

A Causal Approach to Game Theory

author names withheld

Under Review for NExT-Game 2026

Abstract

The tension between rational and irrational behaviors in human decision-making has been acknowledged across a wide range of disciplines, from philosophy to psychology, neuroscience to behavioral economics. Models of multi-agent interactions, such as von Neumann and Morgenstern’s expected utility theory and Nash’s game theory, provide rigorous mathematical frameworks for how agents should behave when rationality is sought. However, the rationality assumption has been extensively challenged, as human decision-making is often irrational, influenced by biases, emotions, and uncertainty, which may even have a positive effect in certain cases. Behavioral economics, for example, attempts to explain such irrational behaviors, including Kahneman’s dual-process theory and Thaler’s nudging concept, and accounts for deviations from rationality. In this paper, we analyze this tension through a causal lens and develop a framework that accounts for rational and irrational decision-making, which we term *Causal Game Theory*. We then introduce a novel notion called counterfactual rationality, which allows agents to make choices leveraging their irrational tendencies. We extend the notion of Nash Equilibrium to counterfactual actions and Pearl Causal Hierarchy (PCH), and show that strategies following counterfactual rationality dominate strategies based on standard game theory. We further develop an algorithm to learn such strategies when not all information about other agents is available.

1. Introduction

Decision-making in multi-agent systems (MAS) underlies problems across economics, social sciences, distributed systems, robotics, and the alignment of AI agents. Game theory has long been the dominant framework for such settings, with solution concepts such as Nash equilibrium [29], cooperative games [37], and Bayesian games [16] prescribing how rational agents should act. Yet human decisions are routinely shaped by cognitive biases, emotions, and unobserved factors that produce seemingly irrational outcomes, and irrational play can occasionally outperform rational play—the so-called *paradox of rationality* [5, 7, 17]. Behavioral economics has cataloged many such deviations [11, 20–22, 28, 40], but the broader question of *when* and *how* an agent can strategically leverage these unobserved biases in MAS remains largely open.

We address this gap by modeling MAS through structural causal models (SCMs) and the Pearl Causal Hierarchy (PCH) [4, 31]. Human decisions are often guided by causal structures [30, 38, 42], and actions can be viewed as interventions [13]. SCMs have been deployed for single-agent decision-making in bandits [3, 44] and reinforcement learning [27, 36], surveyed in Bareinboim et al. [2]. In the *Greedy Casino* [3], players who exploit their biases via counterfactual randomization outperform both their natural play and the rational RCT-based strategy. We extend this perspective to multi-agent settings: agents may act rationally (Nash’s prescription), instinctively, or as a mixture of both. The following example illustrates the resulting dilemma.

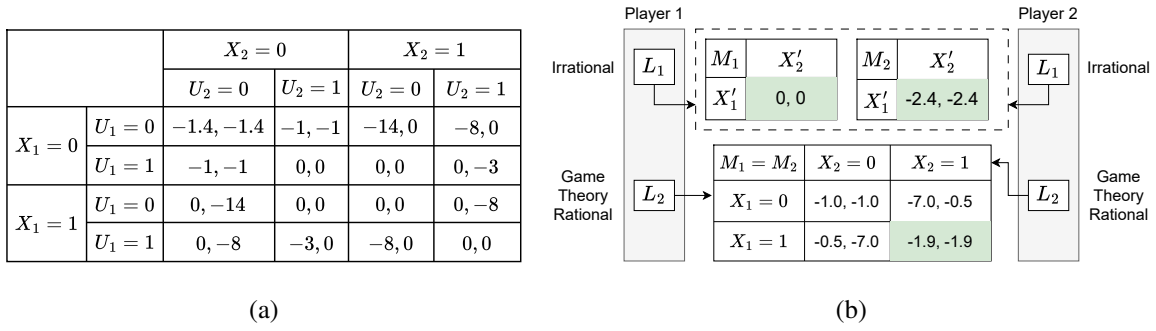


Figure 1: (a) Y_1, Y_2 as a function of U_1, U_2, X_1, X_2 (b) Should the agents be rational or not?

Example 1 (Causal Prisoner’s Dilemma (CPD)) *Two thieves are suspected of a crime, but due to insufficient evidence, they cannot be convicted outright. Now, they have a choice to make – either remain silent (cooperate, C) or betray the other (defect, D). We denote the choices by variables X_1 and X_2 , and cooperation and defection by the values 0 and 1, respectively. The thieves’ decisions are influenced by external circumstances, represented by variables U_1 and U_2 , which capture factors such as the temperament of police officers, the competence of legal defense, new evidence or witnesses emerging, and even the disposition of the judge and the jury. Although these factors cannot be explicitly measured by the prisoners, they may subconsciously shape their decisions.*

Each prisoner has a natural ability to assess their circumstances, denoted by R_1 and R_2 . If prisoner i has an accurate reading of their situation ($R_i = 1$), they choose to cooperate ($X_i = 0$) if the circumstances are favorable ($U_i = 1$), and defect when they are adversarial ($U_i = 0$); conversely, if they have a poor reading of their situation ($R_i = 0$), they defect when circumstances are good, and cooperate when circumstances are bad. For prisoner i , their instinctive or natural choice is modeled as: $X_i \leftarrow f_X(R_i, U_i) = R_i \oplus U_i$, where \oplus is the exclusive-or operator. We note that the variables U_1, U_2, R_1, R_2 and the function f_X are determined by nature and are unknown to the prisoners.

Now, we analyze two scenarios, M_1 and M_2 . In M_1 , the prisoners have a good reading of their situation ($R_1 = R_2 = 1$), while in M_2 , they misjudge their circumstances ($R_1 = R_2 = 0$). In both cases, $P(U_i = 0) = 0.6$ for $i \in \{1, 2\}$. The outcome $\mathbf{Y} = (Y_1, Y_2)$ of their decisions is a function of U_1, U_2, X_1 and X_2 as shown in Fig. 1. For example, when the situation is favorable for both the prisoners ($U_1 = 1, U_2 = 1$) and they cooperate ($X_1 = 0, X_2 = 0$), their payoff is $(0, 0)$. However, if circumstances are favorable for Prisoner 1 and not for Prisoner 2 ($U_1 = 1, U_2 = 0$), and Prisoner 1 defects while Prisoner 2 cooperates ($X_1 = 1, X_2 = 0$), their payoff is $(0, -8)$.

If both prisoners ignore their intuition and search for the optimal strategy, the situation corresponds to the classical Prisoner’s Dilemma, where the payoff for the actions $X_1 = x_1, X_2 = x_2$ is given by:

$$\sum_{u_1, u_2, \mathbf{Y}} \mathbf{Y} \cdot P(u_1, u_2) P(\mathbf{Y} \mid x_1, x_2, u_1, u_2) \quad (1)$$

Notably, both scenarios M_1 and M_2 lead to the same Prisoner’s Dilemma (PD) game, as shown in the 2×2 payoff table at the bottom of Fig. 1. However, if both prisoners rely on their natural instincts, their expected payoff is $(0, 0)$ in M_1 and approximately $(-2.4, -2.4)$ in M_2 . This is illustrated in Fig. 1, where X'_1 and X'_2 denote the players acting based on their natural intuition

(shown in the top row). The situation presents a new dilemma – it is better to follow natural instincts and be irrational in M_1 , whereas it is better to be rational and ignore intuition in M_2 .

This example raises a fundamental question: when should an agent follow intuition, and when should it override it? We explore this tension through a causal lens and address the paradox of rationality (Appendix A). Our contributions are:

1. We formalize a class of games combining rational and irrational behavior (Def. 6) and show it strictly generalizes traditional Normal Form Games (Thm. 7).
2. We develop `Ctf-Nash-Learning` (Alg. 3) that learns the payoff matrix in the counterfactual action space and identifies a Nash Equilibrium, even when other agents’ actions are not fully observed (Thm. 8).

1.1. Preliminaries

We use capital letters for random variables (X) and lowercase for values (x); \mathcal{D}_X is the domain of X , and $|\mathbf{S}|$ is set cardinality. Our framework rests on Structural Causal Models [31]: an SCM M is a tuple $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$, where \mathbf{V} and \mathbf{U} are endogenous and exogenous variables, and \mathcal{F} is a set of mechanisms $V \leftarrow f_V(\mathbf{Pa}(V), \mathbf{U}_V)$ with $\mathbf{Pa}(V) \subseteq \mathbf{V}$ and $\mathbf{U}_V \subseteq \mathbf{U}$. M induces a distribution $P(\mathbf{V})$ over endogenous variables, called the *observational (L_1) distribution*. An intervention $do(\mathbf{x})$ on $\mathbf{X} \subseteq \mathbf{V}$ replaces $\{f_X : X \in \mathbf{X}\}$ with $\mathbf{X} \leftarrow \mathbf{x}$, yielding $M_{\mathbf{x}}$ and *interventional (L_2) distributions* $P_{\mathbf{x}}(\mathbf{Y}) = P(\mathbf{Y}_{\mathbf{x}})$. Counterfactual (L_3) distributions and further background are in Appendix A.1 and Bareinboim et al. [4, Sec. 1.2]; single-agent causal decision-making is reviewed in Bareinboim et al. [2] and Appendix A.5.

2. Causal Normal Form Games

In this section, we model the interaction of multiple agents in a system through the language of SCMs and PCH layers, generalizing the single-agent concepts of Bareinboim et al. [2] to multi-agent settings. We first define action nodes and reward signals together with the underlying SCM.

Definition 1 (Causal Multi-Agent System) *A Causal Multi-Agent System (CMAS) is a tuple $\langle M, N, \mathbf{X}, \mathbf{Y} \rangle$, where (i) $M : \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, \mathbb{P} \rangle$ is an SCM, (ii) N is the set of n agents, (iii) $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ is a tuple of action nodes with disjoint $\mathbf{X}_i, \mathbf{X}_j \subseteq \mathbf{V}$ for $i, j \in [n]$, $i \neq j$, and (iv) $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)$ is the ordered set of reward signals, with $\mathbf{Y}_i \subseteq \mathbf{V} \setminus \mathbf{X}$ for all $i \in [n]$. \square*

A CMAS is essentially an SCM with a set of action nodes \mathbf{X} , each controlled by one of the n agents. In addition, the system includes reward variables, \mathbf{Y} , representing the feedback each agent receives based on their actions and the underlying causal mechanism.

Example 2 *Consider the CPD presented in Ex. 1. The SCM \mathcal{M} corresponding to scenario M_2 is defined as: (i) $\mathbf{U} = \{U_1, U_2, R_1, R_2\}$, $\mathbf{V} = \{X_1, X_2, Y_1, Y_2\}$, (ii) $X_i = R_i \oplus U_i$ for $i \in \{1, 2\}$. Y_1, Y_2 as a function of U_1, U_2, X_1, X_2 are shown in Fig. 1, and (iii) $P(U_i = 1) = 0.4, P(R_i = 0) = 1$ for $i \in \{1, 2\}$. The CMAS can now be defined as $\langle M = \mathcal{M}, N = \{1, 2\}, \mathbf{X} = (\{X_1\}, \{X_2\}), \mathbf{Y} = (\{Y_1\}, \{Y_2\})$. \square*

Now, we define different forms of actions that an agent may take in such a system. First, we define the different action and policy spaces and then explore how the action spaces are related.

Definition 2 (L_1 action) *Given a CMAS $\langle M, N, \mathbf{X}, \mathbf{Y} \rangle$, an L_1 action of an agent i is the one in which the value of their action variables \mathbf{X}_i is determined by the natural mechanism $f_{\mathbf{X}_i} \in \mathcal{F}$. \square*

We will also call such actions *natural actions* and denote them by a_0 . Note that, while performing a_0 , an agent *does not know anything about the underlying SCM* nor do they deliberately change any mechanism of action variable in the system. The L_1 action space is thus $\mathcal{A}^1 = \{a_0\}$ and the L_1 policy space is also a singleton set $\Pi^1 = \{a_0\}$.

An agent can also intervene on the system via atomic interventions (fixed values) [31] or soft interventions (sampling from a distribution) [9]. Next, we define L_2 actions.

Definition 3 (L_2 -action) *Given a CMAS $\langle M, N, \mathbf{X}, \mathbf{Y} \rangle$, L_2 action of an agent i is a hard intervention $do(\mathbf{x})$, where $\mathbf{x} \in \mathcal{D}_{\mathbf{X}_i}$. \square*

Hence, if an agent i performs $do(\mathbf{x}_i)$ in the SCM M , then the natural mechanism $f_{\mathbf{X}_i}$ is replaced by $\mathbf{X}_i \leftarrow \mathbf{x}_i$. The set of L_2 actions is denoted by \mathcal{A}^2 , and an L_2 policy is a distribution over \mathcal{A}^2 .

An agent can also interact through PCH's Layer 3 [3, 4, 33], enabling counterfactual reasoning. For example, in scenario M_2 of Ex. 1, both agents would obtain payoff $(0, 0)$ by playing the opposite of their natural instincts. We now formally define L_3 actions.

Definition 4 (L_3 -action space \mathcal{A}^3) *Given a CMAS $\langle N, M, \mathbf{X}, \mathbf{Y} \rangle$, an L_3 action for agent i is defined as a mapping $h : \mathcal{D}(\mathbf{X}_i) \rightarrow \mathcal{D}(\mathbf{X}_i)$ from intuition to action. \square*

When an agent takes an L_3 action, they first note their natural instinct $\mathbf{X}'_i \leftarrow f_{\mathbf{X}_i}(\mathbf{U}_i)$ and then executes $\mathbf{X}_i \leftarrow h_i(\mathbf{X}'_i)$, where \mathbf{U}_i is the set of unobserved parents of \mathbf{X}_i . If $h(x) = x$, it corresponds to the L_1 action, and if $h(x)$ is constant for all x , it is the L_2 action. Bearing this in mind, we will often denote a_0 as $\mathbf{X} = \mathbf{X}'$, where \mathbf{X} is the action variable and \mathbf{X}' is the intuition.

Concretely, Bareinboim et al. [3] proposes a randomization procedure (Ct f-RCT) in which an agent's would-be choice is perceived as the intuition X' and the executed action is then sampled (or chosen as a function of X'); this enables sampling from counterfactual distributions of the form $P(\mathbf{Y}_x | x')$. We use Ct f-RCT as the data source for the learning algorithm in Sec. 3; full details are deferred to Appendix A.2.

Example 3 (Payoffs across PCH layers in CPD) *For the CMAS in Ex. 2, an L_1 action sets $X_i \leftarrow R_i \oplus U_i$ and yields expected payoff*

$$\sum_{u_1, u_2, x_1, x_2, \mathbf{y}} \mathbf{y} P(u_1, u_2) P(x_1 | u_1) P(x_2 | u_2) P(\mathbf{y} | u_1, u_2, x_1, x_2) \approx (-2.4, -2.4)$$

in M_2 . An L_3 action $X_i \leftarrow g_i(X'_i)$, where $X'_i \leftarrow R_i \oplus U_i$, with $g_1(x) = g_2(x) = 1 - x$ instead yields $(0, 0)$. Standard L_2 payoffs follow by setting g_i constant. The full set of action combinations is shown in Appendix B (Fig. 9). \square

Let $\Delta(A)$ denote the set of distributions over A . Then $\Pi^2 = \Delta(\mathcal{A}^2)$ and $\Pi^3 = \Delta(\mathcal{A}^3)$ are the L_2 and L_3 policy spaces. Next, we define the notion of reward.

