \Omega$, $u_S(a', \omega) \le u_S(\pi(\omega), \omega)$ for all $a' \in A^*$. Suppose, toward a contradiction, that there exists $\omega^* \in \Omega$ and $a' \in A^*$ such that $u_S(a', \omega^*) > u_S(\pi(\omega^*), \omega^*)$. Without loss of generality, we assume $a' > \pi(\omega^*)$; the proof for $a' < \pi(\omega^*)$ follows symmetrically.

Let $a^* \equiv \pi(\omega^*)$ and $\hat{a} \in \min\{\arg\max_{a'>a^*, a'\in A^*} u_S(a', \omega)\}$ denote type ω^*'s smallest optimal action among $\{a'|a' > a^*\}$. Let $\tilde{\omega} = \max\{\omega|\pi(\omega) < \hat{a}\}$ denote the largest type that induces an action smaller than $\hat{a}$. Let $\tilde{a} \equiv \pi(\tilde{\omega}) < \hat{a}$. Since $\hat{a}$ is ω^*'s smallest optimal action among $\{a'|a' > a^*\}$, $u_S(\hat{a}, \omega^*) > u_S(\tilde{a}, \omega^*)$. By definition, $\tilde{\omega} \ge \omega^*$, so by supermodularity,

$$u_S(\hat{a}, \tilde{\omega}) - u_S(\tilde{a}, \tilde{\omega}) \ge u_S(\hat{a}, \omega^*) - u_S(\tilde{a}, \omega^*) > 0. \tag{20}$$

We now construct an alternative outcome distribution $\hat{\pi}$: $\hat{\pi}(\omega) = \pi(\omega)$ for $\omega \neq \tilde{\omega}$, while $\hat{\pi}(\tilde{\omega})$ induces action $\tilde{a}$ with probability $1 - \varepsilon$ and induces action $\hat{a}$ with probability ε. By (20), $\hat{\pi}$ yields a strictly higher value than π. In addition, by (19), for sufficiently small ε, $\hat{\pi}$ remains obedient.

Lastly, we show that $\hat{\pi}$ satisfies comonotonicity. To see this, first note that whether an outcome distribution $\hat{\pi}$ satisfies comonotonicity depends only on its support; that is, the set of (a, ω) such that $\hat{\pi}(a|\omega) > 0$. By construction, the supports of $\hat{\pi}$ and π differ only in that $\hat{\pi}$'s support contains an additional element, $(\tilde{\omega}, \hat{a})$. Since π is comonotone, to establish that $\hat{\pi}$ is comonotone, it suffices to show that: for any $a < \hat{a}$ and $\omega \in \Omega_a$, we have $\omega \le \tilde{\omega}$; for any $a' > \hat{a}$ and $\omega' \in \Omega_{a'}$, we have $\omega' \ge \tilde{\omega}$. To prove the first part, note that $\{\Omega_a\}_{a\in A^*}$ forms a monotone partition of Ω, which implies that for any $a < \hat{a}$ and $\omega \in \Omega_a$, we have $\omega < \pi(\hat{a})$. Recall that $\tilde{\omega} = \max\{\omega|\pi(\omega) < \hat{a}\}$ is largest type that induces an action smaller than $\hat{a}$; therefore, $\omega \le \tilde{\omega}$. To prove the second part, note that since $\{\Omega_a\}_{a\in A^*}$ forming a monotone partition, it follows that that for any $a' > \hat{a}$, and $\omega' \in \Omega_{a'}$, we have $\omega' > \min \Omega_{\hat{a}} > \tilde{\omega}$.

Since $\hat{\pi}$ that satisfies comonotonicity and u_R-obedience, by Lemma 16, there exists a curve equilibrium that yields a strictly higher payoff than (σ, ρ). This contradicts the fact that (σ, ρ) yields the curve payoff. □

B.2 Away from zeros

Theorem 1 tells us that, generically, commitment has zero value if and only if randomization has zero value. A natural question is whether, generically, a small value of commitment implies or is implied by a small value of randomization. This section establishes that the answer is no.

We begin by illustrating the role of the genericity condition for Theorem 1. We present two examples. The first example presents an environment (that violates partitional-unique-response) where commitment is valuable but randomization is not. The second example presents an environment (that violates scant indifferences) where randomization is valuable but commitment is not.

