Role uncertainty and punishment severity in repeated cheap talk: theory and experiment*

By Jinyuqi Huang and Wooyoung Lim

We study how role uncertainty and the severity of grim-trigger punishment via Nash reversion affect strategic information transmission in repeated cheap talk games. Using a baseline with fixed roles and mild punishment, we show theoretically and experimentally that both role uncertainty and stronger punishment independently strengthen incentives for truth-telling. However, their interaction produces more nuanced effects. As predicted, strong punishment under asymmetric role uncertainty reduces communication efficiency. Unexpectedly, strong punishment under symmetric role uncertainty also reduces efficiency, contradicting the predicted null effect. We attribute these inefficiencies to strategic uncertainty and behavioral projection bias.
JEL: C92, D82, D83
Keywords: repeated cheap talk; role uncertainty; punishment level; projection bias; strategic uncertainty

Strategic information transmission is central in many economic interactions, yet it is often impeded by conflicting interests between informed and uninformed parties. In static cheap talk games, such conflicts typically preclude informative communication (Crawford and Sobel, 1982). Repetition offers a potential remedy: players can condition future behavior on past messages, allowing for fully revealing equilibria. However, sustaining such outcomes requires coordination on contingent strategies, which may be difficult to achieve in practice (Wilson and Vespa, 2020). This paper examines two realistic features of repeated cheap talk environments—role uncertainty and punishment severity—that may mitigate coordination challenges and foster truthful long-term communication.

The first feature, **role uncertainty**, captures the dynamic nature of long-term relationships. In many settings, individuals alternate between being the informed sender and the decision-making receiver. This fluidity implies that strategic incentives must account for the possibility of switching roles. The second feature, **punishment severity**, reflects the strength of deterrence provided by grim-trigger strategies, where deviations lead to permanent reversion to the stage-game Nash equilibrium (stage-Nash). The severity is determined by how unfavorable this equilibrium is for the deviating party. We parame-

* Huang: Department of Economics, The Hong Kong University of Science and Technology, Clearwater Bay, Kowloon, Hong Kong, E-mail: jhuangde@connect.ust.hk. Lim: Department of Economics, The Hong Kong University of Science and Technology, Clearwater Bay, Kowloon, Hong Kong, E-mail: wooyoung@ust.hk. We received helpful feedback from Syngjoo Choi, Joyee Deb, Jack Fanning, Nobuyuki Hanaki, Yohanes Eko Riyanto, Can Urgun, Emanuel Vespa, and seminar and conference participants at HKUST, the 2024 Asia-Pacific ESA meeting in Singapore, the 2024 7th World Congress of Game Theory Society in Beijing, China, and the 2025 Asia-Pacific ESA meeting in Osaka, Japan. This study is supported by an RPG student research grant from the Department of Economics at HKUST.

terize punishment intensity via the receiver's share in stage-Nash $b \in (0.5, 1]$: a higher b represents a harsher punishment for the sender.

To motivate our setting, consider the academic job market. A candidate (sender) has private information about her research potential—high or low—and communicates it to a recruitment committee chair (receiver). The receiver then chooses among offering a regular position, a temporary position, or no offer. The candidate values a regular position regardless of type, while the receiver wishes to match offers to true potential. The prior belief about the candidate's type is uniform—that is, both high and low types are equally likely. In a one-shot game, misaligned interests typically lead to pooling on the temporary position. This stage-Nash outcome yields payoffs of $(1-b, b)$ to the sender and receiver, respectively.

In a repeated setting, more informative outcomes become feasible, as players can condition future actions on past behavior. The strategic environment is also richer, reflecting institutional variations in both role uncertainty and punishment severity. Let $\gamma \in [0.5, 1]$ denote the probability that player A is the sender in a given period. In peer relationships, roles are typically symmetric ($\gamma = 0.5$), whereas hierarchical settings often exhibit asymmetry, with one party more frequently in the sender role ($\gamma > 0.5$). Punishment severity, captured by the stage-Nash reversion, also varies across contexts. Its harshness depends on the relative bargaining power of the receiver. For example, a temporary offer may be especially disappointing for a high-potential candidate, implying a higher value of b and thus a more severe punishment.

We analyze how these two factors—individually and jointly—affect the sustainability of informative communication. We formalize this environment as a repeated cheap talk game with two states, two messages and three actions. We focus on two equilibrium constructions. In the Simple Truth-Telling Equilibrium (STE), the sender always tells the truth, and the receiver acts optimally by following the sender's message and extracting the entire information surplus. Only the sender has an incentive to deviate. In the Information Rent Equilibrium (IE), the sender always tells the truth, and the receiver conditions her response on the revealed state. When the message indicates aligned interests, the receiver follows it fully. When it indicates misaligned interests, she plays the stage-game Nash action—partially sharing the surplus with the sender. This latter response, termed the "information rent" action by Wilson and Vespa (2020), gives the sender extra incentive to be truthful. However, it requires the receiver to forgo some surplus and coordinate on a contingent strategy. Both equilibria rely on grim-trigger strategies. Their sustainability depends on whether the discounted cost of future punishment outweighs the short-run gain from deviation.

In our baseline case of fixed roles ($\gamma = 1$) and mild punishment ($b = 0.5$), STE is only sustainable with extremely patient players. Increasing b improves sustainability by raising the cost of deviation. Introducing role uncertainty ($\gamma < 1$) also helps, since both players alternate roles and share the surplus over time, reducing the incentive to misreport.

However, when both role uncertainty and strong punishment are present, their interaction creates opposing forces on the current sender's incentives. Stronger punishment increases the cost of lying when the player expects to be a sender in the future, but simultaneously reduces the expected loss from lying when the player anticipates being a receiver. The net effect depends on the degree of role uncertainty, since the current sender is aware that both roles are possible in the future. With fixed roles, stronger punishment is unambiguously beneficial. With symmetric roles ($\gamma = 0.5$), the positive and negative forces offset each other. With asymmetric roles ($\gamma \in (0.5, 1)$), the effect depends on which player is the binding constraint.

To sustain STE, it is necessary to manage both players' incentives when acting as the sender. When punishment is moderate ($b \in (0.5, 0.75)$), the constraint for the player who is more likely the sender binds, and increasing b helps. When punishment is already strong ($b \in [0.75, 1)$), the constraint for the player who is more likely the receiver binds, and further increases in b may reduce sustainability.

Comparative statics for IE are simpler. Introducing role uncertainty or stronger punishment individually does not affect its sustainability. When $b = 0.5$, the sender and receiver face symmetric incentives, rendering role uncertainty irrelevant. Under fixed roles, the effects of changes in b on punishment severity and the size of the information rent offset each other. However, the interaction between role uncertainty and punishment becomes detrimental: when $b > 0.5$, the sender's incentive to deviate in misaligned states increases as their likelihood of being the sender decreases; when roles are random, the sender's incentive to deviate grows as b increases.

Moreover, implementing IE in practice may be challenging. As noted by Wilson and Vespa (2020), IE relies on coordination devices such as pre-play communication or transfers. In many environments, such devices are unavailable, costly, or legally constrained. By contrast, STE does not require such coordination and is therefore more attainable. This paper contributes by examining factors that promote truthful communication in the STE without relying on information rents.

We experimentally compare six treatments that vary along two dimensions: role uncertainty and punishment severity. Our **Baseline** treatment features fixed roles and mild punishment. Three treatments vary only one factor: **U**nequal-chance **R**andom roles with mild punishment (**UR**), **E**qual-chance **R**andom roles with mild punishment (**ER**), and **H**arsh punishment with fixed roles (**H**). Two treatments vary both factors: **H**arsh punishment with **U**nequal-chance **R**andom roles (**HUR**) and **H**arsh punishment with **E**qual-chance **R**andom roles (**HER**).

Based on our theoretical framework and implemented parameters, we expect information transmission to be more efficient in H, ER, and HER than in the baseline.[1] We expect no significant difference in UR and a decline in HUR. Moreover, we expect similar outcomes in ER and HER, since punishment has no additional effect under symmetric role uncertainty.

Our experimental results provide empirical support for several theoretical predictions. First, we find that introducing either stronger punishment or role uncertainty alone enhances communication efficiency. Specifically, efficiency in treatments H and ER is

[1]The information transmission efficiency is defined as the conditional probability of the receiver's action matching the correct state when the state is of common interest. Detailed explanation is deferred to Section II.D.

consistently higher than in the Baseline. Second, we observe the predicted negative interaction between strong punishment and asymmetric role uncertainty: efficiency in HUR is lower than in the Baseline.

However, two results deviate from theoretical expectations. First, we find that efficiency in UR is higher than in the Baseline, contrary to the predicted null effect. As we show, this seemingly excessive communication is transient. Second, and more notably, efficiency in HER is significantly lower than in both ER and the Baseline, suggesting an unanticipated negative effect of strong punishment under symmetric role uncertainty.

To identify the behavioral drivers behind these results, we decompose the data into two parts. The first focuses on first-round behavior across super-games, capturing how players' initial intentions evolve with experience. The second examines subsequent-round behavior, with particular attention to on-path cooperation conditional on previous coordination.

To reduce noise and clarify trends, we group the first-round data into three blocks: supergames 1–2, 3–5, and 6–7.[2] We compute averages within each block and examine trends across them. Three patterns emerge. First, the efficiency trends in H and ER support the theoretical prediction that each factor individually sustains the STE. Efficiency improves in H, and its decline becomes less pronounced in ER, driven by increased sender cooperation and a characteristic "U"-shaped pattern in receiver behavior. Second, in Baseline, UR, and HUR—where the STE is theoretically unsustainable—efficiency steadily declines, alongside reduced cooperation from both senders and receivers, indicating an erosion of trust and coordination. Third, the trend in HER contradicts theoretical expectations: efficiency deteriorates across blocks, primarily due to declining receiver cooperation, while sender behavior remains relatively stable. This asymmetry suggests that receivers may become increasingly reluctant to follow messages, even when incentives are aligned—a puzzling finding that warrants further investigation.

To better understand these deviations, we apply the concept of strategic uncertainty, using the basin of attraction framework proposed by Blonski and Spagnolo (2001, 2015). A player's basin of attraction for deviation represents the range of beliefs over which her optimal response is to defect. Larger basins indicate greater difficulty in sustaining cooperation due to uncertainty about others' behavior. This framework helps explain both expected and unexpected outcomes. First, H shows the smallest basin of attraction, corresponding to low strategic uncertainty and high efficiency. Second, HUR exhibits the largest basin, consistent with the observed breakdown in cooperation. Third, HER shows higher strategic uncertainty than either H or ER. This is due in part to the sender's relatively indifferent incentives under symmetric roles, which fail to strongly motivate cooperation, and in part to the receiver's large basin, which makes her more sensitive to doubts about the sender's reliability.

In the analysis of subsequent-round behavior, we focus on the receiver's responses following recent successful coordination—specifically, after a "double cooperation" event, where both players previously adhered to the equilibrium path. We examine two cases:

[2] In the experiment, each treatment had four sessions. Each session consisted of seven supergames with stochastic termination.

(1) the probability that the receiver follows the message (i.e., takes the IE or STE equilibrium action) when the message reveals no conflict of interest; and (2) the probability that the receiver pays the information rent (i.e., takes the IE action) when the message reveals a conflict of interest. Two patterns emerge. First, in HUR and HER, receivers exhibit the highest tendency to deviate from message-following even when interests are aligned, undermining the sustainability of informative equilibria and consistent with the low efficiency observed in these treatments. Second, when interests conflict, receivers in the Baseline are least likely to pay the information rent, while those in H, UR, and ER show a greater willingness to cooperate.

While strategic uncertainty provides a compelling explanation for coordination breakdowns in HUR and HER, it does not fully account for the asymmetry in behavior between senders and receivers or the lower-than-expected efficiency in HER. To address this, we turn to projection bias, as proposed by Loewenstein, O'Donoghue and Rabin (2003). Projection bias refers to an individual's tendency to overestimate the relevance of their current preferences or context in future decision-making, underweighting long-term consequences.