Definition 5 (Reward Function) A reward function $\mathcal{R}_i : \mathcal{D}(\mathbf{Y}_i) \rightarrow \mathbb{R}$ of an agent i is a function from outcome \mathbf{Y}_i to real numbers. \square

In Ex. 1, we assume that the reward function is identity, that is, $\mathcal{R}_i(Y_i) = Y_i$ for $i \in \{1, 2\}$. Now that we have all the tools, we are ready to define Normal Form Games in proper causal language.

Definition 6 (Causal Normal Form Game) A tuple $\Gamma = \langle \mathbb{M}, \mathcal{A}, \mathcal{R} \rangle$ is a Causal Normal Form Game (CNFG), where (i) \mathbb{M} is a CMAS $\langle M, N, \mathbf{X}, \mathbf{Y} \rangle$, (ii) $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ is the set of policies for the n agents, $\mathcal{A}_i \subseteq \{\mathcal{A}^1, \mathcal{A}^2, \mathcal{A}^3\}$, and (iii) $\mathcal{R} = (\mathcal{R}_1, \dots, \mathcal{R}_n)$ is the set of reward functions. \square

A CNFG is thus a CMAS, along with the policy space of the n agents and their reward functions. Now we will formally state the result generalizing our observation from CPD (Ex. 1).

Theorem 7 Given a game in normal form, there exist two CNFGs \mathcal{C}_1 and \mathcal{C}_2 with equilibrium payoffs μ_1 and μ_2 under the action space $\mathcal{A}^1 \cup \mathcal{A}^2$, and a Nash Equilibrium (NE) payoff μ_{NE} , such that $\mu_2 < \mu_{NE} < \mu_1$ where $<$ denotes Pareto domination. \square

The theorem implies several observations. CNFGs strictly generalize Normal Form Games (NFGs), capturing instinctive behaviors and counterfactual policies that NFGs cannot naturally express. Although a CNFG can in principle be flattened into an equivalent NFG, causal modeling is not just convenient but necessary: (i) constructing the full payoff matrix requires an SCM, since actions are defined only within that causal structure; (ii) NFGs do not specify how actions are executed or whether agents are capable of executing them, while SCMs provide a concrete notion of agency; (iii) under partial observability, only the SCM exposes the structure between intuitions and executed actions that is needed to learn equilibria (Sec. 3, Alg. 3). See Appendix G for further discussion.


3. Learning Causal Nash Equilibrium

Once agents are allowed to move across PCH layers, a richer solution concept emerges: each agent first chooses which subset of action spaces to operate in, then plays an equilibrium of the induced subgame. We call this the *Causal Nash Equilibrium* (CNE). CNE always exists and admits a payoff-dominance property: under a pure-strategy equilibrium of the layer-selection meta-game, no agent benefits from unilaterally switching PCH layers. The formal definitions, the underlying notion of *counterfactual rationality*, the role of *forgetting*, the existence theorem, and the dominance theorem are deferred to Appendix B.

In practice, payoff matrices are rarely common knowledge. The counterfactual action space introduces an additional difficulty: an agent observes other agents' *executed* actions but not their underlying intuitions, so the rows of the payoff matrix indexed by counterfactual strategies cannot be filled in by direct observation. We address this learning problem next.

3.1. Ctf-Nash-Learning

Suppose both players use Ctf-RCT [3] during exploration (Appendix A.2), generating tuples $(x'_1, x_1, x_2, \mathbf{y})$ from Player 1's perspective: x'_1 is Player 1's intuition, x_1, x_2 are executed actions, and \mathbf{y} is the reward tuple. Player 2's intuition x'_2 is unobserved. For fixed (x'_1, x_1, x_2) , \mathbf{y} is sampled from the mixture $\sum_{x'_2} P(x'_2 | x'_1) P(\mathbf{y}_{x_1, x_2} | x'_2, x'_1)$.



(a) Dynamics in the Causal Prisoner's Dilemma as players move up the layers of PCH (transitions detailed in Appendix B).

(b) Equilibrium payoffs for Player 1 with L_1, L_2, L_3 action spaces under varying $P(R_i = 0) = r$.

Figure 2: Causal Prisoner's Dilemma: dynamics and equilibrium payoffs across PCH layers.

Ctf-Nash-Learning (Alg. 3, Appendix E) operates as follows. First, it estimates the means $R_i(x'_1, x_1, x_2)$ and weights $p_i(x'_1, x_1, x_2)$ of the mixture from samples; under the identifiability conditions of Appendix E, these recover, up to permutation, $E[\mathbf{Y}_{x_1, x_2} | x'_1, x'_2]$ and $P(x'_2 | x'_1)$ (with a degenerate case when intuition does not affect the outcome). Next, it instantiates the agents' L_3 action spaces $\mathcal{F}_1 = \{f: X'_1 \rightarrow X_1\}$ and $\mathcal{F}_2 = \{g: [k] \rightarrow X_2\}$, builds the induced payoff matrix, and returns a Nash Equilibrium of the resulting NFG.

Theorem 8 (Asymptotic correctness of Ctf-Nash-Learning) *Let $\Gamma = \langle \mathbb{M}, (\mathcal{A}_1^3, \mathcal{A}_2^3), \mathcal{R} \rangle$ be a two-player CNFG satisfying the identifiability conditions of Appendix E. Then, as the dataset size $n \rightarrow \infty$, Ctf-Nash-Learning returns a Nash Equilibrium strategy of Γ over the full counterfactual action space, almost surely. \square*

The CPD's learned mixture weights, means, and payoff matrix are reported in Appendix E.4.

4. Experiments

We empirically investigate how game behavior changes as players move across the layers of PCH, using Independent Q-Learning [39]. Fig. 2(a) tracks the dynamics: every 20,000 timesteps, one agent moves up a layer, triggering a payoff change—an experimental realization of the layer transitions discussed in Appendix B. Fig. 2(b) shows equilibrium payoffs as a function of $P(R_i = 0) = r$; for the Causal Prisoner's Dilemma, L_3 uniformly dominates L_2 .

5. Conclusions

We examine the tension between rational and irrational decision-making through a causal lens. We introduce a framework that captures both modes of behavior and strictly generalizes Normal Form Games (Thm. 7). We then develop Ctf-Nash-Learning, which learns the payoff matrix and a Nash Equilibrium in the full counterfactual action space under partial observability, with asymptotic correctness guarantees (Thm. 8). Additional results on counterfactual rationality and the layer-selection-based Causal Nash Equilibrium concept are deferred to the appendix. We hope this framework advances the design of more robust decision-making systems.

References

- [1] Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.
- [2] E. Bareinboim, J. Zhang, and S. Lee. An introduction to causal reinforcement learning. Technical Report R-65, Causal Artificial Intelligence Lab, Columbia University, Dec 2024. <https://causalai.net/r65.pdf>.
- [3] Elias Bareinboim, Andrew Forney, and Judea Pearl. Bandits with unobserved confounders: A causal approach. *Advances in Neural Information Processing Systems*, 28, 2015.
- [4] Elias Bareinboim, Juan D. Correa, Duligur Ibeling, and Thomas Icard. *On Pearl’s Hierarchy and the Foundations of Causal Inference*, page 507–556. Association for Computing Machinery, New York, NY, USA, 1 edition, 2022. ISBN 9781450395861. URL <https://doi.org/10.1145/3501714.3501743>.
- [5] Kaushik Basu. The traveler’s dilemma: Paradoxes of rationality in game theory. *The American Economic Review*, 84(2):391–395, 1994.
- [6] Lawrence Chan, Andrew Critch, and Anca Dragan. Human irrationality: both bad and good for reward inference. *arXiv preprint arXiv:2111.06956*, 2021.
- [7] Andrew M Colman. Cooperation, psychological game theory, and limitations of rationality in social interaction. *Behavioral and brain sciences*, 26(2):139–153, 2003.
- [8] J. Correa, S. Lee, and E. Bareinboim. Nested counterfactual identification from arbitrary surrogate experiments. In *Advances in Neural Information Processing Systems*, volume 34, 2021.
- [9] Juan Correa and Elias Bareinboim. A calculus for stochastic interventions: Causal effect identification and surrogate experiments. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 10093–10100, 2020.
- [10] Tom Everitt, Ryan Carey, Eric D Langlois, Pedro A Ortega, and Shane Legg. Agent incentives: A causal perspective. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 11487–11495, 2021.
- [11] Ernst Fehr and Klaus M Schmidt. A theory of fairness, competition, and cooperation. *The quarterly journal of economics*, 114(3):817–868, 1999.
- [12] Mauricio Gonzalez-Soto, Luis E Sucar, and Hugo J Escalante. Causal games and causal nash equilibrium. *arXiv preprint arXiv:1910.06729*, 2019.
- [13] York Hagmayer and Steven A Sloman. Decision makers conceive of their choices as interventions. *Journal of experimental psychology: General*, 138(1):22, 2009.
- [14] Lewis Hammond, James Fox, Tom Everitt, Alessandro Abate, and Michael Wooldridge. Equilibrium refinements for multi-agent influence diagrams: theory and practice. *arXiv preprint arXiv:2102.05008*, 2021.

- [15] Lewis Hammond, James Fox, Tom Everitt, Ryan Carey, Alessandro Abate, and Michael Wooldridge. Reasoning about causality in games. *Artificial Intelligence*, 320:103919, 2023.
- [16] John C Harsanyi. Games with incomplete information played by “bayesian” players, i–iii part i. the basic model. *Management science*, 14(3):159–182, 1967.
- [17] Nigel Howard. Paradoxes of rationality: theory of metagames and political behavior. (*No Title*), 1971.
- [18] Ronald A Howard, RM Oliver, and JQ Smith. From influence to relevance to knowledge influences diagrams, belief nets and decision analysis. *Eds. RM Oliver and JQ Smith, John Wiley & Sons Ltd*, pages 3–23, 1990.
- [19] Daniel Kahneman. A perspective on judgment and choice: Mapping bounded rationality. *American Psychologist*, 58(9):697–720, 2003. doi: 10.1037/0003-066X.58.9.697.
- [20] Daniel Kahneman. Thinking, fast and slow. *Farrar, Straus and Giroux*, 2011.
- [21] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2):263–291, 1979. ISSN 00129682, 14680262. URL <http://www.jstor.org/stable/1914185>.
- [22] Daniel Kahneman and Amos Tversky. Choices, values, and frames. *American psychologist*, 39(4):341, 1984.
- [23] Daniel Kahneman and Amos Tversky. Prospect theory: An analysis of decision under risk. In *Handbook of the fundamentals of financial decision making: Part I*, pages 99–127. World Scientific, 2013.
- [24] Michael Kearns, Michael L. Littman, and Satinder Singh. Graphical models for game theory. UAI’01, page 253–260, San Francisco, CA, USA, 2001. Morgan Kaufmann Publishers Inc. ISBN 1558608001.
- [25] Daphne Koller and Brian Milch. Multi-agent influence diagrams for representing and solving games. *Games and economic behavior*, 45(1):181–221, 2003.
- [26] Steffen L Lauritzen and Dennis Nilsson. Representing and solving decision problems with limited information. *Management Science*, 47(9):1235–1251, 2001.
- [27] Sanghack Lee and Elias Bareinboim. Characterizing optimal mixed policies: Where to intervene and what to observe. *Advances in neural information processing systems*, 33:8565–8576, 2020.
- [28] George Loewenstein. The role of affect in decision making. *Handbook of Affective Sciences/Oxford University Press*, 2003.
- [29] John F Nash Jr. Equilibrium points in n-person games. *Proceedings of the national academy of sciences*, 36(1):48–49, 1950.
- [30] William Nichols and David Danks. Decision making using learned causal structures. In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 29, 2007.

- [31] Judea Pearl. *Causality*. Cambridge university press, 2009.
- [32] Judea Pearl and Dana Mackenzie. *The book of why: the new science of cause and effect*. Basic books, 2018.
- [33] A. Raghavan and E. Bareinboim. Counterfactual realizability and decision-making. In *The 13th International Conference on Learning Representations*, 2025. forthcoming.
- [34] Arvind Raghavan and Elias Bareinboim. Counterfactual realizability. In *The Thirteenth International Conference on Learning Representations*, 2025. URL <https://openreview.net/forum?id=uuriavczkL>.
- [35] Tim Roughgarden. Algorithmic game theory. *Communications of the ACM*, 53(7):78–86, 2010.
- [36] Kangrui Ruan, Junzhe Zhang, Xuan Di, and Elias Bareinboim. Causal imitation learning via inverse reinforcement learning. In *The Eleventh International Conference on Learning Representations*, 2023.
- [37] Lloyd S Shapley. A value for n-person games. *Contribution to the Theory of Games*, 2, 1953.
- [38] Steven A Sloman and York Hagmayer. The causal psycho-logic of choice. *Trends in cognitive sciences*, 10(9):407–412, 2006.
- [39] Ming Tan. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the tenth international conference on machine learning*, pages 330–337, 1993.
- [40] Amos Tversky and Daniel Kahneman. Judgment under uncertainty: Heuristics and biases: Biases in judgments reveal some heuristics of thinking under uncertainty. *science*, 185(4157): 1124–1131, 1974.
- [41] Amos Tversky and Daniel Kahneman. The framing of decisions and the psychology of choice. *science*, 211(4481):453–458, 1981.
- [42] Amos Tversky and Daniel Kahneman. Causal schemas in judgments under uncertainty. In *Progress in social psychology*, pages 49–72. Psychology Press, 2015.
- [43] Sidney J Yakowitz and John D Spragins. On the identifiability of finite mixtures. *The Annals of Mathematical Statistics*, 39(1):209–214, 1968.
- [44] Junzhe Zhang and Elias Bareinboim. Transfer learning in multi-armed bandit: a causal approach. In *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, pages 1778–1780, 2017.

