Information, Higher-Order Reasoning and Contingent Thinking

Ina Taneva
University of Edinburgh
ina.taneva@ed.ac.uk

Brian Rogers*
WUSTL
brogers@wustl.edu

June 8, 2026

Abstract

We study the role of information complexity on the higher-order rationality of agents in Bayesian games of incomplete information. In a class of dominance solvable Bayesian *chain* games, we find evidence that information environments that rely on more private information, and are therefore more complex in terms of contingent thinking requirements, are associated with lower orders of rationality. The results have implications for the optimal design of information structures, by highlighting a potential source of discrepancy between observed behavior and theoretical predictions.

(*JEL codes*: C72, C92, D80.)

Keywords: Complexity, Contingent Thinking, Incomplete Information, Rationality, Higher-Order Rationality.

*We thank Chiara Aina, Larbi Alaoui, Pierpaolo Battigalli, Andreas Blume, Inga Deimen, Steven Dieterle, Martin Dufwenberg, Françoise Forges, Amanda Friedenberg, Ed Hopkins, Terri Kneeland, Tatiana Kornienko, Ryan Oprea, Antonio Penta, Andreas Steinhauer, Gabriel Ziegler, and Peio Zuazo-Garin for insightful discussions. We also thank several seminar and conference audiences for helpful comments, questions, and suggestions.

1 Introduction

Departures from equilibrium play have been well documented in the experimental literature (see, e.g., Crawford and Iriberri (2007), Georganas (2011)). Understanding the forces and mechanisms underpinning such departures is important from both theoretical and applied perspectives. Indeed, in economics, it is crucial to simultaneously refine our theoretical models to capture the behavioral traits of economic agents, and to gain insights into how we can best design mechanisms, markets, algorithms, and, more broadly, institutions that successfully implement our intended outcomes.

One well-established way to explain departures from equilibrium play is by examining the strategic sophistication of agents as revealed by their strategic choice data. Strategic sophistication is estimated as the extent to which subjects' action choices are in line with the application of iterative reasoning, which relies on rationality and higher-order rationality (or Level-k reasoning). The experimental literature has provided substantial evidence that agents exhibit heterogeneous but limited degrees of strategic sophistication in a variety of different environments (see Beard and Beil (1994), Schotter, Weigelt, and Wilson (1994); Nagel (1995); Costa-Gomes, Crawford, and Broseta (2001); Van Huyck, Wildenthal, and Battalio (2002); Camerer, Ho, and Chong (2004); Costa-Gomes and Weizsäcker (2008); Rey-Biel (2009); Costa-Gomes, Crawford, and Iriberri (2013); Burchardi and Penczynski (2014); Kneeland (2015), among many others). With some exceptions, the focus has been on understanding limited depths of reasoning in games of complete information. Importantly, the documented heterogeneity emerges in unpredictable ways across different strategic environments (Georganas, Healy, and Weber (2015)).

More recently, Friedenberg and Kneeland (2024) have proposed distinguishing between reasoning about rationality and more general forms of strategic reasoning. They provide a novel identification strategy that allows for separate identification of the two bounds — rationality bound and strategic bound — in permuted ring games. Healy (2024) is also concerned with departures from higher-order rationality, but rather than drawing inferences from strategic choice data, they elicit preferences, beliefs about strategies, and beliefs about beliefs in a variety of classic games. Liu and Ziegler (2025) provide a unified theoretical framework that reconciles the structural perspective of Level-k and Cognitive Hierarchy models on the one hand, and the epistemic perspective of reasoning through beliefs and higher-order beliefs on the other hand. In doing so, they provide epistemic foundations for

the Level-k and Cognitive Hierarchy models and are able to distinguish between cognitive limits and belief restrictions. Relatedly, Brandenburger, Friedenberg, and Kneeland (2025) examine to what extent the Level-k analysis captures bounded reasoning in games and provide epistemic assumptions on type structures that guarantee the equivalence of round k rationalizability and the categorization of level k.

It is natural to postulate that the extent to which agents engage in iterative reasoning depends on the complexity of the strategic environment in which they interact. There are, of course, many different ways in which a strategic environment can be complex. Complexity of single-person decision problems is better studied (see, e.g., Oprea (2020), Puri (2025)), while that of strategic interactions is gaining momentum (Cerigioni et al. (2024), Guan and Oprea (2024), Li (2025)). In this paper, we focus on a particular type of complexity: that of the information environment of a Bayesian game. The complexity of an information structure is captured by the extent of contingent thinking it requires, which in turn is determined by the amount of private information in a class of nested information structures we consider.

To be precise, contingent thinking is broadly defined as a form of scenario analysis that takes into account and conditions on the different possible realizations (states) of unobservable or yet unobserved events. In our case, the private signal realizations of one player are not so much *yet unobserved* but are instead *unobservable* to the other players.[1] Pure public signals require no contingent thinking because the signals realizations are common knowledge, while if one of the interacting agents has observed a private signal, that requires a higher level of contingent thinking by his opponents, and even more contingent thinking is required if two of the agents have observed a private signal.

In our experiment, we use information structures that are ordered in terms of the degree of private information they contain, and, therefore, in terms of their contingent thinking complexity. More specifically, we have three treatments: public information, private information to player 1 only, and private information to players 1 and 2. These treatments are ranked according to the extent of contingent thinking that players need to engage in, due to the complexity of higher-order beliefs agents need to formulate. With this design, we

[1] Experimental studies on contingent thinking suggest that subjects find it difficult to perform hypothetical state-by-state reasoning (Charness and Levin (2009), Esponda and Vespa (2014), Ngangoué and Weizsäcker (2021), as well as the overview paper Niederle and Vespa (2023)). In more recent work, Aina, Amelio, and Brütt (2024) experimentally study the impact of contingent thinking on belief updating and show that contingent thinking leads to significant deviations from Bayesian updating.

aim to examine whether the degree of information complexity has a systematic impact on agents' displayed higher-order rationality levels.

To this end, we extend the identification approach of Kneeland (2015), which relies on dominance solvable complete information games to incomplete information games. We introduce a class of incomplete information games called *chain games* to isolate the behavioral implications of different orders of rationality. In a 3-player chain game, player 1's payoff depends only on the realization of the state of nature and his own action choice, player 2's payoff depends on the action of player 1 and her own action choice, and player 3's payoff depends on player 2's action and his own action choice. We then combine the same incomplete information chain game with different information structures, ordered in terms of contingent thinking complexity, which results in different Bayesian chain games. Importantly, we maintain the base payoff environment constant while we vary the information environment of the Bayesian game, thus creating a sequence of dominance solvable Bayesian games. This structured way of varying the information environment of Bayesian games according to their contingent thinking complexity provides a novel approach to isolating a potential determinant of higher-order rationality. As such, the paper provides a methodological contribution to the literature on epistemic experiements.

The opponent structure of chain games, allows us to isolate the behavioral implications of different orders of rationality between subjects and, importantly, compare rationality orders across the different information environments within subjects. It is essential to emphasize that we are not seeking to measure and interpret the exact orders of rationality in each information environment and conclude something about orders of rationality extrapolated to general games. Instead, we are interested in how each subject's rationality orders vary in our classes of Bayesian games as the information complexity changes while the base incomplete information game remains fixed. As such we are able to isolate the effect of increasing contingent thinking information complexity on the subjects' higher-order reasoning about rationality. For each information environment/treatment— public versus private information to player 1 only versus private information to players 1 and 2 — we assign the highest level of reasoning about rationality consistent with the subject's action choices for the different versions of that environment.[2] By assigning the maximum level

[2]In contrast, Kneeland (2015)'s goal was to infer an exact level of reasoning about rationality and they showed how this can be achieved in permuted ring games under two key identifying assumptions: subjects are rational and an exclusion restriction is satisfied. We depart from the exclusion restriction, as the different variations of our Bayesian games induce different hierarchies of beliefs about payoffs. Instead, we identify

of reasoning about rationality consistent with the data to each subject in each information treatment, we provide a conservative lower bound on the effect that the complexity of the information environment can have on the rationality level of agents.

To the best of our knowledge, this is the first paper that studies how strategic sophistication (in its iterative reasoning-about-rationality form) is affected by the contingent thinking complexity of the information structure under which the strategic interaction takes place. We show that the complexity of the information environment can have profound effects on the outcome of the interaction and the within-subject variations in displayed higher-order rationality in strategic environments of incomplete information.

These results have potentially important implications for the design of optimal information structures as they suggest the presence of systematic trade-offs between the nature and complexity of signals, on the one hand, and the depths of reasoning that players display, on the other hand. This is a very relevant consideration for information design, where, for a fixed payoff structure, a designer can manipulate the outcome of a Bayesian game by choosing the information environment in which the interaction takes place. The optimal information structure is chosen to maximize the designer's expected payoff under the usual assumption that players coordinate on the designer's preferred equilibrium when multiplicity exists. For the most part, it has also been assumed that agents are perfectly rational and have unlimited abilities to apply Bayesian updating and higher-order reasoning, irrespective of the nature of information disclosed. Our findings show that the type of information structure plays a pivotal role in the displayed strategic sophistication of agents, through its effect on contingent thinking, which in turn suggests this should be an important design consideration. The recent literature on simplicity in mechanism design addresses separately the concerns about the contingent thinking abilities of agents (Li (2017), Pycia and Troyan (2023), Nagel and Saitto (2025)) and their strategic sophistication (Börgers and Li (2019), Li and Dworczak (2024)).

While our results reveal that information complexity leads to a reduced level of strategic sophistication as measured by the higher-order rationality of agents, we are agnostic about whether this is driven by intrinsic ability bounds on iterative reasoning/strategic sophistication or by the beliefs about the other agents' abilities in iterative reasoning/strategic sophistication. Importantly, both of these effects play a role in determining the displayed strategic sophistication of agents and the resulting departures from theoretical predictions,

only the maximum level of reasoning about rationality consistent with the data.

both according to the theoretical epistemic foundations, as well as the results of experimental studies (see, e.g., Agranov et al. (2012), Georganas, Healy, and Weber (2015), Alaoui and Penta (2016), Alaoui, Janezic, and Penta (2020), Jin (2021)). In future work, it would be therefore important to understand the relative impact of information complexity on each channel. This paper, however, simply aims to shed light on how information complexity impacts these channels in their totality, as captured by the displayed higher-order rationality of agents.

The remainder of the paper is organized as follows. Section 2 presents the motivating framework and explains the intuition behind our Bayesian chain games, which are formally introduced in Section 3. Section 4 describes the experimental design and the specifics of the different treatments. Section 5 presents the data analysis and experimental results. The Appendix contains the experimental instructions.

2 Motivating Framework: The Manager's Problem

Consider a manager who is in charge of two employees collaborating on a project: a supervisor (S) and a worker (W).[3] S and W simultaneously make choices of their own effort levels. S's choice can be interpreted as deciding whether to monitor ($a_S = 1$) or not monitor ($a_S = 0$) W. W can choose to either work ($a_W = 1$) or shirk ($a_W = 0$). The project's difficulty (the state) is uncertain at the time effort choices are made. The project can be either easy ($\theta = 0$) or hard ($\theta = 1$), and is distributed according to the common prior distribution $p_0 := p_0(\theta = 1) = 1/6$, which is common knowledge amongst the manager, S, and W.