In our setting, projection bias manifests when receivers overvalue the immediate gain from defection and undervalue the long-run cost of triggering punishment. The extent of this bias depends on the salience of key game components. In treatments HUR and HER, defection yields a relatively high short-term payoff, amplifying projection bias and encouraging receivers to prioritize immediate gains over future cooperation. In the Baseline treatment, the perceived cost of cooperation is higher compared to H, UR and ER, which also increases projection bias, strengthening the incentive to defect when faced with a potential conflict of interest.

Together, these behavioral mechanisms—strategic uncertainty and projection bias—help explain both the predicted and unexpected patterns in our experimental data. They suggest that even when theoretical conditions support sustainable communication, behavioral frictions may prevent players from coordinating on efficient equilibria. More broadly, our findings highlight that the success of repeated communication hinges not only on institutional design, but also on how individuals perceive incentives and interpret strategic signals. Even subtle asymmetries in roles or payoff timing can distort cooperation and undermine equilibrium selection. This underscores the importance of incorporating behavioral considerations into models of dynamic strategic communication.

In the remainder of the paper, Section I reviews the related literature. Section II presents the theoretical model, comparative statics, and testable implications. Section III describes the experimental design. Section IV presents the empirical findings and behavioral analysis. Section V concludes.

I. Literature

This paper contributes to the literature on effective communication in cheap talk settings. Building on the foundational result of limited information transmission in static interactions between a privately informed sender and a decision-making receiver (Crawford and Sobel, 1982), researchers have explored various conditions that promote communication in static games (see Blume, Lai and Lim (2020) for a review) and examined the potential of repeated interactions to sustain informative outcomes (Golosov et al., 2014; Wilson and Vespa, 2020; Kolotilin and Li, 2021; Kuvalekar, Lipnowski and Ramos, 2022; Best and Quigley, 2024), among others.[3]

Our study contributes to this line of research by examining how realistic institutional features affect information transmission in repeated cheap talk games. Specifically, we focus on the roles of punishment severity and role uncertainty—two features that are both theoretically relevant and empirically observable.

To our knowledge, this is the first study to examine role uncertainty in the context of repeated cheap talk, though it has been explored in related experimental settings such as dictator and trust games. Existing experiments have produced mixed findings on the effect of potential role switching on altruism (Classen, 2005; Iriberri and Rey-Biel, 2011),[4] as well as on reciprocity (Ben-Ner et al., 2004; Herne, Lappalainen and Kestilä-Kekkonen, 2013; Burks, Carpenter and Verhoogen, 2003; Johnson and Mislin, 2011). These studies suggest that role uncertainty tends to reduce altruistic behavior: when the dictator role is not permanent, the current dictator may feel less responsible for the recipient's payoff, making self-interested behavior more normatively acceptable. At the same time, the possibility of switching roles can promote reciprocal thinking, which in turn encourages cooperation. Similarly, our theoretical analysis shows that role uncertainty influences the ex ante distribution of private information, punishment power, and information surplus. We further show that the effect of role uncertainty is context-dependent and interacts with other institutional features.

Our work builds on the framework developed by Wilson and Vespa (2020), extending it in empirically motivated directions. Specifically, we incorporate random role assignments and varying levels of punishment severity—two features that arise naturally in many real-world settings and are likely to affect long-run communication in the absence of explicit coordination mechanisms.

We contribute to the literature by providing both confirming and contrasting empirical evidence. On one hand, we show that each factor—role uncertainty and stronger punishment—can independently enhance information transmission. On the other hand, their interaction introduces additional strategic uncertainty, which we identify as a key source of coordination failure in our experimental setting. Moreover, we find that this interaction can also amplify behavioral biases, particularly projection bias, by increasing the salience of short-term payoffs relative to future consequences. This behavioral channel further undermines the sustainability of informative equilibria, even when theoretical conditions for successful communication are satisfied.

[3]Without observable history, pure relational incentives are generally insufficient to restore full commitment in repeated games (Kuvalekar, Lipnowski and Ramos, 2022). When history is available, fully revealing equilibria can be supported (Golosov et al., 2014). The quality and structure of historical information also influence truth-telling behavior (Best and Quigley, 2024). Additionally, transfers can facilitate communication by acting as signaling devices, as well as mechanisms to discipline the receiver (Kolotilin and Li, 2021) or to incentivize the sender towards full revelation (Wilson and Vespa, 2020).

[4]Mesa-Vázquez, Rodriguez-Lara and Urbano (2021) further show that the effect of role uncertainty on altruism is sensitive to the framing of the game.

II. Theory, Design and Hypotheses

The general repeated cheap talk game involves a fixed pair of players, A and B. The game unfolds over an indefinite number of rounds, with $\delta \in (0,1)$ denoting the probability of continuation. Equivalently, the game can be viewed as having an infinite horizon with a common discount factor δ. In each round, the players engage in the following sequence of actions. First, Nature assigns roles: with probability $\gamma \in [0.5,1]$, A is the sender and B the receiver; with probability $1-\gamma$, A is the receiver and B the sender. Next, Nature draws a state $\theta \in \Theta$, observed only by the sender. The sender sends a message $m \in \mathscr{M}$ to the receiver, who then chooses an action $a \in \mathscr{A}$. Payoffs depend on the state and the receiver's action. Full histories of roles, states, messages, and actions are observed by both players before the next round.

Let t index the rounds. Players are risk-neutral. Player A's objective is to maximize her discounted expected payoff:

$$U_A \equiv \sum_{t=1}^{\infty} \delta^{t-1}\big\{\gamma \mathbb{E}_\theta[u^S(a_t(m_t(\theta)),\theta)] + (1-\gamma)\mathbb{E}_\theta[u^R(a_t(m_t(\theta)),\theta)]\big\},$$

where $u^S(\cdot)$ and $u^R(\cdot)$ are the per-round payoffs when the player acts as sender or receiver, respectively.

Player B maximizes:

$$U_B \equiv \sum_{t=1}^{\infty} \delta^{t-1}\big\{(1-\gamma)\mathbb{E}_\theta[u^S(a_t(m_t(\theta)),\theta))] + \gamma \mathbb{E}_\theta[u^R(a_t(m_t(\theta)),\theta)]\big\}.$$

We simplify the environment to two states, two messages, and three actions.

A. The Stage Game

At the start of each round, roles are assigned and players play a static cheap talk game with the extensive form displayed in Figure 1. The state space is $\Theta = \{L, R\}$, message space $\mathscr{M} = \{\text{“}L\text{”}, \text{“}R\text{”}\}$, and action space $\mathscr{A} = \{L, M, R\}$. The prior is uniform: $\mathbb{P}(\theta = L) = \mathbb{P}(\theta = R) = 0.5$. The sender's payoff is:

$$u^S(a) = \begin{cases} 1 & \text{if } a = L, \\ 1-b & \text{if } a = M, \\ 0 & \text{if } a = R, \end{cases}$$

and the receiver's payoff is:

$$u^R(\theta, a) = \begin{cases} 1 & \text{if } (\theta, a) \in \{(L,L),(R,R)\}, \\ b & \text{if } a = M, \\ 0 & \text{otherwise.} \end{cases}$$

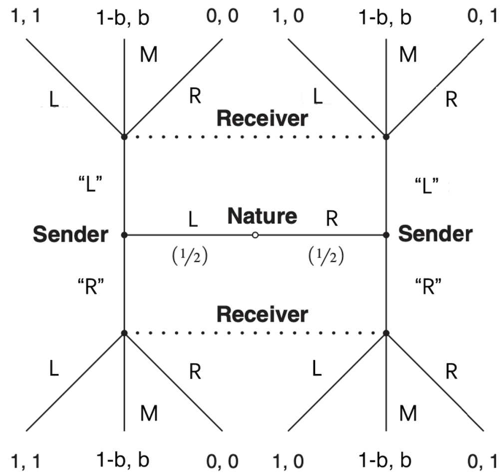

FIGURE 1. THE STAGE GAME.

The sender always prefers action L, while the receiver prefers matching the state ex post and M ex ante (given $b \geq 0.5$). In the static game, this misalignment leads to babbling as the unique equilibrium: the receiver ignores the message and plays M. When $b = 0.5$, she may randomize, a small positive increment ε to b can break the tie; any $b > 0.5$ guarantees M is optimal. This Babbling Equilibrium (BE) yields payoffs of 1 in both states—less than the efficient truth-telling outcome (2 in state L, 1 in state R).

B. *The Simple Truth-Telling Equilibrium*

In repeated play, more informative equilibria become feasible. We focus on the Simple Truth-Telling Equilibrium (STE).[5] In STE, the sender tells the truth, and the receiver follows the message. Any deviation triggers a permanent reversion to babbling. The equilibrium is "simple" in that the receiver always chooses her myopically optimal action given the message, extracting the full surplus without needing explicit reward-sharing. We assume symmetric strategies across players.[6]

On-path, neither player has an incentive to deviate when acting as receiver. The sender, however, may benefit by misreporting in state R, gaining 1 if the receiver follows. Deviation triggers punishment (action M), yielding off-path payoffs of $(1-b, b)$. Thus, larger b implies a harsher punishment for the sender.

STE-BASED COMPARATIVE STATICS. — We define $\hat{\delta}$ as the minimum discount factor required to support STE for a given (γ, b). Smaller $\hat{\delta}$ indicates easier sustainment.

[5] A full analysis is provided in Appendix A.A1.
[6] That is, strategies are independent of identity and the realized role.

To derive $\hat{\delta}$, we examine profitable deviations by senders in state R. The gain from deviation is always 1; the future cost C in each round depends on both b and γ. Figure 2 plots the smaller of the two possible values of C—corresponding to the binding incentive constraint. The deviation-proof condition is $(\delta C)/(1-\delta) \geq 1$, so the binding constraint determines $\hat{\delta}$.

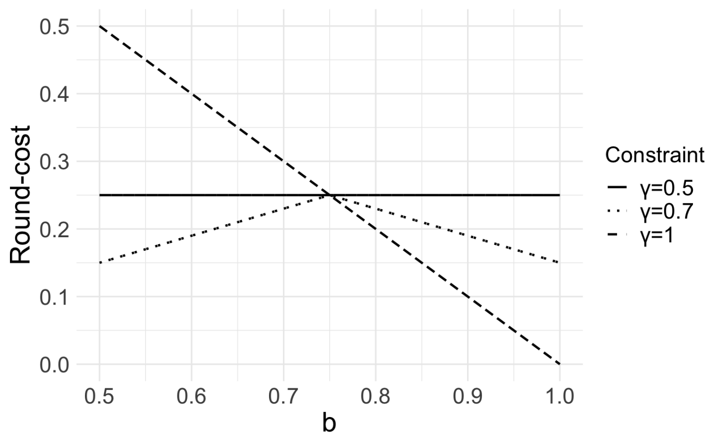

FIGURE 2. ROUND-COST OF DEVIATION UNDER DIFFERENT (γ, b).

Figure 3 shows the resulting values of $\hat{\delta}$—the minimal discount factor required to sustain the STE—across different combinations of punishment severity (b) and role uncertainty (γ). The horizontal axis represents b, and the vertical axis represents $\hat{\delta}$. The point at the top-left corner of the figure—fixed roles ($\gamma = 1$) and mild punishment ($b = 0.5$)—serves as our baseline for comparison.

We highlight four comparative statics:

1. Fixed roles ($\gamma = 1$): Increasing punishment severity (higher b) lowers $\hat{\delta}$, making it easier to sustain the STE.

2. Mild punishment ($b = 0.5$): Increasing role uncertainty (lower γ) also lowers $\hat{\delta}$, facilitating the STE.

3. Symmetric roles ($\gamma = 0.5$): Variation in b has no effect on $\hat{\delta}$—the cost of deviation is constant across punishment levels.

4. Asymmetric roles ($\gamma = 0.7$): The effect of b is non-monotonic. When $b < 0.75$, increasing b lowers $\hat{\delta}$. But when $b \geq 0.75$, increasing b raises $\hat{\delta}$.

Overall, the comparative statics reveal that introducing either stronger punishment or role uncertainty individually facilitates strategic communication relative to the baseline. However, when the two factors interact, their joint effect depends on the degree of role uncertainty. Under symmetric roles, stronger punishment has no additional effect on sustainability. In contrast, under asymmetric roles, the effect of punishment becomes non-monotonic—initially promoting cooperation, but ultimately hindering it when the punishment becomes too severe.