Contents

1	Introduction	1
1.1	Preliminaries	3
2	Causal Normal Form Games	3
3	Learning Causal Nash Equilibrium	5
3.1	Ctf-Nash-Learning	5
4	Experiments	6
5	Conclusions	6
A	Preliminaries and Background	12
A.1	Structural Causal Models and PCH	12
A.2	Counterfactual Randomization	13
A.3	Normal Form Games and Nash Equilibrium	15
A.3.1	What is a Normal-Form Game?	15
A.3.2	Mixed Strategies and Best Responses	16
A.3.3	Nash Equilibrium	16
A.3.4	Example: The Prisoner’s Dilemma	16
A.4	Paradox of Rationality	17
A.4.1	Causal Game Theory and Paradox of Rationality	18
A.5	L_1, L_2 and L_3 actions in single-agent systems: The greedy Casino	18
A.6	Graphical Models and Game Theory	20
A.7	Notes on Hammond et al. [15]	21
B	Causal Nash Equilibrium	25
B.1	Counterfactual rationality and causal strategies	25
B.2	When is a larger action space harmful?	26
B.3	PCH projection and the Layer Selection Game	26
B.4	Causal Nash Equilibrium	27
B.5	Computing CNE: Find-CNE	28
C	Proofs	29
C.1	Proof of Theorem 7	29
C.2	Proof of Theorem 22	30
C.3	Proof of Theorem 23	30
C.4	Proof of Theorem 8	31
D	Discussion of Causal Games & Information sources	32
D.1	Standard versus Causal Normal-Form Games	33
D.2	Correlated Equilibrium	34
D.3	Bayesian Games	37
D.4	CNFG and Correlated Equilibrium	38
D.5	CNFG and Bayesian Games	39

E	Additional Examples and Discussion	41
E.1	SCM for Table 10	41
E.2	Pseudocode for Ctf-Nash-Learning	41
E.3	Assumptions for Alg. 3	41
E.4	Ctf-Nash Learning on Causal Prisoner’s Dilemma	42
E.5	Forgetting in Prisoner’s Dilemma	44
F	PCH Layers and Level-k Rationality: Orthogonal Axes	44
F.1	An Example:	45
F.2	Level-0 Default action vs L_1 action	45
G	FAQ	46

Appendix A. Preliminaries and Background

A.1. Structural Causal Models and PCH

Structural Causal Models is a general class of data-generating models [2, 31] that allows three types of distributions based on three levels of interaction with the system: observational, interventional, and counterfactual. First, we will give the formal definitions of these concepts and the hierarchical relation among them, known as Pearl Causal Hierarchy (PCH). Our presentation mostly follows Bareinboim et al. [4].

Definition 9 (Structural Causal Models) *A structural causal model \mathcal{M} is a 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$, where*

- \mathbf{U} is a set of background variables, also called exogenous variables, that are determined by factors outside the model;
- \mathbf{V} is a set $\{V_1, V_2, \dots, V_n\}$ of variables, called endogenous, that are determined by other variables in the model — that is, variables in $\mathbf{U} \cup \mathbf{V}$.
- \mathcal{F} is the set of functions $\{f_1, f_2, \dots, f_n\}$ such that each f_i is a mapping from (the respective domains of) $U_i \cup Pa_i$ to V_i , where $U_i \subseteq \mathbf{U}$, $Pa_i \subseteq \mathbf{V} \setminus V_i$, and the entire set \mathcal{F} forms a mapping from \mathbf{U} to \mathbf{V} , that is for each $i = 1, 2, \dots, n$, we have $v_i \leftarrow f_i(pa_i, u_i)$;
- $P(\mathbf{U})$ is the distribution over \mathbf{U} .

One way to visualize the dependence among the variables in the SCM is through a causal diagram, formal construction of which is given below (Def. 13, [4]).

Definition 10 (Causal Diagram (Semi-Markovian Models)) *Given an SCM $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$, a causal diagram G of \mathcal{M} is constructed as follows:*

1. add a vertex for every endogenous variable in the set \mathbf{V}
2. add an edge $(V_i \rightarrow V_j)$, for every $V_i, V_j \in \mathbf{V}$ and V_i occurs as an argument in $f_j \in \mathcal{F}$.
3. add a bidirected edge $(V_i \leftarrow \dots \rightarrow V_j)$ for every $V_i, V_j \in \mathbf{V}$ if the corresponding $U_i, U_j \in \mathbf{U}$ are correlated or the corresponding functions f_i, f_j share some $U \in \mathbf{U}$ as an argument.

Next, we define three types of distributions corresponding to distinct modes of interaction with an SCM: the L_1 (observational), L_2 (interventional), and L_3 (counterfactual) distributions (Defs. 2, 5, and 7 in [4]).

Definition 11 (L_1 valuation) *An SCM $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$ defines a joint probability distribution $P^{\mathcal{M}}(\mathbf{V})$ such that for each $\mathbf{Y} \subseteq \mathbf{V}$:*

$$P^{\mathcal{M}}(\mathbf{y}) = \sum_{\mathbf{u} | \mathbf{Y}(\mathbf{u}) = \mathbf{y}} P(\mathbf{u}) \quad (2)$$

Before we define L_2 evaluations, we need to understand interventional SCMs. Let \mathcal{M} be an SCM and \mathbf{x} be an assignment to $\mathbf{X} \subseteq \mathbf{V}$. Then the interventional SCM $\mathcal{M}_{\mathbf{x}}$ is the 4-tuple $\langle \mathbf{U}, \mathbf{V}, \mathcal{F}_{\mathbf{x}}, P(\mathbf{U}) \rangle$, where $\mathcal{F}_{\mathbf{x}} = \{f_i : V_i \notin \mathbf{X}\} \cup \{\mathbf{X} \leftarrow \mathbf{x}\}$. This operation is also known as the $do(\mathbf{x})$ operation.

Definition 12 (L_2 valuation) An SCM $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$ induces a family a joint probability distributions over \mathbf{V} , one for each intervention \mathbf{x} . For each $\mathbf{Y} \subseteq \mathbf{X}$,

$$P^{\mathcal{M}}(\mathbf{y}_{\mathbf{x}}) = \sum_{\mathbf{u} | \mathbf{Y}_{\mathbf{x}}(\mathbf{u}) = \mathbf{y}} P(\mathbf{u}) \quad (3)$$

where $\mathbf{Y}_{\mathbf{x}}(\mathbf{u}) = \mathbf{Y}_{\mathcal{M}_{\mathbf{x}}}(\mathbf{u})$

Such an operation, where the values of random variables \mathbf{X} are set to constant values \mathbf{x} , is known as hard interventions. Conditional or stochastic interventions can be defined similarly [9]. Let $\sigma_{\mathbf{X}} = \{\sigma_X\}_{X \in \mathbf{X}}$ be the set of soft-interventions on the variables $X \in \mathbf{X}$. Given $\sigma_{\mathbf{X}}$, the new model $M_{\sigma_{\mathbf{X}}}$ is defined as $\langle \mathbf{V}, \mathbf{U} \cup \mathbf{U}'_{\mathbf{X}}, \mathcal{F}', P(\mathbf{U} \cup \mathbf{U}'_{\mathbf{X}}) \rangle$, where $\mathbf{U}'_{\mathbf{X}} = \{\mathbf{U}'_X\}_{X \in \mathbf{X}}$ and $\mathcal{F}' = \{\mathcal{F} \setminus \{f_X\}_{X \in \mathbf{X}}\} \cup \{f'_X\}_{X \in \mathbf{X}}$. The distribution $P(\mathbf{V}; \sigma_{\mathbf{X}})$ can then be computed as $P(\mathbf{V})$ in $M_{\sigma_{\mathbf{X}}}$.

Next, we move on to the L_3 distributions, where we ask questions of the form ‘‘Given that the patient died without the treatment, would they be alive if they were given the treatment?’’. The first thing to note here, is that this query is not the same as treatment effect, that is $E[Y_{x=1}] - E[Y_{x=0}]$, where we are taking the difference between the average effects of giving the treatment and not giving the treatment. On the other hand, in the counterfactual question, we are asking the question if it would have helped for the same individual. Now, the difficulty of this problem, lies in the fact, that the patient was already denied treatment and died, and it is not practical to go back in time and give them the treatment. Mathematically, if Y is the variable that denotes whether is the patient is alive and X be the variable that the patient was given the treatment, we can write the above question as $P(Y_{x=1} = 1 \mid X = 0, Y = 0)$. Now, we provide a formal definition on how to compute counterfactual queries, given an SCM.

Definition 13 (L_3 valuation) An SCM $\mathcal{M} = \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P(\mathbf{U}) \rangle$ induces a family of joint distributions over counterfactual events $\mathbf{y}_{\mathbf{x}}, \dots, \mathbf{z}_{\mathbf{w}}$ for $\mathbf{Y}, \mathbf{Z}, \dots, \mathbf{W}, \mathbf{X} \in \mathbf{V}$:

$$P^{\mathcal{M}}(\mathbf{y}_{\mathbf{x}}, \dots, \mathbf{z}_{\mathbf{w}}) = \sum_{\mathbf{u} | \mathbf{Y}_{\mathbf{x}}(\mathbf{u}) = \mathbf{y}, \dots, \mathbf{Z}_{\mathbf{w}}(\mathbf{u}) = \mathbf{z}} P(\mathbf{u}) \quad (4)$$

The collection of observational (L_1), interventional (L_2) and counterfactual (L_3) are together called the PCH.

A.2. Counterfactual Randomization

In practice, interacting through L_3 or counterfactual layer of the Pearl Causal Hierarchy can be extremely nontrivial. To this end, [3] introduces a novel form of randomization to interact through the Layer 3 of PCH. The challenge stems from the observation that agents may consider various alternatives during the deliberation process and change their opinion about the best course of action. For example, if the agent initially considers $X = x$, and then reconsiders and changes to $X = x'$, is this counterfactual action where the natural intuition was x and the performed action would be x' ? What if the agent reconsiders their decision again and changes it to x ; is the agent acting against their intuition? What is the intuition x or x' . Counterfactual randomization [2, 3] addresses this concern.

Algorithm 1: `ctf-RCT`;;
 Counterfactual Randomized Controlled Trials in MAB

Input: domain of actions $\mathcal{D}(X)$, total number of trials $N \in \mathbb{N}$

for $t = 1, 2, \dots$ **do**

- Perceive intended action $X^{(t)}$ and store it;
- if** $t \leq N$ **then**
 - Sample realized action $X'^{(t)} \sim \text{Unif}(\mathcal{D}(X))$;
- else**
 - Set $X'^{(t)} = \arg \max_x \widehat{\mathbb{E}}^{(N)}[Y_{X \leftarrow x} \mid X = X^{(t)}]$;
- end**
- Perform $\text{do}(X'^{(t)})$ and receive reward $Y^{(t)}$;

end




Figure 3: Illustration of decision flow f_X

The main idea is that the agent may consider many options during the deliberation process, but only the final choice matters. Consider the deliberation process shown in Fig. 3: at time step $T = 1$, the agent intends to play $X = x_1$ but reconsiders, thinking it might be sub-optimal, and decides to switch to $X = x_2$ instead, where $x_1 \neq x_2$. As time passes, the agent may realize that $X = x_{t-1}$ was not ideal and switch to an alternative, $X = t$. Ultimately, the final decision defines the intuition of the agent, regardless of the path taken to reach it. In practice, the agent could also in this reasoning process forever without ever reaching a decision.

This challenge calls for novel counterfactual machinery to allow for the counterfactual interaction following layer 3. Bareinboim et al. [3] introduces counterfactual randomization in which an agent is interrupted just before the execution of the choice, the choice being taken as the natural intuition and the final action executed based on this intuition and flip of a coin. Further, Bareinboim et al. [2] also introduces `ctf-RCT`, where an intuition is observed and then an action is chosen at random for execution. This allows us to sample from counterfactual distributions of the form $P(Y_x \mid x') E[Y_x \mid x']$, where Y is the outcome, x is the intervened value and x' is the intuition, measured just before the decision. The algorithm is shown in Alg. 1. For more details on this procedure in the single-agent setting, please refer to Bareinboim et al. [2, Sec. 7]

A.3. Normal Form Games and Nash Equilibrium

In many settings – from economics and political science to computer science and biology – multiple decision-makers interact strategically, each trying to achieve the best possible outcome for themselves. A *normal-form game* provides a compact way to model such one-shot interactions, and the concept of *Nash equilibrium* captures the idea of a stable outcome where no individual can benefit by unilaterally changing their choice. In this section of the appendix, we walk through these ideas step by step, illustrating them with the classical Prisoner’s Dilemma and with the intent of contrasting this later on with other variations and approaches.

A.3.1. WHAT IS A NORMAL-FORM GAME?

Intuitively, a normal-form game asks:

“If each player picks an action simultaneously, how do their combined choices determine everyone’s payoffs?”

This question leads to the following definition of a game:

Definition 14 (Normal-Form Game) *A finite n -player normal-form game is a tuple:*

$$G = \langle N, A, u \rangle$$

where

- $N = \{1, 2, \dots, n\}$ is the set of players.
- $A = A_1 \times A_2 \times \dots \times A_n$, with each A_i a finite set of actions available to player i . An element $a = (a_1, a_2, \dots, a_n)$ is called an action profile.
- $u = (u_1, u_2, \dots, u_n)$ is a collection of payoff functions, one per player:

$$u_i : A \longrightarrow \mathbb{R}, \quad a \mapsto u_i(a).$$

Given a profile a , $u_i(a)$ tells us how much player i “earns” (or how happy they are) under that combination of actions.

To summarize, in Normal Form Games:

- Each player i simultaneously chooses an action $a_i \in A_i$.
- Once all choices $a = (a_1, \dots, a_n)$ are made, each player i receives payoff $u_i(a)$.

Despite their simplicity, Normal-Form Games are extremely powerful in their expressive power and many other richer representations, such as Extensive Form and Bayesian Games, can be reduced to a Normal Form equivalent. Next, we look at how agents can and should behave in a normal form games.

A.3.2. MIXED STRATEGIES AND BEST RESPONSES

Rather than committing to a single action, players may *randomize* over their options. A *mixed strategy* for player i is simply a probability distribution over A_i . Denote by

$$S_i = \Delta(A_i)$$

the set of all such distributions, and by $S = S_1 \times \dots \times S_n$ the collection of all players' mixed strategies.

Given that the other players use some mixed-strategy profile $s_{-i} \in S_{-i}$, player i will choose a distribution $s_i \in S_i$ to maximize their expected payoff

$$u_i(s_i, s_{-i}) = \sum_{a \in A} [s_i(a_i) \times s_{-i}(a_{-i})] u_i(a).$$

Definition 15 (Best Response) A mixed strategy $s_i^* \in S_i$ is a best response to opponents' strategy profile s_{-i} if

$$u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i}) \quad \text{for every } s_i \in S_i.$$

In other words, s_i^* gives player i the highest possible expected payoff, assuming the others stick to strategy s_{-i} .

The notion of best response will play a key role in understanding agent's behavior.

A.3.3. NASH EQUILIBRIUM

A Nash equilibrium is a collection of strategies – one per player – such that each player's choice is a best response to everyone else's. No one can gain by deviating alone. Such a notion gives a concept of stability in a game, or a type of a solution.

Definition 16 (Nash Equilibrium) A mixed-strategy profile $s^* = (s_1^*, \dots, s_n^*)$ is a Nash equilibrium if, for every player i ,

$$u_i(s_i^*, s_{-i}^*) \geq u_i(s_i, s_{-i}^*) \quad \text{for all } s_i \in S_i.$$

A.3.4. EXAMPLE: THE PRISONER'S DILEMMA

To see these definitions in action, consider the *Prisoner's Dilemma*, a two-player game where each must choose either to cooperate (C) or defect (D). The actions of the player i can be written as:

$$A_i = \{C, D\}.$$

Their payoffs are given by the following matrix (first entry is player 1's payoff, second is player 2's):

		P2	
		C	D
P1	C	-1, -1	-7, -0.5
	D	-0.5, -7	-1.9, -1.9

Let's walk through each prisoner's incentives:

- If player 2 cooperates (C), then player 1's payoff is

$$u_1(C, C) = -1 \quad \text{vs.} \quad u_1(D, C) = -0.5.$$

player 1 is better off defecting (D), since $u_1(D, C) > u_1(C, C)$.

- If player 2 defects (D), then player 1's payoff is

$$u_1(C, D) = -7 \quad \text{vs.} \quad u_1(D, D) = -1.9.$$

Again, player 1 prefers to defect (D), since $u_1(D, D) > u_1(C, D)$.

