Bayesian Persuasion and Reciprocity: Theory and Experiment

Pak Hung Au and King King Li

Presenter: Myunghwan Lee (Yonsei University)

Jan 7, 2020

Introduction

- Bayesian persuasion falls into the broad field of information economics.
- The role of information has been widely studied in economics: Information acquisition, asymmetric information, communication, cheap talk, rational inattention ...
- The currently most active area of research in information economics is probably information design.
- Information design investigates the optimal informational environment (who should know what and when), taking as given the preferences of the players and objective function over the players' actions.

Introduction

- Information design can be seen as a parallel to mechanism design.
- In mechanism design problems, the allocation of information (i.e., who knows what) is given and the designer chooses the optimal mechanism.
- In information design problems, the mechanism is given and the designer influences the outcome by specifying the allocation of information.
- Thus, in mechanism design, a mechanism is picked for a given information structure, while in information design, an information structure is picked for a given mechanism.

Introduction

- A planner in information design wants to persuade the receiver who react to information in a rational (Bayesian) manner: thus the term Bayesian persuasion.
- First formerly introduced in Kamenica and Gentzkow (2009), Bayesian persuasion asks how and when the sender with information advantage over receiver can persuade the receiver's behavior via provision of information.
- It focuses on environments where the information designer is motivated by the desire to influence the actions of those who observe the signal realization.

An Illustrative Example

- Consider the example of a prosecutor (Sender) trying to convince a judge (Receiver) that a defendant is guilty.
- There are two states of the world: the defendant is either guilty or innocent. The judge must choose one of two actions: to acquit or convict a defendant.
- The judge gets utility 1 for choosing the just action (convict when guilty and acquit when innocent) and utility 0 for choosing the unjust action (convict when innocent and acquit when guilty).
- The prosecutor gets utility 1 if the judge convicts and utility 0 if the judge acquits, regardless of the state.
- The prosecutor and the judge share a prior belief, $Pr(guilty) = 0.3$.

An Illustrative Example

- If there is no information, the judge always acquits because guilt is less likely than innocence under her prior.
- If the prosecutor chooses a fully informative investigation, one that leaves no uncertainty about the state, the judge convicts 30 percent of the time.
- The prosecutor can still do better with the following signal,

$$\pi(i \mid \textit{innocent}) = \frac{4}{7}, \pi(i \mid \textit{guilty}) = 0$$
$$\pi(g \mid \textit{innocent}) = \frac{3}{7}, \pi(g \mid \textit{guilty}) = 1$$

- Upon seeing g, the judge convicts. Upon seeing i, the judge acquits.
- Even though the judge knows only 30% are guilty and the signal is intended to maximize the probability of conviction in favor of the prosecutor, she convicts 60% of them.

An Illustrative Example

- It can be shown that the signal above is optimal for sender.
- $\pi(i \mid guilty)$ can't be bigger than 0. The prosecutor can increase his payoff by decreasing $\pi(i \mid guilty)$ and increasing $\pi(g \mid guilty)$, which strictly increases the probability of g and the willingness of the judge to convict when she sees g.
- Moreover, when judge convicts, she should be exactly indifferent between two options. If she strictly preferred to convict upon seeing g, the prosecutor can increase his payoff by slightly decreasing $\pi(i \mid innocent)$ and increasing $\pi(g \mid innocent)$, which increases the probability of g and leaves the judge's optimal action given the signal unchanged.

An Illustrative Example

- Intuitively, the sender designs a signal structure that includes noise in guilty signal to increase the probability of materialization, but not too much to induce the receiver to distrust the signal.
- Leaving just as much information as the receiver is induced to choose 'convict' upon seeing g, the sender maximizes the probability of the realization of signal g.
- In terms of posterior beliefs, the sender chooses a signal structure that induces only two posteriors 0 and 0.5.

Motivation

- An implicit premise in most of economics is that behavior is driven by three factors: preferences, technology, and information.
- The prosecutor in the example has no way to affect the judge's payoff through providing additional incentives.
- Instead, she exploits the way the receiver updates her belief according to Bayes's Law by controlling the information structure through which the receiver can learn about the true state.
- Numerous applications explored.