The payoffs are presented in Table 1, which summarizes the strategic interaction between the two employees. We will refer to this as the *base game*. On the one hand, S is only interested in monitoring W if the project is hard, regardless of W's effort choice. We could interpret this as S only being invested in hard projects, as those are the ones that matter for his promotion. On the other hand, W wants to work only if he is being monitored by S, regardless of the state, as his own promotion depends on whether S observes his performance.

Consider a manager whose payoff only depends on W's action, i.e., the manager's payoff

[3] This problem was initially included in an early working paper version of Mathevet, Perego, and Taneva (2020). Taneva is thankful to her coauthors on that paper for allowing her to use this example and build on it.

$\theta = 0$	$a_W = 0$	$a_W = 1$
$a_S = 0$	$1, 1$	$1, 0$
$a_S = 1$	$0, 0$	$0, 1$

$\theta = 1$	$a_W = 0$	$a_W = 1$
$a_S = 0$	$0, 1$	$0, 0$
$a_S = 1$	$1, 0$	$1, 1$

Table 1: Base Game between supervisor (row) and worker (column).

is given by $v(a_S, a_W, \theta) = a_W$ for all $a_S, a_W, \theta \in \{0, 1\}$.[4] That is, she wants to incentivise W to exert high effort as often as possible.

The manager is an information designer, who commits to and announces an information structure under which S and W will make their effort choices. She chooses the information structure that maximizes the expected value of her payoff, $\mathbb{E}v$, assuming that the employees play her preferred Bayes Nash equilibrium (BNE) if there is multiplicity under the chosen information structure.

For this simple setting, the equilibrium actions of S and W can be expressed only in terms of their first- and second-order beliefs, respectively. A first-order belief of S about the state θ is the posterior probability that he attaches to the project being hard and is denoted by $\mu_1 := \Pr_S(\theta = 1)$. Hence, the equilibrium action of S is:[5]

$$a_S^*(\mu_1) = \begin{cases} 0 & \text{if } \mu_1 < 1/2 \\ 1 & \text{if } \mu_1 \geq 1/2. \end{cases}$$

The second-order belief of W about the state can be summarized in the following way: $\lambda_2 := \Pr_W(\mu_1 \geq 1/2)$. Thus, the second-order belief of the worker is defined as the probability he attaches to the supervisor having a first-order belief that is weakly greater than one half. W's equilibrium action is therefore a function of his second-order beliefs:

$$a_W^*(\lambda_2) = \begin{cases} 0 & \text{if } \lambda_2 < 1/2 \\ 1 & \text{if } \lambda_2 \geq 1/2. \end{cases}$$

Consider first two natural benchmarks:

No Information: If the manager does not reveal any information and lets the two em-

[4] This type of objective function is reasonable when the worker is the main productive unit, as in the three-tier hierarchy model of Tirole (1986).

[5] Remember that ties are broken in favor of the manager.

ployees make their effort choices under the prior p_0, it is strictly dominant for S to choose $a_S = 0$ and given that, it is iteratively strictly dominant for W to choose $a_W = 0$. Hence, the expected payoff of the manager under no information is $\mathbb{E}v_n = 0$.

Full Information: If the manager were to always perfectly reveal the true state of the world to the two employees, then for any prior state distribution p her expected payoff will lie on the gray (dotted) line in Figure 1. For the current prior state distribution of $p_0 = 1/6$, the expected payoff of the manager under full information is $\mathbb{E}v_f = 1/6$: with probability 5/6 the project is revealed to be easy, so in turn it is strictly dominant for S to choose $a_S = 0$ and iteratively dominant for W to choose $a_W = 0$, which results in a manager's payoff of 0; and with probability 1/6 the project is revealed to be hard, so in turn it is strictly dominant for S to choose $a_S = 1$ and iteratively dominant for W to choose $a_W = 1$, which results in a manager's payoff of 1.

However, the manager can improve upon these two benchmarks by using more sophisticated information design. She first restricts attention to public information. In this case, the optimal solution can be found using the concavification approach of Kamenica and Gentzkow (2011).

Public Information: If the manager were to use only public signals revealed to both agents simultaneously, then for any prior state distribution p her optimal expected payoff lies on the *solid* blue line in Figure 1. This corresponds to the concavification of the manager's payoff $v(p)$ at each common belief p shared by the two employees, given by the dashed blue lines: the manager has a payoff of 1 whenever S and W both have common knowledge of a common belief $p \geq 1/2$, and a payoff of 0 when $p < 1/2$. At the prior of $p_0 = 1/6$, the manager optimally randomizes between sending a public signal which induces $p' = 1/2$ with probability 1/3, and a public signal which induces $p'' = 0$ with probability 2/3. The weighted average of these posteriors is equal to the prior p_0: $\frac{1}{3} \cdot \frac{1}{3} + \frac{2}{3} \cdot 0 = \frac{1}{6}$, and hence Bayes plausibility is satisfied. By doing this, the manager obtains an expected payoff of $\mathbb{E}v_p = 1/3$.

Optimal Solution: The optimal solution can be found using the techniques developed in Mathevet, Perego, and Taneva (2020) and involves private information revealed to S, which creates additional uncertainty for W, incentivizing him to choose to work more often than when information is public. This corresponds to the concavification of the manager's expected payoff at each common belief p shared by the two employees, where private

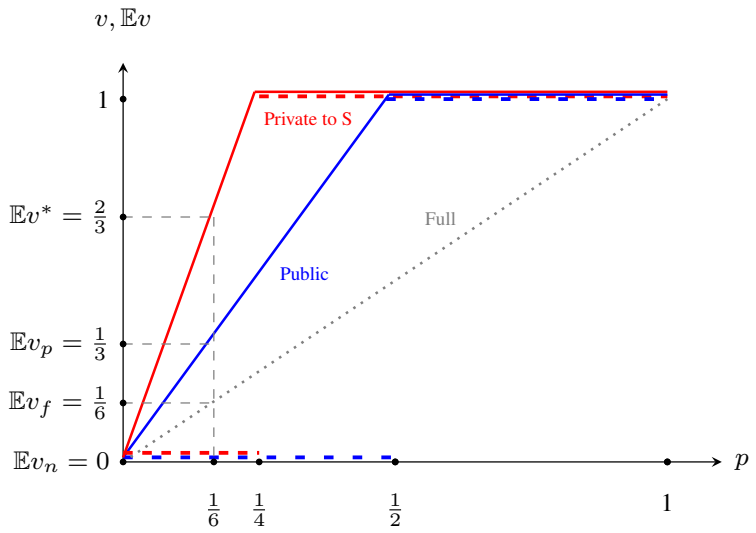


Figure 1: Designer's Expected Payoff under Different Information Modes

signals can be disclosed in addition, given by the dashed red lines:

- for $p \in [\frac{1}{2}, 1]$, use a public signal that establishes the common posterior belief p. Since $p \geq \frac{1}{2}$, this induces $\lambda_2 = 1$, and hence $a_W^* = 1$ and $v^*(p) = 1$.
- for $p \in [\frac{1}{4}, \frac{1}{2})$, use a public signal that establishes the common posterior belief p and a private signal to S which splits her first order belief into $\mu_1' = 0$ (with probability $1 - \lambda_2$) and $\mu_1'' = \frac{1}{2}$ (with probability λ_2), such that $(1 - \lambda_2)\mu_1' + \lambda_2\mu_1'' = p$. This ensures that W's second-order belief λ_2 is consistent with the induced distribution of posteriors. Since $p \geq \frac{1}{4}$, this induces $\lambda_2 \geq \frac{1}{2}$, and hence $a_W^* = 1$ and $v^*(p) = 1$.
- for $p \in [0, \frac{1}{4})$, it must be that $\lambda_2 < \frac{1}{2}$ no matter how the first-order beliefs of S are split. Hence, using a public signal that establishes the common posterior belief p is optimal, which induces $\lambda_2 = 0$, $a_W^* = 0$, and $v^*(p) = 0$.

Concavifying the thus constructed $v^*(p)$ (dashed red lines in Figure 1), gives us the value that can be achieved by the designer when using both public and private signals optimally (solid red graph).

The optimal randomization over posterior belief distributions at the prior $p_0 = 1/6$ gives the designer an expected payoff of $\mathbb{E}v^* = \frac{2}{3}$ and involves the following:

- with probability 2/3 sending a public signal realization (denoted by $P_{1/4}$) which leads to a common posterior belief of $p' = 1/4$. In addition, S receives a private signal,

which with probability 1/2 sends a signal realization (denoted by S_0) that induces a private posterior belief of $\mu_1' = 0$, and with probability 1/2 sends a signal realization (denoted by $S_{\frac{1}{2}}$) that induces a private posterior belief of $\mu_1'' = 1/2$. Note that the private posterior beliefs of S satisfy Bayes plausibility relative to the common posterior belief p': $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 0 = \frac{1}{4}$. Here, the second order belief of W is given by $\lambda' = 1/2$ and hence $a_W^* = 1$.

- with probability 1/3 sending a public signal realization (denoted by P_0) which leads to a common posterior belief of $p'' = 0$. Here, the second order belief of W is given by $\lambda'' = 0$ and hence $a_W^* = 0$.

The distribution of common posterior beliefs is Bayes plausible relative to the common prior p_0: $\frac{2}{3} \cdot \frac{1}{4} + \frac{1}{3} \cdot 0 = \frac{1}{6}$. This ensures there exists an information structure which induces the desired distribution over posterior beliefs. For completeness, we present this information structure in Table 2, where $\pi : \Theta \to \Delta\left(\left\{P_0, \left(P_{\frac{1}{4}}, S_0\right), \left(P_{\frac{1}{4}}, S_{\frac{1}{2}}\right)\right\} \times \left\{P_0, P_{\frac{1}{4}}\right\}\right)$ is the optimal distribution over joint signal realizations for each possible realization of the state.

$\pi(\cdot\|\theta = 0)$	P_0	$P_{\frac{1}{4}}$
P_0	$2/5$	0
$P_{\frac{1}{4}}, S_0$	0	$2/5$
$P_{\frac{1}{4}}, S_{\frac{1}{2}}$	0	$1/5$

$\pi(\cdot\|\theta = 1)$	P_0	$P_{\frac{1}{4}}$
P_0	0	0
$P_{\frac{1}{4}}, S_0$	0	0
$P_{\frac{1}{4}}, S_{\frac{1}{2}}$	0	1

Table 2: Optimal Information Structure

Discussion: In this example, the optimal information structure involves private information to S while leaving W always (weakly) less informed. The splitting of S's first-order beliefs, however, creates the need for contingent thinking by W, which is arguably more complex than in the case of public signals, where S and W make their choices under the same information and a shared posterior belief. In this latter case, W does not need to engage in contingent thinking, as he knows with certainty S's first-order belief and therefore also his optimal choice. Importantly, this contingent thinking complexity associated with the optimal information structure may result in lower orders of displayed rationality and, therefore, less frequent play according to the theoretical predictions. If this is the case, there might be a tradeoff for the designer: while optimal in theory, private information may lead to play that is inconsistent with rationality and higher-order rationality due to the increased contingent thinking complexity; on the other hand, public information, while suboptimal in

theory, may result in play more consistent with the theoretical predictions due to the lower contingent thinking complexity and higher orders of displayed rationality.

