Uncommon Belief in Rationality

Qi Shi, Pavel Naumov

University of Southampton, UK

qi.shi@soton.ac.uk; p.naumov@soton.ac.uk

Abstract

Common knowledge/belief in rationality is the traditional standard assumption in analysing interaction among agents. This paper proposes a graph-based language for capturing significantly more complicated structures of higher-order beliefs that agents might have about the rationality of the other agents. The two main contributions are a solution concept that captures the reasoning process based on a given belief structure and an efficient algorithm for compressing any belief structure into a unique minimal form.

Full version — https://doi.org/10.48550/arXiv.2412.09407

1 Introduction

In the orthodox studies of game theory and game-modelled multiagent systems, the rationality of agents is usually assumed to be *common knowledge* (Aumann 1976). Albeit called "knowledge", it does not have to be the case. This is because, as defined in epistemology, knowledge is something that must be true and justifiable (Steup and Neta 2024). However, from the perspective of each single agent, it is hard to verify that the other agents are indeed rational. As discussed by Lewis (1969), what really matters in the reasoning process is the agents' rationality and ***belief*** about the other agents' rationality, the latter of which, unlike *knowledge*, is not necessarily true or justifiable. Indeed, in the discussion of epistemic game theory (Dekel and Siniscalchi 2015), the assumption of ***rationality and common belief in rationality (RCBR)*** serves as the foundation of the major solution concepts such as Nash equilibrium (Nash 1950), correlated equilibrium (Aumann 1987; Brandenburger and Dekel 1987), and rationalisability (Pearce 1984; Bernheim 1984).

As Lewis (1969) and Schiffer (1972) interpret, RCBR consists of the rationality of all agents and a belief hierarchy that contains all finite sequences in the form that "a_1 believes that a_2 believes … that a_{i-1} believes that a_i is rational", where a_1, a_2, …, a_{i-1}, a_i are (possibly duplicated) agents. Observe that RCBR can be expressed with a complete digraph. For instance, Figure 1(i) illustrates the RCBR among agents a, b, c, where each node represents the

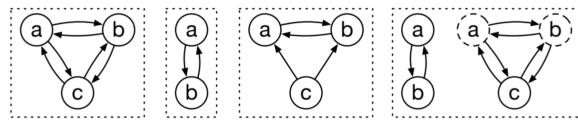

(i) RCBR (ii) (iii) (iv)

Figure 1: Different RBR graphs among agents a, b, c.

rationality of an agent and each path[1] of at least two nodes corresponds to a sequence in the belief hierarchy of RCBR. Specifically, in Figure 1(i), the node labelled with a corresponds to the rationality of agent a; the path labelled with (b, c, b, a) corresponds to the belief sequence "agent b believes that agent c believes that agent b believes that agent a is rational". We denote this belief sequence by the tuple (b, c, b, a) henceforth. Notice that no self-loop exists due to the assumption that *agents do not have introspective beliefs* about their own rationality (Dekel and Siniscalchi 2015).

Yet, the RCBR assumption is too strong, especially in a system consisting of different types of agents. For example, when two adults a and b interact with a child c, who is too young to possess rationality (*i.e.* the ability to do mathematical optimisation), the RCBR assumption is satisfied only between a and b. Then, the belief hierarchy can be captured by Figure 1(ii). Notice that, the irrational agent c is not included in the graph because of the assumption that *irrational agents act arbitrarily and do not possess beliefs about the other agents' rationality*.[2] This assumption also implies that a belief hierarchy about rationality should be "prefix-closed". In other words, if agent a believes that agent b believes that agent c is rational, then agent a must believe that agent b is rational, the latter of which further implies that agent a must be rational. This property makes it possible to illustrate a belief hierarchy with a digraph.

Now we consider the situation that the child c, although young, is a genius. The adults a and b neglect the talent of c. But c not only is rational and believes in the RCBR between

[1] A path refers to a nonempty sequence of connected and possibly duplicated nodes. Hence, a single node forms a singleton path.

[2] Irrational agents' beliefs about others' rationality, even if exist, do not influence strategic behaviours since irrational agents always act arbitrarily. Thus, such beliefs can be disregarded safely when studying the strategic behaviours in a multiagent system.

a and b, but also notices the arrogance of a and b. In this case, there is no RCBR among agents a, b, c, but there is more than RCBR between agents a and b. We illustrate the ***rationality and beliefs in rationality (RBR)*** in this system in Figure 1(iii). We shall call it the "RBR graph" among agents a, b, c. In this system, compared with the RCBR among a, b, c, there is no belief about the rationality of agent c. Meanwhile, compared with the RCBR between a and b, there are agent c's beliefs about the rationality and beliefs of a and b.

Let us now consider another situation that, in the above case, the genius child c fails to identify the arrogance of a and b. Then, in c's belief, the RCBR among a, b, c still exists, which, however, is not the reality. In other words, the agents a and b in c's belief are not the real agents a and b in the system. We call such agents by ***doxastic agents***, who exist in beliefs but are not real. In this case, we visualise the doxastic agents with dashed cycles as Figure 1(iv) shows. In this figure, the solid nodes labelled with a and b represent the real agents a and b between whom the RCBR exists, while the dashed nodes labelled with a and b represent the doxastic agents a and b, with whom agent c believes the RCBR exists. In this case, only the solid nodes represent the rationality of agents; only the paths starting at solid nodes (*i.e.* real agents) correspond to belief sequences in the system.

In this paper, we encode RBR systems with *directed labelled graphs*, as shown in Figure 1, based on which we study the agents' strategic behaviours. Notably, the existing literature on agents' beliefs concentrates on modelling the Bayesian beliefs of agents (*i.e.* agents' subjective probability distributions over the other agents' behaviours). On the contrary, the rationality of agents implies a dynamic "best-response" process (Savage 1954) that cannot be captured by static distributions. Particularly, the *type structure* (Harsanyi 1967; Brandenburger and Dekel 1993) is used in epistemic game theory to implicitly model the hierarchy of Bayesian beliefs. It is a recursive formalisation of strategic behaviours based on the RCBR assumption, rather than a formalisation of beliefs in rationality. Meanwhile, the *influence diagram* (Koller and Milch 2003; Howard and Matheson 2005), a variant of the Bayesian network (Pearl 1988), is used in AI to model a decision system of agents. Due to its strength in modelling the probabilistic uncertainty, the influence diagram is also used to model Bayesian beliefs of agents (Milch and Koller 2000). However, the acyclic nature of the influence graph makes it impossible to model RCBR as the RBR graph does. More literature on modelling agents' beliefs can be seen in the review papers (Albrecht and Stone 2018) and (Doshi, Gmytrasiewicz, and Durfee 2020). As far as we know, no research on explicitly modelling uncommon RBR (*i.e.* not RCBR) of agents exists in the literature.