Motivation

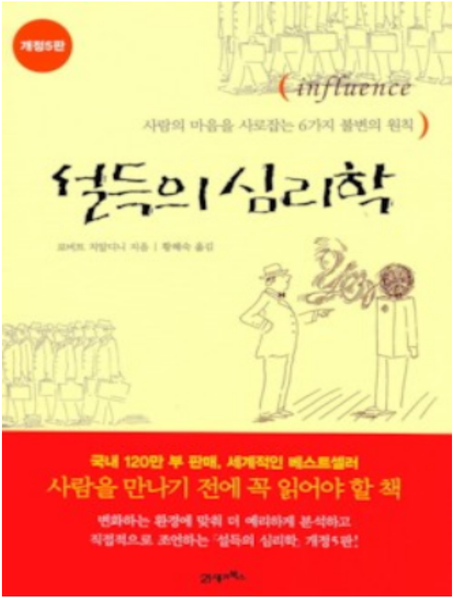
개정5판
influence
사람의 마음을 사로잡는 6가지 불변의 원칙
설득의 심리학
국내 120만 부 판매, 세계적인 베스트셀러
사람을 만나기 전에 꼭 읽어야 할 책
변화하는 환경에 맞춰 더 예리하게 분석하고
직접적으로 조언하는 「설득의 심리학」 개정5판!

Motivation

- However, persuasion in reality works in much complicated manner.
- Receivers' responses may not be as mechanical as the standard model assumes. They make judgments on the sender's behavior which clearly affect their subsequent responses.
- If the receiver perceives that the sender has acted kindly, she would derive utility from returning a favor to the sender.
- Conversely, receivers may not act as the sender induced them to if she perceives the sender's action as unkind, even at the cost of monetary payoff.
- In the example above, while the optimal posterior belief for sender was 0 and 0.5, the receiver's expected monetary payoff is minimized at the posterior belief of 0.5 (too much noise!).

Motivation

- Reciprocation can play a key role in successful persuasion.
- Upon seeing the signal, the receiver perceives it as kind or unkind and react accordingly.
- For example, the receiver's response may not jump sharply at the posterior belief of 50% as the model predicts. She may be willing to punish the sender by choosing the other option at the posterior exceeding 0.5 only marginally.
- The sender rationally expects this in equilibrium, designing more generous(=more informative, less confusing) signal than the standard model predicts.
- This paper proposes a model of Bayesian persuasion that incorporates reciprocity and conduct laboratory experiments to test the model's predictions.

Experimental Design

- Two players: player A (sender) and player B (receiver).
- Player A is given 100 balls in which 100*μ balls are red and 100*(1- μ) balls are black (for some μ between 0 and 1).
- Player A allocates 100 balls into 5 urns (empty urns are allowed). One urn will be randomly drawn, with the probability being the number of balls in the urn divided by 100.
- The composition (i.e., number of red balls and number of black balls) of the drawn urn will be announced to player B.
- One ball will be randomly drawn from the urn, and player B's task is to guess the color of the ball drawn.

Experimental Design

- If her guess is correct, player B receives payoff and zero otherwise. On the other hand, player A receives payoff if player B guesses red regardless of the actual color of the ball drawn.
- In sum, the objective of player A is to design an allocation of balls into urns with the objective of persuading player B to guess red.
- This ball allocation game is strategically equivalent to a Bayesian persuasion game in which the prior of the state is μ.
- Instead of conditional distribution of signals, senders choose a distribution over posteriors.
- Same problem for the sender; Maximize the probability of realizing the signal, so long as the receiver guesses red upon seeing that signal.

Experimental Design

- Experimental studies on Bayesian persuasion are relatively scant. This experimental design avoids abstract language and tests the key predictions of Bayesian persuasion model with simplicity.
- An advantage is that the receiver faces a simple decision that does not involve probability updating: they are given the exact posterior.
- Mistakes in probability calculation or non-Bayesian updating are ruled out as explanations to any behavioral responses by the receiver, leading to more accurate identification.
- The flexibility in the design of information structure (ball allocation) allows us to test whether the prediction of using only two signals holds in the laboratory.

Model

- **Lemma 1** : The unique perfect Bayesian equilibrium is as follows. The receiver chooses R if and only if the posterior belief induced the signal realization is no less than 0.5. If $\mu < \frac{1}{2}$, the sender chooses a signal structure that induces only two posteriors 0 and 0.5. If $\mu = \frac{1}{2}$, the sender chooses a signal structure that induces only a posterior 0.5.
- The equilibrium ball allocation predicted by the standard Bayesian persuasion model is to allocate 100*μ red balls and 100*μ black balls in one urn, and the rest (of black balls) into other urns.
- In the unique perfect Bayesian equilibrium, the receiver's response, given by a simple step function in posterior beliefs, is independent of the prior μ.

Model

- Reciprocity brings significant changes. The key idea is that receiver perceives a posterior belief of 0.5 as unkind, as it is the most uninformative signal but at the same time maximizes the sender's expected payoff.
- Following this idea, this paper develops standard Bayesian persuasion game into psychological game à la Falk and Fischbacher (2006).
- On top of the standard monetary payoffs, a player's utility function consists of terms that capture the kindness of the other player, as well as an appropriate reciprocation to the received treatment.
- This leads to a number of novel predictions, which will be supported by experimental data.