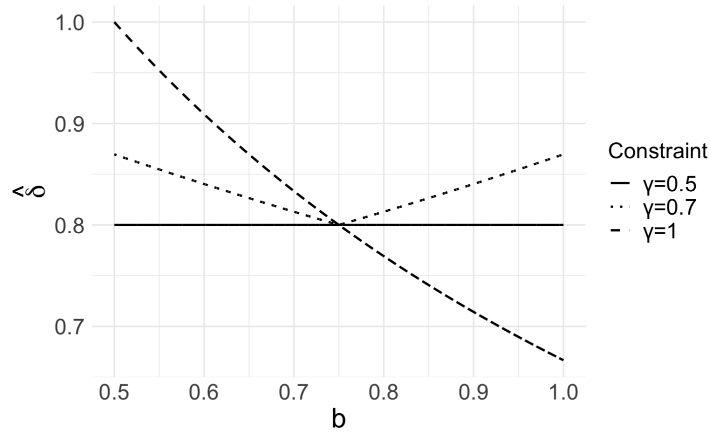

FIGURE 3. COMPARATIVE STATICS FOR STE.

This non-monotonicity reflects a shift in the binding incentive constraint across different values of b. When $b < 0.75$, the constraint binds for player A, who is more likely to act as sender. In this case, harsher punishment increases the expected cost of deviation, thereby supporting the STE. However, when $b \geq 0.75$, the constraint shifts to player B, who is more likely to act as receiver. In that case, further increases in b lower the expected cost of deviation, making cooperation harder to sustain.

C. The Information Rent Equilibrium

In addition to the STE, the Information Rent Equilibrium (IE) can also be sustained in repeated play.[7] In IE, the sender always tells the truth, and the receiver conditions her response on the state revealed by the message. When the message indicates aligned interests, the receiver fully follows the message. When it indicates misaligned interests, she responds with the stage-game Nash action—effectively sharing part of the surplus with the sender. Any deviation triggers a permanent reversion to the babbling equilibrium.

The key feature of IE is the use of information rent: when the state is R, the receiver plays action M in response to message "R." This action, as termed by Wilson and Vespa (2020), gives the sender a positive payoff even in the misaligned state, thereby increasing the sender's incentive to report truthfully. However, this requires the receiver to forgo some surplus, making IE more demanding in terms of coordination. We assume symmetric strategies across players.

On the equilibrium path, the sender may be tempted to lie in state R to secure a higher short-run payoff if the receiver were to blindly follow. Similarly, the receiver may deviate by refusing to pay the information rent and instead choosing her myopically optimal action. As in STE, deviations trigger reversion to punishment (action M), yielding off-

[7] A full analysis is provided in Appendix A.A1.

path payoffs of $(1-b,b)$. Importantly, in IE, the parameter b not only determines the severity of punishment but also affects the size of the information rent, complicating its overall impact.

IE-BASED COMPARATIVE STATICS. — We conduct a parallel analysis for IE and compute the minimum discount factor $\hat{\delta}$ required to support it under different combinations of b and γ. Figure 4 displays the results. As before, the horizontal axis represents the punishment level b, and the vertical axis represents $\hat{\delta}$. The point in the top-left corner—fixed roles ($\gamma = 1$) and mild punishment ($b = 0.5$)—serves as our baseline.

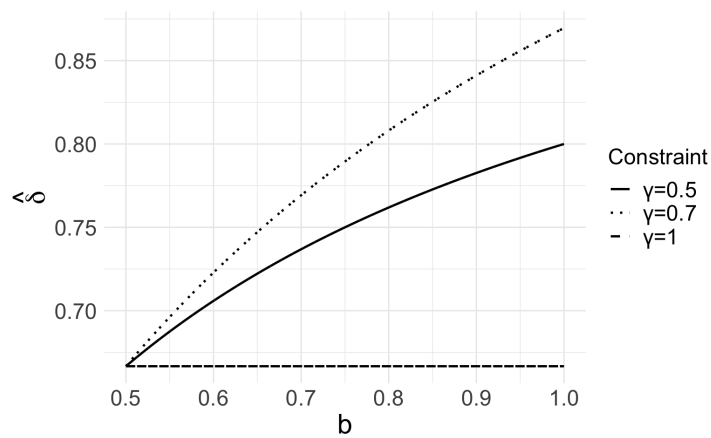

FIGURE 4. COMPARATIVE STATICS FOR IE.

We highlight four comparative statics:

1. Fixed roles ($\gamma = 1$): Increasing b has no effect on $\hat{\delta}$. Effects on punishment severity and the size of information rent are canceled out.

2. Mild punishment ($b = 0.5$): Increasing role uncertainty (lower γ) has no effect on $\hat{\delta}$. Since the sender and receiver face symmetric incentives, role uncertainty does not alter the equilibrium condition.

3. Random roles ($\gamma < 1$): Increasing b raises $\hat{\delta}$. When roles are not fixed, the binding constraint shifts to the sender who is more likely to be in the receiver role in the future. In this case, a larger b increases the temptation to deviate when acting as sender in the misaligned state.

4. Non-mild punishment ($b > 0.5$): Increasing role uncertainty (lower γ) also raises $\hat{\delta}$. As the chance of being the sender declines, the expected future benefit from truth-telling decreases, while the short-run gain from misreporting remains, making deviation more attractive.

Overall, the comparative statics for IE contrast sharply with those for STE. Introducing either stronger punishment or greater role uncertainty individually has no effect relative to the baseline. However, when the two factors interact—i.e., when punishment is harsh and roles are randomized—their joint effect is detrimental. The interaction increases the incentive to deviate in misaligned states, making IE harder to sustain.

D. Design and Predictions

We test the theoretical predictions using six experimental treatments (Table 1), varying along two dimensions: punishment severity (b) and role uncertainty (γ). The **Baseline** features fixed roles and mild punishment. Three treatments vary only one factor: **H** (harsh punishment with fixed roles), **UR** (unequal random roles with mild punishment), and **ER** (equal random roles with mild punishment). The remaining two treatments vary both: **HUR** (harsh punishment and unequal random roles) and **HER** (harsh punishment and equal random roles).

TABLE 1—SIX TREATMENTS.

		Role Uncertainty (γ)		
		Fixed (1)	Unequal Rand. (0.7)	Equal Rand. (0.5)
Punishment (b)	Mild (0.5)	*Baseline*	UR	ER
	Harsh (0.9)	H	HUR	HER

We measure communication efficiency as $\mathbb{P}(a = L \mid \theta = L)$—the probability the receiver takes the correct action when interests are aligned. Efficiency is 1 in any truth-telling equilibrium. The discount factor was fixed at $\delta = 0.8$ across all treatments. If a treatment's $\hat{\hat{\delta}} \leq 0.8$, then STE is theoretically sustainable.

Table 2 summarizes the theoretical predictions based on the minimum required discount factors derived in the comparative statics. According to the STE-based analysis, treatments H, ER, and HER can support the STE. In contrast, the Baseline, UR, and HUR cannot, as their required $\hat{\hat{\delta}}$ exceeds 0.8. The IE-based analysis suggests that all treatments except HUR can support the IE.

TABLE 2—PREDICTIONS UNDER THE IMPLEMENTED TIME DISCOUNT FACTOR OF 0.8.

	Baseline	UR	ER	H	HUR	HER
STE	$\times$	$\times$	✓	✓	$\times$	✓
IE	✓	✓	✓	✓	$\times$	✓

Note: A check mark (✓) indicates that the corresponding equilibrium can be supported. A cross ($\times$) indicates that the equilibrium cannot be supported.

Because the model admits multiple equilibria, we do not aim to test sharp comparative statics. Instead, our goal is to document how subjects behave across institutional settings and assess which environments support more effective communication. The following hypotheses guide our empirical analysis:

Hypothesis 1 (H). *There is no difference in efficiency between the Baseline and H.*

This tests whether harsher punishment improves communication. We expect to reject this hypothesis, as STE is sustainable in H but not in the Baseline.

Hypothesis 2 (UR). *There is no difference in efficiency between the Baseline and UR.*

This tests whether asymmetric role uncertainty alone improves communication. We expect not to reject this hypothesis, as both treatments support IE but not STE.

Hypothesis 3 (ER). *There is no difference in efficiency between the Baseline and ER.*

This tests whether symmetric role uncertainty is sufficient to sustain cooperation. We expect to reject this hypothesis, since ER supports STE, while the Baseline does not.

Hypothesis 4 (HUR). *There is no difference in efficiency between the Baseline and HUR.*

This tests whether harsher punishment under asymmetric roles improves communication. We expect to reject this hypothesis, as HUR supports neither STE nor IE, and is likely to yield lower efficiency than the Baseline.

Hypothesis 5 (HER). *There is no difference in efficiency between ER and HER.*

This tests whether harsher punishment under symmetric roles affects communication. We expect not to reject this hypothesis, as both treatments support STE and IE, implying similar efficiency levels.

III. Experimental Implementation

We conducted a real-time online experiment using oTree (Chen, Schonger and Wickens, 2016) at the Hong Kong University of Science and Technology in November 2023, March 2025, and April 2025.[8] A total of 390 subjects participated, primarily undergraduate students with no prior experience in this experiment. The design included six treatments, each implemented in four sessions. Each session involved 12–20 participants and consisted of seven official super-games. Sessions lasted approximately one hour.

Each unit in Figure 1 corresponded to 40 HKD in the experimental interface. The subject earned on average around 160 HKD (approximately 20 USD), including a 40 HKD show-up fee.

The experiment was conducted in English. Sample instructions are provided in Appendix A.A2. To ensure common knowledge, the experimenter read the instructions aloud. Subjects then completed a quiz to confirm their understanding of the environment. Upon correctly answering all quiz questions, they participated in a non-incentivized practice game to familiarize themselves with the interface before proceeding to the official games.

In treatments with fixed roles, subjects were randomly assigned at the start of each session to be either senders or receivers (half of each). Each sender was randomly matched

[8] IRB protocol no.: HERP-2023-0307. Project title: Random role assignment in repeated cheap talk. Approval period: 22-Sep-2023 to 22-Sep-2026. Institution: Hong Kong University of Science and Technology.

with a receiver at the start of every super-game. Within a super-game, both the pairing and roles remained fixed.

In treatments with random roles, subjects were randomly assigned player IDs A or B at the beginning of the session (again, half of each). Each player A was randomly matched with a player B at the start of each super-game. Within each super-game, the pairing remained fixed, but roles were randomly assigned at the beginning of each round: either A is the sender and B the receiver with probability γ, or the opposite with probability $1-\gamma$.

We visualized the 80% continuation probability to subjects using a spinning wheel interface. At the end of each round, subjects saw a wheel with 80% of its area colored green and 20% red. Pressing a button in the center initiated the spin. If the wheel stopped on green, the game continued; otherwise, it ended. Both players in the same super-game observed the same outcome. Although the number of rounds was pre-determined based on the 80% continuation rule (to standardize session duration), the visual display was designed to provide a more intuitive and engaging representation of stochastic termination.

IV. Empirical Findings

Table 3 reports average communication efficiency across four measures, by treatment. The top two panels show efficiency in the first super-game: the left panel focuses on the first round; the right pools all rounds. The bottom two panels report efficiency across all seven super-games, again separating first-round behavior (left) from all rounds (right). Treatment order follows Table 1 to facilitate comparison.

TABLE 3—EFFICIENCY BY TREATMENT.

First super-game							
First round				All rounds			
$b\backslash\gamma$	1	0.7	0.5	$b\backslash\gamma$	1	0.7	0.5
0.5	0.72	0.71	0.74	0.5	0.48	0.68	0.64
0.9	0.67	0.47	0.54	0.9	0.58	0.35	0.47
All super-games							
First round				All rounds			
$b\backslash\gamma$	1	0.7	0.5	$b\backslash\gamma$	1	0.7	0.5
0.5	0.55	0.59	0.56	0.5	0.47	0.55	0.52
0.9	0.61	0.41	0.46	0.9	0.53	0.27	0.33

Note: In each panel, the corresponding treatments in the first row are (from left to right) Baseline, UR, ER; the corresponding treatments in the second row are (from left to right) H, HUR, HER.

We present five empirical results, each corresponding to one of the five hypotheses.