Contribution and Outline We first discuss how RBR works in agents' strategic reasoning process in Section 2. Then, in section 3, we formally define *RBR graphs*, a graph-based language to capture RBR systems, and *(iterative) rationalisation*, the strategic reasoning process of agents, based on which we propose *doxastic rationalisability*, a solution concept in uncommon RBR systems. After this, we discuss the equivalence of RBR graphs in Section 4, based on which we design an algorithm that can compress any RBR graph to a minimal equivalent form in Section 5. Due to page limits, we retain formal definitions, theorems, and some intuitive discussions in the main text, while relegating all supporting lemmas and tedious formal proofs to the appendix of the full version (Shi and Pavel 2024).

2 Rationality and Rationalisation

As defined by Savage (1954), a rational agent, when faced with uncertainty, first forms a *subjective* probability distribution over all possibilities, and then chooses a strategy that *best responds* (*i.e.* maximises the expected utility) to the subjective probability distribution. This is the commonly acknowledged definition of ***rationality*** in economics, game theory, and multiagent systems. ***Rationalisability*** (Pearce 1984; Bernheim 1984) is a solution concept to "what RCBR exactly implies". In this concept, a strategy is *rationalisable* if it best responds to some subjective probability distribution. Rational agents take only rationalisable strategies. We call the process of finding rationalisable strategies ***rationalisation***. Note that, in a game with compact strategy sets and continuous utility functions, a (mixed) strategy is rationalisable if and only if it is not *strongly dominated*[3] (Gale and Sherman 1950; Pearce 1984). In this situation, rationalisation exactly means the elimination of strongly dominated strategies. In other cases, the elimination of strongly dominated strategies implies[4] rationalisation but not the other way around. For example, Börgers (1993) finds that, if the preference of agents on outcomes is a total order rather than defined as utility functions, then rationalisation is a concept weaker than "eliminating strongly dominated strategies" but stronger than "eliminating weakly dominated strategies".

Since our purpose in this paper is not to discuss the essence of rationality, we simplify the definition of rationalisation to ***the elimination of strongly dominated strategies in pure strategy space***. We say that a rational agent chooses only pure strategies that are *not* strongly dominated by any pure strategy (see Section 3 for formal definitions). The benefit is threefold: first, it avoids the computational intractability in dealing with mixed strategy space and "best response" optimisation; second, it allows us to discuss a more general game frame where the preference of agents is just partial order; third, the simplified definition is more restrictive (but not too much) than those in the literature, so the technical results of this paper can potentially be extended to more general definitions using similar proof techniques but at the expense of mathematical complexity. Henceforth, by saying a strategy is dominated, we mean it is strongly dominated.

To see how rationality interacts with belief, let us consider the following simplified "guess $2/3$ of the average" game (Moulin 1986; Nagel 1995):

> *Each agent chooses an* ***integer*** *in the interval* $[1, 10]$. *The one whose choice is the closest to* $2/3$ *of the average of the* ***other*** *agents' choices wins.*

[3] As standard expressions in game theory, a strategy is strongly dominated if always worse than another, and weakly dominated if never better and sometimes worse.

[4] In the sense that a strongly dominated strategy is never rationalisable (Pearce 1984).

Informally, we say that $2/3$ of the average of the other agents' choices is the *target* which every agent in this game aims to approach. Suppose the agents a, b, c as described in Section 1 play the above game. We analyse the potential choices of each agent in each of the four RBR systems illustrated with the RBR graphs in Figure 1.

We first consider the orthodox case that RCBR exists among agents a, b, c, as depicted in Figure 1(i). According to the description of the game, agent a knows that the choices of agents b and c, denoted by χ_b and χ_c henceforth, lie in the interval $[1, 10]$. Then, the target t_a of agent a lies in the interval $[\frac{2}{3} \cdot \frac{1+1}{2}, \frac{2}{3} \cdot \frac{10+10}{2}] = [\frac{2}{3}, 6\frac{2}{3}]$. Thus, 7 is always closer to the target than any integer greater than it (*i.e.* $8, 9, 10$ are dominated by 7), while every integer in the interval $[1, 7]$ might be the closest to the target. Hence, the rationalisable choice χ_a of agent a should be in the interval $[1, 7]$. Note that, the above analysis works for agents b and c in a symmetric way. Therefore, the rationalisable choices are $\chi_a, \chi_b, \chi_c \in [1, 7]$ after one rationalisation. We mark this observation in column 1^{st} of lines (1-3) in Table 1.

Note that, the RCBR assumption implies the belief sequences (a, b) and (a, c), which represent that "agent a believes that agent b is rational" and "agent a believes that agent c is rational", respectively. Then, agent a, after the above analysis, believes the choices $\chi_b, \chi_c \in [1, 7]$. Thus, a preciser target of agent a becomes $t_a \in [\frac{2}{3}, 4\frac{2}{3}]$. Hence, the rationalisable choice of agent a is $\chi_a \in [1, 5]$ after the second rationalisation. Similarly, $\chi_b \in [1, 5]$ follows from the belief sequences (b, a) and (b, c) and $\chi_c \in [1, 5]$ follows from the belief sequences (c, a) and (c, a). We mark this observation in column 2^{nd} of lines (1-3) in Table 1. Observe that, the RCBR assumption also implies the belief sequences (a, b, c), (a, b, a), (a, c, b), and (a, c, a). Then, agent a, after the above analysis, believes the choices $\chi_b, \chi_c \in [1, 5]$. Hence, $\chi_a \in [1, 3]$ after the third rationalisation, similar for χ_b and χ_c. We mark this observation in column 3^{rd} of lines (1-3) in Table 1. Following the same process, after the fourth rationalisation, $\chi_a, \chi_b, \chi_c \in [1, 2]$ and, after the fifth rationalisation, $\chi_a, \chi_b, \chi_c = 1$. Then, more rationalisation will not eliminate more strategies. We mark these observations in the corresponding columns of lines (1-3) in Table 1. Consequently, given the RCBR assumption, every agent should choose 1. The above analysis follows the standard approach in epistemic game theory (Perea 2012, Example 3.7).