Then, we build on the first example to construct a positive measure of environments where the value of commitment is arbitrarily large but the value of randomization is arbitrarily small. We build on the second example to construct a positive measure of environments where the value of randomization is arbitrarily large but the value of of commitment is arbitrarily small.

Example 1. Consider $\Omega = \{\omega_1, \omega_2\}$ with prior $\mu_0(\omega_1) = \mu_0(\omega_2) = 0.5$ and $A = \{a_1, a_2\}$. Players' payoffs are given in Table 1, where the parameter $k > 0$. Receiver's best response is a_1 when $\mu \equiv \mu(\omega_2) \in [0, 1)$, and she is indifferent between a_1 and a_2 when $\mu = 1$. The concavification of Sender's indirect value function is depicted in Figure 2.

u_S	a_1	a_2
ω_1	0	k
ω_2	0	k

u_R	a_1	a_2
ω_1	1	0
ω_2	1	1

Table 1: Sender and Receiver's payoffs

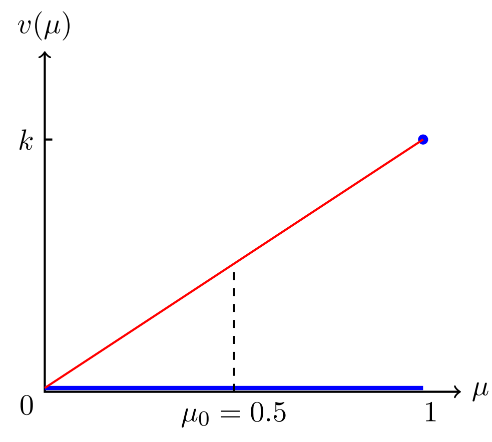

Figure 2: Concavification

Clearly, full revelation is the unique optimal information structure, which yields a payoff of

$k/2$; therefore, Sender does not value randomization. In addition, the only possible cheap-talk equilibrium outcome is babbling, which yields a payoff of 0. Hence, commitment is valuable.

Example 2. Consider $\Omega = \{\omega_1, \omega_2\}$ with prior with prior $\mu_0(\omega_1) = \mu_0(\omega_2) = 0.5$ and $A = \{a_1, a_2, a_3\}$. Players' payoffs are given in Table 2, where the parameter $k > 0$. Receiver's best response is a_1 when $\mu \in [0, 1/3]$, a_2 when $\mu \in [1/3, 2/3]$, and a_3 when $\mu \equiv \mu(\omega_2) \in [2/3, 1]$. This leads to Sender's indirect utility function (blue) and its concave envelope (red) depicted in Figure 3.

u_S	a_1	a_2	a_3
ω_1	0	$2k$	$-2k$
ω_2	$3k$	$-k$	k

u_R	a_1	a_2	a_3
ω_1	1	0	-2
ω_2	-2	0	1

Table 2: Sender's and Receiver's payoffs

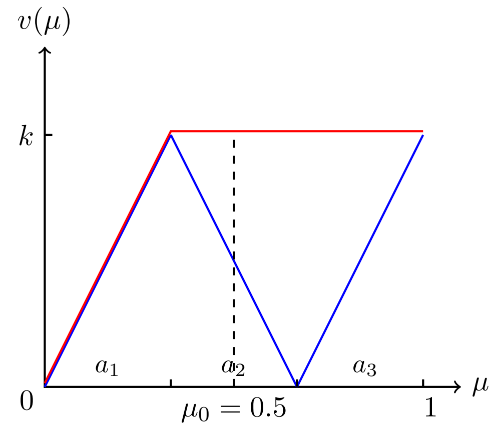

Figure 3: Concavification

Sender values randomization, because the unique optimal information structure that induces a posterior 1/3 cannot be generated by a partitional messaging strategy. Since Sender's indirect utility function is continuous, by Lipnowski (2020), Sender does not value commitment. For the sake of completeness, we construct a cheap talk equilibrium that yields the persuasion payoff k to Sender.

Consider a strategy profile (σ, ρ) with two on-path messages m_1, m_2: $\sigma(m_1|\omega_2) = 1/2$, $\sigma(m_2|\omega_2) = 1/2$, $\sigma(m_1|\omega_1) = 1$; $\rho(a_1|m_1) = 1/2$, $\rho(a_2|m_1) = 1/2$ $\rho(a_3|m_2) = 1$.