First, we consider the comparison between the Baseline ($b=0.5$, $\gamma=1$) and treatment H ($b=0.9$, $\gamma=1$), which relates to Hypothesis 1. Table 3 shows that H achieves higher average efficiency than the Baseline in three out of four measures. The only exception is the first round of the first super-game. While the direction of the effect is consistent with the theoretical prediction, the difference is not statistically significant ($p = 0.34$).[9] Still, behavioral patterns in Sections IV.A and IV.B will suggest a modest and growing positive effect of stronger punishment.

Result 1 (H). *Introducing stronger punishment marginally promotes communication efficiency compared to the Baseline.*

Next, we examine the effect of asymmetric role uncertainty by comparing the Baseline with treatment UR ($b = 0.5$, $\gamma = 0.7$), as stated in Hypothesis 2. Although theory suggests no difference in equilibrium support, UR shows higher communication efficiency across all four measures. The improvement is marginally significant ($p = 0.06$). As will be shown in later sections, the effect diminishes over time.

Result 2 (UR). *Introducing asymmetric role uncertainty marginally promotes communication efficiency compared to the Baseline.*

Turning to Hypothesis 3, we compare the Baseline with treatment ER ($b = 0.5$, $\gamma = 0.5$), which introduces symmetric role uncertainty. ER exhibits higher efficiency than the Baseline across all measures. Although the difference is not statistically significant ($p = 0.23$), the behavioral pattern will suggest a mild but emerging positive effect, especially as subjects gain experience.

Result 3 (ER). *Introducing symmetric role uncertainty marginally promotes communication efficiency compared to the Baseline.*

We now evaluate Hypothesis 4 by comparing the Baseline with treatment HUR ($b = 0.9$, $\gamma = 0.7$), which combines harsh punishment with asymmetric roles. Communication efficiency in HUR is consistently lower than in the Baseline across all four measures. The difference is large and statistically significant ($p = 0.00$). The evidence in Sections IV.A and IV.B will indicate that the negative effect is robust and persistent.

Result 4 (HUR). *The interaction of strong punishment and asymmetric role uncertainty hinders communication efficiency compared to the Baseline.*

Finally, we assess Hypothesis 5 by comparing ER and HER ($b = 0.9$, $\gamma = 0.5$), both of which involve symmetric roles, but differ in punishment severity. HER shows significantly lower efficiency than ER and the Baseline ($p = 0.00$ in both comparisons). This outcome is not predicted by theory, which suggests HER should support both STE and IE. behavioral data, especially from receivers, will point to a low tendency to cooperate in HER.

Result 5 (HER). *The interaction of strong punishment and symmetric role uncertainty yields unexpectedly inefficient communication compared to ER and the Baseline.*

To identify the behavioral drivers behind these results, we decompose the data into two components. Section IV.A will focus on first-round behavior across super-games, highlighting how initial strategies evolve. Section IV.B will examine subsequent-round behavior, with particular attention to on-path cooperation after prior coordination.

[9]Mann-Whitney U tests take supergame-level data as independent observations and use data including all supergames with all rounds. Full results are reported in Appendix A.A3.

A. Initial behavior

In this section, we focus on first-round data, which captures subjects' initial behaviors that can be viewed as largely independent of history. These initial choices set the tone for subsequent interaction and play a critical role in shaping overall efficiency. To reduce noise and better identify patterns, we group the data into three blocks: supergames 1–2, 3–5, and 6–7. We compute averages within each block and examine trends across them.

Figures 5 through 10 plot the sender's message frequencies in the misaligned state ($\theta = R$). The truth-telling frequency is denoted by $\mu(\text{“}R\text{”}|R)$, and the lying frequency by $\mu(\text{“}L\text{”}|R)$. For senders, truth-telling in the misaligned state is crucial for sustaining informative communication and thus a key indicator of equilibrium support.

Figures 11 through 22 report the receiver's action frequencies conditional on receiving message “L” or “R”. The message-following frequencies are $\sigma(L|\text{“}L\text{”})$ and $\sigma(R|\text{“}R\text{”})$. The frequency of choosing the middle action is given by $\sigma(M|\text{“}L\text{”})$ and $\sigma(M|\text{“}R\text{”})$, while the frequency of choosing the opposite action is $\sigma(R|\text{“}L\text{”})$ and $\sigma(L|\text{“}R\text{”})$. For the receiver, we focus on two key metrics: the message-following tendency $\sigma(L|\text{“}L\text{”})$, which is essential for information transmission when the sender is truthful; and the difference $\sigma(M|\text{“}R\text{”}) - \sigma(M|\text{“}L\text{”})$, which partially captures the receiver's willingness to pay information rents in line with the IE. Figure 23 presents first-round efficiencies across supergames by treatment.

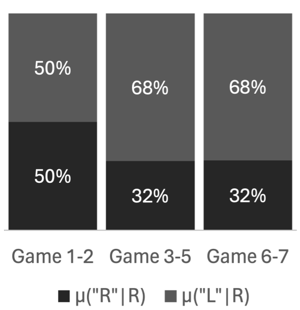

FIGURE 5. ROUND-1 MESSAGE IN STATE R: BASELINE.

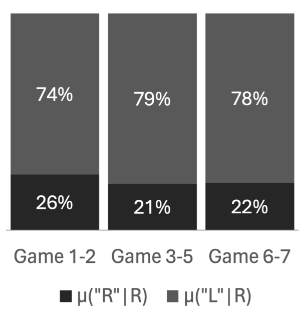

FIGURE 6. ROUND-1 MESSAGE IN STATE R: UR.

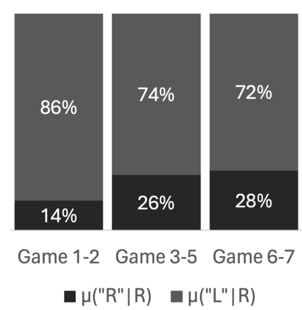

FIGURE 7. ROUND-1 MESSAGE IN STATE R: ER.

We now describe the patterns in initial behavior related to the five results discussed above. While absolute levels are sometimes relevant, the evolving trends are often more informative, as they signal how behavior adapts with experience. In the Baseline treatment, both the sender's truth-telling frequency and the receiver's frequency of following message “L” with action L decline over time, even as the receiver's approximate information rent behavior increases. As a result, Figure 23 shows a downward trend in first-round efficiency in supergames 6–7.

Corresponding to Result 1, on the sender side, with fixed roles, the truth-telling fre-

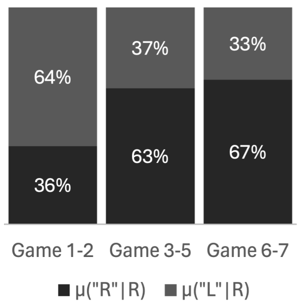

FIGURE 8. ROUND-1 MESSAGE IN STATE R: H.

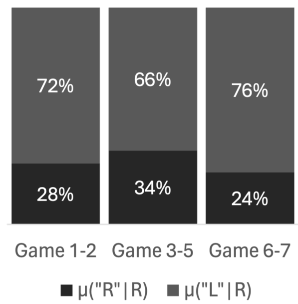

FIGURE 9. ROUND-1 MESSAGE IN STATE R: HUR.

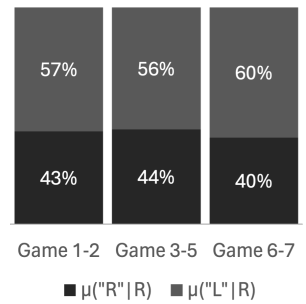

FIGURE 10. ROUND-1 MESSAGE IN STATE R: HER.

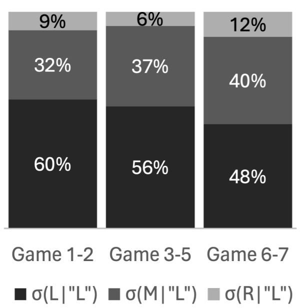

FIGURE 11. ROUND-1 ACTION TO MESSAGE "L": BASELINE.

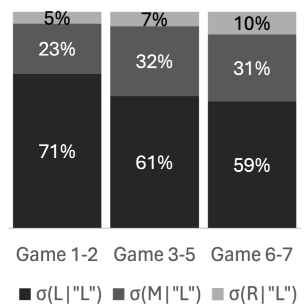

FIGURE 12. ROUND-1 ACTION TO MESSAGE "L": UR.

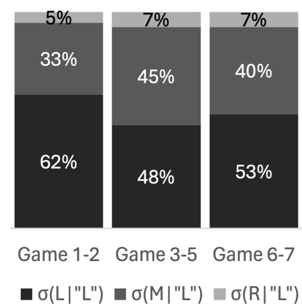

FIGURE 13. ROUND-1 ACTION TO MESSAGE "L": ER.

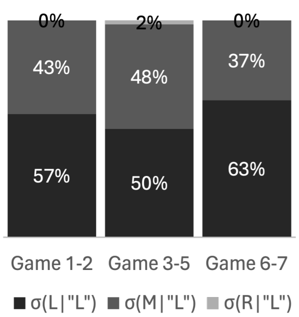

FIGURE 14. ROUND-1 ACTION TO MESSAGE "L": H.

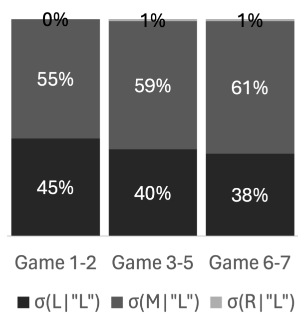

FIGURE 15. ROUND-1 ACTION TO MESSAGE "L": HUR.

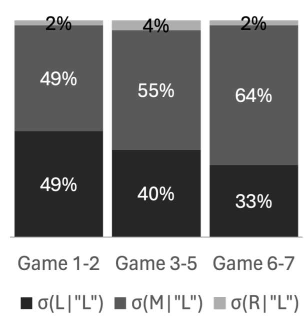

FIGURE 16. ROUND-1 ACTION TO MESSAGE "L": HER.

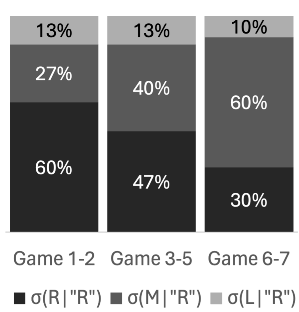

FIGURE 17. ROUND-1 ACTION TO MESSAGE “*R*”: BASELINE.

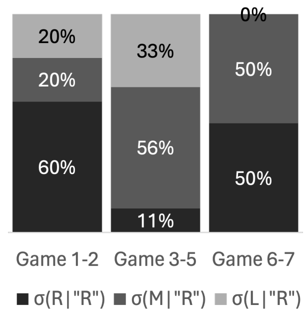

FIGURE 18. ROUND-1 ACTION TO MESSAGE “*R*”: UR.

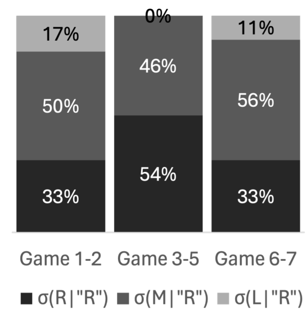

FIGURE 19. ROUND-1 ACTION TO MESSAGE “*R*”: ER.

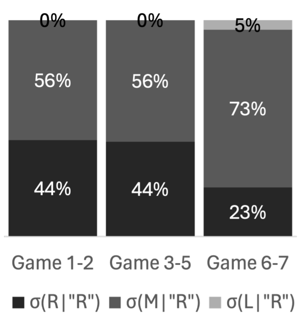

FIGURE 20. ROUND-1 ACTION TO MESSAGE “*R*”: H.

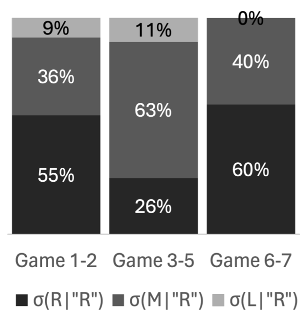

FIGURE 21. ROUND-1 ACTION TO MESSAGE “*R*”: HUR.