Model

- Let $\pi_i(\sigma_S, \sigma_R)$ be the monetary payoff of player $i \in \{S, R\}$ if the sender and receiver play strategy σ_S and σ_R respectively.
- Let $k_S(\sigma_R', \sigma_R'')$ be the sender's kindness *perceived* by the receiver.
- σ_R' is the receiver's first-order belief concerning the sender's strategy, and σ_R'' is the receiver's second-order belief concerning the sender's belief on the receiver's strategy. Then,
$$k_S(\sigma_R', \sigma_R'') = \pi_R(\sigma_R', \sigma_R'') - \pi_S(\sigma_R', \sigma_R'')$$
- This comes from $\pi_R(\sigma_S, \sigma_S') - \pi_S(\sigma_S, \sigma_S')$.
- The higher the payoff that the sender's strategy is *perceived* to bring to the receiver relative to the sender's own payoff, the kinder the receiver thinks the sender is.

Model

- Next, the reciprocation term of the receiver is defined as,
$$\rho_R(\sigma_R, \sigma'_R, \sigma''_R) = \pi_S(\sigma'_R, \sigma_R) - \pi_S(\sigma'_R, \sigma''_R)$$
- This can be interpreted as the alteration in the sender's payoff brought about by the receiver changing her strategy from σ''_R to σ_R.

- The receiver's reciprocity utility is defined as the product of the kindness term and the reciprocation term.
- The idea is that if the receiver perceives the sender to be kind, then she derives a positive utility by returning the sender a favor.
- Conversely, if the receiver perceives the sender to be unkind, then he/she derives a positive utility by taking action to lower the sender's monetary payoff.

Model

- The receiver's overall utility is defined as

$$U_R(\sigma_S, \sigma_R; \sigma'_R, \sigma''_R) \equiv \pi_R(\sigma_S, \sigma_R) + \lambda_R k_s(\sigma'_R, \sigma''_R)\rho_R(\sigma_R, \sigma'_R, \sigma''_R),$$

where λ_R is the receiver's reciprocation parameter.

- λ_R is a positive constant that measures the strength of the reciprocal preference.
- Players' actions as well as beliefs are considered as factors of the receiver's utility function.
- We assume that the sender does not care about reciprocation, i.e., the sender's reciprocation parameter is zero and she maximizes her monetary payoff only.
- This simplification is without loss of generality if the sender's reciprocation parameter is not too large.

Model

- As the players' utilities are assumed to depend on their beliefs, the reciprocity game belongs to the class of psychological games à la Geanakoplos, Pearce, and Stacchetti (1989).
- Traditional game-theoretic models assume that utilities depend only on actions. This is sometimes not sufficient for describing the motivations and choices of decision makers who care about reciprocity, emotions, or social rewards.
- Psychological games allow utilities to depend directly on beliefs (about beliefs) besides which actions are chosen, and they can capture a wider range of motivations.
- The equilibrium notion we adopt is standard in psychological games.
  - Given beliefs, each player maximizes their expected utility.
  - The beliefs match the actual behaviors.

Model

- For simplicity, we restrict senders to provide at most two urns. Feasibility constraint implies that one of which has a fraction of red balls no higher than μ, whereas the other has a fraction of red balls no lower than μ.
- Under this restriction of strategy space, a generic strategy of the sender is a pair of fractions (p, q), with p $\geq \mu \geq q$, describing the respective proportion of red balls in the two urns.
- A generic strategy of the receiver specifies the probability $\sigma_R(p)$ guessing red after an urn with a fraction p of red balls is drawn.
- We assume that given a certain equilibrium sender strategy (p^*, q^*), the receiver evaluates the sender's kindness only by the fraction of red balls in the drawn urn.

Model

- We specify the kindness terms in the equilibrium of our Bayesian persuasion game as follows. There are two cases.
- Suppose the fraction of red balls is $p \geq \mu$ in the drawn urn and let σ''_R the receiver's second-order belief about her own strategy.
- Then the sender's kindness perceived by the receiver is given by,

$$k_s(p, \sigma''_R; (p^*, q^*)) = [\tfrac{p-\mu}{p-q^*}\{q^*\sigma''_R(q^*) + (1-q^*)(1-\sigma''_R(q^*))\} + \tfrac{\mu-q^*}{p-q^*}\{p\sigma''_R(p) + (1-p)(1-\sigma''_R(p))\}] - [\tfrac{p-\mu}{p-q^*}\sigma''_R(q^*) + \tfrac{\mu-q^*}{p-q^*}\sigma''_R(p)]$$