The Three-Player Version: Suppose that the manager could hire another supervisor and augment the interaction as follow: the first supervisor, S_1, only cares about matching the state with his action, the second supervisor, S_2, only cares about matching the action of S_1, and the worker W only cares about matching the action of S_2. In this case, the optimal information structure involves private information to both S_1 and S_2, where S_2 is always (weakly) less informed than S_1, and W is always (weakly) less informed than S_2. In this game, W's equilibrium action is a function solely of his third-order beliefs denoted by $\gamma_3 \coloneqq \Pr_W(\lambda_2 \geq 1/2)$:

$$a^*_W(\gamma_3) = \begin{cases} 0 & \text{if } \gamma_3 < 1/2 \\ 1 & \text{if } \gamma_3 \geq 1/2. \end{cases}$$

More specifically, the optimal randomization over posterior belief distributions at the prior $p_0 = 1/6$ gives the designer an expected payoff of $\mathbb{E}v^* = 1$, since W can be induced to work all of the time, i.e., with probability 1. W receives no signals and is completely uninformed, except for knowing the information structure according to which S_1 and S_2 are informed, which is common knowledge. S_1 and S_2 are informed in the following way:

- with probability 2/3, both S_1 and S_2 receive a signal that leads to a common posterior belief of $p' = 1/4$. In addition, S_1 receives a private signal, which with probability 1/2 induces a private posterior belief of $\mu'_1 = 0$ and with probability 1/2 induces a private posterior belief of $\mu''_1 = 1/2$. Note that the private posterior beliefs of S_1 satisfy Bayes plausibility relative to the common posterior belief p': $\frac{1}{2} \cdot \frac{1}{2} + \frac{1}{2} \cdot 0 = \frac{1}{4}$. Here, the second order belief of S_2 is given by $\lambda'_2 = 1/2$ and hence $a^*_{S_2} = 1$.
- with probability 1/3, both S_1 and S_2 receive a signal that leads to a common posterior belief of $p'' = 0$. Here, the second order belief of S_2 is given by $\lambda''_2 = 0$ and hence $a^*_{S_2} = 0$.

Given this information structure, $\gamma_3 = 2/3$, and hence $a^*_W(\gamma_3) = 1$. The distribution of common posterior beliefs between S_1 and S_2 is Bayes plausible relative to the common prior p_0: $\frac{2}{3} \cdot \frac{1}{4} + \frac{1}{3} \cdot 0 = \frac{1}{6}$, while W's posterior belief about the state remains at the prior. This ensures there exists an information structure which induces the desired distribution over posterior beliefs.

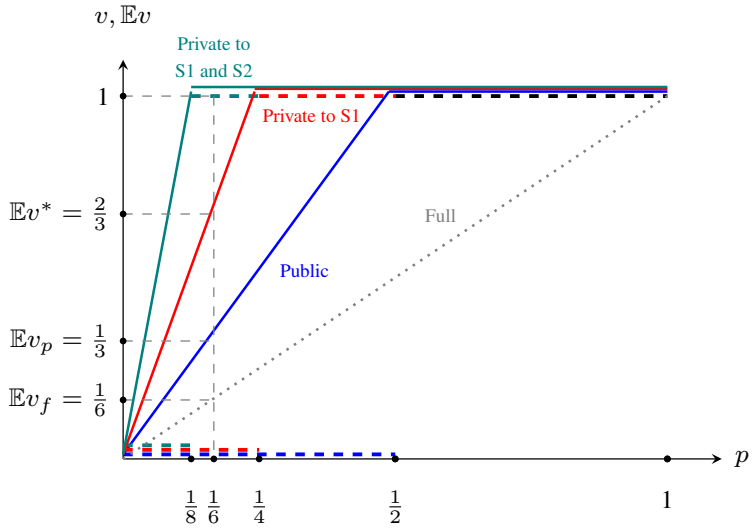


Figure 2: Three-Player Version: Different Information Modes

The three nested information structures, which also form the basis for the different treatments in our experiment, are presented in Figure 2. Under "public", all three players – S_1, S_2, and W – are equally informed; under "private to S_1", S_1 is (weakly) more informed than S_2 and W, who are equally informed; and lastly, under "private to S_1 and S_2", S_1 is always (weakly) more informed than S_2, and S_2 is always (weakly) more informed than W.

An interesting conclusion of this dominance solvable structure is that the manager can guarantee that W works with probability 1 for a larger interval of priors by adding additional supervisors in the chain in a way resembling the payoff incentives presented in our game: the first supervisor only wants to match the state with his action, each additional supervisor only wants to match the action of the previous supervisor, and W only wants to match the action of the last supervisor, where matching each state gives the same positive payoff and not matching the state gives a payoff of zero. For example, with only one supervisor and one worker, the manager can guarantee that the worker is persuaded to exert effort with probability 1 as long as $p_0 \geq 1/4$, while with two supervisors and one worker, this can be achieved for any $p_0 \geq 1/8$. For n supervisors, just by using information design, the manager can guarantee that the worker works with probability one as long as $p_0 \geq 1/2^{n+1}$.[6] This provides a theoretical underpinning for the prevalence of hierarchies in the real world: incentives for performance are created through uncertainty about the information one's superiors have observed and their respectively optimal actions.

[6]We thank Meg Meyer for this observation.

3 Theoretical Framework

In this section we formally define the notion of a Bayesian *chain game*. We start by stating the relevant definition of a game with incomplete information tailored to our environment.[7]

Definition 1. *An n-player incomplete information chain game is a tuple $\Gamma_n = (I, o, (A_i, u_i)_{i \in I}, \Theta)$ where:*

(a) $I = \{1, \ldots, n\}$,

(b) $o : I \to I$ *with* $o(i) = i - 1$ *for* $i = 2, \ldots, n$,

(c) $u_1 : \Theta \times A_1 \to \mathbb{R}$, *and*

(d) $u_i : A_{o(i)} \times A_i \to \mathbb{R}$.

The set Θ denotes a finite set of payoff states, I is the set of players, o is the opponent mapping function designating an opponent $o(i)$ for each player $i = 2, \ldots, n$, A_i is the finite set of feasible actions for each player $i \in I$, and u_i is the payoff function for each player $i \in I$. The payoff and opponent functions ensure that the strategic relationship is one of a chain: player 1's payoff only depends on his own action and the realization of the state; player 2's payoff only depends on her own action and the action of her opponent player 1; and so on, with player n's payoff only depending on his own action and the action of his opponent player $n - 1$.

An information structure which is common knowledge induces a (non-redundant) type space à la Harsanyi with information types $\mathcal{T} = (T_i, \tau_i)_{i \in I}$ where for each player $i \in I$ there is a finite set of types T_i and a belief map $\tau_i : T_i \to \Delta(\Theta \times T_{-i})$ with $T_{-i} = \prod_{j \neq i} T_j$. Hence, each type t_i of player $i \in I$ represents beliefs about the state and about the other players' beliefs.

Definition 2. *An n-player Bayesian chain game is a pair $(\Gamma_n, \mathcal{T})$.*

For complete information games, Bernheim (1984), Pearce (1984), and Tan and da Costa Werlang (1988) establish the following equivalence: a subject who satisfies kth-order rationality (that is, rationality and common belief of order k in rationality) must play a kth-order

[7] To be more specific, this should be categorized as a game with payoff uncertainty as it is a game with complete but imperfect information about an initial chance move rather than a game of genuine incomplete information. The latter class of games is characterized by lack of common knowledge of the mappings from players' actions into their payoffs. Both of these classes are games with asymmetric information (see Battigalli et al. (2011)).

rationalizable action. For incomplete information games, Battigalli et al. (2011) show that the behavioral implications of rationality and common belief in rationality are captured by the concept of *interim correlated rationalizability* (Dekel, Fudenberg, and Morris (2007)). More specifically, when Θ-based hierarchies are as specified by the Harsanyi-type space $\mathcal{T}$, the action profiles of player i of type t_i that are consistent with k-th order rationality are given by the set of **interim correlated k-rationalizable** actions for player i with type t_i, which we formally define below.

Definition 3. *For $i = 1$ and $t_1 \in T_1$, set $ICR_1^0(t_1) = A_1$ and for each $k \geq 1$*

$$ICR_1^k(t_1) \coloneqq \left\{ a_1 \in A_1 \;\middle|\; (i) \quad a_1 \in \arg\max_{a_1' \in A_i} \sum_{\theta} \tau_1(t_1)[\theta] u_i(\theta, a_1',) \right\}.$$

For each $i \geq 2$ and $t_i \in T_i$, set $ICR_i^0(t_i) = A_i$ and for each $k \geq 1$ define recursively,

$$ICR_i^k(t_i) \coloneqq \left\{ a_i \in A_i \;\middle|\; \begin{array}{l} \exists \mu_i \in \Delta(A_{i(o)} \times T_{i(o)} \times \Theta) \text{ s.t.:} \\ (i) \quad a_i \in \arg\max_{a_i' \in A_i} \sum_{t_{i(o)}, a_{i(o)}, \theta} \mu_i(a_{i(o)}, t_{i(o)}, \theta) u_i(a_i', a_{i(o)}, \theta), \\ (ii) \quad \text{If } \mu_i[\{(a_{i(o)}, t_{i(o)})\} \times \Theta] > 0 \text{ then } a_{i(o)} \in ICR_{i(o)}^{k-1}(t_{i(o)}), \\ (iii) \quad \tau_i(t_i)[t_{i(o)}, \theta] = \sum_{a_{i(o)}} \mu_i(a_{i(o)}, t_{i(o)}, \theta) \quad \forall (t_{i(o)}, \theta) \in T_{i(o)} \times \Theta \end{array} \right\}$$

specify the set of **interim correlated k-rationalizable** *actions for player i with type t_i in $(\Gamma_n, \mathcal{T})$. The set of* **interim correlated rationalizable actions** *for player i with type t_i in $(\Gamma_n, \mathcal{T})$ is $ICR_i(t_i) \coloneqq \cap_{k \geq 0} ICR_i^k(t_i)$.*

In our experiment we intentionally use information structures such that for any player in position i and any $t_i \in T_i$, $|ICR_i^i(t_i)| = 1$, that is, there is a unique interim correlated i-rationalizable action for each player-type profile under the correct beliefs $\tau_i(t_i)$

In the experiment, we consider three different classes of information structures, ranked in terms of their contingent thinking complexity – public information (PU), private information to player 1 only (PR1), and private information to players 1 and 2 (PR2). We denote the class of information structure by $c \in \mathcal{C} \coloneqq \{\text{PU, PR1, PR2}\}$.