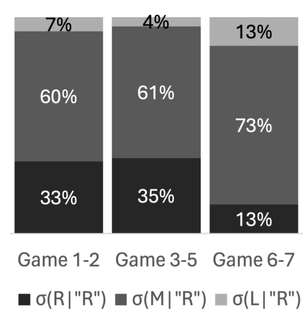

FIGURE 22. ROUND-1 ACTION TO MESSAGE “*R*”: HER.

quency in the misaligned state is significantly higher in H compared to the Baseline. On the receiver side, in supergames 6–7, the frequency of following message “*L*” with action *L* is also higher in H. In addition, the approximate information rent frequency is greater in H than in the Baseline across all supergames. Consequently, Figure 23 shows that first-round efficiency in H increases substantially in the later supergames, following a modest decline in the early ones.

Initial behavior Related to Result 1 (H). *Both the sender's and receiver's initial behaviors become more cooperative in H, resulting in higher initial efficiency in H than the Baseline in supergames 6–7.*

For Result 2, on the sender side, the truth-telling frequency in the misaligned state is consistently lower in UR relative to the Baseline. On the receiver side, the frequency of following message “*L*” with action *L* is initially higher but declines over time. Additionally, the approximate information rent frequency is greater in UR than in the Baseline during supergames 3–5. As shown in Figure 23, first-round efficiency in UR is

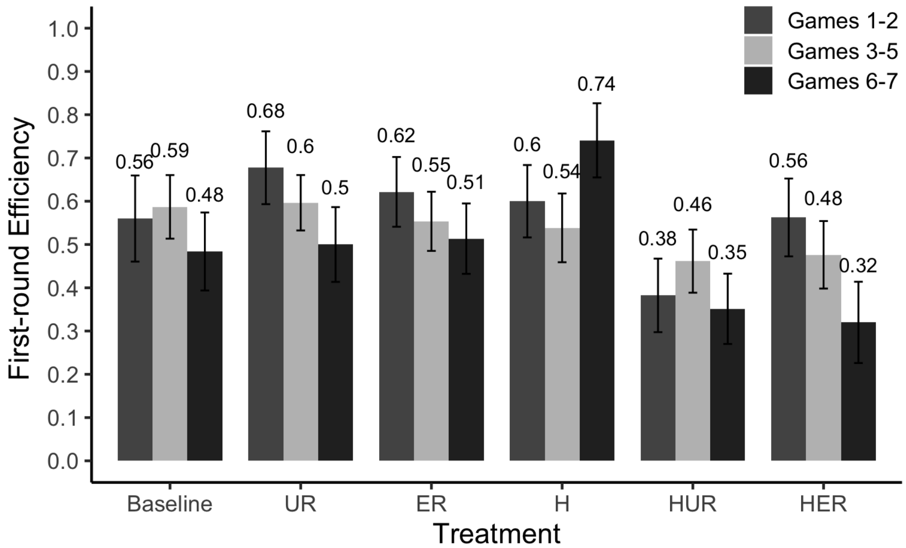

FIGURE 23. ROUND-1 EFFICIENCY BY TREATMENT.

marginally higher than in the Baseline, but the downward trend appears to persist. These patterns suggest that the initially positive effect of asymmetric role uncertainty may be short-lived.

Initial behavior Related to Result 2 (UR). *While the sender's truth-telling frequency is lower in UR than in the Baseline, the receiver's initial behavior is more cooperative, resulting in higher initial efficiency in UR but with a narrowing advantage.*

Corresponding to Result 3, on the sender side, the truth-telling frequency in the misaligned state is lower in ER than in the Baseline initially, but displays a clear upward trend. On the receiver side, the frequency of following message "L" with action L is higher in ER than in the Baseline in supergames 6–7. Additionally, the approximate information rent frequency is higher in ER during supergames 1–2, though slightly lower than in the Baseline in supergames 3–7. As a result, Figure 23 shows that first-round efficiency in ER is higher than in the Baseline during supergames 1–2 and 6–7, and the overall decline in efficiency is less pronounced.

Initial behavior Related to Result 3 (ER). *Both the sender's and receiver's behaviors become more cooperative in ER, resulting in higher initial efficiency in ER by supergames 6–7.*

As for Result 4, on the sender side, the truth-telling frequency in the misaligned state is lower in HUR than in the Baseline. On the receiver side, the frequency of following message "L" with action L is consistently lower and declining in HUR. Furthermore, there is no evidence of information rent behavior in HUR, in line with theoretical predictions that IE cannot be sustained in this environment. Consequently, Figure 23 shows that first-round efficiency in HUR is consistently lower than in the Baseline.

Initial behavior Related to Result 4 (HUR). *Both the sender's and receiver's initial behaviors become less cooperative in HUR, resulting in lower first-round efficiency compared to the Baseline.*

As for Result 5, on the sender side, the truth-telling frequency in the misaligned state is stable and higher in HER than in the Baseline during supergames 3–7. On the receiver side, the frequency of following message "L" with action L declines and remains consistently lower in HER than in the Baseline. Moreover, the approximate information rent frequency in HER is lower than in the Baseline during supergames 6–7. As a result, Figure 23 shows that first-round efficiency in HER declines over time and falls below the Baseline level.

Initial behavior Related to Result 5 (HER). *While sender behavior remains stable in HER, declining receiver cooperation leads to lower first-round efficiency compared to the Baseline.*

To summarize, first, the efficiency trends in treatments H and ER partially support the theoretical prediction that either stronger punishment or symmetric role uncertainty can promote communication. Efficiency improves in H, and its decline becomes less pronounced in ER, driven by increased sender truth-telling and a characteristic "U"-shaped pattern in receiver behavior. Second, in the Baseline, UR, and HUR—where the STE is theoretically unsustainable—efficiency steadily declines, alongside reduced cooperation from both senders and receivers, indicating an erosion of trust and coordination. Third, the trend in HER contradicts theoretical expectations: efficiency deteriorates across blocks, primarily due to declining receiver cooperation, while sender behavior remains relatively stable. This asymmetry suggests that receivers may become increasingly reluctant to follow messages, even when incentives are aligned—a puzzling pattern that warrants further investigation.

Basin-of-attraction Analysis. — To further interpret the variation in first-round behavior across treatments, we examine the role of strategic uncertainty using the concept of the basin of attraction for deviation. This concept, which extends risk dominance (Harsanyi and Selten, 1988), was formalized by Blonski and Spagnolo (2001, 2015) and has been used effectively to explain behavior in repeated games such as the prisoner's dilemma (Dal-Bo and Frechette, 2011, 2018). In our setting, it helps predict whether players are likely to coordinate on cooperative outcomes or fall back to non-cooperative ones.

The basin of attraction reflects the range of beliefs under which a player finds it better to deviate from a cooperative strategy. A larger basin suggests that even small doubts about the other player's behavior can lead to defection, implying greater strategic uncertainty. A smaller basin means that cooperation is more stable and less sensitive to belief variation.

We apply this framework to two equilibrium comparisons in our repeated cheap talk environment: STE vs. BE and IE vs. BE. In both cases, players follow grim-trigger strategies, with deviations punished by permanent reversion to BE. In the STE, the sender is expected to report truthfully and the receiver to follow the message. Deviations include the sender lying when the state is R, or the receiver ignoring the message and choosing M regardless of its content.[10] In the IE, deviations similarly involve sender misreporting or the receiver refusing to pay the information rent—such as choosing R following message "R".

Several features of our experimental design require us to adapt the standard basin analysis. First, since both STE and IE may be sustainable in some treatments, we calculate basins for both comparisons. Second, because sender and receiver face different incentives, we compute basins separately for each role. Third, in some treatments, such as HER, there are multiple contingencies due to random roles and states. We compute basins under each of these and then aggregate.

The aggregation takes two steps. First, for each state, we combine the sender and receiver basins using a multiplicative rule that reflects the joint belief structure necessary for both players to prefer cooperation. Second, we take the union of the resulting basins across states to obtain an integrated measure of strategic uncertainty for each treatment-equilibrium pair. All derivations and multiplication rules are documented in Appendix A.A4.

To illustrate, consider treatment H, which has fixed roles and harsh punishment. In both states, the sender's basin is a null set——she always prefers to cooperate—while the receiver's basins are moderate. Since the sender is known to always cooperate, the receiver's belief requirement becomes irrelevant, and the aggregate basin reduces to the empty set. This implies minimal strategic uncertainty and helps explain the high levels of cooperation observed in H.

In contrast, in treatment HER, where roles are randomly assigned and punishment is harsh, the picture is more complex. When player A is the sender and the state is R, the sender is indifferent between cooperating and defecting, while the receiver's basin stretches to $[0, .45]$. Since the sender may mix between strategies, the receiver's belief becomes pivotal. If the sender's actual truth-telling frequency falls within the receiver's basin—as it does in HER—then the receiver's best response is to defect. This scenario explains the persistent decline in message-following by receivers in HER, despite relatively high sender truth-telling.

The full set of basin results across treatments in Table 4 aligns closely with the behavioral patterns observed in the experiment and offers a coherent explanation for when and why cooperation succeeds or fails. In the Baseline and UR, where the STE is not sustainable, the moderately sized IE-BE basin ($[0, .25]$) suggests intermediate levels of strategic uncertainty. These settings exhibit mixed behavior: some players experiment with harsher punishments, while others shift towards the IE, resulting in modest and declining efficiency. In ER, both STE and IE are sustainable, and the balanced, moderate basin sizes for both comparisons support a smooth transition between equilibria and relatively stable cooperation. In H, the smallest basins—$\emptyset$ for STE-BE and $[0, .09]$ for IE-BE—reflect minimal strategic uncertainty and correspond to the highest rates of sender truth-telling, receiver message-following, and overall efficiency. In

[10]Although the receiver has no incentive to deviate when the message is truthful, we include such a deviation as it captures a generalized distrust of messages.

contrast, HUR shows maximal uncertainty, with both basins equal to $[0,1]$, and a corresponding breakdown in coordination. Finally, in HER, although both equilibria are theoretically supportable, the large STE-BE basin ($[0,.45]$) renders cooperation highly fragile: even moderate doubts about the sender's intentions can push the receiver to defect. This asymmetry—driven not by payoffs but by differential sensitivity to strategic uncertainty—explains the divergence in behavior between roles. The sender's relatively high truth-telling rate (around 40%) still falls within the receiver's basin, making defection the receiver's best response and ultimately undermining communication. Taken together, the basin-of-attraction framework highlights how equilibrium multiplicity, belief heterogeneity, and role asymmetry jointly shape the sustainability of informative communication in repeated interactions.

TABLE 4—AGGREGATED BASINS FOR STE-BE AND IE-BE COMPARISONS.

Treatment	STE-BE Basin	IE-BE Basin
Baseline	$[0,1]$	$[0,0.25]$
UR	$[0,1]$	$[0,0.25]$
ER	$[0,0.25]$	$[0,0.25]$
H	$\emptyset$	$[0,0.09]$
HUR	$[0,1]$	$[0,1]$
HER	$[0,0.45]$	$[0,0.25]$

B. *Subsequent behavior*

We now turn to subjects' ongoing behavior, focusing on how prior cooperation influences subsequent choices. This analysis sheds light on the behavioral mechanisms behind deviations from theoretical predictions.

Our focus lies on two receiver responses: $\sigma(M|\text{“}L\text{”})$ and $\sigma(M|\text{“}R\text{”})$, measured after a prior instance of double cooperation—that is, when both players previously followed the on-path strategies of either STE or IE.[11] The first action, $\sigma(M|\text{“}L\text{”})$, reflects the receiver's tendency to choose the middle action M even when the state is likely favorable and the message indicates aligned interests. This behavior signals a breakdown in cooperation, as it contradicts both STE and IE. The second action, $\sigma(M|\text{“}R\text{”})$, captures the receiver's willingness to pay the information rent when the message suggests a conflict of interest. This response is crucial for sustaining the IE.

Regarding the two responses, we have two pieces of behavioral evidence, one is related to Results 4 and 5 and the other is related to Results 1, 2, 3.

Subsequent behavior Related to Result 4 and 5 (HUR, HER). *The upper panel in Table 5 shows that in treatments HER and HUR, $\sigma(M|\text{“}L\text{”})$ is relatively higher even after*

[11]For state R, message "R", and action M, we classify the outcome as consistent with double cooperation in IE, though it could alternatively be interpreted as partial cooperation under STE.

a round of successful coordination. This suggests that receivers revert to the stage-Nash despite recent evidence of mutual cooperation.