- Hence, it is better for player 1 to defect D , irrespective of what player 2 does. By symmetry, player 2 likewise always prefers D whatever player 1 does.

Thus, each player's unique *best response* to the other is to defect. When both play their best responses, we reach the profile, (D, D) which is the game's *Nash equilibrium*. Ironically, although (C, C) would yield $(-1, -1)$ total payoff (mutual cooperation), rational self-interest drives both to (D, D) , giving only $-1.9 - 1.9 = 3.8$.

A.4. Paradox of Rationality

It has been observed throughout economics and behavioral game theory literature that irrationality can result in better outcomes than rational choices. One such example is the prisoner's dilemma. Both cooperating would be an irrational choice, but it results in a better payoff compared to fully rational players both of whom would choose to confess. Such irrational co-operations have also been observed in practice [7]. There has been several attempts in order to explain such irrationalities observed in human decision-making either through different models of bounded rationality, such as payoff transformations [22, 23, 41] or through alternate forms of reasoning [7].

Consider the example of the Travelers' Dilemma [5], where 2 travelers are asked to write the price of their lost item between \$2-100. One with the lower value receives the lower value + \$2 and one with the higher value receives lower value - \$2. If an agent just tries to maximize their own reward and do not reason over others, both of them will write \$100 and receive that. Now, if they do one step of reasoning, they will think "If I write \$99 and my opponent writes \$100 then, I will get \$101 and my opponent \$97". Hence, both will write \$99 and get \$99. The amount will decrease with more levels of reasoning. Irrational players again get higher payoffs than rational agents.

Basu [5] states that different thought processes lay behind different types of choices that people made playing a version of Traveler's Dilemma with the options ranging from 180 to 300 (pie chart): a spontaneous emotional response (choosing 300), a strategically reasoned choice (295–299) or a random one (181–294). Players making the formal rational choice (180) might have deduced it or known about it in advance. As expected, people making "spontaneous" or "random" selections took the least time to choose (as seen in experiments).




Figure 4: Causal graph for the Markovian Prisoner's Dilemma.

		Convict 2	
		$X_2 = 0$	$X_2 = 1$
Convict 1	$X_1 = 0$	-1, -1	-7, -0.5
	$X_1 = 1$	-0.5, -7	-1.9, -1.9

Figure 5: Payoff matrix for the Prisoner's Dilemma.

A.4.1. CAUSAL GAME THEORY AND PARADOX OF RATIONALITY

Our proposed framework can both model and explain this gap between theory and practice. First, we consider the modeling part through the example of Prisoner's Dilemma, which we will also call the Markovian Prisoners Dilemma. The causal graph for the Markovian PD is shown in Fig. 6. Let X_1 and X_2 be the action variables and their values 0 and 1 correspond to cooperating (C) and defect (D) respectively, and let their natural probability of cooperating be $P(X_i = 0) = 0.9$ for $i \in \{1, 2\}$. If someone is simply collecting observational data, it may happen that the agents are simply playing L_1 and hence the corresponding payoff is higher. Thus our modeling of games can model the irrational tendencies of the agents involved.

Next, comes the explaining part. Consider the scenario M_1 in Ex. 1, and the optimal equilibrium action, which is both the players playing L_1 . Note that, this is infact the best choice from the perspective of a player, who knows how to act in L_1, L_2 or L_3 . However, from an external observers point of view, these players are playing suboptimally, that is, playing C with probability 0.6 and D otherwise. From an external's observers' point of view, if the agents performed RCT, they would have a payoff as shown in Fig. 1 (bottom table), according to which playing D is the optimal strategy with a payoff $(-1.9, -1.9)$. However, the agents, playing seemingly irrationally somehow get a payoff of $(0, 0)$, creating a paradox in the mind of the external experiment designer.

A.5. L_1, L_2 and L_3 actions in single-agent systems: The greedy Casino

The following examples illustrates the limitations of traditional decision-making and how an agent can interact with the system through the three layers of PCH as first introduced by Bareinboim et al. [3]. Consider a casino introducing two new slot machines, denoted 0 and 1. Gamblers choose machines according to two unobserved binary factors: their level of inebriation ($D \in \{0, 1\}$) and whether a machine is blinking ($B \in \{0, 1\}$). Although these factors are hidden from the agent, they influence natural behavior through the rule $X = D \oplus B$, determining the arm $X \in \{0, 1\}$ a gambler is predisposed to choose.

The casino exploits this behavioral pattern by designing reactive slot machines that adjust payouts based on these hidden variables. While ensuring that payout rates meet a government-mandated minimum of 30% when players are assigned arms randomly (e.g., RCT during inspection), the ma-

chines covertly reduce payout rates for players who follow their natural inclinations. The effective payouts are given in Table 1.

	$D = 0$		$D = 1$	
	$B = 0$	$B = 1$	$B = 0$	$B = 1$
$X = 0$	0.10	0.50	0.40	0.20
$X = 1$	0.50	0.10	0.20	0.40

Table 1: Reactive slot machine payouts: bolded entries indicate natural arm choices under the rule $X = D \oplus B$.

Note, that while players are following their natural choice, the payoff is given by

$$E[Y] = \sum_{b,d} y \cdot P(y | X = d \oplus b, b, d) = 0.1 \quad (5)$$

On the other hand, if the inspectors do an RCT, the payoff is given by

$$E[Y | do(X = x)] = \sum_{b,d} y \cdot P(y | X = x, b, d) = 0.3 \quad (6)$$

for any $x \in \{0, 1\}$.

However, if the agents are following opposite of their natural intuition, then the payoff will be given by

$$E[Y_{X=0} | X = 1]P(X = 1) + E[Y_{X=1} | X = 0]P(X = 0) = 0.45 \quad (7)$$

which is significantly higher than the other strategies. In fact, we can now make the following observation:

$$E[Y_x | x'] > E[Y_{x'} | x'] = E[Y | x'] \quad (8)$$

for any $x \neq x'$. The first term corresponds to the scenario when the gambler wants to play x' , but then just before execution they choose x . The payoff of such a strategy is higher than the payoff of just playing their intuition. The term where they simply choose a single machine and play (as in RCT), $E[Y_{x'}]$ lies between these two strategies. In general L_3 strategies will always outperform L_1 and L_2 strategies, since both can be expressed as L_3 strategies.

In fact for single agent systems, it is always better for the agents to follow L_3 policy space, as L_3 space subsumes L_1 and L_2 policies.

$$\max_{\pi} \sum_x E[Y_{X=\pi(x)} | x]P(x) \geq \max\{\max_a E[Y_{X=a}], E[Y]\} \quad (9)$$

For more details, please refer Bareinboim et al. [2, 3].

A.6. Graphical Models and Game Theory

Several works have studied game theory from a graphical models perspective. The main emphasis has been on the computational advantages related to learning equilibria through probabilistic reasoning and corresponding optimization tools [24, 25]. This is a part of the growing and important literature known as algorithmic game theory Roughgarden [35]. Our approach addresses key gaps in existing models, particularly concerning the assumption of Markovianity, issues of irrationality, and multi-agent interactions.

Specifically, Kearns et al. [24] introduced *graphical games* to leverage graph structures for modeling interactions among players, making equilibrium computation more efficient when compared to standard Normal Form Games. Furthermore, Koller and Milch [25] extended influence diagrams [18, 26] to multi-agent settings, where decision nodes represent strategies, and probabilistic dependencies simplify equilibrium computations. Their framework was called Multi-Agent Influence Diagrams (MAIDs). The main goal of these works was connecting graphical models and game theory, and where somewhat silent with respect to how this relate to causality, including interventions and counterfactual reasoning.

The Structural Causal Influence Model by Everitt et al. [10] connects causality with the influence diagrams literature [18, 26]. They study certain notions found in this traditional literature, including value of information, value of control, among others. Their setting focuses on single-agent settings, whereas this paper considers multi-agent interactions, including more equilibrium analysis in scenarios where agents compete in a strategic manner. They also did not consider unobserved confounding, which is one of the key challenges in typical causal settings. Another form of causal games is proposed by Gonzalez-Soto et al. [12], which again focuses on actions as interventions and ignores the other layers of operations by a player.

Hammond et al. [14] extends Koller & Milch’s MAIDs by introducing the concept of MAID subgames and proposing equilibrium refinements such as subgame perfect and trembling hand perfect equilibria. The authors establish equivalence results between MAIDs and Extensive Form Games (EFGs), highlighting the computational advantages of MAIDs in representing and solving certain classes of games. Still, despite its power, this work does not explore causal implications or counterfactual strategies, which are central to our framework. Our model explicitly integrates these aspects for deeper insights into strategic decision-making and the meaning of rationality.

Unlike the Structural Causal Games framework in Hammond et al. [15], which assumes Markovian dynamics, our model handles non-Markovian influences, including unobserved confounding that impact both actions and payoffs. We note that the assumptions required to ascertain Markovianity are inapplicable in our setting, since one of our main goals is to account for irrational behavior – where the agent acts without knowing why. In a Markovian setting, the agent knows the reasons for acting in a particular way. In fact, we model irrationality through the notion of counterfactuals and extend equilibrium concepts beyond purely rational agents, as prescribed by Nash’s framework. A detailed comparison with this work is provided in the next section.

The approach proposed by Chan et al. [6] embeds irrationality in the Bellman equation under a Markovian assumption in a novel way. Our model, however, allows for general irrationality without specifying any functional constraints, which is necessary in a non-Markovian setting. The assumptions required to ascertain Markovianity are inapplicable in our setting, since one of our main goals is to account for irrational behavior – where the agent acts without knowing their reasons. Furthermore, while their focus is on a single-agent environment, ours is on multi-agent, strategic settings.

By bridging these gaps, our model provides a unified view of rational and irrational behaviors through a causal lens and rooted in first principles. It also extends graphical game-theoretic models to multi-agent systems, contributing to a more comprehensive understanding of equilibrium dynamics and rationality. Notably, while our work falls within the realm of causality, it is not primarily focused on its graphical aspects, as evident throughout the main body of the paper. As mentioned earlier, the central issue addressed here concerns the most fundamental decision-making setting and how counterfactual reasoning (and counterfactual randomization Bareinboim et al. [3]) can be leveraged to model and reconcile both irrational and rational behaviors, ultimately resolving the rationality paradox. We believe that the foundational understanding developed in this pervasive setting can be generalized to more complex games, where a graphical model and a more fine-grained structure could play a role, including for computational purposes.

A.7. Notes on Hammond et al. [15]

We note that Hammond et al. [15] also claim to unify causal modeling and game theory through the formalism of causal games and structural causal games (SCGs). In this section, we provide a critical examination of this claim and demonstrate that their approach fails to capture several essential aspects of causal modeling that are fundamental to understanding how agents reason and act in complex systems. Specifically, we identify four key limitations in their formalism: (1) the absence of the causal hierarchy and associated distributions, (2) the neglect of unobserved confounding, (3) the inability to represent L_1 actions even when extended with default functions, and (4) a flawed approach to counterfactual evaluation that misinterprets key semantic and identification issues. Each of these points is discussed in detail below.

1. Absence of the Causal Hierarchy and their distributions: One of the most basic features of causal modeling is the presence of different probability distributions induced by the collection of causal mechanisms, which is organized as three qualitatively different probability distributions – observational, interventional, and counterfactual [4, 32], and which are also known as the Pearl Causal Hierarchy (PCH). These three levels of distributions separate causal models from previously used graphical models such as Bayesian networks, and are considered a novel landmark in evaluation, estimation, representation, and decision-making in complex environments. As a consequence, an agent can interact with the system in three different ways corresponding to the three layers of PCH policies: L_1 , L_2 , and L_3 . However, Hammond et al. [15] collapse this fundamental hierarchy into a single layer by treating all agent actions as (hard and soft) interventions, disregarding the observational (L_1) and counterfactual (L_3) levels of reasoning completely. Recall the key definition of games introduced in this work [15, Def. 22]:

Definition 17 (SCG) *A (Markovian) SCG $M = (G, \theta)$ is a causal game over the exogenous and endogenous variables $\mathbf{E} \cup \mathbf{V}$ such that any deterministic parameterization of the decision variables of CPD π , the induced model with join distribution $P^\pi(V, E)$ is an SCM.*

The authors explain that:

“An SCG can be seen as an SCM without parameters for the decision variables. Given a policy π , we recover an SCM, as we explain in more detail below.”

As a result, an SCG itself does not define a natural distribution (L_1), because the decision variables, say X do not have a natural mechanism f_X and are only determined by the agents or the policy.




Figure 6: Causal Diagram for Markovian Prisoner's Dilemma

This precludes the possibility of modeling agents that interact in a L_1 or L_3 manner, significantly restricting the expressivity of SCGs in modeling how real-world agents reason. Still, in the formalism introduced in this paper, there is a natural distribution of the decision variables (Layer 1 in the PCH), hard/soft interventions (Layer 2), and counterfactual actions (Layer 3). Such fundamental features could not be captured by the models proposed in [15], which is illustrated in the following example.

Example 4 (Markovian Prisoner's Dilemma) *Two thieves are suspected of a crime and are captured. Unfortunately, there is not enough evidence to convict them. They can now cooperate ($X = 0$) or defect ($X = 1$). The payoffs for the actions of the convicts are shown in Fig. 5, where the numbers can be interpreted as the years they have to serve in prison. The Nash equilibrium of this game is when both players defect and the payoffs are $(-2, -2)$, where both players have no incentive to cooperate. The causal diagram for such a scenario is shown in Fig. 6.*

Consider the following scenarios – in the first one, let us call it M_1 , the prisoners are more loyal and their spontaneous instinct is to cooperate ($X = 0$) with a probability of 0.9 (disregarding their utility), and in the second scenario, called M_2 , the instinct is to cooperate ($X = 0$) with a probability of 0.5 and confess otherwise. Now, if the agents follow their natural instincts, their payoffs in the first scenario are

$$\mu_{L_1}^1 = \sum_{x_1, x_2} \mathbf{Y} \cdot P(X_1 = x_1)P(X_2 = x_2)P(\mathbf{Y} | x_1, x_2) = (-1.46, -1.46)$$

and in the second scenario is

$$\mu_{L_1}^1 = \sum_{x_1, x_2} \mathbf{Y} \cdot P(X_1 = x_1)P(X_2 = x_2)P(\mathbf{Y} | x_1, x_2) = (-2.5, -2.5)$$

If μ_{NE} is the NE payoff, we can see that

$$\mu_{L_1}^1 > \mu_{NE} > \mu_{L_1}^2 \quad (10)$$

Now, these natural distributions cannot be represented or modeled by an SCG, where the X_1 and X_2 are determined by interventions. Hence, agents cannot act in L_1 and SCGs cannot capture the subtlety in Eq. 10. In fact, both M_1 and M_2 result in the same SCG in Fig. 6 and mechanized SCG in Fig. 8. In fact M_1 and M_2 are only two instances of infinitely many more scenarios that can happen. Fig. 7 shows how the players' L_1 payoffs change with their probability of cooperating. The scenarios M_1 and M_2 are marked as A and B in the plots. The green line denotes L_1 payoffs equal to the NE payoff. Also, since there is no concept of natural actions, L_3 policies also do not exist in SCGs. For example, in M_2 if both the prisoners decide to act against their natural instinct, then the payoffs are $\mu_{L_3}^2 = (-1.46, -1.46)$ and hence $\mu_{L_3}^2 > \mu_{NE}$.