Now we look at the non-trivial cases where RCBR does not exist. Suppose the RBR among agents a, b, c is as Figure 1(ii) shows. Since agent c is irrational, her choice is always $\chi_c \in [1, 10]$, as depicted in line (6) of Table 1. For agents a and b, the first rationalisation is identical to the case with RCBR. That is $\chi_a, \chi_b \in [1, 7]$, as depicted in column 1^{st} of lines (4-5) in Table 1. Next, note that the belief sequence (a, b) is in the RBR system but (a, c) is not. Then, after the above analysis, agent a believes that $\chi_b \in [1, 7]$ and $\chi_c \in [1, 10]$. Thus, the target of agent a is $t_a \in [\frac{2}{3}, 5\frac{2}{3}]$. Hence, the rationalisable choice of agent a is $\chi_a \in [1, 6]$ after the second rationalisation. Symmetrically, $\chi_b \in [1, 6]$ follows from the belief sequence (b, a) after the second rationalisation. These results are marked in column 2^{nd} of lines (4-5) in Table 1. Then, due to the belief sequences (a, b, a) and (b, a, b), the above analysis implies that the preciser targets of agents a and b are $t_a, t_b \in [\frac{2}{3}, 5\frac{1}{3}]$. Thus, $\chi_a, \chi_b \in [1, 5]$ after the third rationalisation, as marked in column 3^{rd} of lines (4-5) in Table 1. It can be verified that more rationalisation will not eliminate more strategies. Consequently, the choices are $\chi_a, \chi_b \in [1, 5]$ and $\chi_c \in [1, 10]$ in the RBR system depicted in Figure 1(ii).

As the above two cases show, in a given game, RBR implies an *iterative* rationalisation process until a stable state where no more strategy can be eliminated is reached. In essence, the iterative process relies on longer and longer belief sequences. However, since we depict an RBR system with a digraph, if there is an edge from a node labelled with a to a node labelled with b, then, for each belief of agent b, agent a believes that agent b holds this belief. For instance, in Figure 1(i), for the agent b's belief sequence (b, c, a), agent a holds the belief sequence (a, b, c, a) and agent c holds the belief sequence (c, b, c, a). This property allows us to *do iterative rationalisation without explicitly considering the belief sequences.* Instead, if an edge from a node labelled with agent a to a node labelled with agent b exists in an RBR-graph, then agent a does the i^{th} rationalisation based on the result of agent b's $(i-1)^{th}$ rationalisation. In this sense, even if agent b is a doxastic agent, she is treated the same way as a real agent. For convenience of expression, we also refer to the belief sequences of doxastic agents as belief sequences in the RBR system. This is in line with the model of *iterated strategic thinking* (Binmore 1988), which formalises the intuition that "the natural way of looking at game situations ... is not based on circular concepts, but rather on a step-by-step reasoning procedure" (Selten 1998).

It is interesting to observe that, in the RBR system depicted in Figure 1(iii), agents a and b hold the same belief hierarchy as in Figure 1(ii). In other words, agents a and b cannot distinguish the RBR systems in Figure 1(ii) and Figure 1(iii). As a result, the iterative rationalisation process of agents a and b runs identically in Figure 1(ii) and Figure 1(iii), as lines (4-5) in Table 1 show. Contrarily, agent c in Figure 1(iii), different from agent c in Figure 1(ii), is rational and holds a belief hierarchy. In the RBR system depicted in Figure 1(iii), the first rationalisation of agent c, which relies only on the rationality of c, works in the same way as the case in Figure 1(i). From the second rationalisation, agent c does the i^{th} rationalisation based on the results of the $(i-1)^{th}$ rationalisation. We record the process in line (7) of Table 1. For example, in the fourth rationalisation, based on

		1^{st}	2^{nd}	3^{rd}	4^{th}	5^{th}	6^{th}	...
(1)	a	$[1,7]$	$[1,5]$	$[1,3]$	$[1,2]$	1	1	...
(2)	b	$[1,7]$	$[1,5]$	$[1,3]$	$[1,2]$	1	1	...
(3)	c	$[1,7]$	$[1,5]$	$[1,3]$	$[1,2]$	1	1	...
(4)	a	$[1,7]$	$[1,6]$	$[1,5]$	$[1,5]$	...		
(5)	b	$[1,7]$	$[1,6]$	$[1,5]$	$[1,5]$	...		
(6)	c	$[1,10]$	...					
(7)	c	$[1,7]$	$[1,5]$	$[1,4]$	$[1,3]$	$[1,3]$	...	

Table 1: Rationalisations on the RBR graphs in Figure 1.

the analysis of the third rationalisation, agent c believes that the choices $\chi_a, \chi_b \in [1, 5]$, as marked in column 3^{nd} of lines (4-5) in Table 1. Thus, the target $t_c \in [\frac{2}{3}, 3\frac{1}{3}]$. Hence, the rationalisable choice is $\chi_c \in [1, 3]$, as depicted in column 4^{th} of line (7) in Table 1. Also observe that, in the RBR system depicted in Figure 1(iv), agents a and b hold the same belief as in Figure 1(ii), while agent c holds the same belief as in Figure 1(i). As a result, the iterative rationalisation process in Figure 1(iv) is as lines (3-5) of Table 1 show.

It is worth mentioning that, our RBR-graph can be used to model *bounded rationality* (Simon 1955, 1957) of agents. In particular, the *cognitive hierarchy model* (Camerer, Ho, and Chong 2004; Chong, Ho, and Camerer 2016) is a mathematical model of bounded rationality that limits the length of belief sequences in a belief hierarchy. The model assumes every agent believes herself to be smarter than everyone else and thus performs deeper reasoning. Particularly, an i-step reasoner believes that every other agent is a j-step reasoner with some probability, where $j < i$, and a 0-step reasoner is an irrational agent. Because of this assumption, the statement that "an agent cannot reason too deeply" is reduced to the statement that "the agent believes that the other agents cannot reason too deeply". The former is the essence of bounded rationality, while the latter makes it possible to depict a bounded rationality RBR system with an acyclic RBR-graph. Moreover, the iterative rationalisation process in bounded rationality RBR graphs is exactly the step-by-step reasoning process in the cognitive hierarchy model. However, the cognitive hierarchy model works implicitly in a probabilistic approach and is not used to model the explicit RBR systems as the RBR graphs do.

3 Terminology and Solution Concept

In preparation for the study of the iterative rationalisation process in games among agents with uncommon RBR, we next formalise the concepts informally introduced in the previous sections. Throughout this paper, unless specified otherwise, we assume a fixed nonempty set $\mathcal{A}$ of agents. We start with a general definition of (strategic) games that uses partial orders (Osborne and Rubinstein 1994).

