Strategic Ignorance and Information Design

Ina Taneva
U of Edinburgh and CEPR

Thomas Wiseman
U of Texas at Austin

ESEM

Bocconi University

August 23, 2022

MOTIVATION

The **standard setting** of information design:

MOTIVATION

The **standard setting** of information design:

A designer **commits** to disclosing information about an uncertain payoff relevant state to a set of interacting agents.

Motivation

The **standard setting** of information design:

- A designer **commits** to disclosing information about an uncertain payoff relevant state to a set of interacting agents.
- Information structure: joint distributions of (private) signals conditional on each possible realization of the state.

MOTIVATION

The **standard setting** of information design:

- A designer **commits** to disclosing information about an uncertain payoff relevant state to a set of interacting agents.
- Information structure: joint distributions of (private) signals conditional on each possible realization of the state.
- Information incentivizes agents to take actions that benefit the designer in expectation.

MOTIVATION

The **standard setting** of information design:

- A designer **commits** to disclosing information about an uncertain payoff relevant state to a set of interacting agents.
- Information structure: joint distributions of (private) signals conditional on each possible realization of the state.
- Information incentivizes agents to take actions that benefit the designer in expectation.
- **Implicit assumption**: players agree to get informed according to the chosen information structure.

MOTIVATION

The **standard setting** of information design:

- A designer **commits** to disclosing information about an uncertain payoff relevant state to a set of interacting agents.
- Information structure: joint distributions of (private) signals conditional on each possible realization of the state.
- Information incentivizes agents to take actions that benefit the designer in expectation.
- **Implicit assumption**: players agree to get informed according to the chosen information structure.

Our setting: Agents can publicly commit not to observe their private signals.

Motivation

Unlike in standard single-receiver environments, **players** may benefit from refusing to get informed, if they can credibly signal this choice to the other players.

MOTIVATION

Unlike in standard single-receiver environments, **players** may benefit from refusing to get informed, if they can credibly signal this choice to the other players.

- **Cost:** reduced ability to tailor own action to the state (or to opponents' actions)

Motivation

Unlike in standard single-receiver environments, **players** may benefit from refusing to get informed, if they can credibly signal this choice to the other players.

- **Cost:** reduced ability to tailor own action to the state (or to opponents' actions)
- **Gain:** the opponents' changed behaviour in response to own ignorance

MOTIVATION

Unlike in standard single-receiver environments, **players** may benefit from refusing to get informed, if they can credibly signal this choice to the other players.

- **Cost:** reduced ability to tailor own action to the state (or to opponents' actions)
- **Gain:** the opponents' changed behaviour in response to own ignorance

Many economic settings where committing to ignorance is valuable.

Motivation

Unlike in standard single-receiver environments, **players** may benefit from refusing to get informed, if they can credibly signal this choice to the other players.

- **Cost:** reduced ability to tailor own action to the state (or to opponents' actions)
- **Gain:** the opponents' changed behaviour in response to own ignorance

Many economic settings where committing to ignorance is valuable.

Unreasonable to assume players can be always induced to play under the designer chosen information structure.

Motivation

For **robustness** considerations it is important to model players' incentives to **accept** information.

Motivation

For **robustness** considerations it is important to model players' incentives to **accept** information.

$\Rightarrow$ Different from the issue of **equilibrium selection**.

Motivation

For **robustness** considerations it is important to model players' incentives to **accept** information.

$\Rightarrow$ Different from the issue of **equilibrium selection**.

The requirement that agents be incentivized to view their signals imposes new constraints on the designer's problem: **"Look constraints"**.

Motivation

For **robustness** considerations it is important to model players' incentives to **accept** information.

$\Rightarrow$ Different from the issue of **equilibrium selection**.

The requirement that agents be incentivized to view their signals imposes new constraints on the designer's problem: **"Look constraints"**.

$\Rightarrow$ Similar to **participation constraints** in mechanism design.

MOTIVATION

For **robustness** considerations it is important to model players' incentives to **accept** information.

$\Rightarrow$ Different from the issue of **equilibrium selection**.

The requirement that agents be incentivized to view their signals imposes new constraints on the designer's problem: **"Look constraints"**.

$\Rightarrow$ Similar to **participation constraints** in mechanism design.