The profile satisfies R-BR because the posterior upon observing m_1 is $1/3$ and when observing m_2 is 1. We next show that the profile also satisfies S-BR. For type ω_2 Sender, the expected payoff of sending message m_2 is k and the expected payoff of sending message m_1 is $\frac{1}{2} \cdot 3k + \frac{1}{2} \cdot (-k) = k$, so type ω_2 Sender is indifferent and has no incentive to deviate. For type ω_1 Sender, the expected payoff

u_S	a_1	a_2
ω_1	$0 + s_{11}$	$k + s_{12}$
ω_2	$0 + s_{21}$	$k + s_{22}$

u_R	a_1	a_2
ω_1	$1 + r_{11}$	r_{12}
ω_2	$1 + r_{21}$	$1 + r_{22}$

Table 3: Sender's and Receiver's payoffs

of sending message m_2 is $-2k$ and the expected payoff of sending message m_1 is $\frac{1}{2} \cdot 0 + \frac{1}{2} \cdot 2k = k$, so he strictly prefer sending message m_1. Hence, (σ, ρ) is a cheap-talk equilibrium that yields the persuasion payoff.

B.2.1 Large value of commitment, small value of randomization

We will construct a positive measure of environments where the value of commitment is arbitrarily large but the value of randomization is arbitrarily small. Formally, given an environment (u_S, u_R), let $\Delta_{\mathcal{R}}(u_S, u_R) =$ *Persuasion Payoff* $-$ *Partitional Persuasion Payoff* and $\Delta_{\mathcal{C}}(u_S, u_R) =$ *Persuasion Payoff* $-$ *Cheap-Talk Payoff* denote the value of randomization and the value of commitment in this environment. We will establish that for any $\varepsilon > 0$ and $B > 0$, there exist Ω, A, and μ_0, and a positive measure set of environments $E \subseteq \mathbb{R}^{2\,|A|\,|\Omega|}$ such that, for any $(u_S, u_R) \in E$, we have $\Delta_{\mathcal{R}}(u_S, u_R) < \varepsilon$ and $\Delta_{\mathcal{C}}(u_S, u_R) > B$.

Fix any $\varepsilon > 0$ and $B > 0$, we perturb players' payoffs in Example 1 to construct a positive measure of environments in which $\Delta_{\mathcal{R}} < \varepsilon$ and $\Delta_{\mathcal{C}} > B$. The idea is that for sufficiently small perturbations in the proper direction, the changes to the persuasion payoff, partitional persuasion payoff, and cheap-talk payoff will also be small. Since in Example 1, the value of randomization is zero and the value of commitment can be arbitrarily large (when scaling up k), we obtain a positive measure of environments with small value of randomization and large value of commitment.

Players' payoffs are perturbed as in Table 3, where s_{ij}, r_{ij} are small perturbations to Sender's and Receiver's payoffs, respectively, when action a_j is taken in state ω_i.

Let $s_{ij} \in [0, \delta]$, $r_{11}, r_{21} \in [-\delta, 0]$, and $r_{12}, r_{22} \in [0, \delta]$, where $\delta > 0$. These perturbations generate a positive measure set of preferences, denoted by E^k_δ. We will establish that for $k = 2B + 2$ and

$\delta < \min\{\frac{1}{4}, \frac{\varepsilon}{2(B+2+\varepsilon)}\}$, for any $(u_S, u_R) \in E_\delta^k$, the value of commitment is greater than B and the value of randomization is less than ε.

Consider any $(u_S, u_R) \in E_\delta^k$. Since $\delta < 1/4$, Receiver's best response is a_2 iff $\mu \geq \mu^* \equiv \frac{1+r_{11}-r_{12}}{1+r_{11}-r_{21}-r_{12}+r_{22}} \in [1-2\delta, 1]$. Sender's indirect value function is thus

$$v(\mu) = \begin{cases} s_{11} + (s_{21} - s_{11})\mu & \text{if } \mu < \mu^* \\ 2B + 2 + s_{12} + (s_{22} - s_{12})\mu & \text{if } \mu \geq \mu^*. \end{cases}$$

By inducing beliefs $\mu = 0$ and $\mu = \mu^*$, Sender achieves her persuasion payoff

$$\frac{1}{2\mu^*}[2B + 2 + s_{12} + (s_{22} - s_{12})\mu^*] + (1 - \frac{1}{2\mu^*})s_{11}.$$