Definition 1 *A* ***game*** *is a tuple* $(\Delta, \preceq)$ *such that*

1. $\Delta = \{\Delta_a\}_{a\in\mathcal{A}}$*, where* $\Delta_a \neq \varnothing$ *is a finite strategy space for each agent* $a \in \mathcal{A}$*;*
2. $\preceq = \{\preceq_a\}_{a\in\mathcal{A}}$*, where* $\preceq_a$ *is a partial order on the Cartesian product* $\prod_{b\in\mathcal{A}} \Delta_b$*.*

An element $s_a \in \Delta_a$ is called a *strategy* of agent a. An *outcome* is a tuple $\boldsymbol{s} \in \prod_{a\in\mathcal{A}} \Delta_a$ consisting of a strategy for each agent $a \in \mathcal{A}$. Binary relation $\preceq_a$ shows agent a's preference over the outcomes. For two outcomes $\boldsymbol{s}$ and $\boldsymbol{s}'$, if $\boldsymbol{s} \preceq_a \boldsymbol{s}'$ and $\boldsymbol{s}' \npreceq_a \boldsymbol{s}$, then we write $\boldsymbol{s} \prec_a \boldsymbol{s}'$ and say that agent a *strictly prefers* outcome $\boldsymbol{s}'$ to $\boldsymbol{s}$. For example, in our simplified "2/3 game" in Section 2, the strategy space of each agent is all integers in the interval $[1, 10]$; an outcome is a collection of every agent's choice. An agent strictly prefers the outcomes where her choice is closer to her target.

Note that, the commonly used definition of games where preference is defined via utility functions is a special case of Definition 1 in which preference $\preceq_a$ is a total order for each agent $a \in \mathcal{A}$. In particular, $\boldsymbol{s} \preceq_a \boldsymbol{s}'$ if $u_a(\boldsymbol{s}) \leq u_a(\boldsymbol{s}')$, where $u_a(\boldsymbol{s})$ and $u_a(\boldsymbol{s}')$ are utilities of agent a toward outcome $\boldsymbol{s}$ and $\boldsymbol{s}'$, respectively.

Definition 2 *For any agent* $a \in \mathcal{A}$*, a* ***reasoning scene*** Θ_a *in the game* $(\Delta, \preceq)$ *is a Cartesian product* $\prod_{b\in\mathcal{A}\setminus\{a\}} \Theta_a^b$ *such that* $\varnothing \subsetneq \Theta_a^b \subseteq \Delta_b$ *for each agent* $b \neq a$*.*

A reasoning scene describes a static context in which a rational agent rationalises. It captures the uncertainty of an agent toward the other agents' strategies. In other words, given that every other agent b may choose a strategy from set Θ_a^b, agent a reasons about which strategies of her own are rationalisable. Recall that, when agents a, b, c with RBR in Figure 1(ii) play the simplified "2/3 game", in the third rationalisation, agent a believes $\chi_b \in [1, 6]$ and $\chi_c \in [1, 10]$, as column 2^{nd} of lines (5-6) in Table 1 shows. In this situation, we say that agent a rationalises in the reasoning scene Θ_a such that Θ_a^b consists of all integers in the interval $[1, 6]$ and Θ_a^c consists of all integers in the interval $[1, 10]$. A tuple $\boldsymbol{s}_{-a} \in \Theta_a$ is a combination of all agents' strategies except agent a and thus $(\boldsymbol{s}_{-a}, s_a)$ is an outcome.

The next definition formalises the notion of dominance as discussed in Section 2.

Definition 3 *For a given reasoning scene* Θ_a *of agent* a *and any strategies* $s_a, s'_a \in \Delta_a$*, strategy* s'_a ***dominates*** *strategy* s_a *if* $(\boldsymbol{s}_{-a}, s_a) \prec_a (\boldsymbol{s}_{-a}, s'_a)$ *for each tuple* $\boldsymbol{s}_{-a} \in \Theta_a$*.*

We write $s_a \lhd_{\Theta_a} s'_a$ if strategy s_a is dominated by strategy s'_a in the reasoning scene Θ_a. Note that, dominance relation is asymmetric (*i.e.* if $s_a \lhd_{\Theta_a} s'_a$, then $s'_a \not\lhd_{\Theta_a} s_a$).

The next definition formalises the result of rationalisation in a given reasoning scene. It is in line with our simplified definition of rationalisation (*i.e.* the elimination of dominated strategies) as discussed in Section 2.

Definition 4 *In reasoning scene* Θ_a*, the* ***rational response*** *of agent* a *is a set of strategies*

$$\Re_a(\Theta_a) := \{s_a \in \Delta_a \mid \neg\exists s'_a \in \Delta_a(s_a \lhd_{\Theta_a} s'_a)\}.$$

In other words, the set $\Re_a(\Theta_a)$ consists of all rationalisable (*i.e.* not dominated) strategies of agent a in the reasoning scene Θ_a. Note that, $\Re_a(\Theta_a) \neq \varnothing$ for any agent a and any reasoning scene Θ_a due to the asymmetry of dominance relation $\lhd_{\Theta_a}$. Next, we formally define RBR graphs. In the rest of the paper, notation nEm is short for $(n, m) \in E$, notation ℓ_n is short for $\ell(n)$, and notation π_a is short for $\pi(a)$.

Definition 5 *An* ***RBR graph*** *is a tuple* (N, E, ℓ, π) *where:*

1. (N, E) *is a finite digraph with set* N *of the nodes and set* $E \subseteq N \times N$ *of the directed edges;*
2. $\ell : N \to \mathcal{A}$ *is a labelling function such that for each node* $n, m_1, m_2 \in N$*,*
   (a) *if* nEm_1*, then* $\ell_n \neq \ell_{m_1}$*;*
   (b) *if* nEm_1*,* nEm_2*, and* $\ell_{m_1} = \ell_{m_2}$*, then* $m_1 = m_2$*;*
3. $\pi : \mathcal{A} \to N$ *is a partial designating function such that for each agent* a*, if* π_a *is defined, then* $\ell_{\pi_a} = a$*;*
4. *for each node* $n \in N$*, there is an agent* $a \in \mathcal{A}$ *and a path from node* π_a *to* n*.*

An RBR graph defined above represents an RBR system among all agents in set $\mathcal{A}$. Note that the nodes in RBR graphs are not agents but just labelled by agents. This is because multiple nodes may represent the same agent when doxastic agents exist, as shown in Figure 1(iv). In particular, item 1 defines the finite digraph structure (N, E) of an RBR graph. Each node in set N represents either a real agent or a doxastic agent. A sequence of nodes connected by edges in set E forms a path corresponding to a belief sequence in the RBR system. Item 2 defines the labelling function ℓ such that ℓ_n is the agent that node n represents. Item 2a formalises the assumption that agents do not have introspective beliefs. Item 2b captures the intuition that an agent has only one identity in another agent's belief, thus preventing belief conflicts. Item 3 defines the designating function π on the set of all rational agents such that π_a is the node representing the real agent a (captured by the solid nodes in Figure 1). Note that, each node n *not* in the *image* of function π (captured by the dashed nodes in Figure 1) represents a doxastic agent ℓ_n. Item 4 requires that each node in an RBR graph should be reachable from (*i.e.* relevant to) a real agent so that every object satisfying Definition 5 represents an RBR system.