For each of these classes, we consider different variations of the information structure within the class, which we call precisions and denote by $\pi^c \in \Pi^c$. Each precision is an information structure that gives rise to a Harsanyi type space $\mathcal{T}^{\pi_c}$. We next describe our identification approach, which assigns the maximum rationality level consistent with action

choices for each class of information structures.

Identification: Given an information class $c \in \mathcal{C}$ and $k \geq 1$, a subject who satisfies kth-order rationality plays (the unique) $a_i \in ICR_i^i(t_i^{\pi_c})$ for all games $(\Gamma_n, \mathcal{T}_{\pi_c})$ with $\pi_c \in \Pi_c$ when in player positions $i = 1, \ldots, k$. We assign to each subject the maximum k such that this holds for the given $c \in \mathcal{C}$

Hence, the identification strategy has a nested structure: a subject can be categorized as kth-order rational only if he is also rational of all orders smaller than k. Additionally, the above implies that if a subject satisfies kth-order rationality but not $(k+1)$th-order rationality for information class $c \in \mathcal{C}$, he must be playing an action $\tilde{a}_i \notin ICR_i^{k+1}(t_i)$ for some game $(\Gamma_n, \mathcal{T}_{\pi_c})$ with $\pi_c \in \Pi^c$ when in player position $i = k+1$. Note further that an n-player chain game is able to identify rationality of up to order n, where a subject identified as nth-order rational is at least nth-order rational, but potentially satisfying a higher rationality order.

4 Experimental Design

We use the preceding framework and results to empirically identify from strategic choice data the orders of rationality that each subject exhibits under the different information treatments. We use the three-player version of the manager's problem in a context-free environment where the state is the color of a ball drawn from a jar. The information structures in our experiment are exogenously set, so there is no role for the manager as an information designer. Instead, we are interested in the effect of different information environments on the interacting agents' displayed orders of rationality, which may be relevant from an information design perspective.

We presented the base game to the subjects in the following way:

Role 1's payoffs

		Ball draw: **P**	Ball draw: **G**
Role 1's actions	e	1	0
	f	0	1

Role 2's payoffs

		Role 1's actions: e	Role 1's actions: f
Role 2's actions	j	1	0
	k	0	1

Role 3's payoffs

		Role 2's actions: j	Role 2's actions: k
Role 3's actions	s	1	0
	t	0	1

Table 3: Three-Player Chain Game

Role 1's payoff only depends on whether he matches the color of the ball with his actions. If the player in Role 1 is rational, then he will play e if he believes the ball is more likely

to be purple (**P**), and f if he believes it is more likely to be green (**G**). Role 2's payoff only depends on whether she matches the action of Role 1 with her action. If the player in Role 2 is second-order rational (is rational and believes that the player in Role 1 is rational), then she must play j if she believes that the player in Role 1 believes the ball is more likely to be purple (and hence Role 1 will play e), and she must play k if she believes that the player in Role 1 believes the ball is more likely to be green (and hence Role 1 will play f). Role 3's payoff only depends on whether he matches the action of Role 2 with his action. If the player in Role 3 is third-order rational (is rational, and believes the player in Role 2 is second-order rational), then he must play s if he believes that the player in Role 2 believes that the player in Role 1 believes the ball is more likely to be purple (and hence Role 2 will play j), and he must play t if he believes that the player in Role 2 believes that the player in Role 1 believes the ball is more likely to be green (and hence Role 2 will play k).

This constitutes a three-player chain game of incomplete information, which permits identification of up to third order of rationality. These payoffs were visible on the screen at all times, presented stacked vertically, starting with the left-most matrix, instead of horizontally (see Appendix A: Instructions). The use of binary payoffs guarantees that risk-attitude considerations do not play a role. We clearly explained the structure of the game during the instructions, as we are not interested in the absolute values of the subjects' rationality orders, but in their relative change under the different information regimes. Therefore, it was important that subjects understood the structure of the base game well, as the different information structures introduce additional complexity. We are interested in whether that additional complexity plays a role for the subjects' rationality orders.

For each treatment, depending on their play in each role under each precision level of the corresponding treatment, we assign each subject the maximum rationality level consistent with their action choices. Hence, for each treatment, we separate subjects into four rationality categories, irrational (R0), rational but not second-order rational (R1), second-order rational but not third-order rational (R2), and (at least) third-order rational (R3). The precise classification rules for each treatment are presented in the following subsections.

4.1 Treatments

Each subject's order of rationality is estimated from choice data three times, in each of the three *treatments*, which represent different types of information environments: public information (PU), private information to Role 1 (PR1), and private information to both Role

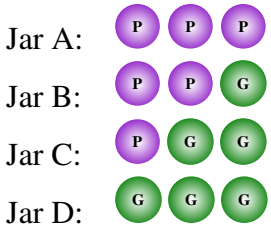


Figure 3: The Different Jars in the Experiment/Precisions in PU Treatment

1 and Role 2 (PR2). These correspond to the dark blue, red, and light blue graphs in Figure 2, respectively. Each treatment has three or four variants of the information structure, with different informativeness about the state and/or the other players' information, which we call *precisions*. Public information has four precisions, which include two cases of complete information, while the other two treatments have three precisions each. The details of each treatment are presented next.

4.1.1 Public Information (PU) Treatment

In the PU treatment, the jar from which a ball (either purple (**P**) or green (**G**)) will be drawn at random is announced to all three roles, and this is common knowledge. The chosen jar establishes the common belief that all three roles share about the ball (i.e., the state). The four jars labeled A, B, C, and D are presented in Figure 3 and these labels are maintained across all treatments.[8]

Each jar's composition represents a given *precision* in the PU treatment, as it establishes a different common posterior belief about the ball color. For example, if it is announced that the ball will be drawn from Jar A, and the players in all roles are told that everyone else has the same information, then the common belief is that the ball will be purple with probability 1, which is also common knowledge. On the other hand, if it is announced that the ball will be drawn from Jar B, and the players in all roles are told that everyone else has the same information, then the players in all roles assign probability $2/3$ to the ball being purple, and $1/3$ to the ball being green, and that is common knowledge.

We separate subjects into four rationality categories depending on their play in each role under each precision level of the PU treatment, i.e. according to their *action choices*.[9] For each role, we call the vector of action choices across the different precisions a *choice*

[8] In the experiment, the balls were presented only in their respective color — purple or green — but in the paper we have added letters P and G to avoid ambiguity under black-and-white print.

[9] Each precision-role combination is a game.

vector. The choice vectors consistent with each rationality type R1–R3 are presented in Table 4, where for each role the four entries of each vector correspond to the action choice in each of the four precisions of this treatment. For example $(e; e; f; f)$ means that the subject in Role 1 plays e when it is announced that the ball will be drawn from Jar A, e from Jar B, f from Jar C, and f from Jar D.

If a subject is categorized as R3, she must always play the rationalizable choice vector in all three roles: $(e; e; f; f)$ when she plays in Role 1, $(j; j; k; k)$ when she plays in Role 2, and $(s; s; t; t)$ when she plays in Role 3. If a subject is categorized as R2, she must play the rationalizable choice vector in Roles 1 and 2, but she plays any other choice vector except the rationalizable one in Role 3. If a subject is categorized as R1, she must play the rationalizable choice vector in Role 1, but she can play any choice vector in Role 2 except the rationalizable one and any choice vector in Role 3. Subjects who are not categorized under one of the three categories R1-R3, are classified as R0.

PU	Role		
Type	Role 1	Role 2	Role 3
R1	$(e; e; f; f)$	$\neg(j; j; k; k)$	
R2	$(e; e; f; f)$	$(j; j; k; k)$	$\neg(s; s; t; t)$
R3	$(e; e; f; f)$	$(j; j; k; k)$	$(s; s; t; t)$

Table 4: Action Choices under Rationality in PU Treatment

We follow the same process for assigning the rationality types in the other two treatments, which, however, have three precisions instead of four. We made this choice in order to impose a higher threshold for assigning types in the PU treatment, due to it being the simplest information environment. However, most subjects (103 out of 115) are exact R3 matches even with choice vectors of length four. We would have gotten even more R3 matches if we considered only three precisions in this treatment (choice vectors of length three), which would have only strengthened our finding of very high rationality in the PU treatment.

4.1.2 Private Information to Role 1 (PR1) Treatment

In the PR1 treatment we introduce a box of cards labeled either A, B, C or D. The composition of the card box is announced to all three roles, which is common knowledge. A card is then drawn at random from the box and revealed only to the person in Role 1, while

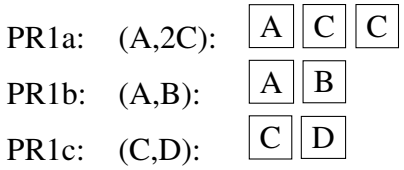


Figure 4: Precisions in PR1 Treatment

Roles 2 and 3 remain uninformed about the realization of the card draw. This information structure is common knowledge. The letter on the card determines which jar the ball will be drawn from, i.e., if the drawn card has letter A on it, then the ball will be randomly drawn from Jar A, and similarly for all other letters. Therefore, only the player in Role 1 will be certain about the jar that is going to be used for the ball draw, while players in Roles 2 and 3 will have a distribution over the possible jars, determined by the box composition. The composition of the card box is the public signal which creates a common belief for all three players, while the private signal — the revelation of the drawn card to Role 1 — creates a splitting of Role 1's first-order beliefs.

The three precisions of the PR1 treatment are presented in Figure 4. Consider, for example, the precision PR1a, where the box consists of one A card and two C cards. Under this precision, the common belief about the ball (state) created by the public signal (box composition) is that with probability 5/9 the ball is purple and will probability 4/9 it is green. However, the player in Role 1 will either hold the belief μ_1' that the ball is purple with probability 1 (if he observes the private signal "card A"), in which case he will take action e if he is rational, or the belief μ_1'' that the ball is purple with probability 1/3 (if he observes the private signal "card C"), in which case he will take action f if he is rational. The players in Role 2 and 3 in turn know that μ_1' occurs with probability $1/3$, while μ_1'' occurs with probability 2/3. Hence, if the player in Role 2 is second-order rational, she would choose to play k as she believes it more likely that Role 1 will play f. In turn, if the player in Role 3 is third-order rational, he would choose to play t as he believes that Role 2 will play k with certainty. Note that it is really important for Roles 2 and 3 to engage in the correct type of contingent thinking, where they consider what Role 1 would do under each of the possible contingencies, that is, depending on each possible realization of the private signal, and to also consider the probabilities of those contingencies. If instead Roles 2 and 3 only consider what Role 1 would do under the common belief which assigns probability 5/9 to the ball being purple, then they would expect Role 1 to choose e, which would lead to a reversal to j and s of their own choices, respectively.

Every time a subject is asked to make an action choice, a new card is drawn from the card box of the current precision and a new ball is drawn from the corresponding jar, so the players in Roles 2 and 3 know that the card they might have observed when they played in Role 1, need not be the same when they make their choices as players in Roles 2 and 3. Everyone is informed of this through the instructions and the software also makes this salient.