WE ASK:

Motivation

For **robustness** considerations it is important to model players' incentives to **accept** information.

$\Rightarrow$ Different from the issue of **equilibrium selection**.

The requirement that agents be incentivized to view their signals imposes new constraints on the designer's problem: **"Look constraints"**.

$\Rightarrow$ Similar to **participation constraints** in mechanism design.

We Ask:

How is the set of implementable outcomes impacted by the presence of strategic ignorance?

MOTIVATION

For **robustness** considerations it is important to model players' incentives to **accept** information.

$\Rightarrow$ Different from the issue of **equilibrium selection**.

The requirement that agents be incentivized to view their signals imposes new constraints on the designer's problem: **"Look constraints"**.

$\Rightarrow$ Similar to **participation constraints** in mechanism design.

WE ASK:

How is the set of implementable outcomes impacted by the presence of strategic ignorance?

How does strategic ignorance limit the scope of information design?

MOTIVATING EXAMPLE

Build up from complete information case:

	H	L
Y	$3, 0$	$1, 1$
X	$2, 2$	$0, 0$

Y is strictly dominant.

L is the unique best response to Y.

Outcome: (Y, L), payoff $(1, 1)$.

MOTIVATING EXAMPLE

Add random state $\omega \in \{a, b\}$, equally likely.

$\omega = a$:

	H	L
Y_a	$3, 0$	$1, 1$
X	$2, 2$	$0, 0$
Y_b	$0, 0$	$-2, 1$

$\omega = b$:

	H	L
Y_a	$0, 0$	$-2, 1$
X	$2, 2$	$0, 0$
Y_b	$3, 0$	$1, 1$

If it is common knowledge that player 1 knows ω, then outcome us (Y_ω, L), payoff $(1, 1)$.

MOTIVATING EXAMPLE

	H	L
Y_a	$3, 0$	$1, 1$
X	$2, 2$	$0, 0$
Y_b	$0, 0$	$-2, 1$

$\omega = a$

	H	L
Y_a	$0, 0$	$-2, 1$
X	$2, 2$	$0, 0$
Y_b	$3, 0$	$1, 1$

$\omega = b$

	H	L
Y_a	$1.5, 0$	$-0.5, 1$
X	$2, 2$	$0, 0$
Y_b	$1.5, 0$	$-0.5, 1$

$\mu(a) = \mu(b) = 1/2$

At the prior, X is strictly dominant.

If it is common knowledge that player 1 does not know ω, then outcome is (X, H), payoff $(2, 2)$.

Player 1 gains from ignorance.

If designer sends message to reveal ω, player 1 would choose to publicly ignore it.

PREVIEW OF RESULTS

Characterization of the implementable outcome distributions under strategic ignorance in **general finite environments**.

Appropriate definition of **correlated equilibrium** for our environment.

Show it is without loss to consider **direct contingent** information structures:

- Of higher dimension than in standard information design: single "on-path" action recommendations no longer sufficient.
- Individual messages are vectors of pure action recommendations, one for each possible choice of the other players in the pre-play stage.

Showcase the impact of strategic ignorance in **traditional economic applications**: investment game and currency attack.

Preview of Results

Necessity of strategic ignorance: restricting attention to equilibria where everyone chooses to look at signals is with loss.

Robustness to strategic ignorance undoes two standard qualitative results:

- **Multistage communication** may be worse for the designer than providing all information at once.
- Allowing for **communication between players** may be beneficial to the designer.

Harm of strategic ignorance: relevance for environments without a designer, where players try to coordinate on the best information structure ex ante.

Relation to Literature

Strategic ignorance:

- Survey: **Goldman, Hagmann, and Loewenstein (2017)**.
- private values in second-price auctions: **McAdams (2012)**.
- buyer valuations in bilateral trade: **Roesler and Szentes (2017)**.

Sequential and multi-stage information design:

- **Doval and Ely (2020), de Oliveira and Lamba (2019), Makris and Renou (2021)**.
- main differences: designer provides information once and extensive form fixed

Most closely related: **Arcuri (2021)**.