Model

$$k_s(p, \sigma''_R; (p^*, q^*)) = [\frac{p-\mu}{p-q^*}\{q^*\sigma''_R(q^*) + (1-q^*)(1-\sigma''_R(q^*)\} + \frac{\mu-q^*}{p-q^*}\{p\sigma''_R(p) + (1-p)(1-\sigma''_R(p)\}] - [\frac{p-\mu}{p-q^*}\sigma''_R(q^*) + \frac{\mu-q^*}{p-q^*}\sigma''_R(p)]$$

- In the specification above, even though the receiver does not observe the composition of the undrawn urn, he assumes that its fraction of red balls stays at the equilibrium value of q^*.
- If the sender's strategy is (p, q^*), then with probability $\frac{p-\mu}{p-q^*}$, urn with fraction $\mu \geq q$ of red balls is drawn. In this case, the payoffs of the receiver and the sender are $q^*\sigma''_R(q^*) + (1-q^*)(1-\sigma''_R(q^*))$ and $\sigma''_R(q^*)$ respectively.
- With complementary probability $\frac{\mu-q^*}{p-q^*}$, urn with a fraction $p \geq \mu$ of red balls is drawn. In this case, the expected payoffs of the receiver and the sender are $p\sigma''_R(p) + (1-p)(1-\sigma''_R(p))$ and $\sigma''_R(p)$ respectively.

Model

- Similarly, if the drawn urn has a fraction $\mu \geq q$ red balls, the sender's perceived kindness is,

$$k_s(q, \sigma''_R; (p^*, q^*)) = [\tfrac{p^*-\mu}{p^*-q}\{q\sigma''_R(q) + (1-q)(1-\sigma''_R(q)\} + \tfrac{\mu-q}{p^*-q}\{p^*\sigma''_R(p^*) + (1-p^*)(1-\sigma''_R(p^*)\}] - [\tfrac{p^*-\mu}{p^*-q}\sigma''_R(q) + \tfrac{\mu-q}{p^*-q}\sigma''_R(p^*)]$$

Equilibrium

Reciprocity equilibrium corresponds to subgame perfect psychological Nash equilibrium. A pair of strategy $((p^*, q^*), \sigma_R(\cdot))$ constitutes a reciprocity equilibrium if and only if

- The sender's strategy maximizes utility given her belief $\sigma_R(\cdot)$,

$(p^*, q^*) \in arg\ max_{\{(p',q'):p' \geq \mu \geq q'\}} \frac{p'-\mu}{p'-q'} \sigma_R(q') + \frac{\mu - q'}{p'-q'} \sigma_R(p')$

- If an urn with a fraction $p \geq \mu$ of red balls is realized, the receiver maximizes her utility given beliefs, i.e., $\sigma_R(p) \in$

$arg\ max_{\sigma \in [0,1]} p\sigma + (1-p)(1-\sigma) + \lambda_R k_s(p, \sigma_R; (p^*, q^*))[\sigma - \sigma_R(p)]$

- If an urn with a fraction $\mu \geq q$ if red balls is realized, the receiver maximizes her utility given beliefs, i.e., $\sigma_R(q) \in$

$arg\ max_{\sigma \in [0,1]} q\sigma + (1-q)(1-\sigma) + \lambda_R k_s(q, \sigma_R; (p^*, q^*))[\sigma - \sigma_R(q)]$

Equilibrium

- **Proposition 1** : Let $\mu \leq \frac{1}{2}$. The reciprocity equilibrium (in which the sender offers two urns) is unique.
- The sender chooses (p^*,0) where
$p^* = \frac{1}{4}[(1-\lambda_R(1+\mu)) + \sqrt{(1-\lambda_R(1+\mu))^2 + 16\mu\lambda_R}]$.
- Upon observing an urn with a fraction p of red balls, the receiver's probability of guessing red is given by $\sigma_R(p) =$

$$\begin{cases} \max\{0, \frac{\frac{(2p-1)(p^*-p)}{\lambda_R} + (1-\mu)(p^*-p) - 2(\mu-p)(1-p^*)}{2(p^*-\mu)(1-p)}\}, & \text{if } p < \mu, \\ \min\{\frac{\frac{(2p-1)p}{\lambda_R} + (1-\mu)p}{2\mu(1-p)}, 1\}, & \text{if } p \geq \mu, \end{cases}$$