PR1	Role		
Type	Role 1	Role 2	Role 3
R1	$(e\|A, f\|C; e\|A, e\|B; f\|C, f\|D)$	$\neg(k; j; k)$	
R2	$(e\|A, f\|C; e\|A, e\|B; f\|C, f\|D)$	$(k; j; k)$	$\neg(t; s; t)$
R3	$(e\|A, f\|C; e\|A, e\|B; f\|C, f\|D)$	$(k; j; k)$	$(t; s; t)$

Table 5: Action Choices under Rationality in PR1 Treatment

We separate subjects into the four rationality categories depending on their action choices in each role under each precision level of the PR1 treatment, i.e. according to their choice vectors. The choice vectors consistent with each rationality type R1–R3 are presented in Table 5. Again, all subjects who are not categorized under one of the three categories R1-R3, are classified as R0.

Note that for Role 1 the choice vector is contingent on the card realization to reflect that the optimal action choice of Role 1 may differ depending on the private signal (card letter) they observed. Therefore, for Role 1, the choice vector consists of six entries: two entries for each of the three precision levels, one entry for each possible realization of the card. Indeed, Role 1's rationalizable profile is contingent on the private signal realization (card letter) he observes. The strategy profile $(e|A, f|C; e|A, e|B; f|C, f|D)$ has the strategies for each precision level separated by semi-colons: the first two entries correspond to precision PR1a and specify that e is chosen if the card observed is A, and f is chosen if the card observed is C. Of course, in the data we only observe one choice for Role 1 for each precision, depending on which card is drawn and revealed to him, i.e., depending on the realization of the private signal.

In contrast, the choice vectors for Roles 2 and 3 are not contingent, as players on those roles do not observe private signals. Each entry in the vector corresponds to the action choice for a given precision for the respective role, and hence they each have three entries. For example, $(k; j; k)$ means that Role 2 chooses k under precision PR1a, j under PR1b, and k

under PR1c.

4.1.3 Private Information to Roles 1 and 2 (PR2) Treatment

In the PR2 treatment, there are two boxes of cards — a Primary Box and a Secondary Box — which contain cards labeled either A, B, C, D or X. The Secondary Box does not contain any X cards. The composition of the two card boxes is announced to all three roles, which is common knowledge. A card is then drawn at random from the Primary Box and revealed only to the players in Roles 1 and 2, while Role 3 remains uninformed about the realization of the draw. If the card has an A, B, C, or D, then that card will determine the jar from which the ball will be drawn, and so both Role 1 and Role 2 will know which jar is going to be used, and that will be common knowledge between the two of them. Alternatively, if the card has an X on it, a card from the Secondary Box will be drawn and revealed only to Role 1, while Roles 2 and 3 remain uninformed about the realization of the draw from the Secondary Box. In this case, the letter of the card drawn from the Secondary Box will determine the jar from which the ball will be drawn and so only Role 1 will know which jar is going to be used; Role 2 will know that the first card was an X and that the jar will be determined by the draw from the Secondary Box; while Role 3 will not observe any card draws.

The three precisions of the PR2 treatment are presented in Figure 5. Consider, for example, the precision PR2b, where the Primary Box consists of one C card and one X card, while the Secondary Box consists of one A card and two C cards. Under this precision, the common belief about the ball (state) created by the public signal (the composition of the two boxes) is that with probability 1/2 the ball is purple and will probability 1/2 it is green. However, the player in Role 1 will either hold the belief μ_1' that the ball is purple with probability 1 (if he observes the private signal "card A"), in which case he will take action e if he is rational, or the belief μ_1'' that the ball is purple with probability 1/3 (if he observes the signal "card C" – either common to both Roles 1 and 2 as the draw from the Primary Box, or private to Role 1 as the draw from the Secondary Box), in which case he will take action f if he is rational. The player in Role 2 would either know with certainty that μ_1'' has realized, when the card from the Primary Box is C, or she would believe that μ_1' occurs with probability 1/3 while μ_1'' occurs with probability 2/3, when the card from the Primary Box is X. Hence, if the player in Role 2 is second-order rational, she would choose to play k if she observes "card C" as in this case she believes that Role 1 will play f with certainty, and she would

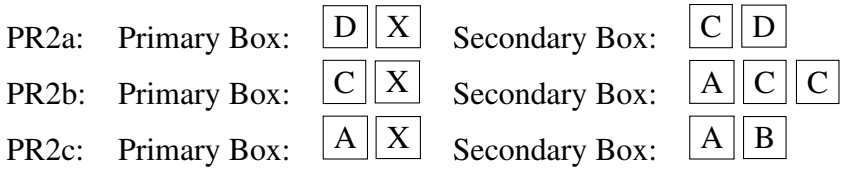


Figure 5: Precisions in PR2 treatment

choose to also play k if she observes "card X" as in this case she believes Role 1 is more likely to play f. In turn, if the player in Role 3 is third-order rational, he would choose to play t as he believes that Role 2 will play k with certainty. Note that it is essential for Roles 2 and 3 to engage in the correct type of contingent thinking, where they consider what Role 1 would do under each of the possible contingencies, that is, depending on each possible realization of the signals which are private to Role 1, or common to Roles 1 and 2, and to also consider the probabilities of those contingencies. For example, if instead Role 2, upon observing "card X", only considers what Role 1 would do under the induced common belief (shared by Roles 1 and 2) which assigns probability 5/9 to the ball being purple, then she would expect Role 1 to play e, which in turn would change her own action choice from k to j.

Every time a subject is asked to make an action choice, a new card is drawn from the Primary Box of the current precision, and, if necessary, a new card is drawn from the Secondary Box of the current precision. Hence, all players know that any cards they may have observed in their previous roles, are irrelevant for their current choice. Also, a new ball is drawn from the newly selected jar for every new action choice. Everyone is informed of this through the instructions and the software also makes this salient.

PR2	Role		
Type	Role 1	Role 2	Role 3
R1	$(f\|C, f\|D; e\|A, f\|C; e\|A, e\|B)$	$\neg(k\|D, k\|X; k\|C, k\|X; j\|A, j\|X)$	
R2	$(f\|C, f\|D; e\|A, f\|C; e\|A, e\|B)$	$(k\|D, k\|X; k\|C, k\|X; j\|A, j\|X)$	$\neg(t; t; s)$
R3	$(f\|C, f\|D; e\|A, f\|C; e\|A, e\|B)$	$(k\|D, k\|X; k\|C, k\|X; j\|A, j\|X)$	$(t; t; s)$

Table 6: Action Choices under Rationality in PR2

We separate subjects into the four rationality categories depending on their action choices in each role under each precision level of the PR1 treatment, i.e. according to their choice vectors. The choice vectors consistent with each rationality type R1–R3 are presented in Table 6. All subjects who are not categorized under one of the three categories R1-R3, are

classified as R0.

Note that for Roles 1 and 2 the choice vector is contingent on the card realization to reflect that the optimal action choice of those roles may differ depending on the private signal they observed (the Primary Box card letter, which is common to Roles 1 and 2, and the Secondary Box letter, which is private to Role 1). For Role 1, the choice vector consists of six entries: two entries for each of the three precision levels, one entry for each possible realization of the card — either from the Primary or the Secondary Box — which determines the Jar. Indeed, Role 1's rationalizable action choice vector is contingent on the signal realization (card letter), either private and shared with Role 2, that he observes. The choice vector $(f|C, f|D; e|A, f|C; e|A, e|B)$ has the action choices for each precision level separated by semi-colons: the first two entries correspond to precision PR2a and specify that f is chosen if the card observed is C, and f is chosen if the card observed is D.

For Role 2, the choice vector also consists of six entries: two entries for each of the three precision levels, one entry for each possible realization of the card from the Primary Box — either a card that determines the jar, or card X. Consider the strategy profile $(k|D, k|X; k|C, k|X; j|A, j|X)$, which presents the action choices for each precision level separated by semi-colons: the first two entries correspond to precision PR2a and specify that k is chosen if the card observed is D, and k is chosen if the card observed is X. Of course, in the data we only observe one choice for Roles 1 and 2 for each precision, depending on which Primary and Secondary Box cards are drawn and revealed to them, i.e., depending on the realization of the card signals.

In contrast, the choice vectors for Role 3 are not contingent, as players in that Role do not observe any private signals. Each entry in the vector corresponds to the action choice for a given precision for the respective role, and hence they each have three entries. For example, $(t; t; s)$ means that Role 3 chooses t under precision PR2a, t under PR2b, and s under PR2c.

4.2 Laboratory Implementation

The experimental sessions were conducted in the Laboratory for the Integration of Theory and Experiments (LITE) with undergraduate students at the University of California Santa Barbara. Each subject participated in only one session and made decisions through a software interface. The subjects were not allowed to interact directly with one another during the experiment.

There were six sessions in total, and each session implemented one of the possible sequences of the three treatments.[10] The precisions also appeared in random order within each treatment. Each subject played in all three roles in all treatments and under all precisions, for a total of 30 strategic choices. The possibility to revise one's choices for all three roles was provided at every precision level. Once subjects had confirmed their action choices, feedback about the computer choices (ball and card draws) and a reminder of one's own action choice was given. Feedback about the other subjects' action choices was not provided. Subjects were then rematched randomly and anonymously in another group of three players, and played the game under a new precision level.

One game was randomly selected for payment at the end of the experiment. Payments were in US dollars. Each subject received the value of their payoff in the selected game times \$20; that is, they received a fixed payment of \$20 if their payoff was 1, and \$0 if their payoff was 0. All subjects received a showup fee of \$5. The average payment for all active subjects was \$21.87.

Instructions were read aloud by the same experimenter at the beginning of each session and before each treatment. Instructions to all subjects were the same, but the order in which they were read depended on the particular session's order of treatments. The instructions can be found in the Appendix.

Three of the sessions consisted of 18 subjects, and the other three sessions consisted of 21 subjects. There were 115 subjects which remained active until the end of the respective session, with 2 subjects choosing to leave prematurely. One of the experimenters had to take the place of the missing subject in the two affected sessions as the software was programmed in a way that could not progress for the affected 3-player groups without the choices of the missing subject. The data of the missing subjects was discarded, however the data of the subjects they interacted with was not. As subjects were rematched at each precision level, any effects on individual subject's behavior would be minimal. The analysis also confirms that there are no significant session effects on the displayed rationality levels (the dummies for all sessions, including the two affected sessions 3 and 5, are insignificant).

[10] For example, the first session's sequence was PU-PR1-PR2, which means that it started with the public information treatment, then private information to role 1, and concluded with private information to roles 1 and 2. The second session's sequence was PR1-PU-PR2, and so on.