- "Hear-no-evil" BCE motivated by the same question
- weaker form of robustness to strategic ignorance
- an outcome could be "hear-no-evil" BCE even if a player prefers his worst BNE outcome after choosing to remain uninformed

Model

Environment

States: $\omega \in \Omega$, finite

- commonly known prior $\mu \in \Delta(\Omega)$

Agents: $i \in \mathcal{I}$ with $|\mathcal{I}| = N > 1$

- A_i finite
- $u_i : A \times \Omega \to \mathbb{R}$

Basic game: $G = \langle (A_i, u_i)_{i \in \mathcal{I}}, \mu \rangle$

Designer: $u^D : A \times \Omega \to \mathbb{R}$

Information structure: (T, P)

- T_i finite, $T \equiv T_1 \times \cdots \times T_N$
- $P : \Omega \to \Delta(T)$

Informational Environment: $(T_{\mathcal{L}}, P_{\mathcal{L}})$

- common knowledge that
- all $i \in \mathcal{L}$ are informed according to (T, P)
- all $i \in \mathcal{G} := \mathcal{I} \setminus \mathcal{L}$ are uninformed

TIMING

1. Designer publicly commits to (T, P)
2. State $\omega \in \Omega$ is realized according to μ
3. Vector of signals drawn according to $P(\cdot|\omega)$
4. **Look-Ignore Stage**: $s_i \in S_i \equiv \{\ell, g\}$; $s \in S \equiv \{\ell, g\}^N$
   - public and simultaneous choices of whether to
   - *Look* $(s_i = \ell)$ and learn the signal realization t_i, or to
   - *Ignore* $(s_i = g)$ and remain uninformed
5. **Action Stage**: Bayesian game $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$
   - each i chooses a_i
6. Payoffs are realized

Some Definitions

Dynamic Game $G^*(T, P)$: given (T, P), base game G augmented by the Look-Ignore and action stages.

Strategy in G^*: $\left(\gamma, (\tilde{\beta}^s)_s\right)$ with

- Look-Ignore strategy: $\gamma_i \in \Delta(\{\ell, g\})$
- post-Look strategy: $\tilde{\beta}_i^s : T_i \to \Delta(A_i)$ if $i \in \mathcal{L}(s)$
- post-Ignore strategy: $\tilde{\beta}_i^s \in \Delta(A_i)$ if $i \in \mathcal{G}(s)$

SOLUTION CONCEPT

PBE*=PBE with no signaling what you don't know refinement

Strategy profile $\left(\gamma, (\tilde{\beta}^s)_s\right)$ is a **PBE*** of $G^*(T, P)$ if:

1. for each $s \in S$, $\tilde{\beta}^s$ is a Bayes Nash Equilibrium of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$, and BNE
2. for each $i \in \mathcal{I}$ and $s_i \in \{\ell, g\}$ with $\gamma(s_i) > 0$,

$$\sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^{s_i, s_{-i}})(a, \omega) u_i(a_i, a_{-i}, \omega)$$
$$\geq \sum_{s_{-i}, a, \omega} \prod_{j \neq i} \gamma_j(s_j) v(\tilde{\beta}^{s'_i, s_{-i}})(a, \omega) u_i(a_i, a_{-i}, \omega) \quad \text{for all } s'_i \in \{\ell, g\}.$$

PBE* **outcome**: $v(a, \omega) := \sum_{s \in S} \prod_{i \in \mathcal{I}} \gamma_i(s_i) v(\tilde{\beta}^s)(a, \omega)$ for all $a \in A, \omega \in \Omega$.

Set of PBE* outcomes: $PBE^*(G^*(T, P))$.

Designer's Problem

Designer chooses (T, P) to solve

$$\max_{(T,P)} \sum_{a,\omega} u^D(a,\omega)v(a,\omega) \quad \text{s.t.} \quad v \in PBE^* \left(G^* \left(T, P\right)\right).$$

Essentially, the designer maximizes over the set $\cup_{(T,P)} PBE^* \left(G^* \left(T, P\right)\right)$ by choosing the information structure (T, P).

CHARACTERIZATION

Definition of Correlated Equilibrium

The designer can provide correlation of strategies with the state and with the strategies of other players **only** at the **action stage** and **only** for players who choose **"Look"**.

- Look-Ignore-stage choices must be independent of each other and of ω.
- Action-stage choices for $s_i = g$ must be independent of ω and the actions of $-i$, but can depend on s_{-i}.