Equilibrium

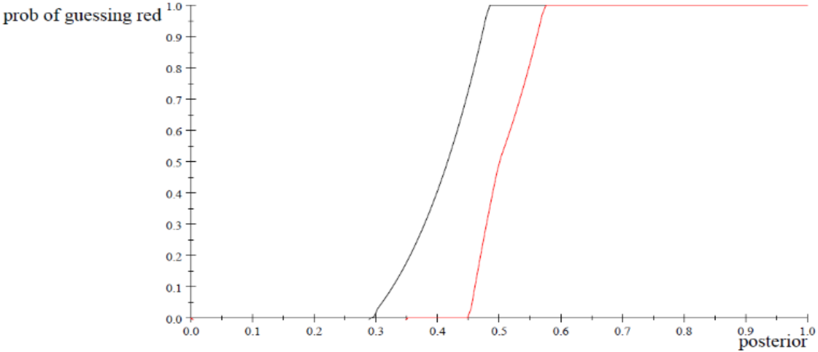

Figure 1: Receiver's equilibrium responses (Black: $\mu = 0.3$; *red* : $\mu = 0.5$)

- This illustrates the receiver's equilibrium strategy $\sigma_R(p)$ for the case $\lambda_R = 0.6$.

Equilibrium

- The receiver's probability of guessing red $\sigma_R(p)$ is not a step-function in the fraction p, as predicted in Lemma 1. If λ_R is positive, $\sigma_R(p)$ increases continuously over an interval of p.
- The receiver's responses and hence the sender's optimal information structure depends on the prior belief μ, which is (ex-post) payoff-irrelevant factor.
- As μ increases, the sender's expected payoff goes up with as it is more likely that the realized posterior is favorable for the state being red. Consequently, the receiver's perception of the sender's kindness goes down, making them less willing to be persuaded to choose R.

Hypotheses

- **Hypothesis 1** : The probability that the receiver guesses red increases continuously in the fraction of red balls in the drawn urn around the 50% mark.

- Proposition 1 also implies that an increase in μ would shift $\sigma_R(p)$ downwards and raise the p^* of the sender's equilibrium strategy.
- **Hypothesis 2** : An increase in the initial fraction of red balls weakly lowers the probability that the receiver guesses red for any urn composition.
- **Hypothesis 3** : When the initial number of red balls increases, the sender makes the fraction of red balls in the urns more extreme (i.e., closer to 0 and 1).

Hypotheses

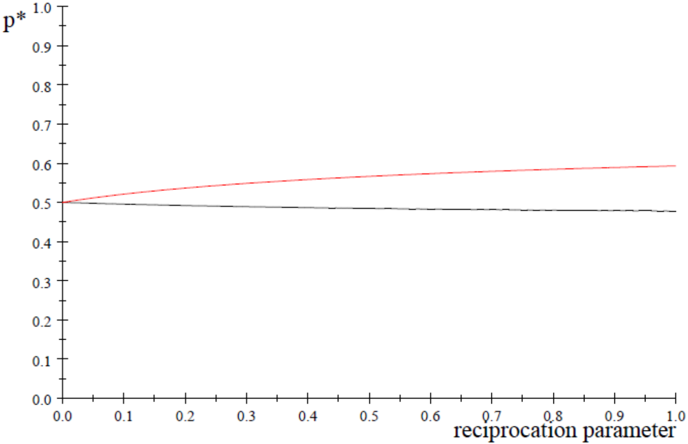

Figure 2: Sender's equilibrium choice of p^* against λ_R

(Black: $\mu = 0.3$; red: $\mu = 0.5$)

Hypotheses

- The underlying intuition is analogous to that of an ultimatum game. The receiver is willing to sacrifice monetary payoffs to punish the sender if she perceives the sender to be unkind.
- The punishment takes the form of guessing black (lowering his own expected payoff in order to make the sender getting a zero payoff).
- The standard Bayesian persuasion model without reciprocation incentives predicts, in sharp contrast to the hypotheses above, that the receiver's response does not depend on the posterior belief as long as it exceeds 50%.
- It also predicts that neither the receiver's strategy $\sigma_R(p)$ nor the sender's strategy $(p^*, 0)$ would be affected by the prior belief.

Hypotheses

- Last hypothesis concerns the expected payoff of the players.
- In a Bayesian persuasion model without reciprocity, an increase in the prior could strictly harm the receiver.
- The receiver gets the positive payoff for certain if the drawn urn has 0% red ball, and gets the positive payoff with probability 50% if the drawn urn has 50% red ball.

$$(1 - \frac{\mu}{0.5}) \times 1 + \frac{\mu}{0.5} \times 0.5 = 1 - \mu.$$

- **Corollary 2** : Suppose $\mu < \mu' \leq \frac{1}{2}$. The sender's expected monetary payoff is higher if the initial fraction of red balls is μ' than when it is μ. The same is true for the receiver's expected monetary payoff if the reciprocation parameter λ_R is sufficiently large.
- **Hypothesis 4** : When the initial number of red balls increases, both the expected monetary payoffs of both sender and receiver increase.