In an RBR graph (N, E, ℓ, π), a path $p = (n_1, \ldots, n_k)$ where $k \geq 1$ is a sequence of k nodes such that $n_i E n_{i+1}$ for each integer $i < k$. We call the sequence $\sigma = (\ell_{n_1}, \ldots, \ell_{n_k})$ of agents the ***belief sequence*** of path p and read it as "agent ℓ_{n_1} believes that ... believes that agent ℓ_{n_k} is rational". For a finite sequence σ of (possibly duplicated) agents, denote by $|\sigma|$ the length of sequence σ. For any agent $a \in \mathcal{A}$ and any sequence σ of agents, $a :: \sigma$ is the sequence obtained by attaching agent a at the beginning of sequence σ. A sequence σ is called an ***alternating sequence*** if every two consecutive agents in σ are not equal.

Definition 6 *For each node n in an RBR graph (N, E, ℓ, π), each integer $i \geq 1$, and each integer $j \geq 0$,*

$$\Pi_n^i := \begin{cases} \{\ell_n\}, & \text{if } i = 1; \\ \{\ell_n :: \sigma \mid nEm, \sigma \in \Pi_m^{i-1}\}, & \text{if } i \geq 2; \end{cases} \quad (1)$$

$$\Psi_n^j := \bigcup_{0 < i \leq j} \Pi_n^i; \quad (2)$$

$$\Psi_n^* := \bigcup_{i > 0} \Pi_n^i. \quad (3)$$

Informally, for a node n in an RBR graph, set Π_n^i consists of the belief sequences of all paths starting at node n and of length i; set Ψ_n^j consists of the belief sequences of all paths starting at node n and of length at most j; set Ψ_n^* consists of the belief sequences of all paths starting at node n. Intuitively, Ψ_n^* denotes the belief hierarchy of the (real or doxastic) agent represented by node n.

In the rest of this section, we consider solution concepts of games among agents with (possibly) uncommon RBR.

Definition 7 *A **solution** S of the game $(\Delta, \preceq)$ on the RBR graph (N, E, ℓ, π) is a family of sets $\{S_n\}_{n \in N}$ such that $\varnothing \subsetneq S_n \subseteq \Delta_{\ell_n}$ for each node $n \in N$.*

Informally, solution S describes a type of uncertainty in the choice of each (real or doxastic) agent in an RBR system: the (real or doxastic) agent represented by node n chooses only strategies in set S_n. We denote the solution $\{S_n\}_{n \in N}$ by S if it causes no ambiguity. Specifically, let

$$S^\Delta := \{\Delta_{\ell_n}\}_{n \in N} \quad (4)$$

be the solution corresponding to the whole strategy space.

Note that a solution does not have to be "reasonable". For instance, consider the "2/3 game" in Section 2 and the RBR graph depicted in Figure 1(iv). A solution S could be such that S_n is the set of all *prime* integers in the interval $[1, 10]$ for all five nodes n in the RBR graph, which is obviously unreasonable. Next, we consider the rationalisation of solutions. To do this, we have the next auxiliary definition.

Definition 8 *For any solution S of the game $(\Delta, \preceq)$ on the RBR graph (N, E, ℓ, π), the **belief scene** $\tilde{\Theta}_n(S)$ of any node $n \in N$ is the reasoning scene of agent ℓ_n such that for each agent $b \neq \ell_n$,*

$$\tilde{\Theta}_n^b(S) := \begin{cases} S_{n'}, & \text{if } \exists n' \in N (nEn' \text{ and } \ell_{n'} = b); \\ \Delta_b, & \text{otherwise.} \end{cases} \quad (5)$$

Specifically, for any solution S and any node n in the RBR graph, consider the (real or doxastic) agent ℓ_n denoted by node n. For each agent $b \neq \ell_n$, if agent ℓ_n believes b is rational, then there must be a node n' labelled with b such that nEn' in the belief graph. Moreover, node n' captures the agent b in agent ℓ_n's belief. In this sense, given the solution S, agent ℓ_n believes that agent b chooses a strategy from set S_n. On the other hand, if agent ℓ_n believes b is irrational, then no node n' exists such that $\ell_{n'} = b$ and nEn' and, in agent ℓ_n's belief, agent b choose any strategy from set Δ_b. Hence, given solution S, the (real or doxastic) agent ℓ_n denoted by node n believes that she rationalises in the reasoning scene in statement (5). We refer to such a reasoning scene as the belief scene of node n. Then, the rationalisation of solution S is such that every (real or doxastic) agent in an RBR graph rationalises (*i.e.* rational response) in her belief scene, as formally defined below.

Definition 9 *The **rationalisation** of any solution S of the game $(\Delta, \preceq)$ on the RBR graph (N, E, ℓ, π) is the solution $\mathbb{R}(S) = \{\mathbb{R}(S)_n\}_{n \in N}$ such that $\mathbb{R}(S)_n := \Re_{\ell_n}(\tilde{\Theta}_n(S))$ for each node $n \in N$.*

Note that, for any solution S, the rationalisation $\mathbb{R}(S)$ is a solution of the same game on the same RBR graph. It captures one turn of the iterative rationalisation process as discussed in Section 2, which is formally defined below.

Definition 10 *The i^{th} **rationalisation** on solution S is*

$$\mathbb{R}^i(S) := \begin{cases} S, & i = 0; \\ \mathbb{R}(\mathbb{R}^{i-1}(S)), & i \geq 1. \end{cases}$$

Recall that, the i^{th} column of Table 1 shows the result of the i^{th} rationalisation for our "2/3 game" in the RBR systems denoted in Figure 1. As shown there, the iterative rationalisation process may lead to a stable state. We call it *stable solution* and define it as follows.