Robust Correlated Equilibrium

The object of interest is an element

$$(\gamma, \beta^g, \tilde{v}) \in \times_i \left(\Delta\{\ell, g\} \times \left(\times_{s_{-i}} \Delta A_i\right)\right) \times \Delta(\mathscr{A} \times \Omega), \quad \text{where}$$

$\gamma \in \times_i (\Delta\{\ell, g\})$: **(public) Look-Ignore recommendations**

$\beta^g \in \times_i \left(\times_{s_{-i}} \Delta A_i\right)$: **post-Ignore recommendations**

$\tilde{v} \in \Delta(\mathscr{A} \times \Omega)$: **post-Look recommendations**

$\mathscr{A}_i \equiv A_i^{|S_{-i}|}$: set of agent i's (pure) mappings from S_{-i} to A_i, $\mathscr{A} \equiv \mathscr{A}_1 \times \cdots \times \mathscr{A}_N$

$m_i \in \mathscr{A}_i$: **message**

$m_i(s_{-i}) \in A_i$: **action recommendation** after combination s_{-i} of other agents' Look-Ignore choices

Equivalence Result

RCE outcome: $v(\gamma, \beta^g, \tilde{v}) \in \Delta(A \times \Omega)$

RCE outcome

Set of RCE outcomes: $RCE\,(G^*)$.

Theorem

$\cup_{(T,P)} PBE\,(G^*\,(T,P)) = RCE\,(G^*)$.

Direct Contingent Recommendations

The designer maximizes over the set of all information structures (T, P).

$\Rightarrow$ a very large space

(T, P) is an information structure with **direct contingent recommendations** if $T_i = \mathscr{A}_i$ for each agent i.

An outcome $v \in \Delta(A \times \Omega)$ is **implementable with direct contingent recommendations** if there exists a conditional message distribution $P : \Omega \rightarrow \Delta(\mathscr{A})$ such that $v \in PBE^* (G^* (\mathscr{A}, P))$.

Theorem

An outcome $v \in \Delta(A \times \Omega)$ is a PBE^* outcome if and ***only if*** *it is implementable with direct contingent recommendations.*

Next Steps

Allow agents general garbling of designer messages, rather than all or nothing.

Implications for optimal monitoring in repeated games.

Different possible extensive forms – sequential moves at the Look-Ignore stage.

CONCLUSION

Ability of agents to publicly refuse information has important effects for information design in strategic settings:

- Significantly alters implementable outcomes and optimal information structures in many settings.
- Undoes standard qualitative results from the information design literature.

Findings also relevant for settings where agents seek to coordinate on what pre-play information to gather.

- Agreement that maximizes expected payoffs ex ante may not be sustainable (harm of ignorance).

Thank You!

INVESTMENT GAME

Basic Game

	A	B
A	c, c	$d, 0$
B	$0, d$	$0,0$

$\omega = a$

	A	B
A	$0,0$	$0, d$
B	$d, 0$	c, c

$\omega = b$

$c > d > 0$

Two symmetric firms seeking to coordinate on the right project.

Profitability increases with total investment.

Unknown state determines which project has potential: $\mu(a) = \mu(b) = \frac{1}{2}$.

DESIGNER

The designer wants the firms to fail: $u^D(A, A, b) = u^D(B, B, a) = 1$ and 0 otherwise.

Without loss restricting attention to **symmetric** BCE outcome distributions.

	A	B
A	r	$q - r$
B	$q - r$	$1 - 2q + r$

$\omega = a$

	A	B
A	$1 - 2q + r$	$q - r$
B	$q - r$	r

$\omega = b$

$\mathbb{E}(u^D) = 1 - 2q + r$

Baseline: No Strategic Ignorance

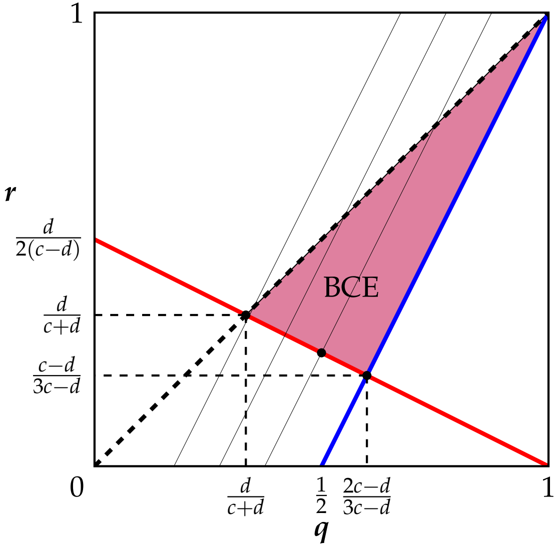

Baseline: No Strategic Ignorance

Optimal direct information structure:

	A	B
A	$\frac{d}{c+d}$	0
B	0	$\frac{c}{c+d}$

$\omega = a$

	A	B
A	$\frac{c}{c+d}$	0
B	0	$\frac{d}{c+d}$

$\omega = b$

Designer sends a **public** signal.

Exploits coordination incentive: right project recommended with prob $\frac{d}{c+d} < \frac{1}{2}$.

Each firm is just willing to obey the recommendation given the other one will.

$\mathbb{E}(u^D) = 1 - 2q + r = \frac{c}{c+d} > \frac{1}{2}$ $\qquad$ $\mathbb{E}(u_i) = cr + d(q - r) = \frac{cd}{c+d}$

WITH STRATEGIC IGNORANCE

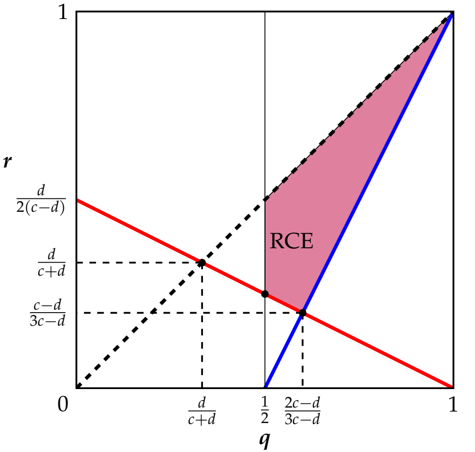

After a deviation to Ignore:

Case 1. If $q \geq \frac{1}{2}$, opponent continues to follow action recommendation.

$\Rightarrow$ **No gain** from ignoring.

WITH STRATEGIC IGNORANCE

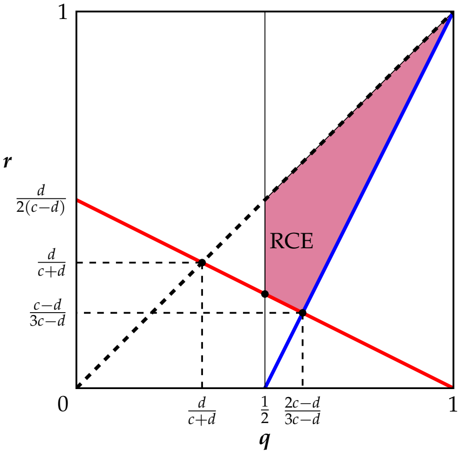

After a deviation to Ignore:

Case 2: If $q < \frac{1}{2}$, opponent plays the opposite of action recommendation.
$\Rightarrow$ **Potential gain** from ignoring.

Payoff from BNE post-Ignore: $\frac{c}{2} - \frac{c-d}{2}q$

Payoff from BNE post-Look: $cr + d(q - r)$

Gain from ignorance: $r < \frac{c}{2(c-d)} - \frac{c+d}{2(c-d)}q$

WITH STRATEGIC IGNORANCE

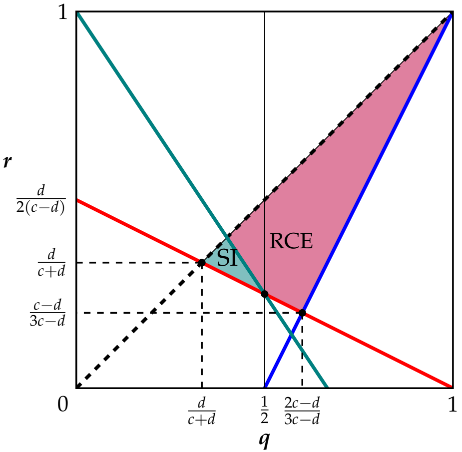

After a deviation to Ignore:

Case 2. If $q < \frac{1}{2}$, opponent plays the opposite of action recommendation.