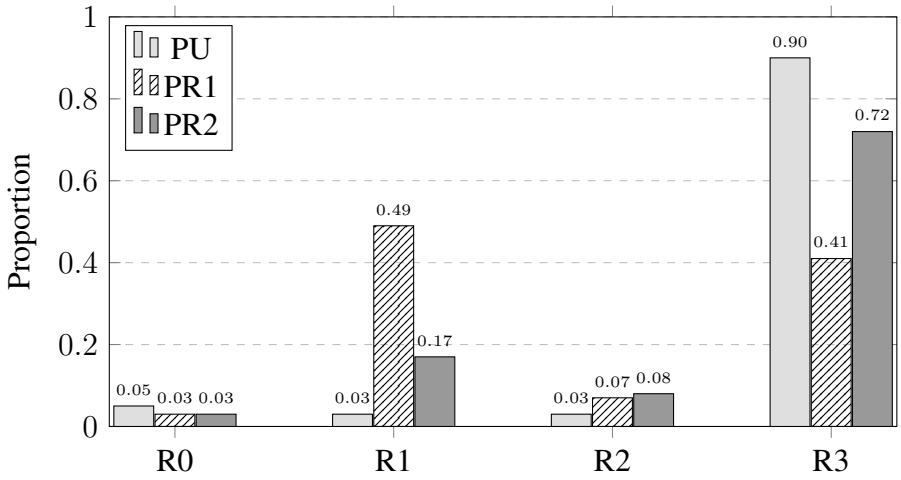


Figure 6: Order of Rationality Classification by Treatment

5 Data Analysis and Results

Rationality Type Frequencies. The proportions of subjects classified as each rationality type R0-R3 in each of the three treatments are presented in Figure 6. For each of the rationality type categories, the bars are ordered in sequence according to the treatments as follows: PU (light solid), PR1 (hatched), PR2 (dark solid). The results are reassuring, as they provide evidence of subjects maintaining first-order rationality consistently throughout the three treatments, with most subjects being classified as R1 or higher in all treatments. The higher levels of rationality, however, vary significantly across the three treatments.

In the PU treatment, most subjects (103 of 115, or $\approx 90\%$) are classified as R3 types. Only 3 of 115 subjects, or $\approx 3\%$, are classified as R1 or R2 types. 6 of 115 subjects, or $\approx 5\%$, are classified as R0. The fact that most subjects are classified as R3 types in the PU treatment provides reassurance that the subjects understood the dominance solvable nature of the base incomplete information chain game well, and that any variations in the rationality orders are due to the complexity of the information environment.

Things change dramatically when we introduce private information to player 1 in the chain game, and thus the necessity for contingent thinking for roles 2 and 3. There is a substantial drop in the displayed levels of rationality in the PR1 treatment compared to the PU treatment. In the PR1 treatment, only 47 of 115 subjects, or $\approx 41\%$, are classified as R3, which is a drop of 31 percentage points relative to the PU treatment. There are 8 of 115 subjects, or $\approx 7\%$, classified as R2. Almost half of the subjects, 57 of 115, or $\approx 49\%$ are classified

as R1. Only 3 of 115 subjects, or $\approx 3\%$, are classified as R0.

Something surprising happens when we introduce even more private information, by providing it to both players 1 and 2, in the PR2 treatment. There is an increase in the proportion of subjects classified as R3 types relative to the PR1 treatment (but not relative to the PU treatment): 83 of 115, or $\approx 72\%$. This constitutes an increase of 31 percentage points of R3 subjects as compared to the PR1 treatment. 9 of 115 or $\approx 8\%$ are classified as R2 in the PR2 treatment, which is very similar to the number of R2 subjects in the PR1 treatment. 20 of 115, or $\approx 17\%$, are classified as R1. And the same number of subjects — 3 of 115, or $\approx 3\%$ — are classified as R0.

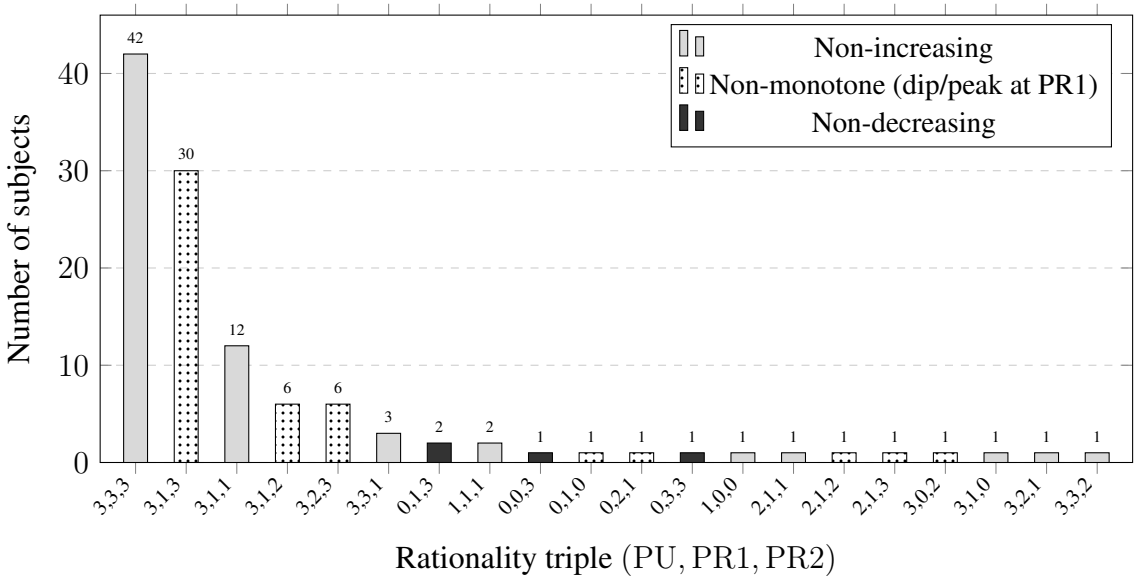


Figure 7: Overall Classification by Rationality Triples

In Figure 7, we present the classification of agents according to their rationality levels in the three treatments combined, so each agent is assigned a rationality triple listing the agent's revealed rationality levels in PU, PR1, and PR2, in that order. Overall, 64 of 115 subjects have non-increasing orders of rationality across the three treatments, which would be in line with the conjecture that the treatments are ranked in increasing order of complexity, at least weakly so. Only 4 of 115 subjects display increasing orders of rationality across the three treatments. However, there are 47 of 115 subjects who display a non-monotone pattern to their rationality levels, almost all of whom (except for 2 subjects) actually exhibit a dip of their rationality level in PR1 relative to the other two treatments. In particular, there are 30 subjects who are classified as R3 in both PU and PR2, but as R1 in PR1. We further

explore this and potential explanations in the subsequent analysis.

Where Mistakes Happen. One plausible explanation for the increase in the higher-order reasoning of agents in the PR2 treatment relative to the PR1 treatment, despite the higher complexity and contingent thinking requirements of the former relative to the latter, is that in PR2 agents in Role 2 do not need to employ contingent thinking when they observe a Primary Box card other than X, because in that case they have the same information as Role 1. In contrast, in PR1, Roles 2 and 3 need to employ contingent thinking, as they are equally informed, and strictly less informed than Role 1. This suggests that the higher-order reasoning in PR2 relative to PR1 may be due to beliefs (in the opponent's ability) rather than (own) ability: in PR2, Role 2 is strictly more informed than Role 3 and half the time does not need to use contingent thinking, which could be the reason why Role 3 believes in Role 2's second-order rationality more often than in the PR1 treatment where Role 2 is equally informed to Role 3 and needs to use contingent thinking all of the time. We think this is an interesting channel to explore in future work.

Mistakes	Precision		
Treatment	a	b	c
PR1	**101 (82%)**	10 (8%)	12 (10%)
PR2	10 (18%)	**39 (68%)**	8 (14%)

Table 7: Frequency of Mistakes across Roles for each Precision within PR1 and PR2

In Table 7 we present the frequency of mistakes that agents make across different roles for the different precisions within the PR1 and PR2 treatments.[11] The precisions (information structures) that have the highest contingent thinking complexity within each treatment (precisions PR1a and PR2b) are also associated with the greatest incidence of mistakes, which negatively impact the rationality order. This aligns with our overall hypothesis that this form of information complexity plays a role in determining the displayed higher-order reasoning of agents.

Regression Analysis. To determine the effect of information complexity on the order of rationality while utilizing the within-subject design of our experiment, we estimate an

[11] Agents hardly make any mistakes in the PU treatment, which is why it is omitted from this table.

ordered logit regression, where the dependent variable is the observed rationality type (R0-R3) for each individual in each of the three treatments (i.e., three rationality types per person, for a total of 345 observations). The main independent variables are two treatment dummies — PR1 and PR2 — with the baseline category being the PU treatment. The results are presented in Table 8. In Column II we also include session dummies, with Session 1 as the baseline category.

	I	II
PR1	-2.419^{***}	-2.463^{***}
	(0.352)	(0.359)
PR2	-1.143^{***}	-1.151^{***}
	(0.338)	(0.342)
Session 2		-0.588
		(0.543)
Session 3		-0.418
		(0.545)
Session 4		-0.139
		(0.594)
Session 5		-0.833
		(0.586)
Session 6		-0.007
		(0.555)
Observations	345	345
Clusters	115	115
Pseudo R^2	0.104	0.115

Table 8: Regression Analysis

Notes: Ordered Logit Estimates. Standard errors clustered at the individual level are in parenthesis. $*p < 0.10$, $**p < 0.05$, $***p < 0.01$.

The regression coefficients reveal a significant negative relationship between the rationality type of an agent and the inclusion of private information in the information structure. Both the coefficient on the PR1 treatment dummy and that on the PR2 treatment dummy are negative and significant (at the 1% level of significance), with the former being larger in absolute value than the latter. Given the baseline treatment of public information (PU), the negative sign of the coefficients is in line with our hypothesis that the inclusion of private

information results in greater contingent thinking complexity and impacts the observed rationality level negatively. In particular, the negative sign of the coefficients shows that the probability of being classified as R0 is higher and the probability of being classified as R3 is lower in the PR1 and PR2 treatments than in the PU treatment. To determine the direction of the effects for the intermediate rationality levels, we need to consider the estimated probabilities of being classified as a given rationality type, which are presented in Table 9 using the estimated coefficients from regression I. We see that the estimated probability of being classified as R1 or R2 increases in the PR1 and PR2 treatments relative to the PU treatment. Importantly, the estimated distribution of rationality types in the PU treatment first-order stochastically dominates the estimated distributions in the PR1 and PR2 treatments, which confirms our hypothesis that the inclusion of private information negatively impacts higher-order rationality.