Experiment Procedures

- In each round, subjects are randomly matched into pairs. In each matched pair, one subject is assigned to the role of player A and the other the role of player B.
- Each subject's role remained fixed throughout the experiment and subjects' decisions were anonymous.
- The game has two stages.
  - Stage 1: Player A allocates all of 100 colored balls into 5 urns. One urn is randomly drawn. The total number of balls in each urn determines the probability that the urn is drawn.
  - Stage 2: Player B receives a message on the screen specifying the number of red balls and number of black balls in the urn. Then, a ball is randomly drawn from the urn and Player B is asked to guess whether the ball drawn is red or black.
- Repeat 10 independent rounds.

Experimental Procedures

Urn	1	2	3	4	5
Red					
Black					

Figure 3: Sender's choice

- Player A receives 40 HKD if player B guesses that the ball is red and receives 0 if player B guesses that the ball is black, regardless of the actual color of the ball drawn.
- Player B receives 40 HKD if her guess coincides with the ball drawn and receives 0 if her guess is incorrect.
- At the end of each round, subjects are informed about (i) the urn drawn, (ii) the guess made by player B, (iii) the color of the ball drawn, (iv) earning.

Experimental Procedures

- Ran two treatments with $\mu = 0.3$ and $\mu = 0.5$ to identify the effect of prior belief on players' behaviors.
- In total, 162 subjects participated in 8 sessions, 4 sessions for each treatment.
- The subjects were undergraduate students in a major university in Hong Kong, recruited through electronic recruitment system. The experiment was programmed with z-tree (Fischbacher, 2007).
- Before the beginning of the 10 decision-making rounds, subjects were given one practice round. Out of the 10 rounds, one round was randomly drawn for payment.
- Subjects received a show-up fee of 40 HKD, in addition to the earnings from the randomly-drawn round.

Receivers' Responses

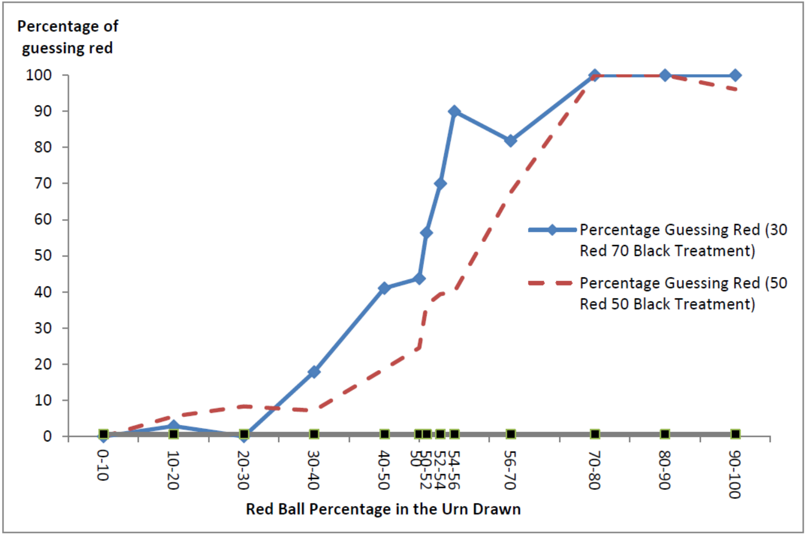

Figure 4: Percentage of guessing red against urn composition

Receivers' Responses

- The probability that player B guesses red is significantly positive when presented with urns with 40-50% of red balls, even though such a choice gives her a negative expected payoff (*p-value* = *0.00*).
- Although the probability of guessing red spikes at 50% mark, it falls short of 100% by a large margin (*p-value* = *0.00*). These probabilities are even lower in $\mu = 0.5$ treatment.
- These behaviors are not completely consistent with expected utility maximization, which would call for guessing red whenever the proportion of red balls exceeds 50%.
- **Result 1**: Player B does not guess red for sure even if the urn contains more than 50% of red ball. For urns with fractions of red balls between 50% and 70%, the probability that Player B guesses red is significantly lower than 1.

Receivers' Responses

	Dependent variable: Guessing Red	
	(1) Full Sample	(2) More than 50 Percentage of Red Ball in the Drawn Urn
Percentage of Red Ball	0.02***	0.02***
	(0.001)	(0.01)
50-red-50-black Treatment	-0.14***	-0.16***
	(0.03)	(0.17)
Number of Observations	810	349
Pseudo R-square	0.50	0.18

Figure 5: Determinants of Probability of Guessing Red

- Marginal effect coefficient estimates of Probit regression reported. 50-black-50-red treatment is a dummy variable which equals to 1 for the $\mu = 0.5$ treatment, zero otherwise.