$\Rightarrow$ **Potential gain** from ignoring.

Payoff from BNE post-Ignore: $\frac{c}{2} - \frac{c-d}{2}q$

Payoff from BNE post-Look: $cr + d(q - r)$

Gain from ignorance: $r < \frac{c}{2(c-d)} - \frac{c+d}{2(c-d)}q$

WITH STRATEGIC IGNORANCE

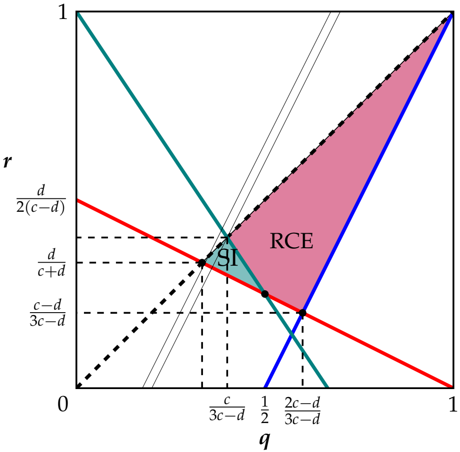

At the initial information structure, **no equilibrium** where both choose Look.

The designer needs to reduce the probability of recommending the wrong project: from $\frac{c}{c+d}$ to $\frac{2c-d}{3c-d}$ (still $> \frac{1}{2}$).

This probability is equal to the designer's payoff, but it also increases the payoff from the worst post-Ignore BNE – **fundamental trade-off**.

Firms are better off than before:
$\mathbb{E}(u_i) = \frac{c^2}{3c-d} > \frac{cd}{c+d}$

WITH STRATEGIC IGNORANCE

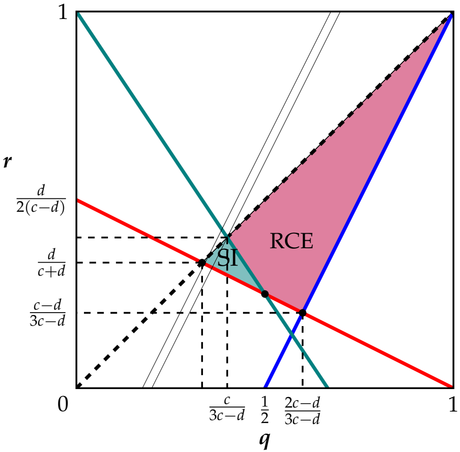

Note: Unlike BCE, The set of RCE is **non-monotone** in the amount of exogenous information that players have.

Appendix

FURTHER RESULTS

The Necessity of Ignorance

Theorem

A PBE outcome v may be implementable only if $\gamma_i(g) > 0$ for some $i \in \mathcal{I}$.*

2-player example where the designer does **strictly** better in a PBE* with $\gamma_2(g) > 0$.

1's Look-constraint needs the low punishment that exists only if it is **common knowledge** that 2 is uninformed.

For the designer's objective it is important that 2 tailors his action to the state, so has to be informed.

The optimal solution is a **compromise** where 2 mixes between Look and Ignore.

Note: Pure Look equilibria not enough even of we **enlarge the message space** to indicate whether the others' messages are informative or uninformative.

Communication between Players May Be Better

Unlike in standard information design, designer can be **strictly better off** if players can communicate.

Suppose 2 is willing to punish 1 effectively only when uninformed.

2 must be informed on path to play the designer's state-contingent desired action.

Suppose players can communicate and 1 has **incentives to share information**.

The designer can give the relevant information to 1, who now has an incentive to Look and will subsequently share it with 2.

In situations where 1 should not know the state on path, coded messages that only when combined reveal the state can be used.

MULTISTAGE COMMUNICATION MAY BE WORSE

In general, the designer is hurt by having to provide recommendations for all possible combinations of Look-Ignore choices at once.

The information about the state that those recommendations imply may interfere with the standard obedience constraint (without strategic ignorance).

Suggestive that **recommendations conditional on Look-Ignore choices**, so that only relevant part of the message is observed, would be (weakly) beneficial.

The extra information needed for punishing a deviator does not disturb on-path obedience constraints.

This turns out to be **wrong** – the designer can be **strictly worse off** in this case.

Multistage Communication May Be Worse

The designer starts out by giving on-path Look recommendations only.