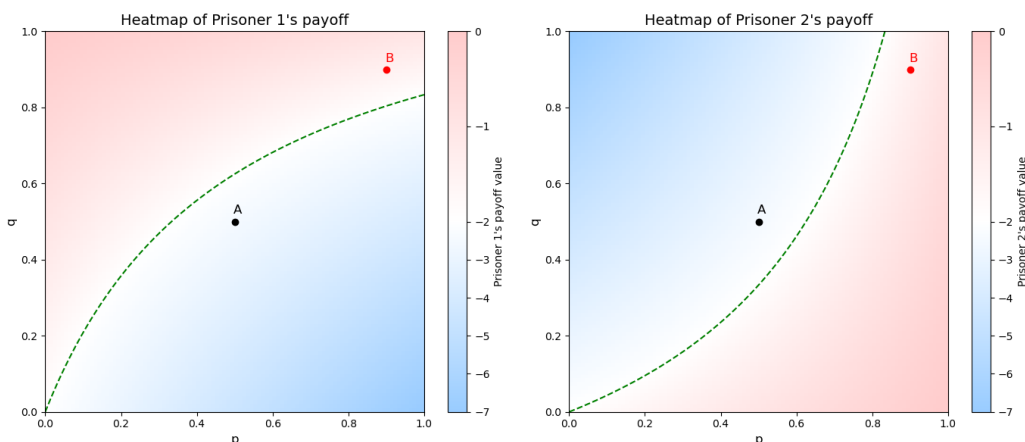


Figure 7: Payoffs of the two players as a function of their lying probability. A and B denote the payoffs when the agents follow their natural instincts in scenario M_1 and M_2 , respectively.




Figure 8: m-MAID is same for both M_1 and M_2 in Ex. 1

2. Unobserved Confounders: One of the major challenges in real-world settings that causal inference is concerned with is the existence of unobserved confounding, variables that cannot be measured but influence both decisions and outcomes. The classic saying that “causation is not association” comes precisely because of the existence of such confounders. However, Hammond et al. [15, Sec. 2.1] does not take unobserved confounders into account.

“In this paper, we make the simplifying assumption that all SCMs are Markovian, meaning that each variable V has exactly one exogenous parent E_V and the exogenous variables are independent.”

The problem becomes even more fundamental in the context of causal game theory, as highlighted in the Causal Prisoner Dilemma (Ex. 1). In words, both M_1 and M_2 imply the same SCG and mechanized-MAID (Fig. 8), but entail entirely different causal analysis and decision-making strategies. For example, in M_1 , following the L_1 strategies is the equilibrium strategy, and in M_2 following the L_3 strategies are the equilibrium strategies, and in both cases the equilibrium is better than the equilibrium payoffs of L_2 . However, this observation is not an idiosyncrasy but rather part of a broader phenomenon, as highlighted by the following proposition.

Theorem 18 (CNFG representation of SCG) *Given an SCG corresponding to a Normal Form Game, there exist two CNFGs \mathcal{C}_1 and \mathcal{C}_2 , with equilibrium payoffs μ_1 and μ_2 under L_1 and L_2 actions and Nash Equilibrium payoffs μ_{NE} , such that*

$$\mu_1 < \mu_{NE} < \mu_2 \quad (11)$$

Proof This result follows from Thm. 7. Recall Def. 22 from [15], and Ex. 1 of this work. Hence, if we follow the proposed construction, we see there is a bijection between a normal form game and its SCG representation, which implies that given a normal form game there is a single SCG which represents the same game as the normal form game. However, Thm. 7 implies the existence of two CNFGs satisfying Eq. 11. Hence, CNFG represents a larger class of models than SCGs for normal form games. ■

Formally, CNFGs is a strictly larger class than normal form games represented as SCG.

3. Why can't an SCG with default actions represent L_1 actions?

One may surmise that a Strategic Causal Game (SCG) could be extended to include a default action to represent an L_1 action. However, this would face two fundamental challenges:

1. **Defining the Default Action:** In CNFGs, L_1 actions are determined by nature (the SCM) through mechanisms that are entirely unknown to the agent. For instance, in Example 1, suppose we attempt to proxy an L_1 action in an SCG by defining a default function such as $X = U \oplus R$. But why should this specific function be chosen? It could just as well be $U \cdot R$, $U \vee R$, or any arbitrary combination of U and R . Since the agent has no knowledge of these variables or the functional form governing them, they cannot meaningfully prefer one default over another. This renders the notion of a default action indeterminate from the agent's perspective.
2. **Dependence on Unobserved Variables:** A default action that depends on unobserved variables (like U or R) inherently contradicts the assumption that these variables are unobservable to the agent. If the agent is able to use these variables as inputs in its decision-making process, then, by definition, they are no longer unobserved. This undermines the epistemic foundations on which the model is built. Put differently, this denies the possibility of unobserved confounding, a common challenge in causal inference.

These issues highlight a deeper point: SCGs and CNFGs represent fundamentally different models of agent-environment interactions. Attempting to equip SCGs with mechanisms to mimic CNFG behavior would ultimately collapse the SCG framework into that of CNFG.

4. Counterfactual Evaluation: Next, we look at the evaluation of counterfactual queries and the rationale behind computing such queries, as proposed in Hammond et al. [15]. In words, they want to answer “If we have evidence that the equilibrium π was played in the actual world, how and to what extent should that inform us of the equilibrium π' played in the counterfactual world where the values of some mechanism variables may have changed?” They claim to compute the quantity $P^{\pi'}(\mathbf{x}_I | z_\pi)$ through the following procedure:

- For every actual rational outcome $\pi \in \mathcal{R}(\mathcal{M} | z)$, update $P(\mathbf{u})$ to $P(\mathbf{u} | z)$ (‘abduction’)
- Apply the intervention I , on variables \mathbf{Y} , recomputing any rational responses to form π' and adding new exogenous variables as required (‘action’)

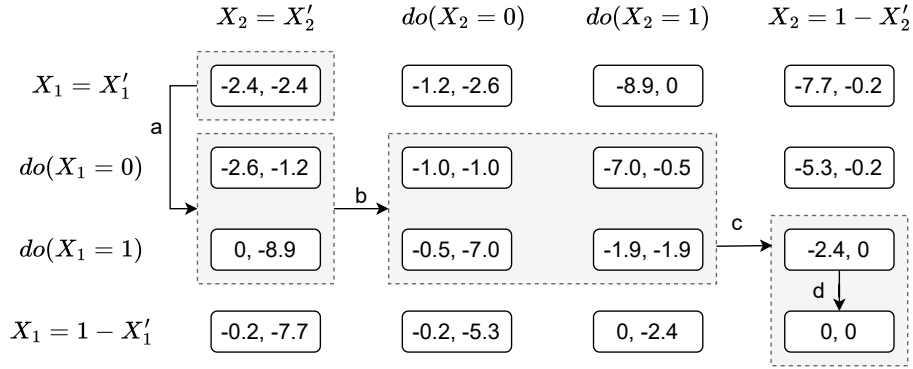


Figure 9: Change of Equilibrium with change of policies in the Causal Prisoner's Dilemma.

- Return each marginal distribution $\int_{D_{\mathbf{u}}} P\pi'(\mathbf{x} | \mathbf{u}')P(\mathbf{u}')d\mathbf{u}'$ in the modified model for each counterfactual rational outcome π' ('prediction')

Note that even though this follows Pearl's algorithm in theory, it misses a fundamental practical point: neither \mathbf{u} is observed, nor is $P(\mathbf{u})$ known, which makes the above evaluation impossible in most cases [31]. Formally, this procedure provides clear semantics for counterfactuals [4, Sec. 1.2], but it does not immediately imply their identification. Methods to overcome this impossibility of directly evaluating counterfactual quantities are known in the literature (see, e.g., [8, 33]). For example, one counterfactual quantity that [15] would be interested in, within our context from Example 1, is

$$P(Y_{X_1=1, X_2=0} = y | Y_{X_1=0, X_2=0} = y'), \quad (12)$$

which is known not to be identifiable from observational or interventional data without further assumptions. We note that their counterfactual evaluations do not correspond to the counterfactual actions studied in this paper. In contrast, we compute the quantity

$$\sum_x P(Y_{X_1=x', X_2=0} | X_1 = x)P(X_1 = x), \quad (13)$$

which SCGs cannot represent, since they do not define a natural distribution. Moreover, this quantity is counterfactually realizable and can be used in practice through counterfactual randomization [33].

Appendix B. Causal Nash Equilibrium

This appendix expands on the brief Causal Nash Equilibrium (CNE) discussion of Sec. 3. We introduce counterfactual rationality, the Layer Selection Game, the role of forgetting, and prove the existence and dominance theorems.

B.1. Counterfactual rationality and causal strategies

Allowing agents to transition between layers of the PCH leads to a two-step decision process. First, the agent determines which layers to operate in – instinct-based (L_1), classical rationality (L_2), or counterfactual reasoning (L_3). Second, the agent must decide which action to take within the chosen layer. We refer to this two-step process as a *causal strategy*. An agent is *counterfactually rational*

if it seeks to maximize its expected payoff using causal strategies, given that other agents are also counterfactually rational.

Example 5 (Equilibria in CPD) Consider Ex. 1 (M_2) where we analyze how the payoffs and equilibria evolve as agents move across the layers of the PCH, from instinct-based L_1 policies to counterfactual-based L_3 policies. Fig. 9 shows the payoff of the prisoners in this larger action space. If both prisoners follow their natural choices, playing L_1 , their payoffs are $(-2.4, -2.4)$.

Now, suppose prisoner 1 starts thinking rationally, ignoring their natural instincts, which results in transition (a) in Fig. 9. Prisoner 1 eventually defects, meaning they play the action $do(X_1 = 1)$, while prisoner 2 still follows their instinct, $X'_2 = X_2$. As a result, the payoffs become $(0, -8.9)$, where prisoner 1 benefits while prisoner 2 suffers. Eventually, prisoner 2 also learns to think rationally, leading to transition (b). In this case, both prisoners enter the realm of Standard Game Theory (SGT), each choosing to defect, playing the actions $(do(X_1 = 1), do(X_2 = 1))$. This results in NE with payoffs of $(-1.9, -1.9)$. First, the scope of SGT is highlighted in the four central cells of Fig. 9. Second, the equilibrium in SGT is worse than when both agents act irrationally (L_1). The SGT analysis stops at this point, but our framework suggests that strategic thinking may continue.

Over time, prisoner 2 introspects and contemplates counterfactual decisions, as highlighted in transition (c). They realize that their natural instincts provide insights that can be leveraged, and they should choose to act opposite to their natural choices, $X_1 = 1 - X'_1$. This yields payoffs of $(-2.4, 0)$, improving their baseline and hurting prisoner 1. Eventually, prisoner 1 also reaches L_3 , leading to transition (d). Both players, now operating under Causal Game Theory (CGT), settle on actions against their natural instincts, $X_1 = 1 - X'_1, X_2 = 1 - X'_2$, achieving payoffs of $(0, 0)$. This is the final state, where no unilateral deviation can increase payoffs. \square

The game in this example reflects an increasingly refined form of human rationality, tracing its evolution from primitive instincts based on raw intuition (L_1) to a notion of rationality based on game theory, where the intuition is ignored (L_2), and going to advanced strategic thinking leveraging both rational and irrational aspects of human cognition (L_3).

B.2. When is a larger action space harmful?

A natural question is whether it is always better to consider the full payoff table, since it provides the largest action space. Consider the example shown in Fig. 10. The full game specification is given in Appendix E. If Player 1's action space is limited to L_1 and L_2 , then the equilibrium payoff is $(0, 0)$ (marked in blue). However, if the action space L_3 is considered, the last row in the table is also considered (gray), and the equilibrium payoff decreases to $(-1, -1)$. Hence, regardless of what the other player does, Player 1's mere consideration of a larger action space harms them. Broadly, deciding which action space to follow is non-trivial.

B.3. PCH projection and the Layer Selection Game

Definition 19 (PCH Projection) Given a CNFG $\Gamma = \langle \mathbb{M}, \mathcal{A}, \mathcal{R} \rangle$, the PCH projection of Γ , denoted by $\Gamma(A_1, \dots, A_n)$, is the subgame of Γ where the action space of agent i is constrained to a subset $\mathcal{A}_i \supseteq A_i \in \{\mathcal{A}_i^1, \mathcal{A}_i^2, \mathcal{A}_i^1 \cup \mathcal{A}_i^2, \mathcal{A}_i^3\}$. \square

This projection captures how a game evolves when agents operate within a restricted subset of available strategies corresponding to different levels of reasoning within the PCH. The problem now

P1 \ P2	$X_2 = X'_2$	$X_2 = 0$	$X_2 = 1$
$X_1 = X'_1$	-2, 2	-2, -2	-2, -2
$X_1 = 0$	0, 0	-1.5, -1.5	-1.5, -1.5
$X_1 = 1$	0, 0	-1.5, -1.5	-1.5, -1.5
$X_1 = 1 - X'_1$	2, -2	-1, -1	-1, -1

Figure 10: It is not always optimal to jump to L_3 policy.

is to find a projection from where agents have no incentive to unilaterally deviate to a different layer of the PCH. To address this, we introduce a strategic layer selection game, a meta-game, where agents choose which layer of PCH to operate at.

Definition 20 (Layer Selection Game) *Given a CNFG $\Gamma = \langle \mathbb{M}, \mathcal{A}, \mathcal{R} \rangle$, its Layer Selection Game L_Γ is the NFG with (i) the same set of agents N , (ii) action space $A = A_1 \times \dots \times A_n$, where $\mathcal{A}_i \supseteq A_i \in \{\mathcal{A}_i^1, \mathcal{A}_i^2, \mathcal{A}_i^1 \cup \mathcal{A}_i^2, \mathcal{A}_i^3\}$ and (iii) utility $u(A) = NE(\Gamma(A_1, \dots, A_n))$, where $NE(\Gamma(A_1, \dots, A_n))$ is a Nash Equilibrium payoff of the CNFG Γ when actions spaces are restricted to A_1, \dots, A_n . \square*

This metagame represents a higher-level decision process, where each cell in the payoff matrix corresponds to a PCH projection of Γ , and its equilibrium will determine the layer of reasoning in which the agents should operate. We assume such counterfactual rationality is common knowledge: both players are aware that the other player can forget a part of their action space and choose the PCH layers in which they operate. Let s_i^* be the NE strategy for player i in the layer selection game. Let $\text{supp}(s_i^*)$ denote the support of s_i^* – the set of action spaces with non-zero probability in s_i^* . In particular, if $\mathcal{A}_i^j \notin \text{supp}(s_i^*)$, then the agent can ignore, or “forget” about this action space, and instead play a PCH projection of Γ that excludes \mathcal{A}_i^j . For instance, in Fig. 10, if Player 1 is able to forget that it can play L_3 , the payoff for the agent is (0, 0), which is higher than the payoff with playing L_3 , (-1, -1).