Probability	Treatment		
Type	PU	PR1	PR2
R0	0.01	0.07	0.02
R1	0.07	0.42	0.19
R2	0.03	0.09	0.06
R3	0.89	0.42	0.73

Table 9: Ordered Logit Estimated Probabilities Using Regression I Estimated Coefficients

Note that the absolute value of the coefficients suggests a more negative effect on the observed rationality level when private information is given only to Role 1 (PR1 treatment) than when it is given to both Roles 1 and 2 (PR2 treatment), despite the (arguably) higher contingent thinking complexity of the latter. This is also confirmed by the estimated distributions in Table 9, where the PR2 distribution first-order stochastically dominates the PR1 distribution. As discussed above, one possible explanation for this is that the differing levels of informedness between Roles 2 and 3 in the PR2 treatment, which imply that Role 2 needs to use contingent thinking less often than in the PR1 treatment. This in turn could be affecting Role 3's belief in the rationality of Role 2 in a positive way in the PR2 treatment versus the PR1 treatment, leading to the higher levels of rationality observed in the former relative to the latter.

6 Conclusion

When thinking about the complexity of a strategic environment under incomplete information, we argue it useful to consider the impact that information complexity might play in determining the displayed orders of agents' strategic sophistication. We identify one possible such channel of complexity, which has to do with the commonality of information and the contingent thinking requirements it imposes on formulating higher-order beliefs. In our experimental study we show that the complexity of the information environment can have profound effects on the outcome of the interaction and the within-subject variations in displayed higher-order rationality in strategic environments of incomplete information. Our results reveal a significant negative relationship between the rationality type of an agent, which we use as a measure of strategic sophistication, and the inclusion of private information in the information structure.

One promising direction for future research is to better understand the implications of our results for the design of optimal information structures. In particular, it would be interesting to uncover the features of the strategic environment which determine whether the designer faces a tradeoff between optimality and contingent-thinking simplicity of the information structure. To this end, a rigorous and general definition of contingent-thinking complexity would be necessary, one that applies in more general environments than the nested information structures we use in our experimental design.

While our results indicate that the inclusion of private information leads to a reduced level of strategic sophistication as measured by the higher-order rationality of agents, we think it would be both important and interesting to disentangle how much of this effect is due to intrinsic ability bounds versus the belief about other agents' abilities. Understanding the nuanced impact of information complexity on each of these channels may provide an underpinning to our finding that strategic sophistication is not necessarily monotone in the degree of information complexity beyond the initial inclusion of private information.

References

Agranov, Marina, Elizabeth Potamites, Andrew Schotter, and Chloe Tergiman. 2012. "Beliefs and Endogenous Cognitive Levels: An Experimental Study." *Games and Economic Behavior* 75:449–463. 1

Aina, Chiara, Andrea Amelio, and Katharina Brütt. 2024. "Contingent Belief Updating." Working Paper. 1

Alaoui, Larbi, Katharina A. Janezic, and Antonio Penta. 2020. "Reasoning about Others' Reasoning." *Journal of Economic Theory* 189:105091. 1

Alaoui, Larbi and Antonio Penta. 2016. "Endogenous Depth of Reasoning." *Review of Economic Studies* 83 (4):1297–1333. 1

Battigalli, Pierpaolo, Alfredo Di Tillio, Edoardo Grillo, and Antonio Penta. 2011. "Interactive Epistemology and Solution Concepts for Games with Asymmetric Information." *The B.E. Journal of Theoretical Economics (Advances)* 11:Article 6. 7, 3

Beard, T. Randolph and Richard Beil. 1994. "Do People Rely on the Self-Interested Maximization of Others? An Experimental Test." *Management Science* 40 (2):252–262. 1

Bernheim, B. Douglas. 1984. "Rationalizable Strategic Behavior." *Econometrica* 52 (4):1007–1028. 3

Börgers, Tilman and Jiangtao Li. 2019. "Strategically Simple Mechanisms." *Econometrica* 87 (6):2003–2035. 1

Brandenburger, Adam, Amanda Friedenberg, and Terri Kneeland. 2025. "From Anchors to Hierarchies: Identifying Levels of Reasoning in Games." Working paper. 1

Burchardi, Konrad B. and Stefan P. Penczynski. 2014. "Out of Your Mind: Eliciting Individual Reasoning in One Shot Games." *Games and Economic Behavior* 84:39–57. 1

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. 2004. "A Cognitive Hierarchy Model of Games." *The Quarterly Journal of Economics* 119 (3):861–898. 1

Cerigioni, Francesco, Fabrizio Germano, Pedro Rey-Biel, and Peio Zuazo-Garin. 2024. "Higher Orders of Rationality and the Structure of Games." Working Paper. 1

Charness, Gary and Dan Levin. 2009. "The Origin of the Winner's Curse: A Laboratory Study." *American Economic Journal: Microeconomics* 1 (1):207–236. 1

Costa-Gomes, Miguel, Vincent P. Crawford, and Bruno Broseta. 2001. "Cognition and Behavior in Normal-Form Games: An Experimental Study." *Econometrica* 69 (5):1193–1235. 1

Costa-Gomes, Miguel, Vincent P. Crawford, and Nagore Iriberri. 2013. "Structural Models of Nonequilibrium Strategic Thinking: Theory, Evidence, and Applications." *Journal of Economic Literature* 51:5–62. 1

Costa-Gomes, Miguel and Georg Weizsäcker. 2008. "Stated Beliefs and Play in Normal-Form Games." *Review of Economic Studies* 75:729–762. 1

Crawford, Vincent P. and Nagore Iriberri. 2007. "Level-k Auctions: Can a Nonequilibrium Model of Strategic Thinking Explain the Winner's Curse and Overbidding in Private-Value Auctions?" *Econometrica* 75 (6):1721–1770. 1

Dekel, Eddie, Drew Fudenberg, and Stephen Morris. 2007. "Interim Correlated Rationalizability." *Theoretical Economics* 2:15–40. 3

Esponda, Ignacio and Emanuel Vespa. 2014. "Hypothetical Thinking and Information Extraction in the Laboratory." *American Economic Journal: Microeconomics* 6 (4):180–202. 1

Friedenberg, Amanda and Terri Kneeland. 2024. "Beyond Reasoning about Rationality: Evidence of Strategic Reasoning." Working paper. 1

Georganas, Sotiris. 2011. "English Auctions with Resale: An Experimental Study." *Games and Economic Behavior* 73:147–166. 1

Georganas, Sotiris, Paul J. Healy, and Roberto A. Weber. 2015. "On the Persistence of Strategic Sophistication." *Journal of Economic Theory* 159:369–400. 1

Guan, Menglong and Ryan Oprea. 2024. "Three Faces of Complexity in Strategic Choice." Working Paper. 1

Healy, Paul J. 2024. "Epistemic Experiments: Utilities, Beliefs, and Irrational Play." Working paper. 1

Jin, Ye. 2021. "Does Level-*k* Behavior Imply Level-*k* Thinking?" *Experimental Economics* 24 (1):330–353. 1

Kamenica, Emir and Matthew Gentzkow. 2011. "Bayesian Persuasion." *American Economic Review* 101 (6):2590–2615. 2

Kneeland, Terri. 2015. "Identifying Higher-Order Rationality." *Econometrica* 83 (5):2065–2079. 1, 2

Li, Jiangtao and Piotr Dworczak. 2024. "Are Simple Mechanisms Optimal when Agents are Unsophisticated?" Working paper. 1

Li, Jiaying. 2025. "Complexity and Higher Order Rationality: An Experimental Study." Working Paper. 1

Li, Shengwu. 2017. "Obviously Strategy-Proof Mechanisms." *American Economic Review* 107 (11):3257–3287. 1

Liu, Shuige and Gabriel Ziegler. 2025. "Reasoning about Bounded Reasoning." Working paper. 1

Mathevet, Laurent, Jacopo Perego, and Ina Taneva. 2020. "On Information Design in Games." *Journal of Political Economy* 128 (4):1370–1404. 3, 2

Nagel, Lea and Roberto Saitto. 2025. "A Measure of Complexity for Strategy-Proof Mechanisms." Working paper. 1

Nagel, Rosemarie. 1995. "Unraveling in Guessing Games: An Experimental Study." *The American Economic Review* 85 (5):1313–1326. 1

Ngangoué, M. Kathleen and Georg Weizsäcker. 2021. "Learning from Unrealized versus Realized Prices." *American Economic Journal: Microeconomics* 13 (2):174–201. 1

Niederle, Muriel and Emanuel Vespa. 2023. "Cognitive Limitations: Failures of Contingent Thinking." *Annual Review of Economics* 15:307–328. 1

Oprea, Ryan. 2020. "What Makes a Rule Complex?" *American Economic Review* 110 (12):3913–3951. 1

Pearce, David G. 1984. "Rationalizable Strategic Behavior and the Problem of Perfection." *Econometrica* 52:1029–1050. 3

Puri, Indira. 2025. "Simplicity and Risk." *The Journal of Finance* 80 (2):1029–1080. 1

Pycia, Marek and Peter Troyan. 2023. "A Theory of Simplicity in Games and Mechanism Design." *Econometrica* 91 (4):1495–1526. 1

Rey-Biel, Pedro. 2009. "Equilibrium Play and Best Response to (Stated) Beliefs in Normal Form Games." *Games and Economic Behavior* 65 (2):572–585. 1

Schotter, Andrew, Keith Weigelt, and Charles Wilson. 1994. "A Laboratory Investigation of Multiperson Rationality and Presentation Effects." *Games and Economic Behavior* 6:445–468. 1

Tan, Tommy Chin-Chiu and Sérgio Ribeiro da Costa Werlang. 1988. "The Bayesian Foundations of Solution Concepts of Games." *Journal of Economic Theory* 45 (2):370–391. 3

Tirole, Jean. 1986. "Procurement and Renegotiation." *Journal of Political Economy* 94 (2):235–259. 4

Van Huyck, John, John Wildenthal, and Ray Battalio. 2002. "Tacit Coordination games, Strategic Uncertainty, and Coordination Failure: Evidence from Repeated Dominance Solvable Games." *Games and Economic Behavior* 38:156–175. 1

Appendix

Appendix A: Instructions

Welcome to the experiment!

Welcome to the experiment.
Please wait for instructions and do not touch the computer until you are instructed to do so. Please put away all personal belongings and silence your phone. We need your full attention for the entire experiment.

Instructions

The experiment you will be participating in today is an experiment in decision making. At the end of the experiment you will be paid for your participation in cash. Each of you may earn different amounts. The amount you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. You will be using the computer for the experiment, and all choices will be made through the computer. DO NOT socialize or talk during the experiment.

If you have any questions, please raise your hand and your question will be answered out loud so everyone can hear.

No Deception

All instructions and descriptions that you will be given in this experiment are accurate and truthful. The instructions are intended to convey a clear and precise description of the choices we will ask you to make. Our interest is in observing how people make decisions. In accordance with the policy of this lab, at no point will we attempt to deceive you in any way.

The Experiment

We will now explain the basic idea of the experiment. You will be randomly placed into groups of 3 through the computers and asked to make a set of three decisions. You will

then be randomly placed into a new 3 person group and will make a new set of decisions, and this process will repeat. Each time you are regrouped, the decisions will be similar but the setting will change, so that you will be in a new version of the decision making environment. It is important to always pay attention to the details of the current version when making your decisions.

In each version, you will be randomly matched in a group with two other participants currently in the room. The people in your group will be assigned three roles, which we call simply Role 1, Role 2 and Role 3. Each one of you will play once in each role. You will always know which role you are currently in, but the people in the other two roles will always remain anonymous.