If an opponent deviates to Ignore, the other player gets the recommendation of what to play in that instance **only after** the deviation has occurred.

However, the **Look-constraint** needs to be satisfied **every time** new recommendations are sent.

By giving both recommendations in one message at the start, the designer can satisfy the second Look-constraint for free.

Easy to construct examples where the designer does benefit from multistage communication.

THE HARM OF IGNORANCE

Not surprisingly, option to Ignore ends up being harmful to agents in some cases.

We provide such an example which has the flavour of a **prisoner's dilemma**.

A perfectly informative information structure maximizes the players' payoffs.

However, **Ignore is strictly dominant** after plugging in the unique BNE after each Look-Ignore choice as continuation payoffs:

	ℓ	g
ℓ	$1.11, 1.11$	$0.14, 1.12$
g	$1.12, 0.14$	$0.15, 0.15$

$\Rightarrow$ A designer who wants to maximize the total expected payoff must provide less than perfect information.

$\Rightarrow$ Players worse off than in the case of automatically observed messages.

Bayes Nash Equilibrium

Given (T, P) and $s \in S$, $\tilde{\beta}^s$ is a **Bayes Nash Equilibrium** of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$ if:

1. for each $i \in \mathcal{L}(s)$, $t_i \in T_i$, and $a_i \in A_i$ with $\tilde{\beta}^s_i(a_i|t_i) > 0$, we have

$$\sum_{a_{-i}, t_{\mathcal{L}(s)\setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_i, t_{\mathcal{L}(s)\setminus i}|\omega) \left(\prod_{j \in \mathcal{L}(s)\setminus i} \tilde{\beta}^s_j(a_j|t_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}^s_k(a_k) \right) u_i(a_i, a_{-i}, \omega)$$
$$\geq \sum_{a_{-i}, t_{\mathcal{L}(s)\setminus i}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_i, t_{\mathcal{L}(s)\setminus i}|\omega) \left(\prod_{j \in \mathcal{L}(s)\setminus i} \tilde{\beta}^s_j(a_j|t_j) \prod_{k \in \mathcal{G}(s)} \tilde{\beta}^s_k(a_k) \right) u_i(a'_i, a_{-i}, \omega) \text{ for all } a'_i \in A_i,$$

and

back

2. for each $i \in \mathcal{G}(s)$ and $a_i \in A_i$ with $\tilde{\beta}^s_i(a_i) > 0$, we have

$$\sum_{a_{-i}, t_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}^s_j(a_j|t_j) \prod_{k \in \mathcal{G}(s)\setminus i} \tilde{\beta}^s_k(a_k) \right) u_i(a_i, a_{-i}, \omega)$$
$$\geq \sum_{a_{-i}, t_{\mathcal{L}(s)}, \omega} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}^s_j(a_j|t_j) \prod_{k \in \mathcal{G}(s)\setminus i} \tilde{\beta}^s_k(a_k) \right) u_i(a'_i, a_{-i}, \omega) \text{ for all } a'_i \in A_i.$$

BAYES NASH EQUILIBRIUM

Then $v(\tilde{\beta}^s) \in \Delta(A \times \Omega)$ defined as

$$v(\tilde{\beta}^s)(a,\omega) := \sum_{t_{\mathcal{L}(s)}} \mu(\omega) P_{\mathcal{L}(s)}(t_{\mathcal{L}(s)}|\omega) \left(\prod_{j \in \mathcal{L}(s)} \tilde{\beta}^s_j(a_j|t_j) \prod_{i \in \mathcal{G}(s)} \tilde{\beta}^s_i(a_i) \right)$$

for all $a \in A$ and $\omega \in \Omega$ is a BNE outcome of $G(T_{\mathcal{L}(s)}, P_{\mathcal{L}(s)})$.