In practice, agents can limit their reasoning layers by restricting their capabilities: (i) at L_1 , agents act instinctively without requiring sampling mechanisms, (ii) at L_2 , agents may need access to randomization (e.g., coin flips) for mixed strategies, and (iii) at L_3 , agents must introspect, observe their intuition, and then decide how to act based on it, through more sophisticated procedures, such as ctf-randomization. Refusing to observe intuition renders L_3 inaccessible. Forgetting part of the action space may not always be beneficial: in the simple prisoner’s dilemma, agents who play only $\{C\}$ get payoff (-1, -1), but a deviator playing $\{C, D\}$ can defect for payoff (-0.5, -7.0), so forgetting defection is not in the agent’s interest (Appendix E).

B.4. Causal Nash Equilibrium

Definition 21 (Causal Nash Equilibrium, or CNE) *Let Γ be a CNFG and L_Γ be its corresponding layer selection game with NE strategy s^* . A strategy profile π^* is called CNE if π^* is the Nash Equilibrium of $\Gamma(A^*)$, where $A^* = A_1 \times \dots \times A_n$, and $A_i = \cup_{\mathcal{A} \in \text{supp}(s_i^*)} \mathcal{A}$. \square*

Theorem 22 (Existence of CNE) *For any CNFG, CNE always exists. \square*

P1 \ P2	\mathcal{A}^1	\mathcal{A}^2	$\mathcal{A}^1 \cup \mathcal{A}^2$
\mathcal{A}^1	-2, 2	-2, -2	-2, 2
\mathcal{A}^2	0, 0	-1.5, -1.5	0, 0
$\mathcal{A}^1 \cup \mathcal{A}^2$	0, 0	-1.5, -1.5	0, 0
\mathcal{A}^3	2, -2	-1, -1	-1, -1

Figure 11: Layer selection game for the example in Fig. 10.

P1 \ P2	\mathcal{A}^1	\mathcal{A}^2	$\mathcal{A}^1 \cup \mathcal{A}^2$	\mathcal{A}^3
\mathcal{A}^1	-2.4, -2.4	-8.9, 0	-8.9, 0	-8.9, 0
\mathcal{A}^2	0, -8.9	-1.9, -1.9	-1.9, -1.9	-2.4, 0
$\mathcal{A}^1 \cup \mathcal{A}^2$	0, -8.9	-1.9, -1.9	-1.9, -1.9	-2.4, 0
\mathcal{A}^3	0, -8.9	0, -2.4	0, -2.4	0, 0

Figure 12: Layer selection game for the Causal Prisoner's Dilemma.

If playing L_2 is a pure strategy NE of the layer selection game L_Γ , then the CNE of Γ in CGT and the NE of the normal form game induced by Γ coincide. Note that it is possible for a CNFG to have multiple layer selection games and CNEs. We write $NE(\Gamma(A); L_\Gamma)$ for the NE payoff with action space A as chosen in L_Γ .

Theorem 23 (Dominance of causal strategies) *Let Γ be a CNFG with CNE payoff μ^* and L_Γ be its layer selection game with NE strategy s^* . If s^* is a pure strategy NE and $A_i^* = \text{supp}(s_i^*)$, $\mu^* \geq NE(\Gamma(A_i, A_{-i}^*); L_\Gamma)$ for all $A_i \in \{A_i^1, A_i^2, A_i^1 \cup A_i^2, A_i^3\}$ and $i \in [n]$. \square*

In other words, Thm. 23 guarantees that if the layer selection game L_Γ admits a pure strategy NE, no agent benefits by unilaterally switching to a different PCH reasoning layer. Consider Fig. 11, which shows the layer selection game for the game in Fig. 10: if Player 1 follows L_2 policies and Player 2 follows L_1 and L_2 policies, neither has an incentive to switch to a different layer of PCH. Thus, the CNE payoff of Γ is $(0, 0)$, while the NE payoff of Γ with L_3 actions is $(-1.5, -1.5)$ and that with interventions is $(0, 0)$. In contrast, Fig. 12, corresponding to the CPD in Ex. 5, has a pure strategy NE at $(\mathcal{A}^3, \mathcal{A}^3)$, indicating both players should adopt L_3 policies. This is consistent with Fig. 9, resulting in a payoff $(0, 0)$ while the NE payoff in L_2 is $(-1.9, -1.9)$.

B.5. Computing CNE: Find-CNE

When the action spaces and corresponding payoffs of the CNFG Γ are known to both agents (as in SGT), Find-CNE (Alg. 2) implements the procedure above. The algorithm first constructs the layer selection game L_Γ corresponding to Γ and then computes its NE strategy. Any action space that occurs with nonzero probability in the NE strategy is used for CNE; the rest are discarded. The final step computes the NE of the projection of Γ with the restricted action space.

Algorithm 2: Find-CNE

Input: PCH projections of CNFG $\Gamma = \langle \mathbb{M}, \mathcal{A}, \mathcal{R} \rangle$
Output: CNE strategies π^*

 Construct the Layer Selection Game, L_Γ : For all $A = A_1 \times \dots \times A_n$, such that

 $\mathcal{A}_i \supseteq A_i \in \{\mathcal{A}^1, \mathcal{A}^2, \mathcal{A}^1 \cup \mathcal{A}^2, \mathcal{A}^3\}$, $u(A) \in \text{NE}(\Gamma(A_1, \dots, A_n))$;

 Let s^* be the NE strategy of L_Γ and $A^* = A_1^* \times \dots \times A_n^*$, where $A_i^* = \bigcup_{A \in \text{supp}(s_i^*)} A$;

return NE strategies of $\Gamma(A^*)$;

Appendix C. Proofs

C.1. Proof of Theorem 7

Consider a normal form game \mathcal{G} with the action space $A = A_1 \times \dots \times A_n$ and the utility function $r = (r_1, \dots, r_n)$. Assume all the utilities are finite. Suppose s^* is the NE strategy and μ_{NE} is the NE payoff.

Next, we will construct an SCM that induces the same Normal Form Game under L_2 actions as follows:

- $\mathbf{U} = \{U_1, \dots, U_n\}$, $D_{U_i} = A_i$ for all $i \in [n]$
- $\mathbf{X} = \{X_1, \dots, X_n\}$, $D_{X_i} = A_i$ for all $i \in [n]$
- $X_i = U_i$ for all $i \in [n]$
- $P(U_i = a_i^j) = s_i^*(a_i^j)$ where $s_i^*(a_i^j)$ is the probability of playing a_i^j by agent i in the NE strategy s_i^* .
- Now, we will define the observed variable the payoff of the i -th agent for $i \in [n]$. For $i \neq 1$, we define $Y_i(a, \mathbf{u}) = r_i(a)$, for all $a \in D_{\mathbf{X}}$. For $i = 1$, we have

$$Y_1(a, \mathbf{u}) = r_1(a) + M \cdot (\mathbb{1}\{U_1 = a_1\} - s_1^*(a_1)) \quad (14)$$

Now it is easy to see that for all $i \neq 1$,

$$\mathbb{E}[Y_i \mid do(a)] = \sum_{\mathbf{u}} Y_i(a, \mathbf{u}) P(\mathbf{u}) = r_i(a) \sum_{\mathbf{u}} P(\mathbf{u}) = r_i(a) \quad (15)$$

For $i = 1$ and for all $a \in D_{\mathbf{X}}$,

$$\mathbb{E}[Y_1 \mid do(a)] = \sum_{\mathbf{u}} Y_1(a, \mathbf{u}) P(\mathbf{u}) \quad (16)$$

$$= \sum_{\mathbf{u}} (r_1(a) + \mathbb{1}\{U_1 = a_1\} \cdot M \cdot (|A_1| - 1) - \mathbb{1}\{U_1 \neq a_1\} \cdot M) P(\mathbf{u}) \quad (17)$$

$$= r_1(a) + M \cdot (P(U_1 = a_1) - s_1^*(a_1)) \quad (18)$$

$$= r_1(a) \quad (19)$$

Hence, for L_2 action space, the SCM induces the same normal form game \mathcal{G} .

For Γ_1 , suppose M is a significantly large positive number and for Γ_2 , let M be a significantly large negative number, then we have for agent 1, the L_1 payoff is higher in Γ_1 and lower in Γ_2 than the NE payoff; for all other agents the payoff is the same. Hence, it follows, that

$$\mu_2 < \mu_{NE} < \mu_1$$

□

C.2. Proof of Theorem 22

Let Γ be a CNFG and consider the associated Layer Selection Game (LSG) L_Γ . By construction, L_Γ is a finite normal form game, since each agent chooses a reasoning layer from a finite set, and the payoffs are well-defined as the NE payoffs of the induced subgames. By Nash's theorem, L_Γ admits at least one mixed strategy Nash equilibrium. Let this equilibrium strategy profile be $s^* = (s_1^*, \dots, s_n^*)$.

For each agent i , let

$$A_i^* = \text{supp}(s_i^*) \quad (20)$$

the set of action spaces played with positive probability in s_i^* . Define the restricted action space

$$A^* = A_1^* \times \dots \times A_n^* \quad (21)$$

The PCH projection of Γ with action space A^* is then a subgame of Γ . This subgame can be represented in normal form, where each player's strategy space is finite. Hence, by Nash's theorem again, this subgame admits at least one Nash equilibrium.

The resulting equilibrium constitutes a Causal Nash Equilibrium (CNE) of the original CNFG Γ . Therefore, a CNE exists for every CNFG. Moreover, just as normal form games may have multiple Nash equilibria, CNFGs may admit multiple CNEs.

C.3. Proof of Theorem 23

First note that $\mu^* = NE(\Gamma(A^*))$. Suppose an agent is able to change the action space from A_i^* to A_i' and improve their payoff. However, if that was true, then $NE(\Gamma(A_i', A_{-i}^*)) > NE(\Gamma(A_i^*, A_{-i}^*))$, which implies in the PCH-LSG L_Γ , agent i would be able to improve the payoff moving from A_i^* to A_i' . However, by our assumption A^* is the pure strategy NE of L_Γ , hence no such deviations are incentivised - a contradiction. Hence $\mu^* \geq NE(\Gamma(A_i', A_{-i}^*))$ for all

Let Γ be a CNFG and let L_Γ be its layer-selection game. Assume L_Γ admits a pure-strategy Nash equilibrium, denoted $A^* = (A_1^*, \dots, A_n^*)$, where A_i^* is the action-space (PCH layer set) chosen by player i . Let $\mu^* = NE(\Gamma(A^*))$ be the Nash-equilibrium payoff vector of the PCH projection of Γ in which the action spaces are restricted to A^* .

Fix a player i and any alternative admissible action-space A_i' for player i , $A_i' \in \{A_i^1, A_i^2, A_i^1 \cup A_i^2, A_i^3\}$. Suppose, that deviating to A_i' would strictly improve player i 's payoff, that is,

$$NE_i(\Gamma(A_i', A_{-i}^*)) > NE_i(\Gamma(A^*)) = \mu_i^*.$$

By the definition of L_Γ , the payoff to player i from choosing an action-space is precisely the equilibrium payoff of the corresponding PCH-projected game. Hence, holding A_{-i}^* fixed, player i would obtain a strictly higher payoff in L_Γ by switching from A_i^* to A_i' . This contradicts the fact that A^* is a Nash equilibrium of L_Γ .

Therefore, for every player i and every admissible unilateral deviation A'_i ,

$$\mu_i^* \geq \text{NE}_i(\Gamma(A'_i, A_{-i}^*)).$$

Hence, the theorem follows.

C.4. Proof of Theorem 8

First, we will show that the payoff matrix learned is a permutation of the true payoff matrix, and then find out why L_2 or L_3 payoffs will be properly learned. First note that, since the mixture is identifiable, we recover the $D_{X_2} = k$ distributions, where each of them corresponds to a value of X'_2 . However the deduced values $\{\hat{x}_2^1, \dots, \hat{x}_2^k\}$ are arranged in decreasing order of distribution, that is

$$p(\hat{x}_2^1) > \dots > p(\hat{x}_2^k)$$

In reality, the original values of X'_2 may not be so well arranged and hence $\{\hat{x}_2^1, \dots, \hat{x}_2^k\} = h(\{x_2^1, \dots, x_2^k\})$ where h is a permutation function.

Now, L_3 action space consists of all the functions from natural intuition X'_2 to X_2 . Hence the values of X'_2 are essentially irrelevant and we can learn the whole table upto a permutation of the action of the second player. Since NE of Player 1 and the NE payoff remains same even with the permutation of the action space, we have that $\text{NE}(\Gamma(\mathcal{A}_1^3, \mathcal{A}_1^3))$ will be properly learned.

Now, L_2 action spaces are constant functions and remain invariant to permutations of X'_2 . Hence, in a similar manner $\text{NE}(\Gamma(\mathcal{A}_1^2, \mathcal{A}_2^2))$ will be correctly learned, as will $\text{NE}(\Gamma(\mathcal{A}_1^3, \mathcal{A}_2^2))$ and $\text{NE}(\Gamma(\mathcal{A}_1^2, \mathcal{A}_2^3))$, and so on. By our assumption, the NE strategy of the PCH-LSG for the other agent spans over \mathcal{A}_2^2 and \mathcal{A}_2^3 . Hence, the NE strategy of PCH-LSG lies on the space $\{\mathcal{A}_1^1, \mathcal{A}_1^2, \mathcal{A}_1^1 \cup \mathcal{A}_1^2, \mathcal{A}_1^3\} \times \{\mathcal{A}_2^2, \mathcal{A}_2^3\}$. Since, we are able to learn NE corresponding to each of these policies, we can correctly identify the CNE strategy.

Let $\Gamma = \langle M, (\mathcal{A}_1^3, \mathcal{A}_2^3), R \rangle$ be a two-player CNFG and let the data be generated by `Ctf-RCT`. For fixed (x'_1, x_1, x_2) the outcome distribution of Y is a finite mixture whose components correspond to the latent values of the opponent's intuition X'_2 ; by the identifiability assumption of the mixture, Algorithm 3 recovers the $k = |D_{X'_2}|$ component means and weights, but only up to a permutation of the component index. Formally, there exists a permutation π of $\{1, \dots, k\}$ such that the learned pairs $(\hat{p}_j(x'_1, x_1, x_2), \hat{R}_j(x'_1, x_1, x_2))$ satisfy

$$\hat{p}_j(x'_1, x_1, x_2) = p_{\pi(j)}(x'_1, x_1, x_2), \quad \hat{R}_j(x'_1, x_1, x_2) = R_{\pi(j)}(x'_1, x_1, x_2), \quad (22)$$

for all $j = 1, \dots, k$ and all (x'_1, x_1, x_2) . Hence the payoff table that Algorithm 3 constructs for the L_3 action game is isomorphic to the true payoff table via a relabeling of the columns indexed by the opponent's intuition values.