TABLE 5—CONDITIONAL $\sigma(M|$“L”$)$ BY TREATMENT.

Prev. Event	State	Mess.	Action	Baseline	UR	ER	H	HUR	HER
	L	“L”	L	0.21	0.18	0.25	0.11	0.35	0.27
Double-coop.	R	“R”	R	0.18	0.14	0.18	0.20	0.34	0.54
	R	“R”	M	0.47	0.18	0.29	0.44	0.75	0.47
Partial-coop.	L	“L”	M/R	0.48	0.45	0.51	0.82	0.82	0.76
	R	“L”	L	0.31	0.39	0.45	0.75	0.83	0.76
Non-coop.	R	“L”	M/R	0.51	0.47	0.54	0.85	0.92	0.84

Several factors contribute to this pattern. First, in environments with role uncertainty, successful cooperation may fail to resolve strategic doubts about future behavior—especially when players anticipate switching roles and facing different incentives. Second, the high punishment severity makes *M* particularly attractive for the receiver, as it secures a large and safe payoff regardless of the sender's message. Third, these incentive structures amplify projection bias (Loewenstein, O'Donoghue and Rabin, 2003): receivers overweight their current circumstances and undervalue the long-run consequences of defection. Specifically, when the immediate gain from deviating is salient—as it is in HER and HUR—players are more likely to act myopically, favoring short-term advantage over sustained cooperation. The combination of role uncertainty and harsh punishment thus undermines trust, even in the presence of recent coordination, and contributes to the persistent inefficiency observed in these treatments.

Subsequent behavior Related to Result 1, 2 and 3 (H, UR and ER). *The upper panel of Table 6 shows that, $\sigma(M|$“R”$)$ is lowest in the Baseline after a round of successful coordination among the Baseline, H, UR and ER, all of which are theoretically predicted to support the IE. Even when the previous round involved the receiver paying the information rent (i.e., choosing M following message* “*R*”*), receivers in the Baseline are least likely to continue doing so, with* $\sigma(M|$“R”$) = 0.68$.

This discrepancy reflects a distinct form of projection bias. In the Baseline, fixed roles imply that current receivers expect to remain in that role, reducing the salience of future gains from cooperation. At the same time, the cost of paying the information rent—giving up the short-run optimal action—feels more immediate and burdensome. This asymmetry increases the temptation to deviate. By contrast, when roles are randomized or information rent is smaller, receivers anticipate switching roles or bearing a lower information rent burden, making long-run incentives more appealing. Receivers in these environments are thus more willing to cooperate, even when doing so requires sacrificing immediate payoff.

Together, these findings demonstrate how recent cooperative history interacts with institutional features to shape ongoing behavior. While theory predicts that double co-

TABLE 6—CONDITIONAL $\sigma(M|``R")$ BY TREATMENT.

Prev. Event	State	Mess.	Action	Baseline	UR	ER	H	HUR	HER
	L	"L"	L	0.35	0.49	0.53	0.47	0.31	0.82
Double-coop.	R	"R"	R	0.08	0.25	0.07	0.15	0.08	0.24
	R	"R"	M	0.68	0.95	0.75	0.91	0.82	0.80
Partial-coop.	L	"L"	M/R	0.30	0.33	0.64	0.67	0.80	0.52
	R	"L"	L	0.18	0.46	0.54	0.78	0.77	0.83
Non-coop.	R	"L"	M/R	0.31	0.39	0.25	0.84	0.76	0.68

operation should reinforce trust and facilitate continued coordination, behavioral frictions—particularly projection bias—can erode this effect. In treatments like HER and HUR, where the immediate benefits of defection are high and future incentives are uncertain, cooperation deteriorates even after success. In the Baseline, the fixed-role structure diminishes the perceived value of future cooperation, weakening receivers' willingness to pay the information rent. These patterns help explain the divergence between theoretical sustainability and observed outcomes, highlighting the importance of behavioral considerations in repeated strategic communication.

V. Conclusion

This paper studies how role uncertainty and punishment severity shape strategic information transmission in repeated cheap talk settings. Through a combination of theory and experiment, we identify when these factors enhance or hinder the sustainability of truthful communication in long-term relationships.

When roles are fixed, stronger punishment has a clear and positive effect—both theoretically and empirically. A higher punishment level not only deters sender deviation but also reduces the burden of information rent on the receiver. This encourages cooperation even without explicit coordination mechanisms. In such environments, combining strong punishment with low information rent functions effectively as a "stick-and-carrot" mechanism, promoting efficient communication.

With random roles, the dynamics become more nuanced. Under mild punishment, theory predicts that role uncertainty facilitates cooperation, but the experimental support is limited. This may reflect our setting being close to the threshold of theoretical sustainability, where incentives are not sufficiently salient within the experimental time frame. Still, treatments with random roles rarely underperform the baseline. We expect the benefits of role uncertainty under mild punishment to become more visible over longer horizons. When punishment is strong, the predicted and observed negative effect of asymmetric role uncertainty holds. More surprisingly, even symmetric role uncertainty—where theory predicts no additional effect—leads to efficiency losses in the lab. This suggests that the interaction of strong punishment and random roles increases strategic uncertainty, making coordination more fragile. Reducing strategic uncertainty remains a practical challenge. Possible interventions include manipulating beliefs through pre-play communication or focal points, or reshaping incentive structures to alter the basin of attraction and reinforce cooperation.

Beyond strategic considerations, our findings also highlight the role of behavioral biases—particularly projection bias—in undermining cooperation. When players overweight immediate payoffs and underappreciate future consequences, even theoretically sustainable equilibria may unravel. Projection bias is especially pronounced in environments where defection yields high short-term gains or where future benefits are less salient. These distortions interact with institutional features to shape behavior in ways that theory alone may not fully anticipate.

In sum, while role uncertainty and punishment severity individually can each promote strategic communication, their interaction may generate unintended frictions. Effective communication design in long-term relationships must account not only for theoretical incentive structures but also for behavioral responses to uncertainty and potentially biased preferences.

REFERENCES

Ben-Ner, Avner, Louis Putterman, Fanmin Kong, and Dan Magan. 2004. "Reciprocity in a two-part dictator game." *Journal of Economic Behavior and Organization*, 53(3): 333–352.

Best, James, and Daniel Quigley. 2024. "Persuasion for the Long Run." *Journal of Political Economy*, 132(5): 1391–1792.

Blonski, Matthias, and Giancarlo Spagnolo. 2001. "Prisoners' other dilemma." *Stockholm School of Economics Economic Research Institute Working Paper 437*.

Blonski, Matthias, and Giancarlo Spagnolo. 2015. "Prisoners' other dilemma." *International Journal of Game Theory*, 44: 61–81.

Blume, Andreas, Ernest K. Lai, and Wooyoung Lim. 2020. "Strategic Information Transmission: A Survey of Experiments and Theoretical Foundations." *Handbook of Experimental Game Theory*, C. Monica Capra, Rachel Croson, Mary Rigdon and Tanya Rosenblat (eds)(Cheltenham, UK and Northampton, MA, USA: Edward Elgar Publishing).

Burks, S., J. Carpenter, and E. Verhoogen. 2003. "Playing both roles in the trust game." *Journal of Economic Behavior and organization*, 51: 195–216.

Chen, D.L., M. Schonger, and C. Wickens. 2016. "oTree - An open-source platform for laboratory, online and field experiments." *Journal of Behavioral and Experimental Finance*, 9: 88–97.

Classen, Timothy J. 2005. "Role Switching in Multiple Round Dictator Game Experiments." *Working paper*, https://www.tclass1.sites.luc.edu/DictatorExperiment/RoleSwitchingDictators_Aug05.pdf.

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

Dal-Bo, Pedro, and Guillaume R Frechette. 2011. "The evolution of cooperation in infinitely repeated games: Experimental evidence." *American Economic Review*, 101: 411–429.

Dal-Bo, Pedro, and Guillaume R Frechette. 2018. "On the determinants of cooperation in infinitely repeated games: A survey." *Journal of Economic Literature*, 56: 60–114.

Golosov, Mikhail, Vasiliki Skreta, Aleh Tsyvinski, and Andrea Wilson. 2014. "Dynamic strategic information transmission." *Journal of Economic Theory*, 151: 304–341.

Harsanyi, John C., and Reinhard Selten. 1988. "A General Theory of Equilibrium Selection in Games." *Cambridge, MA: MIT Press*.

Herne, Kaisa, Olli Lappalainen, and Elina Kestilä-Kekkonen. 2013. "Experimental comparison of direct, general, and indirect reciprocity." *The Journal of Socio-Economics*, 45: 38–46.

Iriberri, Nagore, and Pedro Rey-Biel. 2011. "The role of role uncertainty in modified dictator games." *Experimental Economics*, 14: 160–180.

Johnson, Noel D., and Alexandra A. Mislin. 2011. "Trust games: A meta-analysis." *Journal of Economic Psychology*, 32(5): 865–889.

Kolotilin, Anton, and Hongyi Li. 2021. "Relational communication." *Theoretical Economics*, 16(4): 1391–1430.

Kuvalekar, Aditya, Elliot Lipnowski, and João Ramos. 2022. "Goodwill in communication." *Journal of Economic Theory*, 203: 105467.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin. 2003. "Projection Bias in Predicting Future Utility." *The Quarterly Journal of Economics*, 118(4): 1209–1248.

Mesa-Vázquez, Ernesto, Ismael Rodriguez-Lara, and Amparo Urbano. 2021. "Standard vs random dictator games: On the effects of role uncertainty and framing on generosity." *Economics Letters*, 206: 109981.

Wilson, Alistair J., and Emanuel Vespa. 2020. "Information Transmission under the Shadow of the Future: An Experiment." *American Economic Journal: Microeconomics*, 12(4): 75–98.

APPENDIX

A1. Derivation of the Minimal Required Time Discount Factor $\hat{\delta}$

We consider a generalized version of the Information Rent Equilibrium (IE), which may not always be feasible but is theoretically well-defined. In this setting, the on-path strategy is: *truth-telling as sender*, and *message-following with information rent r when the truthful message is* "*R*". We analyze four possible one-shot deviations.

CASE 1. — Player A (as the current sender) deviates from truth-telling to lying in state *R*.

- No deviation:

$$r+\frac{\delta}{1-\delta}\left(\gamma\frac{1+r}{2}+(1-\gamma)\frac{2-r}{2}\right)$$

- Deviation:

$$1+\frac{\delta}{1-\delta}\left(\gamma(1-b)+(1-\gamma)b\right)$$

- Incentive compatibility condition:

$$\delta\geq\frac{2-2r}{4-3r-2b-3\gamma+2\gamma r+4\gamma b}$$

CASE 2. — Player B (as the current sender) deviates from truth-telling to lying in state *R*.

- No deviation:

$$r+\frac{\delta}{1-\delta}\left((1-\gamma)\frac{1+r}{2}+\gamma\frac{2-r}{2}\right)$$

- Deviation:

$$1+\frac{\delta}{1-\delta}\left((1-\gamma)(1-b)+\gamma b\right)$$

- Incentive compatibility condition:

$$\delta\geq\frac{2-2r}{1-r+3\gamma+2b-4\gamma b-2\gamma r}$$

CASE 3. — Player A (as the current receiver) deviates by refusing to pay the rent on a truthful message "*R*".

- No deviation:

$$1-r+\frac{\delta}{1-\delta}\left(\gamma\frac{1+r}{2}+(1-\gamma)\frac{2-r}{2}\right)$$

- Deviation:

$$1+\frac{\delta}{1-\delta}\left(\gamma(1-b)+(1-\gamma)b\right)$$

- Incentive compatibility condition:

$$\delta \geq \frac{2r}{2+r-3\gamma-2b+2\gamma r+4\gamma b}$$

CASE 4. — Player B (as the current receiver) deviates by refusing to pay the rent on a truthful message "*R*".