Definition 11 *A **stable solution** S is such that $\mathbb{R}(S) = S$.*

Note that, without extra assumptions (*e.g.* the inaccessibility of strategies), the iterative rationalisation should start at the solution S^{Δ} where every agent chooses from the whole strategy space. Moreover, without extra assumptions (*e.g.* the limited mental capacity), the iterative rationalisation would continue forever because no agent has the incentive to stop it. The definition below captures this idea.

Definition 12 *The* ***rational solution*** $\mathbb{S}$ *is* $\lim_{i\to\infty}\mathbb{R}^i(S^{\Delta})$.[5]

However, as the following theorem shows, this process does not have to continue forever because a stable state will be achieved after finite iterations of rationalisation.

Theorem 1 *There is an integer* $i \geq 0$ *such that* $\mathbb{S} = \mathbb{R}^j(S^{\Delta})$ *for each integer* $j \geq i$.

This theorem shows the rational solution is well-defined. The intuition behind it is that, with the initial solution S^{Δ}, the iterative rationalisation process eliminates more and more but not all strategies (Lemma 13 in Appendix B of the full version). Since the strategy space is finite (item 1 of Definition 1), the elimination process has to reach a stable solution in finite steps and stays there forever (Lemma 14 in Appendix B of the full version). Such a stable solution is the rational solution by Definition 12.

In a sense, rational solution is a solution concept in uncommon RBR systems. It predicts the strategic behaviours of both real and doxastic agents in a game. Indeed, what matters are the real agents. The next definition extracts the elements of the real agents in a rational solution.

Definition 13 *The* ***doxastic rationalisability*** *of the game* $(\Delta, \preceq)$ *on the RBR graph* (N, E, ℓ, π) *is a family of sets* $\mathfrak{R} = \{\mathfrak{R}_a\}_{a\in\mathcal{A}}$ *such that*

$$\mathfrak{R}_a := \begin{cases} \mathbb{S}_{\pi_a}, & \text{if } \pi_a \text{ is defined;} \\ \Delta_a, & \text{otherwise;} \end{cases}$$

where $\mathbb{S}$ *is the rational solution of the same game on the same RBR graph.*

Doxastic rationalisability, the proposed solution concept, is the exact extension of *rationalisability* (Pearce 1984; Bernheim 1984) into uncommon RBR systems. In other words, without any other assumption than RBR among the agents, doxastic rationalisability is the unique reasonable prediction[6] of the agents' strategic behaviours in a game.

4 Equivalence in RBR Graphs

Intuitively, if there is no other assumption than uncommon RBR, then an agent's strategic behaviour is only affected by her own belief. In this sense, if a (real or doxastic) agent always has the same strategic behaviour in two RBR systems, then we say that the agent has *equivalent* beliefs in these RBR systems. Recall that a (real or doxastic) agent is denoted by a node in RBR graphs. For simplicity, we say that two nodes are *doxastically equivalent* if the agents denoted by them have the same strategic behaviour in every game.

Formally, we consider nodes n and n' in (possibly equal) RBR graphs $B = (N, E, \ell, \pi)$ and $B' = (N', E', \ell', \pi')$. For any game $G = (\Delta, \preceq)$, denote its rational solutions on the RBR graphs B and B' by $\mathbb{S}(G) = \{\mathbb{S}(G)_m\}_{m\in N}$ and $\mathbb{S}'(G) = \{\mathbb{S}'(G)_{m'}\}_{m'\in N'}$, respectively. Then, doxastic equivalence between nodes is defined as follows.

Definition 14 *The nodes* n *and* n' *are* ***doxastically equivalent*** *if* $\mathbb{S}(G)_n = \mathbb{S}'(G)_{n'}$ *for each game* G.

Recall that, as discussed in Section 1, agent a in Figure 1(ii) and in Figure 1(iii) have the same belief: either of them believes that there is RCBR with agent b and agent c is irrational. The same goes with agent b in Figure 1(ii) and in Figure 1(iii), and agent c in Figure 1(i) and in Figure 1(iv). Note that, by "the same belief" we mean that the belief hierarchy (*i.e.* the set of belief sequences) is the same. In this sense, by saying that the agents denoted by nodes n and n' have the same belief, we mean $\Psi^*_n = \Psi^*_{n'}$.

We find that beliefs are equivalent if and only if they are the same, as formally stated in the following theorem.

Theorem 2 *The nodes* n *and* n' *are doxastically equivalent if and only if* $\Psi^*_n = \Psi^*_{n'}$.

For the "only if" part of Theorem 2, note that $\Psi^*_n \neq \Psi^*_{n'}$ implies the existence of an integer k such that $\Psi^k_n \neq \Psi^k_{n'}$ by Definition 6. Then, we show the existence of a parameterised game G_k such that $\mathbb{S}(G_k)_n \neq \mathbb{S}'(G_k)_{n'}$ (Definition 18 and Lemma 20 in Appendix C.1 of the full version). For the "if" part of Theorem 2, we prove by induction that, for each integer $i \geq 1$, after the i^{th} rationalisation, $\mathbb{R}^i_B(S^{\Delta})_n = \mathbb{R}^i_{B'}(S^{\Delta})_{n'}$ in every game (Lemma 25 in Appendix C.2 of the full version). Then, the "if" part statement of Theorem 2 follows from Theorem 1.

Now, we consider the equivalence of RBR systems. Recall that an RBR system is a collection of the belief hierarchies of all real agents. The real agents are whom we care about. In this sense, we say that two RBR systems are equivalent if no real agents can distinguish them. In other words, every real agent should always have the same strategic behaviour in two equivalent RBR systems. Note that we use RBR graphs to denote RBR systems. Formally, we consider the equivalence of two arbitrary RBR graphs B and B'. For any game G, denote by $\mathfrak{R}(G) = \{\mathfrak{R}(G)_a\}_{a\in\mathcal{A}}$ and $\mathfrak{R}'(G) = \{\mathfrak{R}'(G)_a\}_{a\in\mathcal{A}}$ the doxastic rationalisabilities of the game G on the RBR graphs B and B', respectively.

Definition 15 *The RBR graphs* B *and* B' *are* ***equivalent*** *if* $\mathfrak{R}(G)_a = \mathfrak{R}'(G)_a$ *for each agent* $a \in \mathcal{A}$ *and each game* G.