Consider first the projection $\Gamma(\mathcal{A}_1^3, \mathcal{A}_2^3)$. Player 1's L_3 action space is $F_1 = \{f : D_{X'_1} \rightarrow D_{X_1}\}$ and Player 2's L_3 action space is $F_2 = \{g : \{1, \dots, k\} \rightarrow D_{X_2}\}$. Because the learned table differs from the true table only by the permutation π of the intuition labels, replacing each $g \in F_2$ by $g \circ \pi$ yields a bijection between the learned and true games that preserves payoffs. Therefore the two normal form games $\Gamma(\mathcal{A}_1^3, \mathcal{A}_2^3)$ (true) and $\hat{\Gamma}(\mathcal{A}_1^3, \mathcal{A}_2^3)$ (learned) are strategically equivalent; in particular, they have the same set of Nash equilibria and the same equilibrium payoffs. Thus Algorithm 3 learns $\text{NE}(\Gamma(\mathcal{A}_1^3, \mathcal{A}_2^3))$ correctly.

Next consider projections that use L_2 actions for Player 2. An L_2 action for Player 2 is a constant function $c : \{1, \dots, k\} \rightarrow D_{X_2}$. Such constants are invariant under any permutation of the intuition labels; hence the learned payoff table for $\Gamma(A_1^3, A_2^2)$ (and, symmetrically, for $\Gamma(A_1^2, A_2^2)$ and $\Gamma(A_1^2, A_2^3)$) coincides with the true one. Consequently the corresponding Nash equilibrium payoffs are learned correctly.

By the hypothesis of the theorem, in the layer selection game L_Γ the equilibrium action space of Player 2 is $A_2 \in \{A_2^2, A_2^3\}$. From the previous paragraphs, for every action space A_1 of Player 1 that appears in L_Γ the value $\text{NE}(\Gamma(A_1, A_2))$ is learned correctly by Algorithm 2. Therefore the learned layer selection game coincides (up to relabeling) with the true one, and applying `FIND-CNE` to the learned payoff matrix returns the same best response set for Player 1 as in the true game. Hence the CNE strategy for Player 1 is correctly identified.

Appendix D. Discussion of Causal Games & Information sources

While the Causal Prisoner’s Dilemma highlights the difficulty of cooperation – and shows that counterfactual reasoning can improve upon standard Nash-like outcomes – other strands of the literature explore strategic interactions from orthogonal perspectives, including frameworks such as Correlated Equilibrium and Bayesian Games.

The concept of Correlated Equilibrium (CE), introduced by Aumann [1], generalizes Nash Equilibrium by allowing players to coordinate their strategies through signals from an external correlation device. Unlike in Nash equilibria, where each player optimizes independently, correlated equilibria permit coordinated play, which can yield higher social welfare. This framework is particularly effective in settings where cooperation can be facilitated by signals or mediators without direct communication. In addition, CE is often easier to compute and achieves better efficiency in certain games, especially when compared to independently derived strategies. Applications include traffic routing, bargaining, multi-agent learning, and regret minimization.

Another important extension is Bayesian Games, introduced by Harsanyi [16], which address strategic interactions under incomplete information. Here, players possess private information about their own types (e.g., preferences, available actions, or payoffs) but maintain beliefs about others’ types, often represented by probability distributions. This framework allows players to form and update strategies based on their beliefs. Bayesian Games naturally model scenarios involving uncertainty about the environment or about other agents, such as auctions, signaling games, contract theory, and mechanism design. From a computational standpoint, they introduce additional complexity due to the structure of beliefs and type spaces.

Comparing these frameworks to the standard Nash equilibrium setting shows how they enrich strategic analysis by incorporating coordination (in CE) and information asymmetry (in Bayesian Games). In this section, we offer a preliminary discussion on how these concepts relate to causal reasoning, and how counterfactual thinking may further enhance strategic decision-making in complex environments.

Specifically, we argue that the notion of information, as traditionally understood in the literature, is orthogonal to the causal structure captured in Causal Normal Form Games (CNFGs). This means that the causal framework can be naturally extended to incorporate sources of information available to agents. We first will revisit the standard definitions of normal form games (Sec. D.1), correlated equilibrium Sec. D.2, and Bayesian games (Sec. D.3) using formal causal language. We then in-

roduce CNFGs with information and compare them to these classical game-theoretic frameworks (Secs. D.4 and D.5), highlighting their expressive differences and extensions.

D.1. Standard versus Causal Normal-Form Games

To ground the discussion and establish a common denominator, we start with the definition discussed earlier, first of a causal normal form game (Def. 6):

Definition 24 (Causal Normal Form Game) A tuple $\Gamma = \langle \mathbb{M}, \mathcal{A}, \mathcal{R} \rangle$ is a Causal Normal Form Game (CNFG, for short), where

- \mathbb{M} is a CMAS $\langle M, N, \mathbf{X}, \mathbf{Y} \rangle$,
- $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ is the set of policies for the n agents, where $\mathcal{A}_i \in \{\mathcal{A}^1, \mathcal{A}^2, \mathcal{A}^1 \cup \mathcal{A}^2, \mathcal{A}^3\}$,
- $\mathcal{R} = (\mathcal{R}_1, \dots, \mathcal{R}_n)$ is the set of reward functions.

Next, we provide the standard definition of a normal form game.

Definition 25 A Normal-Form Game is defined as a 3-tuple $G = (N, A, u)$ where

- N : set of players.
- $A = A_1 \times \dots \times A_n$: action space, where A_i is the set of actions available to player i .
- $u = (u_1, \dots, u_n)$: utility functions, where $u_i : A \rightarrow \mathbb{R}$.

Whenever the policy space is constrained for the interventional layer (L_2), CNFG reduces to an NFG.

Theorem 26 A CNFG with L_2 actions can be converted to an NFG and vice versa with the same action space \mathcal{A} .

Proof The proof is constructive. CNFG with L_2 actions can be converted to a normal form game, simply by having the L_2 interventions as actions and $\mathbb{E}[\mathcal{R}_i(\mathbf{Y}_i(\mathbf{x}))]$ as the utility. The other way can be done as follows: define an action variable X_i for each of the n agents, where $D_{X_i} = A_i$. Then define $Y_i(x) = u_i(x)$ and \mathcal{R} as a set of identity functions. ■

Hence, a definition of a Normal Form Games in the causal terms would be:

Definition 27 A Normal-Form Game is defined as a 3-tuple $\Gamma = \langle \mathbb{M}, \mathcal{A}, \mathcal{R} \rangle$ where

- \mathbb{M} is a CMAS $\langle M, N, \mathbf{X}, \mathbf{Y} \rangle$,
- $\mathcal{A} = (\mathcal{A}_1^2, \dots, \mathcal{A}_n^2)$ is the set of L_2 policies for the n agents,
- $\mathcal{R} = (\mathcal{R}_1, \dots, \mathcal{R}_n)$ is the set of reward functions.

This gives an intuitive understanding of why and how CNFGs generalize NFGs – a claim made in Thm. 7. For concreteness, consider the following example.

		Player 2	
		$X_2 = B$	$X_2 = F$
Player 1	$X_1 = B$	2, 1	0, 0
	$X_1 = F$	0, 0	1, 2

Figure 13: Payoff matrix for Battle of Sexes

		P2		
		L_1	$X_2 = B$	$X_2 = F$
P1	L_1	1.875, 1.875	0.875, 1.375	1.375, 0.875
	$X_1 = B$	1.375, 0.875	2, 1	0, 0
	$X_1 = F$	0.875, 1.375	0, 0	1, 2

Figure 14: Payoff matrix for Battle of Sexes with L_1 and L_2 actions




Figure 15: Battle of Sexes with Information from a coin toss




Figure 16: Causal Diagram for Battle of Sexes with unobserved confounding

Example 6 (Prisoner’s Dilemma) Consider the classical Prisoner’s Dilemma game with payoffs as shown in Fig. 5. To represent this game as a CNFG, define X_1 and X_2 as the action variables in a CMAS, and $\mathbf{Y} = Y_1, Y_2$ as the corresponding reward signals. Assume X_1, X_2 are binary, where 0 represents cooperation (C) and 1 represents defection (D). Define Y_1 and Y_2 as deterministic functions of X_1 and X_2 , based on the given payoff table. For completeness, let $X_1 = U_1$ and $X_2 = U_2$, where the prior distribution over exogenous variables is $P(U_1 = 1) = P(U_2 = 1) = 0.5$. Note that this distribution is defined for formal completeness but does not affect the L_2 payoffs.

To convert this CNFG into an NFG, we focus on L_2 actions. The L_2 policy space \mathcal{A}^2 maps to the action values of X_1 and X_2 , i.e., either 0 (C) or 1 (D). The payoff for any joint action (x_1, x_2) is given by $\mathbb{E}[\mathbf{Y} \mid do(x_1, x_2)]$, aligning exactly with the utility function in the standard Normal Form representation. \square

As discussed earlier, Thm. 7 noted that a CNFGs is strictly more expressive than NFGs by showing two CNFGs that agree on the interventional layer but may have different equilibriums, when other layers of PCH are considered.

D.2. Correlated Equilibrium

In this section, we investigate how CNFGs can be extended to systems with information – an important step toward modeling more realistic decision-making scenarios. We begin by examining Correlated Equilibrium through a classical example known as the Battle of the Sexes.

Example 7 (Battle of Sexes) A couple of agents want to spend time together in the evening. Agent 1 wants to go to the ballet, while Agent 2 prefers a football match. Their payoffs, based on whether they go to the ballet or football, are shown in Tab. 13. The symmetric Nash equilibrium for this game occurs when both agents go to their preferred location two-thirds of the time, yielding a joint

payoff of (0.75, 0.75). We can represent this Normal Form Game as a CNFG with L_2 actions. The causal diagram corresponding to the SCM is shown in Fig. 15. The actions of the agent i in this SCM is given by $do(X_i = B)$ and $do(X_i = F)$ for $i \in \{1, 2\}$.

Suppose the players now have access to a coin and observe the outcome of the toss, H, T , making decisions accordingly. Assume they follow the strategy $\{H \rightarrow B, T \rightarrow F\}$ – that is, if the coin shows heads, they go to the ballet; if it shows tails, they go to the football match. Note that this is an equilibrium strategy, as neither player has an incentive to deviate from it. The resulting equilibrium payoff is (1.5, 1.5). Such an equilibrium is called a correlated equilibrium, and it is superior to the Nash equilibrium.

We can also represent this graphically in the causal diagram. In Fig. 15, a new variable C can be introduced, which takes two values $\{H, T\}$, each with probability 0.5. Now, since the outcome of the coin is available to the two agents, they can condition their policies on the outcome of this coin. So, the new policy space of the agents would be a mapping from outcome of coin to the show they want to attend, that is $A_i : \{H, T\} \rightarrow \{B, F\}$. If we are talking of mixed strategy, the policy of the agent is given by the distribution $\pi_i(\cdot | C)$ over the values $\{F, B\}$. Thus, we can simply represent additions of random variables in a correlated equilibrium as new variables in the causal model (also known as confounders).

We now formally define Correlated Equilibrium using causal language.

Definition 28 (Correlated Equilibrium) Given a CNFG $\Gamma = \langle \mathbb{M}, \mathcal{A}^2, \mathcal{R} \rangle$ with policy space restricted to L_2 , a correlated equilibrium $(\mathbf{S}, P(\mathbf{S}), \pi)$ is a tuple, where $\mathbf{S} = (S_1, \dots, S_n)$ is a tuple of random variables with distribution $P(\mathbf{S})$ and $\pi = (\pi_1, \dots, \pi_n)$ is the set of mappings $\pi_i : D_{S_i} \rightarrow \mathcal{A}_i^2$ and for each agent i and every other mapping π'_i ,

$$\begin{aligned} & \sum_{\mathbf{s} \in D_{\mathbf{S}}} P(\mathbf{s}) \mathcal{R}_i(\mathbf{Y}_{i[X_1=\pi_1(S_1), \dots, X_i=\pi_i(S_i), \dots, X_n=\pi_n(S_n)]}) \\ & \geq \sum_{\mathbf{s} \in D_{\mathbf{S}}} P(\mathbf{s}) \mathcal{R}_i(\mathbf{Y}_{i[X_1=\pi_1(S_1), \dots, X_i=\pi'_i(S_i), \dots, X_n=\pi_n(S_n)]}) \end{aligned} \quad (23)$$

As observed in the previous example, it is possible to represent the correlated equilibrium as a NE of a CNFG with information. Next, we formally define CNFG where the agents have access to pieces of information (possibly shared) before acting.

Definition 29 (CMAS with States) A Causal Multi-Agent System (CMAS) with states is a tuple $\langle M, N, \mathbf{X}, \mathbf{S}, \mathbf{Y} \rangle$, where $M : \langle \mathbf{U}, \mathbf{V}, \mathcal{F}, P \rangle$ is an SCM and

- N is the set of n agents,
- $\mathbf{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)$ is the ordered set of action nodes with $\mathbf{X}_i, \mathbf{X}_j \subset \mathbf{V}$ for $i, j \in [n]$ and $\mathbf{X}_i \cap \mathbf{X}_j = \emptyset$ if $i \neq j$,
- $\mathbf{S} = (\mathbf{S}_1, \dots, \mathbf{S}_n)$ is the ordered set of context nodes $\mathbf{S}_i \subset \mathbf{V}$ for the agent i for $i \in [n]$, and for all $i \in [n]$, $\mathbf{S}_i \notin De(\mathbf{X}_i)$
- $\mathbf{Y} = (\mathbf{Y}_1, \dots, \mathbf{Y}_n)$ is the ordered set of reward signals, with $\mathbf{Y}_i \subseteq \mathbf{V}$ for all $i \in [n]$. \square

Once we have introduced the notion of a CMAS to model the environment, we can consider the information available to the agents about the states.

		Suspect	
		$S = 1$	$S = 0$
Police	$P = 1$	0, 0	2, -2
	$P = 0$	-2, -1	-1, 1

Figure 17: Payoff when suspect is criminal (that is $T = 1$).

		Suspect	
		$S = 1$	$S = 0$
Police	$P = 1$	-3, -1	-1, -2
	$P = 0$	-2, -1	0, 0

Figure 18: Payoff when suspect is civilian (that is $T = 0$).

Definition 30 (CNFG with Information) A tuple $\Gamma = \langle \mathbb{M}, \mathcal{A}, \mathcal{R}, \mathcal{I} \rangle$ is a Causal Normal Form Game (CNFG), where

- \mathbb{M} is a CMAS with states $\langle N, M, \mathbf{X}, \mathbf{S}, \mathbf{Y} \rangle$,
- $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ is the set of policies for the n agents, where $\mathcal{A}_i \in \{\mathcal{A}^1, \mathcal{A}^2, \mathcal{A}^1 \cup \mathcal{A}^2, \mathcal{A}^3\}$,
- $\mathcal{R} = (\mathcal{R}_1, \dots, \mathcal{R}_n)$ is the set of reward functions.
- \mathcal{I} is the information available to the agents. □

The information \mathcal{I} can take many forms and is introduced to make the definition more complete and general. For example, one form of information might be the distribution over states, $P(\mathbf{S})$, which helps illustrate the relationship between Correlated Equilibrium and equilibrium concepts in CNFGs with states. Other forms of information available to the agents could include interventional or counterfactual distributions. While this is outside the scope of the present paper, it presents a promising direction for future work.