- No deviation:

$$1-r+\frac{\delta}{1-\delta}\left((1-\gamma)\frac{1+r}{2}+\gamma\frac{2-r}{2}\right)$$

- Deviation:

$$1+\frac{\delta}{1-\delta}\left((1-\gamma)(1-b)+\gamma b\right)$$

- Incentive compatibility condition:

$$\delta \geq \frac{2r}{-1+3r+3\gamma+2b-2\gamma r-4\gamma b}$$

These four inequalities jointly determine the minimal discount factor $\hat{\delta}$ needed to sustain the generalized IE. The two equilibrium types discussed in the main text—IE and STE—are special cases of this framework:

- **IE:** $r=1-b$ (receiver pays the information rent)
- **STE:** $r=0$ (receiver always plays the myopic best response)

SPECIAL CASE: IE ($r=1-b$)The inequalities simplify as follows:

- Case 1:

$$\delta \geq \frac{2b}{1+b-\gamma+2\gamma b}$$

- Case 2:

$$\delta \geq \frac{2b}{3b+\gamma-2\gamma b}$$

- Case 3:

$$\delta \geq \frac{2-2b}{3-3b-\gamma+2\gamma b}$$

- Case 4:

$$\delta \geq \frac{2-2b}{2+\gamma-b-2\gamma b}$$

SPECIAL CASE: STE ($r = 0$)In this case, the receiver has no incentive to deviate, as they always choose the myopic best response. Only the sender's constraints matter:

- Case 1:
$$\delta \geq \frac{2}{4 - 2b - 3\gamma + 4\gamma b}$$
- Case 2:
$$\delta \geq \frac{2}{1 + 2b + 3\gamma - 4\gamma b}$$

COMPUTING $\hat{\delta}$For any given (γ, b) pair, the maximum of these inequalities yields the minimal required discount factor $\hat{\delta}$. The resulting values for STE and IE are illustrated in Figure 3 and Figure 4, respectively.

A2. *The Sample Instructions: HUR*

Welcome to the experiment! This experiment studies decision-making in pairs of individuals. Please read the instruction carefully. Please do not communicate with other participants during the experiment. The experiment will last about 60 minutes. The payment you will receive from this experiment depends on your decisions.

1. Overview of this experiment

You will go through the following modules:

- **Instruction:** you will learn about the basic environment of this experiment.
- **Quiz:** it will give you an opportunity to learn whether your understanding about the instruction is correct or not.
- **One Practice Game:** it will allow you to be familiar with the interface of the game.
- **Seven Official Games:** the formal part of the experiment.

2. Your session and pair

There are XX participants in this session, including yourself. At the beginning of the experiment, one-half of the participants will be assigned with a player ID "A" and the other half with a player ID "B". Your player ID will remain fixed throughout the experiment.

The experiment consists of 7 official games. At the beginning of each game, one player A and one player B are randomly paired. The pair is fixed within each game. After each game, participants will be reshuffled to form new pairs. Everyone won't know the identity of the paired participant.

3. Overview of each game: roles, states and choices

Each game consists of at least 1 round, possibly multiple rounds of interactions.

At the beginning of each round, there are two possibilities of the role assignment: "A is the Sender, B is the Receiver" (with 70% chance) or "B is the Sender, A is the Receiver" (with 30% chance). This means that the player A is more likely to be the Sender in each round.

At the beginning of each round, there are two possibilities of the state: Left or Right. One of the states will be randomly determined by the computer. They are equally likely.

If your role in the current round is Sender, you will observe the state. Then you will be asked to choose a message to send to the Receiver: "State is left" or "State is right". You are allowed to send any of these two messages regardless of the actual state being Left or Right.

If your role in the current round is Receiver, you will not observe the state. Instead, you will receive a message from the Sender in your pair. Then you will be asked to choose an action: Left, Middle, or Right.

Main takeaway: you will possibly play multiple rounds in each game. You will interact with the same participant within the game. But the state (Left or Right) and your roles (Sender or Receiver) will be random in each round. The participant with player ID "A" will have more chance to be the Sender in each round.

4. Overview of each game: round-reward

The reward you can get in each round depends on the current state and the current Receiver's action, which can be summarized in the table below. Each cell corresponds to the reward for a given state and action. Remember that in each cell, no matter whether you are player A or B, the first element denotes the current Sender's round-reward, while the second element denotes the current Receiver's round-reward. One unit denotes 1HKD.

For example, in the cell located by "State is Right, Receiver's Action is Left", the Sender's round-reward is 40, and the Receiver's round-reward is 0.

If you are the current Sender in a round, your round-reward will be 40 if the Receiver's action is Left, no matter the state. It will be 4 if the Receiver's action is Middle, no matter the state. It will be 0 if the Receiver's action is Right, no matter the state.

If you are the current Receiver in a round, your round-reward will be 40 if your action correctly matches the state. It will be 36 if your action is Middle, no matter the state. It will be 0 if your action is the opposite to the state.

TABLE A1—PAYOFF DISTRIBUTION.

		Receiver's action		
		Left	Middle	Right
State	Left	40,40	4,36	0,0
	Right	40,0	4,36	0,40

When the game terminates, your total reward of each game is the sum of all round-rewards.

5. Overview of each game: information feedback

At the end of each round, you will be provided with a summary of what happened in this round, including the state, the Sender's chosen message, the Receiver's chosen action, and your reward in this round.

Within each game, during each round, no matter you are Sender or Receiver, you can see the history information about the previous rounds (if exist), including the past role assignments, past states, past messages and the past actions.

The following is a constructed example of the available history at round 4: we have the information about the past three rounds.

For instance, we can know from the list that in round 2, the realized state was Left, the role assignment was "A was the Receiver, B was the Sender", the message from the Sender (played by B) was "State is Left", and the action from the Receiver (played by A) was Left.

TABLE A2—EXAMPLE OF HISTORY INFORMATION.

Round	State	Sender	Message	Action
1	Right	A	"State is Right"	Right
2	Left	B	"State is Left"	Left
3	Right	A	"State is Right"	Right

6. Random termination of each game

As mentioned, each game has at least one round, possibly multiple rounds. After each round, there are two possibilities: the game terminates, or it continues to the next round.

The probability of going to the next round is 80%, while the probability of termination is 20%. You will see a spinning wheel that consists of red area (20%) and green area (80%) as illustrated below. Once you click the "Spin" button, the wheel starts spinning. If the spinning wheel stops at the green area, the game continues. Otherwise, the game terminates. Both participants in the same pair will see the same outcome from the wheel for each round.

Note that no matter how many rounds the game has been played, the probability of continuation is always 80%.

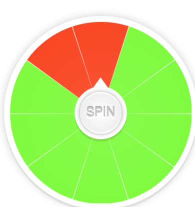

7. Summarizing the uncertainties in each game

In the previous pages, you have learnt the details of the game, including your possible roles, available choices and the corresponding rewards. Basically, you will play 7 official games. In each game, you may have multiple rounds. In each round, the current Sender sends a message after observing the State, and the current Receiver takes an action after receiving the message.

In each round of the game, there are three aspects of uncertainty: random role assignment, random state and random termination.

At the beginning of each round, roles (Sender or Receiver) are randomly assigned with an unequal chance (if you are player A, you have 70% chance to be the Sender; if you are player B, you have 30% chance to be the Sender), and a state (Left or Right) is randomly determined with a 50-50 chance. At the end of each round, the game will continue to the next round with 80% chance.

8. Your cash payment

The practice game won't affect your cash payment. After completing the last official game, the computer will randomly select one game out of 7 games for your payment. Every game has an equal chance to be selected for your payment, so it is in your best interest to take each game equally seriously.

For example: Suppose the randomly selected game ended at the 5th round and the round-rewards for you were 36, 40, 40, 40, 0, then your final cash payment is 40HKD show-up fee plus the total reward of the selected game 156HKD, which in total is 196HKD.

9. Quiz and practice game

Next, you will first finish a short quiz. You can only proceed with all correct answers. Then, you will participate in a practice game. The official games will follow afterward.

If you have any confusion about the process of the experiment, you can review the instruction. If you are ready for the quiz, please press "NEXT PAGE".

Quiz

Please answer the following 8 questions:

1) During the whole experiment, you will be paired with the same participant in all games. (True or False)

2) In each game, you may be the Sender for several rounds and be the Receiver for other rounds, depending on the computer's randomization. If you are player B, you are more likely to be the Sender than your paired participant. (True or False)

3) In each game, it is possible to have only one round. (True or False)

4) The Sender, but not the Receiver, will learn the current state determined by the computer. (True or False)

5) Within a game, before making decision, you will know the past decisions in previous rounds, but not the past states. (True or False)

6) Your final cash payment may depend on your paired participant's decision. (True or False)

7) Look at the round-reward table. If the state is Left and the Receiver's action is Right, what will be the Sender's round-reward? (40, 36, 4, 0)

8) Your future roles (Sender or Receiver) may affect your total reward within each game. (True or False)

A3. Test Statistics

Table A3 presents the *p*-value of Mann-Whitney U tests using four scales of data. For example, at the all-game all-round scale, we use all the data. We use super-game level data as independent observations. The corresponding *p*-values indicate whether the difference in average efficiency at the all-game all-round scale between two treatments is significant or not.

TABLE A3—*p*-VALUES OF MANN-WHITNEY U TESTS.

Difference	Game 1 round 1	p	Game 1 all rounds	p	All games round 1	p	All games all rounds	p
ER− Baseline	+	0.937	+	0.130	+	0.865	+	0.227
ER− H	+	0.659	+	0.641	−	0.425	−	0.925
ER− UR	+	0.906	−	0.643	−	0.647	−	0.451
ER− HER	+	0.263	+	0.114	+	0.124	+	0.000
ER− HUR	+	0.111	+	0.022	+	0.014	+	0.000
Baseline− H	+	0.736	−	0.427	−	0.354	−	0.340
Baseline− UR	+	0.981	−	0.018	−	0.544	−	0.055
Baseline− HER	+	0.311	+	0.621	+	0.186	+	0.000
Baseline− HUR	+	0.140	+	0.228	+	0.028	+	0.000
H− UR	−	0.795	−	0.413	+	0.721	−	0.470
H− HER	+	0.492	+	0.363	+	0.029	+	0.000
H− HUR	+	0.256	+	0.153	+	0.002	+	0.000
UR− HER	+	0.369	+	0.024	+	0.053	+	0.000
UR− HUR	+	0.186	+	0.001	+	0.004	+	0.000
HER− HUR	+	0.735	+	0.638	+	0.465	+	0.074

A4. The Basin-of-attraction Analysis

DERIVATION OF THE PAYOFF MATRICES. — We derive the compressed payoff matrix of STE-BE for UR for example. UR is the most general case which has four contingencies and asymmetric role randomness.

The four contingencies are: case 1 - current state is L, current sender (receiver) is A (B); case 2 - current state is R, current sender (receiver) is A (B); case 3 - current state is L, current sender (receiver) is B (A); case 4 - current state is R, current sender (receiver) is B (A).

In case 1, if both player A and B cooperate on STE-STE, then player A's payoff is

$$1+\frac{0.8}{1-0.8}(0.7\times\frac{1+0}{2}+0.3\times\frac{1+1}{2})=3.6$$

where the first term is A's current payoff as sender in state L, the first term in the bracket is the sender's expected round-payoff weighted by the probability of A as sender, the second term in the bracket is the receiver's expected round-payoff weighted by the probability of A as receiver.

Similarly, player B's STE-STE payoff is

$$1+\frac{0.8}{1-0.8}(0.3\times\frac{1+0}{2}+0.7\times\frac{1+1}{2})=4.4$$

If player A chooses STE and B chooses BE, then player A's payoff is

$$0.5+\frac{0.8}{1-0.8}0.5=2.5$$

where the first term is current sender's payoff when current receiver takes BE action M, and the second term is the time discounted total payoff in BE no matter sender or receiver.

Similarly, player B's STE-BE payoff is

$$0.5+\frac{0.8}{1-0.8}0.5=2.5$$

If both player A and B choose BE-BE, then both player A and B's payoff is

$$0.5+\frac{0.8}{1-0.8}0.5=2.5$$

If player A chooses BE and player B chooses STE, this is a bit more complicated. Since choosing BE as current sender means always sending message "L" no matter the state, when the state is revealed at the end of the current round, in player B's view as current receiver, player A did not lie. Since player B is playing the grim-trigger strategy, she will continue doing so until she identifies deviation. Note that if the state is R, then the current receiver can immediately tell the current sender lied once the state is revealed, and the payoffs in that case can be computed. So, suppose player A's payoff is x, we have the following equation

$$x=1+0.8\times(0.7\times0.5\times x+0.7\times0.5\times3+0.3\times0.5\times2.5+0.3\times0.5\times2.5)$$

where the first term is the payoff of current sender, the first term in the bracket is the payoff in case 1 for the sender weighted by the probability of case 1, the second term in the numerator is the payoff in case 2 for the sender weighted by the probability of case 2, and so on. Solving for x, we get

$$x = 3.4$$

Similarly, the equation for player B is

$$y = 1 + 0.8 \times (0.7 \times 0.5 \times y + 0.7 \times 0.5 \times 2 + 0.3 \times 0.5 \times 2.5 + 0.3 \times 0.5 \times 2.5)$$

where the solution is

$$y = 3$$

The compressed payoff matrices in Table A4 and A5 are for the STE-BE discussion. The compressed payoff matrices in Table A6 and A7 are for the IE-BE discussion.