The next theorem shows the necessary and sufficient condition for two RBR systems to be equivalent. That is, for each agent a, either a is irrational in both systems, or a is rational and has equivalent beliefs in both systems. In the RBR graphs, the former means that agent a is not in the domain of definition of the designating functions. The latter is formally

[5] Technically, $\mathbb{S} = \{\mathbb{S}_n\}_{n\in N}$ such that $\mathbb{S}_n = \lim_{i\to\infty}\mathbb{R}^i_n(S^{\Delta})$, where $\mathbb{R}^i_n(S^{\Delta})$ is a set of strategies for each integer i and the limit of a sequence of sets is defined in the standard way (Resnick 1998, Section 1.3). In particular, if a sequence of sets stabilises after some element, then the limit is equal to the stable value.

[6] In the sense that (1) every strategy not in the doxastic rationalisability is believed to be dominated by another strategy, so no agent would like to choose it; (2) more assumptions/beliefs are required to eliminate a strategy in the doxastic rationalisability.

expressed with the doxastic equivalence between the nodes denoting agent a. Proofs of the theorem and its corollary below can be found in Appendix D of the full version.

Theorem 3 *RBR graphs* (N, E, ℓ, π) *and* (N', E', ℓ', π') *are equivalent if and only if, for each agent* $a \in \mathcal{A}$, ***either*** *both* π_a *and* π'_a *are not defined,* ***or*** π_a *and* π'_a *are both defined and doxastically equivalent.*

The next corollary follows directly from Theorem 3 and Theorem 2. It shows that the core of the equivalence of two RBR systems is the belief hierarchy (*i.e.* the set of all belief sequences, $\Psi^*_{\pi_a}$) of each real agent. This property is used in the next section for minimising an RBR graph.

Corollary 1 *RBR graphs* (N, E, ℓ, π) *and* (N', E', ℓ', π') *are equivalent if and only if* π *and* π' *have the same domain* $\mathcal{D}$ *of definition,* ***and*** $\Psi^*_{\pi_a} = \Psi^*_{\pi'_a}$ *for each agent* $a \in \mathcal{D}$.

5 Minimisation of RBR Graphs

So far we have assumed that RBR systems are given. However, this is not the case in most situations. Researchers in behaviour economics study how to *elicit* the belief of a single agent (Schotter and Trevino 2014; Charness, Gneezy, and Rasocha 2021; Danz, Vesterlund, and Wilson 2022). To get an RBR system, we probably need to elicit the belief hierarchy of each agent and combine them as a whole. For instance, to get the RBR system depicted in Figure 1(iii), we first know that agent a believes RCBR exists between herself and agent b, agent b believes RCBR exists between herself and agent a, and agent c believes that RCBR exists agents a and b. Then, we depict each agent's belief with a graph and combine all of them as a whole, as shown in Figure 2. In other words, an RBR system is a collection of the belief hierarchy of all (real) agents in it. Recall that, in Definition 5, we never require an RBR graph to be a connected graph.

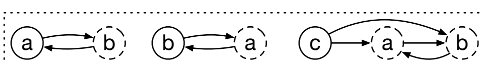

Figure 2: The collection of individual beliefs in Figure 1(iii).

It is easily observable and verifiable using Corollary 1 that the RBR systems denoted in Figure 1(iii) and Figure 2 are indeed equivalent. However, the RBR graph in Figure 2 has more than twice the number of nodes than that in Figure 1(iii). Note that, by Definition 9, the time complexity of rationalisation is proportional to the number of nodes in an RBR graph. Minimising an RBR graph is to find an equivalent RBR graph with the fewest nodes. On the one hand, it reduces the time complexity of computing the doxastic rationalisability. On the other hand, it helps to find the most condensed expression of an RBR system. Formally, we have the next definition for "the most condensed expression".

Definition 16 *An RBR graph* (N, E, ℓ, π) *is* ***canonical*** *if* $\Psi^*_n \neq \Psi^*_{n'}$ *for all distinct nodes* $n, n' \in N$.

Intuitively, Ψ^*_n denotes the belief hierarchy of the agent denoted by node n. Then, an RBR graph is canonical if the nodes represent agents with different belief hierarchies. It is proved in Appendix E.1 of the full version that *a canonical RBR graph is not equivalent to any RBR graph with fewer nodes* (Lemma 29) and *two equivalent canonical RBR graphs must be isomorphic* (Definition 19 and Theorem 5). Due to these properties, to minimise an RBR graph, we only need to compute an equivalent canonical RBR graph.

Algorithm 1: Minimise an RBR graph

Input: RBR graph (N, E, ℓ, π)
Output: RBR graph (N', E', ℓ', π')
$\mathbb{P} \leftarrow \big\{\{n' \in N \mid \ell_{n'} = \ell_n\} \,\big|\, n \in N\big\}$; `//`$\Psi^1_n$ `equivalence`
$stable \leftarrow false$;
while *not stable* **do** `//`Ψ^i_n `equiv.` $\to \Psi^{i+1}_n$ `equiv.`
  $stable \leftarrow true$;
  $\mathbb{P}' \leftarrow \varnothing$;
  for *each set* $P \in \mathbb{P}$ **do**
    for *each node* $n \in P$ **do**
      $type(n) \leftarrow \{P' \in \mathbb{P} \mid nEn', n' \in P'\}$;
    $\mathbb{Q} \leftarrow \big\{\{n' \in P \mid type(n') = type(n)\} \mid n \in P\big\}$;
    $\mathbb{P}' \leftarrow \mathbb{P}' \cup \mathbb{Q}$;
    if $|\mathbb{Q}| > 1$ **then**
      $stable \leftarrow false$;
  $\mathbb{P} \leftarrow \mathbb{P}'$;
$N' \leftarrow \mathbb{P}$; `//equivalent classes as nodes`
$E' \leftarrow \{(P, Q) \mid (n, m) \in E, n \in P, m \in Q\}$;
for *each node* $P \in N'$ **do**
  pick an arbitrary node $n \in P$ and $\ell'_P \leftarrow \ell_n$;
for *each agent* $a \in \mathcal{A}$ **do**
  $\pi'_a \leftarrow P : \pi_a \in P$ if π_a is defined;
return (N', E', ℓ', π');

Technically, Ψ^*_n is the set of "labelling sequences" of all paths starting at node n. This reveals a similarity between RBR graphs and automata. Inspired by Myhill–Nerode theorem (Myhill 1957; Nerode 1958) and Hopcroft's algorithm (Hopcroft 1971), we design Algorithm 1 that works based on *partition refinement* and outputs an equivalent canonical RBR graph of the input RBR graph, which is unique up to isomorphism. A detailed explanation and formal proof of its correctness can be found in Appendices E.2 and E.3 of the full version. The time complexity of Algorithm 1 is $O(|\mathcal{A}| \cdot |N|^2 \cdot log|N|)$, where $|\mathcal{A}|$ is the number of agents and $|N|$ is the numbers of nodes in the input RBR graph. See Appendix E.4 of the full version for the a detailed analysis.