Meanwhile, full revelation yields a payoff of

$$B + 1 + \frac{s_{11} + s_{22}}{2}.$$

Therefore,

$$\begin{aligned} \Delta_{\mathcal{R}}(u_S, u_R) \leq & \frac{1}{2\mu^*}[2B + 2 + s_{12} + (s_{22} - s_{12})\mu^*] + (1 - \frac{1}{2\mu^*})s_{11} - B - 1 - \frac{s_{11} + s_{22}}{2} \\ = & \frac{1-\mu^*}{\mu^*}(B + 1 + \frac{s_{12} - s_{11}}{2}) \\ \leq & \frac{2\delta}{1-2\delta}(B + 1 + \delta) \\ < & \varepsilon \end{aligned}$$

where the third line follows from $\mu^* \geq 1 - 2\delta$ and $s_{ij} \in [0, \delta]$, and the last line follows from $\delta < 1$ and $\delta < \frac{\varepsilon}{2(B+2+\varepsilon)}$.

In addition, since $\delta < 1/4$, a_2 is Sender's preferred action regardless of the states. It follows that any cheap-talk equilibrium outcome must be the babbling outcome where Receiver takes action a_1 with probability 1. Therefore, Sender's cheap-talk payoff is $\frac{s_{11}+s_{21}}{2}$. Hence, $\Delta_{\mathcal{C}}(u_S, u_R) =$ *Persuasion Payoff* $-$ *Cheap-Talk Payoff* $\geq B + 1 + \frac{s_{11}+s_{22}}{2} - \frac{s_{11}+s_{21}}{2} > B$, where the weak inequality follows from the fact that persuasion payoff is greater than the payoff from full revelation, and the

u_S	a_1	a_2	a_3
ω_1	$0 + s_{11}$	$2k + s_{12}$	$-2k + s_{13}$
ω_2	$3k + s_{21}$	$-k + s_{22}$	$k + s_{23}$

u_R	a_1	a_2	a_3
ω_1	$1 + r_{11}$	$0 + r_{12}$	$-2 + r_{13}$
ω_2	$-2 + r_{21}$	$0 + r_{22}$	$1 + r_{23}$

Table 4: Sender's and Receiver's payoffs

strict inequality follows from $\frac{s_{22}-s_{21}}{2} \geq \frac{-\delta}{2} > -1$.

B.2.2 Small value of commitment, large value of randomization

We will establish that for any $\varepsilon > 0$ and $B > 0$, there exist finite spaces Ω, A, a prior μ_0, and a positive measure set of environments $E \subseteq \mathbb{R}^{2|A||\Omega|}$ such that, for any $(u_S, u_R) \in E$, we have $\Delta_{\mathcal{C}}(u_S, u_R) < \varepsilon$ and $\Delta_{\mathcal{R}}(u_S, u_R) > B$.

Fix any $\varepsilon > 0$ and $B > 0$, we perturb players' payoffs in Example 2 to construct a positive measure of environments in which $\Delta_{\mathcal{C}} < \varepsilon$ and $\Delta_{\mathcal{R}} > B$. Similar to Section B.2.1, the idea is to show that changes to the persuasion payoff, partitional persuasion payoff, and cheap-talk payoff are small under small perturbations. Since in Example 2, the value of commitment is zero and the value of randomization can be arbitrarily large (when scaling up k), we obtain a positive measure of environments with small value of commitment and large value of randomization.

Players' payoffs are perturbed as in Table 4, where s_{ij}, r_{ij} are small perturbations to Sender's and Receiver's payoffs, respectively, when action a_j is taken in state ω_i.

Let $s_{ij}, r_{ij} \in [0, \delta]$, where $\delta > 0$. These perturbations generate a positive measure set of preferences, denoted by E_δ^k. We will establish that there exists $\delta^* > 0$ such that for $k = 2B + 2\varepsilon$ and $\delta < \delta^*$, for any $(u_S, u_R) \in E_\delta^k$, the value of randomization is greater than B and the value of commitment is less than ε.