TABLE A4—COMPRESSED PAYOFF MATRICES FOR STE-BE DISCUSSION (PART 1).

Baseline					
State $= L$			State $= R$		
S \R	STE	BE	S \R	STE	BE
STE	(3, 5)	(2.5, 2.5)	STE	(2, 5)	(2.5, 2.5)
BE	(3.7, 3)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)
UR					
State $= L$			State $= R$		
AS \BR	STE	BE	AS \BR	STE	BE
STE	(3.6, 4.4)	(2.5, 2.5)	STE	(2.6, 4.4)	(2.5, 2.5)
BE	(3.4, 3)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)
AR \BS	STE	BE	AR \BS	STE	BE
STE	(3.6, 4.4)	(3, 3.1)	STE	(3.6, 3.4)	(2, 3)
BE	(2.5, 2.5)	(2.5, 2.5)	BE	(2.5, 2.5)	(2.5, 2.5)
ER					
State $= L$			State $= R$		
AS \BR	STE	BE	AS \BR	STE	BE
STE	(4, 4)	(2.5, 2.5)	STE	(3, 4)	(2.5, 2.5)
BE	(3.3, 3)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)
AR \BS	STE	BE	AR \BS	STE	BE
STE	(4, 4)	(3, 3.3)	STE	(4, 3)	(2, 3)
BE	(2.5, 2.5)	(2.5, 2.5)	BE	(2.5, 2.5)	(2.5, 2.5)

TABLE A5—COMPRESSED PAYOFF MATRICES FOR STE-BE DISCUSSION (PART 2).

		H			
State $= L$			State $= R$		
S \R	STE	BE	S \R	STE	BE
STE	(3, 5)	(0.5, 4.5)	STE	(2, 5)	(0.5, 4.5)
BE	(0.9, 2.4)	(0.5, 4.5)	BE	(1.4, 3.6)	(0.5, 4.5)
		HUR			
State $= L$			State $= R$		
AS \BR	STE	BE	AS \BR	STE	BE
STE	(3.6, 4.4)	(1.46, 3.54)	STE	(2.6, 4.4)	(1.46, 3.54)
BE	(3.06, 3.33)	(1.46, 3.54)	BE	(2.36, 2.64)	(1.46, 3.54)
AR \BS	STE	BE	AR \BS	STE	BE
STE	(3.6, 4.4)	(2.25, 3.89)	STE	(3.6, 3.4)	(1.36, 3.64)
BE	(2.26, 2.74)	(2.26, 2.74)	BE	(2.26, 2.74)	(2.26, 2.74)
		HER			
State $= L$			State $= R$		
AS \BR	STE	BE	AS \BR	STE	BE
STE	(4, 4)	(2.1, 2.9)	STE	(3, 4)	(2.1, 2.9)
BE	(3.5, 2.8)	(2.1, 2.9)	BE	(3, 2)	(2.1, 2.9)
AR \BS	STE	BE	AR \BS	STE	BE
STE	(4, 4)	(2.8, 3.5)	STE	(4, 3)	(2, 3)
BE	(2.9, 2.1)	(2.9, 2.1)	BE	(2.9, 2.1)	(2.9, 2.1)

BASIN COMPUTATION. — Given the payoff matrices, we can compute the basin of attraction for BE in each case for each treatment. For the column player, denote the probability of playing the cooperative equilibrium (STE in STE-BE discussion and IE in IE-BE discussion) by p and denote the probability of playing the non-cooperative BE by $1-p$. For the row player, denote the probability of playing the cooperative equilibrium (STE in STE-BE discussion and IE in IE-BE discussion) by q and denote the probability of playing the non-cooperative BE by $1-q$.

For example, consider the STE-BE discussion for treatment UR. In case 1, for the player A as current sender to think of BE as best response to player B's mixed strategy $p\cdot$STE $+(1-p)\cdot$BE, we need

$$3.6 \times p + 2.5 \times (1-p) \leq 3.4 \times p + 2.5 \times (1-p)$$

leading to $p \in \emptyset$.

Similarly, we can conduct this computation for all contingencies in all treatments. Table A8 includes the basins for the STE-BE discussion. Table A9 includes the basins for the IE-BE discussion.

TABLE A6—COMPRESSED PAYOFF MATRICES FOR IE-BE DISCUSSION (PART 1).

Baseline					
	State $= L$			State $= R$	
S \R	IE	BE	S \R	IE	BE
IE	(4, 4)	(2.5, 2.5)	IE	(3.5, 3.5)	(2, 3)
BE	(3.7, 3)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)
UR					
	State $= L$			State $= R$	
AS \BR	IE	BE	AS \BR	IE	BE
IE	(4, 4)	(2.5, 2.5)	IE	(3.5, 3.5)	(2, 3)
BE	(3.6, 2.9)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)
AR \BS	IE	BE	AR \BS	IE	BE
IE	(4, 4)	(2.8, 3.3)	IE	(3.5, 3.5)	(2, 3)
BE	(2.5, 2.5)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)
ER					
	State $= L$			State $= R$	
AS \BR	IE	BE	AS \BR	IE	BE
IE	(4, 4)	(2.5, 2.5)	IE	(3.5, 3.5)	(2, 3)
BE	(3.4, 2.9)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)
AR \BS	IE	BE	AR \BS	IE	BE
IE	(4, 4)	(2.9, 3.4)	IE	(3.5, 3.5)	(2, 3)
BE	(2.5, 2.5)	(2.5, 2.5)	BE	(3, 2)	(2.5, 2.5)

ORGANIZING THE BASINS. — We apply a specific multiplication rule: (1) For intervals like $[0,a]$ where $0 < a < 1$, the multiplication of two such intervals results in an interval whose upper bound is the product of the upper bounds of the original intervals; (2) An interval $[0,a]$ with $0 < a < 1$ multiplied with $\emptyset$ results in $\emptyset$; (3) An interval $[0,a]$ with $0 < a < 1$ multiplied with $[0,1]$ results in $[0,1]$; (4) An interval $[0,a]$ with $0 < a < 1$, multiplied with "$=$" results in $[0,a]$ itself; (5) The interval $[0,1]$ multiplied with $\emptyset$ results in $[0,1]$; (6) The interval $[0,1]$ multiplied with "$=$" results in $[0,1]$; (7) The null set $\emptyset$ multiplied with "$=$" results in "$=$". The aggregated basins for the STE-BE and the IE-BE discussions are displayed in Table A10.

Table A7—Compressed payoff matrices for IE-BE discussion (part 2).

H					
State $= L$			State $= R$		
S \R	IE	BE	S \R	IE	BE
IE	(3.2, 4.8)	(0.5, 4.5)	IE	(2.3, 4.7)	(0.4, 4.6)
BE	(0.9, 2.4)	(0.5, 4.5)	BE	(1.4, 3.6)	(0.5, 4.5)
HUR					
State $= L$			State $= R$		
AS \BR	IE	BE	AS \BR	IE	BE
IE	(3.68, 4.32)	(1.46, 3.54)	IE	(2.78, 4.22)	(1.36, 3.64)
BE	(3.08, 3.31)	(1.46, 3.54)	BE	(2.36, 2.64)	(1.46, 3.54)
AR \BS	IE	BE	AR \BS	IE	BE
IE	(3.68, 4.32)	(2.22, 3.92)	IE	(3.58, 3.42)	(1.36, 3.64)
BE	(2.26, 2.74)	(2.26, 2.74)	BE	(2.36, 2.64)	(2.26, 2.74)
HER					
State $= L$			State $= R$		
AS \BR	IE	BE	AS \BR	IE	BE
IE	(4, 4)	(2.1, 2.9)	IE	(3.1, 3.9)	(2, 3)
BE	(3.5, 2.8)	(2.1, 2.9)	BE	(3, 2)	(2.1, 2.9)
AR \BS	IE	BE	AR \BS	IE	BE
IE	(4, 4)	(2.8, 3.5)	IE	(3.9, 3.1)	(2, 3)
BE	(2.9, 2.1)	(2.9, 2.1)	BE	(3, 2)	(2.9, 2.1)

TABLE A8—BASINS OF ATTRACTION FOR EACH CONTINGENCY (STE-BE).

State	Role	Basin of Sender	Basin of Receiver
		Baseline	
L	AS, BR	$[0,1]$	$\emptyset$
R	AS, BR	$[0,1]$	$[0,0.17]$
		UR	
L	AS, BR	$\emptyset$	$\emptyset$
R	AS, BR	$[0,1]$	$[0,0.21]$
L	AR, BS	$\emptyset$	$\emptyset$
R	AR, BS	$[0,0.31]$	$\emptyset$
		ER	
L	AS, BR	$\emptyset$	$\emptyset$
R	AS, BR	$=$	$[0,0.25]$
L	AR, BS	$\emptyset$	$\emptyset$
R	AR, BS	$=$	$[0,0.25]$
		H	
L	AS, BR	$\emptyset$	$[0,0.81]$
R	AS, BR	$\emptyset$	$[0,0.64]$
		HUR	
L	AS, BR	$\emptyset$	$[0,0.20]$
R	AS, BR	$[0,0.01]$	$[0,0.50]$
L	AR, BS	$\emptyset$	$\emptyset$
R	AR, BS	$[0,0.40]$	$[0,1]$
		HER	
L	AS, BR	$\emptyset$	$[0,0.08]$
R	AS, BR	$=$	$[0,0.45]$
L	AR, BS	$\emptyset$	$[0,0.08]$
R	AR, BS	$=$	$[0,0.45]$

TABLE A9—BASINS OF ATTRACTION FOR EACH CONTINGENCY (IE-BE).

State	Role	Basin of Sender	Basin of Receiver
		Baseline	
L	AS, BR	$\emptyset$	$\emptyset$
R	AS, BR	$[0,0.5]$	$[0,0.5]$
		UR	
L	AS, BR	$\emptyset$	$\emptyset$
R	AS, BR	$[0,0.5]$	$[0,0.5]$
L	AR, BS	$\emptyset$	$\emptyset$
R	AR, BS	$[0,0.5]$	$[0,0.5]$
		ER	
L	AS, BR	$\emptyset$	$\emptyset$
R	AS, BR	$[0,0.5]$	$[0,0.5]$
L	AR, BS	$\emptyset$	$\emptyset$
R	AR, BS	$[0,0.5]$	$[0,0.5]$
		H	
L	AS, BR	$\emptyset$	$[0,0.88]$
R	AS, BR	$[0,0.1]$	$[0,0.9]$
		HUR	
L	AS, BR	$\emptyset$	$[0,0.23]$
R	AS, BR	$[0,0.19]$	$[0,0.61]$
L	AR, BS	$[0,0.03]$	$\emptyset$
R	AR, BS	$[0,0.42]$	$[0,1]$
		HER	
L	AS, BR	$\emptyset$	$[0,0.08]$
R	AS, BR	$[0,0.5]$	$[0,0.5]$
L	AR, BS	$\emptyset$	$[0,0.08]$
R	AR, BS	$[0,0.5]$	$[0,0.5]$

TABLE A10—AGGREGATED BASINS FOR STE-BE AND IE-BE COMPARISONS.

Treatment	STE-BE Basin	IE-BE Basin
Baseline	$[0,1]$	$[0,0.25]$
UR	$[0,1]$	$[0,0.25]$
ER	$[0,0.25]$	$[0,0.25]$
H	$\emptyset$	$[0,0.09]$
HUR	$[0,1]$	$[0,1]$
HER	$[0,0.45]$	$[0,0.25]$