Receivers' Responses

- Sharp contrast to the prediction of the standard Bayesian persuasion model that player B's response is a step function.
- Column 1 of Table 1 reports the marginal effect coefficient estimations of the probability of guessing red. The higher the percentage of red balls is, the more likely that player B will guess red.
- Column 2 of Table 1 reports the same regression for the subsample where the percentage of red ball in the drawn urn is higher than 50%. The probability of guessing red still increases with the percentage of red balls. Note that Bayesian persuasion predicts that the coefficient to be zero.
- **Result 2**: As the percentage of the red balls in the drawn urn increases from 30% to 70%, the probability of guessing red by player B increases continuously.
- Overall, these results support Hypothesis 1.

Receivers' Responses

- Next, compare player B's responses between two treatments.
- **Result 3**: An increase in the initial fraction of red balls weakly lowers the probability that the receiver guesses red for any urn composition, and strictly so when the fraction of red balls is in the range of 30% to 70%.
- This result supports Hypothesis 2. Facing an urn with 50% of red ball, player B in the $\mu = 0.3$ treatment may interpret that player A is relatively kind than the case of $\mu = 0.5$ treatment where the percentage of red ball is the same as the initial condition.

Receivers' Responses

- The subjects' responses in the questionnaire answered at the end of the experiment sessions shed light on the role of reciprocation concerns in their decisions.
- "When the number of the balls is similar or red balls are only 2-3 balls more, I choose black as my answer, as the probability if winning is approximately 1/2 only. Therefore, even if I am wrong, player A cannot get his pay."
- "If the ratio of red and black ball is 1:1. I would choose black ball so that player A can't get additional payoff."
- It is clear from the responses that player B's behavior was driven by reciprocation concerns, not by mistaken probabilistic beliefs.

Receivers' Behaviors

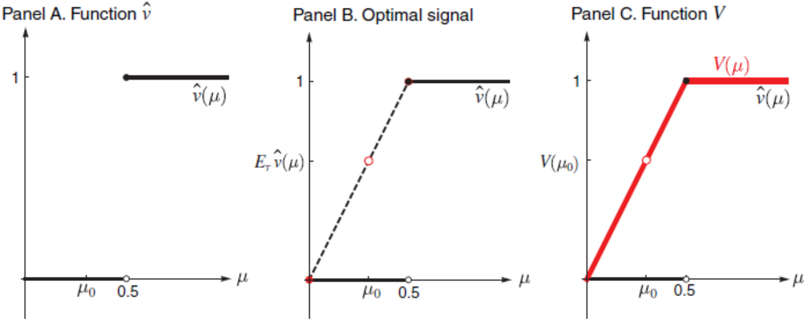

Figure 2. The Motivating Example

Receivers' Behaviors

- Given player B's responses differ from the prediction of standard Bayesian persuasion model, the empirically optimal ball allocation differs from the theoretically optimal urn.
- Apply the technique of Kamenica and Gentzkow (2011) to Figure 3 and look for the concave closure of the payoff function in posteriors. For $\mu = 0.3$ treatment, the empirically optimal ball allocation is (0.55, 0). For $\mu = 0.5$ treatment, the empirically optimal ball allocation is (0.75, 0).
- The corresponding expected payoff to Player A is $(0.3/0.55) \times 0.9 \times 40 = 19.64$ vs $(0.3/0.5) \times 0.6 \times 40 = 14.4$, $(0.5/0.75) \times 1 \times 40 = 26.67$ vs $(0.5/0.5) \times 0.25 \times 40 = 10$.
- **Result 4**: The empirically optimal ball allocation is different from the prediction of the standard Bayesian persuasion model. In particular, it involves providing an urn with a fraction of red balls significantly exceeding 50%.