back

DEFINITION OF CORRELATED EQUILIBRIUM

$(\gamma, \beta^g, \tilde{v})$ is a **Robust Correlated Equilibrium** of G^* if

1. **(Consistency with the prior)** $\tilde{v}(\mathscr{A} \times \{\omega\}) = \mu(\omega)$ for all $\omega \in \Omega$;
2. **(Obedience for agent *i* recommended Look)** for every $s \in S$, $i \in \mathcal{L}(s)$, $m_i \in \mathscr{A}_i$, and $a_i' \in A_i$, agent i (weakly) prefers following $m_i(s_{-i})$ than deviating to a_i'. 2
3. **(Obedience for agent *i* recommended Ignore)** for every $s \in S$, $i \in \mathcal{G}(s)$, and $a_i, a_i' \in A_i$ such that $\beta_i^g(a_i|s_{-i}) > 0$, agent i (weakly) prefers following a_i than deviating to a_i'. 3
4. **(Obedience for agent *i* at the Look-Ignore stage)** for every $i \in \mathcal{I}$, s_i such that $\gamma_i(s_i) > 0$, and $s_i' \in S_i$, agent i (weakly) prefers following s_i than deviating to s_i'. 4

RCE CONDITION 2

Obedience for agent *i* who chooses Look: For every $s \in S$, $i \in \mathcal{L}(s)$, $m_i \in \mathscr{A}_i$, and $a_i' \in A_i$

$$\sum_{m_{\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega} v(m_i, m_{\mathcal{L}(s)\setminus i}, \omega) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k|s_{-k}) u_i(m_i(s_{-i}), (m_j(s_{-j}))_{j \in \mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega)$$
$$\geq \sum_{m_{\mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega} v(m_i, m_{\mathcal{L}(s)\setminus i}, \omega) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k|s_{-k}) u_i(a_i', (m_j(s_{-j}))_{j \in \mathcal{L}(s)\setminus i}, a_{\mathcal{G}(s)}, \omega).$$

back

RCE CONDITION 3

Obedience for agent i who chooses Ignore: For every $s \in S$, $i \in \mathcal{G}(s)$, and $a_i, a_i' \in A_i$ such that $\beta_i^g(a_i|s_{-i}) > 0$

$$\sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega} v(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s)\setminus i} \beta_k^g(a_k|s_{-k}) u_i(a_i, (m_j(s_{-j}))_{j \in \mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega)$$
$$\geq \sum_{m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega} v(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s)\setminus i} \beta_k^g(a_k|s_{-k}) u_i(a_i', (m_j(s_{-j}))_{j \in \mathcal{L}(s)}, a_{\mathcal{G}(s)\setminus i}, \omega).$$

back

RCE CONDITION 4

Obedience for agent *i* at the Look-Ignore stage: For every $i \in \mathcal{I}$, s_i such that $\gamma_i(s_i) > 0$, and $s_i' \in S_i$

$$\sum_{s_{-i}, m_{\mathcal{L}(s)}, a_{\mathcal{G}(s)}, \omega} \prod_{j \neq i} \gamma_j(s_j) v(m_{\mathcal{L}(s)}, \omega) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k | s_{-k}) u_i((m_j(s_{-j}))_{j \in \mathcal{L}(s)}, a_{\mathcal{G}(s)}, \omega)$$
$$\geq \sum_{s_{-i}, m_{\mathcal{L}(s')}, a_{\mathcal{G}(s')}, \omega} \prod_{j \neq i} \gamma_j(s_j) v(m_{\mathcal{L}(s')}, \omega) \prod_{k \in \mathcal{G}(s')} \beta_k^g(a_k | s'_{-k}) u_i((m_j(s'_{-j}))_{j \in \mathcal{L}(s')}, a_{\mathcal{G}(s')}, \omega)$$

where $s \equiv (s_i, s_{-i})$ and $s' \equiv \left(s_i', s_{-i}\right)$.

back

RCE OUTCOME

Given a RBC $(\gamma, \beta^g, \tilde{v})$, the associated outcome distribution $v \in \Delta(A \times \Omega)$ is defined as

$$v(\gamma, \beta^g, \tilde{v})(a, \omega) := \sum_{s \in S} \prod_{i \in \mathcal{I}} \gamma_i(s_i) \left(\sum_{m_{\mathcal{L}(s)} : (m_j(s_{-j}))_{j \in \mathcal{L}(s)} = a_{\mathcal{L}(s)}} \tilde{v}\left(m_{\mathcal{L}(s)}, \omega\right) \right) \prod_{k \in \mathcal{G}(s)} \beta_k^g(a_k | s_{-k})$$

for all $a \in A$ and $\omega \in \Omega$.

back