Consider any $(u_S, u_R) \in E^k_\delta$ for $\delta < \min\{\frac{1}{4}, k\}$. Since $\delta < 1/4$, Receiver's best response is

$$a_R(\mu) = \begin{cases} a_1 & \text{if } \mu \in [0, \mu_{12}] \\ a_2 & \text{if } \mu \in [\mu_{12}, \mu_{23}] \\ a_3 & \text{if } \mu \in [\mu_{23}, 1] \end{cases}$$

where $\mu_{12} = \frac{1}{3+r_{12}+r_{22}-r_{12}-r_{21}}$ and $\mu_{23} = \frac{2+r_{12}-r_{13}}{3+r_{23}-r_{13}+r_{12}-r_{22}}$. Similar to Example 2 , the optimal information structure induces beliefs μ_{12} and 1, yielding a value

$$\frac{1}{2(1-\mu_{12})}\max\{\mu_{12}(3k+s_{21})+(1-\mu_{12})s_{11}, \mu_{12}(-k+s_{22})+(1-\mu_{12})(2k+s_{12})\}+\frac{1-2\mu_{12}}{2(1-\mu_{12})}(k+s_{23}).$$

Since $s_{ij} \in [0, \delta]$, taking $\delta \to 0$, we have $\mu_{12} \to 1/3$, and the above value approaches k. By continuity, there exists $\delta^1 > 0$ such that for any $\delta < \delta^1$, the persuasion value lies within the interval $[k - \frac{\varepsilon}{2}, k + \frac{\varepsilon}{2}]$.

Meanwhile, full revelation yields a payoff of $\frac{k+s_{11}+s_{23}}{2}$ and providing no information yields a payoff of $\frac{k+s_{12}+s_{22}}{2}$. Both values approach $k/2$ when $\delta \to 0$. By continuity, there exists $\delta^2 > 0$ such that for any $\delta < \delta^2$, the partitional persuasion value is less than $\frac{k+\varepsilon}{2}$.

Next, we will construct a cheap-talk equilibrium that yields a payoff close to k for small δ. Consider a strategy profile (σ, ρ) with two on-path messages m_1, m_2: $\sigma(m_1|\omega_2) = \frac{\mu_{12}}{1-\mu_{12}}$, $\sigma(m_2|\omega_2) = \frac{1-2\mu_{12}}{1-\mu_{12}}$, $\sigma(m_1|\omega_1) = 1$; $\rho(a_1|m_1) = p$, $\rho(a_2|m_1) = 1-p$, $\rho(a_3|m_2) = 1$, where $p = \frac{2k+s_{23}-s_{22}}{4k+s_{21}-s_{22}}$. Since $\delta < k$, $p \in (0, 1)$ is a well defined probability.

The strategy profile satisfies R-BR because the posterior upon observing m_1 is $1 - \mu_{12}$, and upon observing m_2 is 1. We now show that the profile also satisfies S-BR. For type ω_2 Sender, the expected payoff of sending message m_2 is $k + s_{23}$ and the expected payoff of sending message m_1 is $p(3k + s_{21}) + (1 - p)(-k + s_{22}) = k + s_{23}$, so type ω_2 Sender is indifferent and has no incentive to deviate. For type ω_1 Sender, the expected payoff of sending message m_2 is $-2k + s_{13}$ and the expected payoff of sending message m_1 is $p(s_{11}) + (1 - p)(2k + s_{12})$. As $\delta \to 0$, $p(s_{11}) + (1 - p)(2k + s_{12}) \to k$ and $-2k + s_{13} \to -2k$, so type ω_1 Sender strictly prefers to send message m_1. By continuity, there exists $\delta^3 > 0$ such that for any $\delta < \delta^3$, the strategy profile (σ, ρ) is a cheap-talk equilibrium, and the cheap-talk value is $\frac{1}{2}(k + s_{23}) + \frac{1}{2}[p(s_{11}) + (1-p)(2k + s_{12})] > k - \frac{\varepsilon}{2}$.

Therefore, for $k = 2B + 2\varepsilon$, $\delta < \delta^* \equiv \min\{\delta^1, \delta^2, \delta^3, \frac{1}{4}, k\}$, and for any $(u_S, u_R) \in E^k_\delta$, $\Delta_{\mathcal{C}}(u_S, u_R) < (k + \frac{\varepsilon}{2}) - (k - \frac{\varepsilon}{2}) = \varepsilon$, and $\Delta_{\mathcal{R}}(u_S, u_R) > (k - \frac{\varepsilon}{2}) - \frac{k+\varepsilon}{2} = \frac{k}{2} - \varepsilon = B$.