Senders' Behaviors

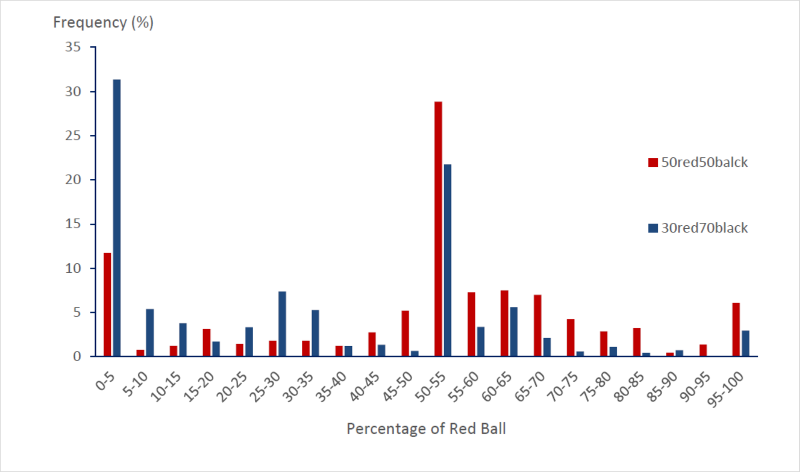

Figure 6: Distribution of Urns

Senders' Behaviors

- In either treatment, urns with 0% red balls and 50-55% red balls are the most frequently offered urns.
- Urns with more than 50% red balls are offered quite often. The proportions of urns with more than 55% of red balls are about 16.83% and 39.99% in $\mu = 0.3$ treatment and $\mu = 0.5$ treatment respectively. Both proportions contradict the prediction of the standard Bayesian persuasion model.
- More importantly, the proportion is significantly higher in $\mu = 0.5$ treatment than in $\mu = 0.3$ treatment (*p-value* = *0.00*).

Senders' Behaviors

- **Result 5**: Player As choose urns with more than 50% of red balls quite often. The empirical frequency of urns with more than 50% of red balls is much higher in $\mu = 0.5$ treatment than $\mu = 0.3$ treatment.
- This result supports Hypothesis 3. In particular, Player A provides urns with composition that are more favorable to player B in $\mu = 0.5$ treatment. This finding indicates that player As in our experiments respond to the difference in player B's behavior across the two treatments.

Senders' Behaviors

- As Player As make more "generous" offers in $\mu = 0.5$ treatment, Player Bs indeed benefit a lot from a higher initial fraction of red balls. The average monetary payoff of Player Bs is 25.8 HKD and 17.3 HKD in $\mu = 0.5$ treatment and $\mu = 0.3$ treatment respectively (*p-value = 0.00*).
- Similarly, Player A also benefit from an increase in the initial fraction of red balls. Their average monetary payoffs of Player As are 16.9 HKD and 12.1 HKD in $\mu = 0.5$ treatment and $\mu = 0.3$ treatment respectively (*p-value = 0.00*).
- **Result 6**: The average monetary payoff of both Player A and Player B are much higher in the 50-red50-black treatment than the 30-red-70-black treatment.
- This result supports the Hypothesis 4 that both players share the benefit of having a higher initial fraction of red balls.

Senders' Behaviors

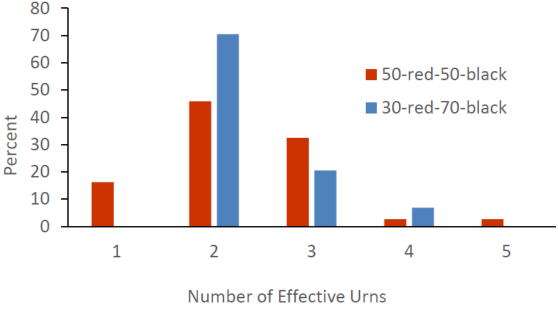

Figure 7: Frequency of Number of Effective Urns

- Both the standard Bayesian persuasion model and our reciprocity model predict that Player A uses, effectively, only two urns. The number of effective urns in our experiment is close to two most of the time.

Concluding Remarks

- This paper studies how reciprocation incentives affect people's behavior and economic outcome in the context of Bayesian persuasion, both theoretically and experimentally.
- The results are analogous to the finding in ultimatum game experiments.
- Receivers are persuaded to take the sender's preferred action only if the sender is willing to offer sufficiently informative signals. Choosing an information structure that is too opaque leaves too little rent to the receiver, who may punish the sender for being too "stingy" by refusing to be persuaded.
- Furthermore, when the prior belief of the state is more favorable, the sender is deemed to have a higher (ex-ante) expected payoff, so the receiver would be more demanding in the sender's information revelation.

Concluding Remarks

- The concept of reciprocity per se is not new in economics. However, its role in the context of persuasion has not been explored before.
- The effect of reciprocity concern in other communication settings, such as cheap talk and disclosure of verifiable information, can be considered. It is interesting to investigate whether reciprocity concern improves or worsens the quality of information transmission in these settings.
- The simplicity of the experiment design makes it possible to investigate persuasion behaviors in settings with multiple senders or multiple receivers.
- As receivers are likely to suffer bias in information processing, it is interesting to study how the optimal dynamic persuasion technique may exploit these biases.