6 Concluding Discussion

Generally speaking, the RBR graph proposed in this paper is a *syntactic* presentation of an RBR system. It is in line with our linguistic intuition about beliefs (*i.e.* the correspondence between the label sequence of a path and a belief sequence in the hierarchy). In a sense, doxastic rationalisability, the solution concept we propose, is a *semantic* interpretation of an RBR system in games. From this perspective, Theorem 2 and Corollary 1 reveal the correlation between syntax and semantics of our graph-based language, based on which, we design an efficient algorithm that computes the most condensed syntactic expression of an RBR system.

Acknowledgements

The research is funded by the China Scholarship Council (CSC No.202206070014).

References

Albrecht, S. V.; and Stone, P. 2018. Autonomous agents modelling other agents: A comprehensive survey and open problems. *Artificial Intelligence*, 258: 66–95.

Angluin, D. 1980. Local and global properties in networks of processors. In *Proceedings of the Twelfth Annual ACM Symposium on Theory of Computing*, 82–93.

Aumann, R. J. 1976. Agreeing to disagree. *The Annals of Statistics*, 4(6): 1236–1239.

Aumann, R. J. 1987. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica: Journal of the Econometric Society*, 1–18.

Bernheim, B. D. 1984. Rationalizable strategic behavior. *Econometrica: Journal of the Econometric Society*, 1007–1028.

Binmore, K. 1988. Modeling rational players: Part II. *Economics & Philosophy*, 4(1): 9–55.

Börgers, T. 1993. Pure strategy dominance. *Econometrica: Journal of the Econometric Society*, 423–430.

Brandenburger, A.; and Dekel, E. 1987. Rationalizability and correlated equilibria. *Econometrica: Journal of the Econometric Society*, 1391–1402.

Brandenburger, A.; and Dekel, E. 1993. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory*, 59(1): 189–198.

Camerer, C. F.; Ho, T.-H.; and Chong, J.-K. 2004. A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, 119(3): 861–898.

Charness, G.; Gneezy, U.; and Rasocha, V. 2021. Experimental methods: Eliciting beliefs. *Journal of Economic Behavior & Organization*, 189: 234–256.

Chong, J.-K.; Ho, T.-H.; and Camerer, C. 2016. A generalized cognitive hierarchy model of games. *Games and Economic Behavior*, 99: 257–274.

Danz, D.; Vesterlund, L.; and Wilson, A. J. 2022. Belief elicitation and behavioral incentive compatibility. *American Economic Review*, 112(9): 2851–2883.

Dekel, E.; and Siniscalchi, M. 2015. Epistemic game theory. In *Handbook of Game Theory with Economic Applications*, volume 4, 619–702. Elsevier.

Doshi, P.; Gmytrasiewicz, P.; and Durfee, E. 2020. Recursively modeling other agents for decision making: A research perspective. *Artificial Intelligence*, 279: 103202.

Gale, D.; and Sherman, S. 1950. Solutions of finite two-person games. In *Contributions to the Theory of Games (AM-24), Volume I*, volume 1, 37–50. Princeton University Press.

Harsanyi, J. C. 1967. Games with incomplete information played by "Bayesian" players, I–III Part I. The basic model. *Management Science*, 14(3): 159–182.

Hopcroft, J. 1971. An n log n algorithm for minimizing states in a finite automaton. In *Theory of Machines and Computations*, 189–196. Elsevier.

Howard, R. A.; and Matheson, J. E. 2005. Influence diagrams. *Decision Analysis*, 2(3): 127–143.

Hsieh, S.-M.; Hsu, C.-C.; and Hsu, L.-F. 2006. Efficient method to perform isomorphism testing of labeled graphs. In *Proceedings of the Sixth International Conference on Computational Science and Its Applications*, 422–431. Springer.

Koller, D.; and Milch, B. 2003. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 45(1): 181–221.

Lewis, D. 1969. *Convention: A philosophical study*. Harvard University Press.

Milch, B.; and Koller, D. 2000. Probabilistic models for agents' beliefs and decisions. In *Proceedings of the 16th Conference on Uncertainty in Artificial Intelligence*, 389–396.

Moulin, H. 1986. *Game Theory for the Social Sciences*. NYU Press.

Myhill, J. 1957. Finite automata and the representation of events. *WADD Technical Report*, 57: 112–137.

Nagel, R. 1995. Unraveling in guessing games: An experimental study. *The American Economic Review*, 85(5): 1313–1326.

Nash, J. F. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences*, 36(1): 48–49.

Nerode, A. 1958. Linear automaton transformations. *Proceedings of the American Mathematical Society*, 9(4): 541–544.

Osborne, M. J.; and Rubinstein, A. 1994. *A Course in Game Theory*. MIT press.

Pearce, D. G. 1984. Rationalizable strategic behavior and the problem of perfection. *Econometrica: Journal of the Econometric Society*, 1029–1050.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. San Francisco: Morgan Kaufmann.

Perea, A. 2012. *Epistemic Game Theory: Reasoning and Choice*. Cambridge University Press.

Resnick, S. 1998. *A Probability Path*. Springer.

Savage, L. J. 1954. *The Foundation of Statistics*. New York: John Wiley and Sons.

Schiffer, S. R. 1972. *Meaning*. Oxford University Press.

Schotter, A.; and Trevino, I. 2014. Belief elicitation in the laboratory. *Annual Review of Economics*, 6(1): 103–128.

Selten, R. 1998. Features of experimentally observed bounded rationality. *European Economic Review*, 42(3-5): 413–436.

Shi, Q.; and Pavel, N. 2024. Uncommon belief in rationality. *arXiv preprint arXiv:2412.09407*.

Simon, H. A. 1955. A behavioral model of rational choice. *The Quarterly Journal of Economics*, 99–118.

Simon, H. A. 1957. *Models of Man: Social and Rational*. John Wiley.

Steup, M.; and Neta, R. 2024. Epistemology. In Zalta, E. N.; and Nodelman, U., eds., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.