Each decision works in the following way. The computer screen will show you several Jars. Each Jar contains a certain number of Green and Purple balls. The computer will choose a Jar and then draw a ball from that Jar. Different Roles may be given different information about which Jar is chosen, in ways that we will explain later. We will always be explicit and accurate when describing what information is given to the different roles. Then, before seeing which color ball is drawn, each person in your group will make their choice.

Your earnings are described in the tables on your screen. The earnings for Role 1 are determined by her choice and the color of the ball. These are displayed in the top table. If the ball is **P**urple, Role 1 earns 1 point by choosing e and 0 points by choosing f. If the ball is **G**reen, then Role 1 earns 1 point by choosing f and 0 points by choosing e. So, to maximize the chance of earning a point as Role 1, you should choose e if you think a Purple ball is more likely and you should choose f if you think a Green ball is more likely.

Next, the earnings for Role 2 depend on her choice and the choice by Role 1. These are given by the middle table. If Role 1 chooses e then Role 2 earns 1 point by choosing j and 0 points by choosing k. If Role 1 chooses f then Role 2 earns 1 point by choosing k and 0 points by choosing j. So, to maximize the chance of earning a point as Role 2, you should choose j if you think Role 1 is more likely to choose e than f, and you should choose k if you think Role 1 is more likely to choose f than e.

Finally, the earnings for Role 3 depend on her choice and the choice by Role 2. These are given by the bottom table. If Role 2 chooses j, then Role 3 earns 1 point by choosing s and 0 points by choosing t. If Role 2 chooses k then Role 3 earns 1 point by choosing t and zero points by choosing s. So, to maximize the chance of earning a point as Role 3,

you should choose s if you think Role 2 is more likely to choose j than k, and you should choose t if you think Role 2 is more likely to choose k than j.

Role 1

	P	**G**
e	1	0
f	0	1

Role 2

	e	f
j	1	0
k	0	1

Role 3

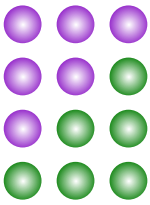
	j	k
s	1	0
t	0	1

These tables will always be listed in the same order and will appear on the righthand side of your computer screen for your reference.

Jars

So far, we have described how your earnings depend on the color of the ball and the choices of your group members. That structure will remain the same throughout the entire experiment. Also, throughout the experiment, the computer will always draw a ball from 1 of 4 jars. These jars are displayed on your screen and are labeled A, B, C, and D. As you can see, Jar A contains 3 Purple balls, Jar B contains 2 Purple and 1 Green ball, Jar C contains 1 Purple and 2 Green balls, and Jar D contains 3 Green balls.

Jar A:
Jar B:
Jar C:
Jar D:



So, for example, if the Jar is A, then the ball will be Purple with 100% chance. If the Jar is B, then there is a 2 out of 3 chance the ball will be Purple, and a 1 out of 3 chance the ball will be Green, and so forth.

Even though the 4 Jars always stay the same, you will not always know exactly which Jar the computer is drawing from, and the different Roles may have different information

about which Jar is being used. Before you make any choices, we will provide a complete description of how the computer chooses the Jar, and what information you and your group members are given about that jar.

Structure

In each group you will make 3 choices: once as Role 1, once as Role 2, and once as Role 3. For each choice the computer randomly draws a new ball from one of the Jars. So the color of the ball may change from choice to choice.

After you make a choice in each of the 3 Roles, you will be given a screen that allows you to revise any of your choices. Once you are satisfied with your choices, you can confirm your final choices. Once that is done, you will be randomly regrouped before proceeding with the next set of decisions.

Earnings

You will earn a show-up payment of $5 for arriving to the experiment on time. In addition to the show-up payment, at the conclusion of the experiment, once you have made all your choices, the computer will randomly select exactly one of these choices and your additional payment will be based on the outcome of this one choice. Specifically, a point is worth $20. Thus if you earned 1 point in the randomly selected decision, you will earn a total of $25 ($5+$20). If you earned zero points in that decision, you will earn a total of $5.

Since any of your choices could be the one selected at the end, you should treat each choice as if it is the one determining your earnings. For the selected choice, you will receive $20 if you earned a point in that choice.

Feedback

The identity of your randomly chosen group members will not be revealed, and so your choices are kept anonymous.

During the experiment, we will summarize the outcomes at various stages to help you keep track of what has happened. However, we will not reveal information about other people's choices until the end of the experiment. The software records all choices and outcomes as

they happen and all of this information will be presented in full on your computer screen before the end of the experiment.

Once you see all of the choice data from your groups in the experiment, the computer will then reveal which choice was randomly chosen for your payment, what the outcome in your group was for that choice, and your resulting earnings. You will be given information only about the groups in which you were a member.

We will now start the software.
Please click on the OK button at the bottom of your screen when it appears in a moment.
Then wait for further instructions before proceeding.

[[Always call this version 1, but insert PU, PR1, or PR2 as appropriate.]]

In Version 1,

.....**PU treatment**.....

all roles will have exactly the same information. Specifically, all three roles will be told which jar the computer is going to use. For example, suppose you are told that the current jar is Jar B (containing 2 purple (**P**) balls and one green (**G**) ball) and that your current role is Role 2. Then you know that the computer will draw from Jar B and you also know that the people in Role 1 and Role 3 will also have the same information and know that the current jar is Jar B.

Remember, every time you make a choice, the computer draws a ball from a jar that contains all three balls.

.....**end PU treatment**.....

Are there any questions [pause]?

We are now going to begin the first part of the experiment. Remember to pay close attention to the information on your screen. Take your time and enter your choices carefully. To begin, please click the OK button at the bottom of your screen and then follow the prompts.

.....

We have now completed the first version and I will now describe the second version. Please pay careful attention as the differences between the versions are important to understand. Do not click OK until instructed.

[[Always call this version 2, but insert PU, PR1, or PR2 as appropriate.]]

In Version 2,

.....**PR1 treatment**.....

there will be a Box with cards. Each card will be labeled either A, B, C or D. There may be more than one card with the same letter in the Box, and there may not be a card with each letter. Everyone in your group will be told exactly which cards are in the Current Box. The computer will then randomly draw a card from the Box. The letter on that card will determine which Jar the computer will use for the current decision. For example, if the card drawn has a letter A on it, this means that the ball will be drawn from Jar A, and similarly for all other letters.

However, and this very important, only the person in Role 1 will observe which card is drawn. So only Role 1 will know for sure which jar is going to be used. The people in Roles 2 and 3 will only know which cards are in the current Box; they will not see which card is drawn.

For example, suppose the Current Box contains one C card and one D card. Then all three group members are given this information, and so each person knows the current box contains one C card and one D card. The computer will then draw a card from that box. If your current role is Role 1, you will observe the drawn card: if the card drawn is C, you will know that the jar is going to be Jar C; if the card drawn is D, you will know that the jar is going to be Jar D. However, if your current role is Role 2 or Role 3, you will not observe the draw of the card. But notice that you do know that the current jar is going to be either Jar C or Jar D, each with equal chance, because there is one C card and one D card in the box.

Remember, every time you make a choice, the computer will draw a new card from the Current Box, so the Jar can change for each choice you make. And for each choice, the computer draws a new ball from a jar that contains all three balls.

.....**end PR1 treatment**.....

Are there any questions [pause]?

As before, remember to pay close attention to the information on your screen. Take your time and enter your choices carefully. Now please click the OK button at the bottom of your screen to proceed with the next part of the experiment.

.....

We have now completed the second version and I will now describe the third version. Again, please pay careful attention as the differences between the versions are important to understand. Do not click OK until instructed.

[[Always call this Version 3, but insert PU, PR1, or PR2 as appropriate.]]

In Version 3,

.....**PR2 treatment**.....

there will be two boxes filled with cards — we call these the Primary Box and the Secondary Box.

Each card will be labeled either A, B, C, D or X. There may be more than one card with the same letter in the Box, and there may not be a card with each letter. The Secondary Box will not contain any X cards. Everyone in your group will be told exactly which cards are in both Boxes.

The computer will then randomly draw a card from the Primary Box. However, and this is very important, only Role 1 and Role 2 will observe which card is drawn from the Primary Box; Role 3 will not observe the draw.

If the card has an A, B, C, or D, then that card will determine the jar for the next decision and both Role 1 and Role 2 will know which jar is going to be used. For example, if the card drawn has a letter A on it, this means that the ball will be drawn from Jar A, and similarly for all other letters.

But, if the card has an X on it, then the computer will randomly draw a card from the Secondary Box to determine the draw. In this case,and again this is very important, only Role 1 will observe the card from the Secondary Box; Role 2 and Role 3 will not observe the draw.

So in the case that the card from the Primary Box is X, then Role 1 will see the card from the Secondary Box and so Role 1 will know which Jar is going to be used. Role 2 will know that the first card was an X, and so Role 2 knows that the Jar will be determined by the draw from the Secondary Box, but will not observe that draw. Role 3 will not observe either card draws.

For example, suppose that the current Primary Box contains one B and one X card, and the current Secondary Box contains one A card and two C cards. Then all 3 group members will be given this information. The computer will then draw a card from that Primary Box. If your current role is Role 1 or Role 2, you will observe the card drawn from the Primary

Box. If the card drawn from the Primary Box is, for example, B, and your current role is Role 1 or Role 2, then you will know that the current jar is going to be Jar B.

If your current role is Role 3, even though you won't observe any card draws, you will know that the current jar is Jar A, Jar B, or Jar C; it cannot be Jar D because neither box contains a D card. If the card drawn from the Primary Box is X, then the person in Role 1 will observe the draw from the Secondary Box. If the card drawn from the Secondary Box is A (C), and your current role is Role 1, you will know that the current jar is Jar A (C); if your current role is Role 2, you will know that the current jar is either Jar A or Jar C, as determined by the draw from the Secondary Box; again if your current role is Role 3, you will know that the current jar is either Jar A, Jar B, or Jar C, as determined by the draw(s) from the box(es).

Remember, every time you make a choice, the computer will draw a new card from the Current Primary Box and, if necessary, a new card from the current Secondary Box, so the Jar can change for each choice you make. And for each choice, the computer draws a new ball from a jar that contains all three balls.

.....**end PR2 treatment**.....

Are there any questions [pause]?

As before, remember to pay close attention to the information on your screen. Take your time and enter your choices carefully. Now please click the OK button at the bottom of your screen to proceed with the last part of the experiment.

.....

We have now completed the experiment. The next 3 screens will summarize all of your choices, the choices of others, and all of the cards and balls that were drawn. You will now be able to see your points earned for each choice. Please take as much time as you like to review this information. Proceed through the summary until you reach the last page, which will tell you which choice was randomly selected and your corresponding earnings. Please record your ID number and your earnings information on your payment receipt. We need to collect your signed receipts before you leave. We will now call you one at a time to the front of the lab to receive your payments. Please collect your belongings as you leave your desks.