Theorem 31 If $(\mathbf{S}, P(\mathbf{S}), \pi)$ is a correlated equilibrium of NFG Γ , then, π is a NE of the CNFG with Information designed as follows:

1. $\mathbf{X} = (X_1, \dots, X_n)$ are the actions and $\mathbf{Y} = u(x)$ for each combination are obtained from NFG Γ . \mathcal{R} is identity for all agents.
2. Introduce variables $\mathbf{S} = (S_1, \dots, S_n)$ with distribution $P(\mathbf{S})$.
3. Define the L_2 action space \mathcal{A}_i^2 for the agent i as soft interventions $\pi(X_i | S_i)$
4. $P(\mathbf{S})$ is available to the agents and the expected payoff for policies π is given by:

Proof The proof follows from the definition of correlated equilibrium. Since π is the policy in the correlated equilibrium, it is also the best response as per Eq. 23. Hence, if every agent is playing the best response given \mathbf{S} , we have each agent is playing NE policy (from Def. 16) in the game with policies conditioned on \mathbf{S} . ■




Figure 19: Causal Diagram for Sheriff's Dilemma with L_2 policy space.




Figure 20: Causal Diagram for Sheriff's Dilemma with L_1 policy space.

D.3. Bayesian Games

Here, we introduce an additional layer of complexity in the information structure through the concept of Bayesian Games. Before presenting the formal definition, we begin with an example.

Example 8 (Sheriff's Dilemma) *A police officer faces an armed suspect, and they must simultaneously decide whether to shoot. The suspect could be either a criminal or a civilian, but the officer is unaware of the suspect's true identity. It is preferable for the suspect to shoot if they are a criminal and not to shoot if they are a civilian. However, in hindsight, it is better for the officer to shoot if the suspect shoots – but in reality, they must act simultaneously.*

Depending on the type of the player, criminal or civilian, the payoffs corresponding to the actions of the players are shown in Fig. 17 and Fig. 18 respectively. Now, the question is how do we compute the best policy for the agents.

Let's start with the suspect's actions. If the suspect is a criminal, then shooting is a dominant strategy and if the suspect is a civilian, not shooting is the dominant strategy. However, from the perspective of the policemen, things are not so simple, since they do not know the type of the suspect. If they know that the suspect is highly likely to be criminal, then it is better for them to shoot, and otherwise not. Hence, in order to form a decision, they need to have a belief over the likelihood of someone being a criminal.

The scenarios can be represented in a corresponding causal graph shown in Fig. 19. The variable T represents the type of the suspect: $T = 0$ indicates a civilian, and $T = 1$ indicates a criminal. The variable P captures the officer's decision to shoot or not, while S denotes whether the suspect chooses to shoot. If we consider L_2 actions only, then the policy space is a soft intervention over the values of P . Finally, $Y = (Y_1, Y_2)$ represent the utilities of the officer and the suspect, respectively. The value of Y as a function of P, T and S is shown in Fig. 18 and Fig. 17. In addition, the agent should also have a belief of the likelihood of someone being criminal, that is $P(T)$.

This class of Games where there is an uncertainty about the nature of the game is called Bayesian Game. Formally, Bayesian Games can be defined as follows [16]:

Definition 32 (Bayesian Games) *A Bayesian Game is a tuple $\langle N, A, \Theta, p, u \rangle$, where*

- N is the set of n players indexed by i ;
- $A = A_1 \times \dots \times A_n$, where A_i is the action set available to player i ;
- $\Theta = \Theta_1 \times \dots \times \Theta_n$ where Θ_i is the type space for player i ;

		$X_2 = B$		$X_2 = F$	
		$U_2 = 0$	$U_2 = 1$	$U_2 = 0$	$U_2 = 1$
$X_1 = B$	$U_1 = 0$	3,0	3,3	0,0	0,0
	$U_1 = 1$	0,0	2,1	0,0	0,0
$X_1 = F$	$U_1 = 0$	0,0	0,0	0,3	0,0
	$U_1 = 1$	0,0	0,0	3,3	1,2

Table 2: Battle of Sexes with Unobserved Confounding

- $p : \Theta \rightarrow [0, 1]$ is a common prior over types;
- $u = (u_1, \dots, u_n)$ where $u_i : A \times \Theta \rightarrow \mathbb{R}$ is the utility function for player i

Now, we formalize the idea of Bayesian Games in the causal framework.

Theorem 33 (Bayesian Games in Causal Framework) *A Bayesian Game is a CNFG with information $\Gamma = \langle M, \mathcal{A}^2, \mathcal{R}, \mathcal{I} \rangle$, where M is a CMAS with states.*

Proof The construction follows in the same way as Normal Form Games, but now with introduction of the type variables. The CMAS contains the variables \mathbf{X} corresponding to the actions A , state variables \mathbf{S} corresponding to the type Θ , $\mathbf{Y}(x) = u(x)$ and \mathcal{R} is the identity function. The information available to the agents is $P(\mathbf{S})$. ■

Now, a best response π_i^* to a strategy profile π_{-i} in a Bayesian Game is defined as

$$BR_i(\pi_{-i}) = \arg \max_{\pi_i'} EU(\pi_i', \pi_{-i}) \quad (24)$$

Definition 34 (Bayes-Nash Equilibrium) *A Bayes Nash Equilibrium is a strategy profile π such that for all i , $\pi_i \in BR(\pi_{-i})$.*

D.4. CNFG and Correlated Equilibrium

In Sec. D.2, we saw how we can represent correlated equilibrium with L_2 actions. However, SCMs can inherently represent two more layers of distribution L_1 and L_3 . Using the more general action space is not only a matter of choice, but can be essential in obtaining a better payoff (that is finding the corresponding equilibrium), as illustrated in the following example.

Example 9 (Causal Battle of Sexes) *Considering Ex.7, we note that, in reality, the decision to go to the ballet or the football game may be influenced by several external factors. For example, when a new ballet is released, agent 1 – who generally prefers football – may also want to attend the ballet. Conversely, if there is a major football event, such as the Super Bowl, agent 2 may also be happy to join agent 1 for the match. These unobserved factors can therefore influence the preferences of both players. The corresponding causal graph is shown in Fig.16.*

In addition, the importance of the event may also affect the payoffs. Let $U_1 = 0$ ($U_2 = 0$) indicate that there is a major football match (ballet performance), and $U_1 = 1$ ($U_2 = 1$) that the football match (ballet) is not particularly significant. Now, both agents' intuitions are given by:

$$X_i = \begin{cases} F & \text{if } U_1 = 0, U_2 = 1 \\ B & \text{if } U_1 = 1, U_2 = 0 \\ F \text{ or } B \text{ with equal probability} & \text{otherwise} \end{cases} \quad (25)$$

This means that if either the ballet or football performance is particularly good, the agents choose to attend that event. If both events are equally good or equally unappealing, they make the decision randomly. The payoff for such an L_1 action is $(1.875, 1.875)$. In contrast, if they instead base their decisions on signals from a coin toss, the resulting payoff is $(1.5, 1.5)$ – lower than what they would receive by following their natural intuitions.

We now define a correlated equilibrium over a general CNFG, where the agents' action spaces may correspond to L_1 , L_2 , or L_3 policies.

Definition 35 (Causal Correlated Equilibrium) *Given a CNFG $\Gamma = \langle \mathbf{M}, \mathcal{A}, \mathcal{R} \rangle$, a causal correlated equilibrium $(\mathbf{S}, P(\mathbf{S}), \pi)$ is a tuple, where $\mathbf{S} = (S_1, \dots, S_n)$ is a tuple of random variables with distribution $P(\mathbf{S})$ and $\pi = (\pi_1, \dots, \pi_n)$ is the set of mappings $\pi_i : D_{S_i} \rightarrow \mathcal{A}_i$ and for each agent i and every other mapping π'_i ,*

$$\begin{aligned} & \sum_{\mathbf{s} \in D_{\mathbf{S}}} P(\mathbf{s}) \mathcal{R}_i(\mathbf{Y}_{i[X_1=\pi_1(S_1), \dots, X_i=\pi_i(S_i), \dots, X_n=\pi_n(S_n)]}) \\ & \geq \sum_{\mathbf{s} \in D_{\mathbf{S}}} P(\mathbf{s}) \mathcal{R}_i(\mathbf{Y}_{i[X_1=\pi_1(S_1), \dots, X_i=\pi'_i(S_i), \dots, X_n=\pi_n(S_n)]}) \end{aligned} \quad (26)$$

The definition is similar to that of a classical correlated equilibrium, except that the equilibrium is for a CNFG with action spaces that can span across any layer of the PCH.

D.5. CNFG and Bayesian Games

Similarly, in Bayesian Games, agents need not be restricted to layer L_2 , as L_1 and L_3 policies may potentially yield higher rewards, as illustrated next.

Example 10 (Causal Sheriffs Dilemma) *Consider Ex. 8. In reality, the situation may not be so simple or well-defined. Unobserved factors might influence both the officer's assessment of the suspect and the suspect's decision to shoot. For instance, a suspect's background might affect both their likelihood of being a criminal and their behavior. A well-trained officer might intuitively discern such subtle cues and make a quick judgment about whether to shoot. An untrained officer, on the other hand, may lack this ability and be more prone to error. This creates unobserved confounding between the suspect's identity and the officer's tendency to shoot. In other words, the officer may not be able to articulate why they want – or do not want – to shoot, but their instinct carries information about their internal state.*

Consider two scenarios, M_1 and M_2 , that induce the same causal diagram shown in Fig. 20. In M_1 , the officers are well-trained; in M_2 , they are not. In both scenarios, let an adverse background be denoted by the variable $U_T = 1$, with $P(U_T = 1) = 0.1$. Suppose the suspect is a criminal,

Action Space	Payoff depends on agents' actions	Agents act based on a signal	Agents act based on type and has belief about the types
L_2	Normal Form Games (Def. C.2)	Correlated Equilibrium (Def. C.7)	Bayesian Games (Def. C.13)
L_1, L_2, L_3	Causal Normal Form Games (Def. 2.10)	Causal Correlated Equilibrium (Def. C.16)	Causal Bayesian Games (Def. C.18)

Figure 21: Comparison of different representations with information and action spaces

that is, $T = 1$ if and only if they come from an adverse background. This background may influence the suspect's behavior, which in turn can influence the officer's decision to shoot. In the causal diagram, this pathways are represented by the dashed arrows.

Further, in scenario M_1 , the officer is able to pick up on these non-verbal cues, and their probability of shooting is given by $P(P = 1 \mid U_T = 1) = 0.9$ and $P(P = 0 \mid U_T = 0) = 0.9$. In the second scenario, M_2 , the officer almost always makes mistakes, and their probability of shooting is given by $P(P = U_T) = 0.1$. The payoffs $\mathbf{Y} = (Y_1, Y_2)$, as a function of P , T , and S , are shown in Tables 17 and 18.

Now suppose Congress wants to recommend a new policy by passing a law that determines whether officers should shoot or not. They disregard the officers' natural intuitions entirely and compute the Bayesian Nash Equilibrium (BNE) of the game induced by the model, concluding that it is better if the officer does not shoot at all. The expected payoff for the officer under the BNE is therefore given by:

$$\mu_{BE} = -2 \cdot 0.1 = -0.2 \quad (27)$$

However, if the law is not implemented, then in M_1 and M_2 , the expected L_1 -payoff of the policeman are respectively

$$\mu_1 = -0.11, \quad \mu_2 = -0.99 \quad (28)$$

This implies $\mu_2 < \mu_{BE} < \mu_1$, indicating that, even though both SCMs induce the same Bayesian game, implementing the law would be harmful in M_1 , while beneficial in M_2 .

In essence, this is similar to the scenarios illustrated in Ex. 1. Now, we can rewrite the definition of Causal Bayesian Games without restricting ourselves to L_2 layer.

Definition 36 (Causal Bayesian Games) A Bayesian Game is a CNFG with information $\Gamma = \langle M, \mathcal{A}, \mathcal{R}, \mathcal{I} \rangle$, where M is a CMAS with states.

Note that even if two CMASs coincide on their L_2 distributions, they may differ in their L_1 and L_3 distributions – particularly in the corresponding L_1 and L_3 actions. This is formalized below.

Theorem 37 Given a Bayesian Game, there exists two Causal Bayesian Games Γ_1 and Γ_2 with expected L_1 payoffs μ_1 and μ_2 and Bayes-Nash Equilibrium (BE) payoffs μ_{BE} , such that

$$\mu_2 \leq \mu_{BE} \leq \mu_1 \quad (29)$$

Proof The construction is similar to the one in the proof of Thm. 7. ■

The above discussion illustrates the fact that information structure and actions based on layers of PCH are orthogonal to each other. We can have one without the other and even both of them in the same model, without compromising the other. The summary of the axis can be shown in Fig. 21.

Appendix E. Additional Examples and Discussion

E.1. SCM for Table 10

Consider the SCM with $\mathbf{U} = \{U_1, U_2\}$, $\mathbf{X} = \{X_1, X_2\}$ and $\mathbf{Y} = \{Y_1, Y_2\}$. The domains of U_1, U_2, X_1 and X_2 are $\{0, 1\}$. $P(U_1 = 0) = P(U_2 = 0) = 0.5$. $X_1 = U_1$ and $X_2 = U_2$. \mathbf{Y} as a function of U_1, U_2, X_1, X_2 are shown in Table 3.

The action space available to Player 1 and Player 2 are \mathcal{A}^3 and $\mathcal{A}^1 \cup \mathcal{A}^2$ respectively.

E.2. Pseudocode for Ctf-Nash-Learning

Algorithm 3: Ctf-Nash-Learning

Input: Dataset from Ctf-RCT: $(x'_1, x_1, x_2, \mathbf{y})$

Output: Nash Equilibrium strategy f^* over the counterfactual action space

For each (x'_1, x_1, x_2) , estimate the means and weights of the mixture over \mathbf{y} . Let the means be

$R_1(x'_1, x_1, x_2), \dots, R_k(x'_1, x_1, x_2)$ with weights $p_1(x'_1, x_1, x_2), \dots, p_k(x'_1, x_1, x_2)$ in descending order;

If k distributions cannot be identified, treat them as a single distribution: set $R_i(x'_1, x_1, x_2)$ to its mean and $p_i(x'_1, x_1, x_2) = p_i(x'_1, \bar{x}_1, \bar{x}_2)$ where $(x_1, x_2) \neq (\bar{x}_1, \bar{x}_2)$. If this fails, set $p_i = 1/k$;

Define the action spaces: $\mathcal{F}_1 = \{f: X'_1 \rightarrow X_1\}$, $\mathcal{F}_2 = \{g: [k] \rightarrow X_2\}$;

Build the payoff matrix: for each $(f, g) \in \mathcal{F}_1 \times \mathcal{F}_2$, compute

$$\sum_{x'_1, i} P(x'_1) p_i(x'_1, f(x'_1), g(i)) R_i(x'_1, f(x'_1), g(i));$$

Return the Nash Equilibrium strategy f^* of the resulting NFG;

E.3. Assumptions for Alg. 3

For the algorithm to work, we will make the following assumptions. Assume that the learning is from the perspective of Player 1.

Assumption 38 (Identifiability of Mixture) Let $\mathbf{Y}_{x_1, x_2} \mid x'_1, x'_{2i} \sim \phi_i$, for $i \in k$, where ϕ_i is a distribution dependent on x'_1, x_1, x_2 and $k = |D(X)|$. We assume that the distributions are such that their mean and weights are identifiable from their mixture upto a permutation of the i 's:

$$\sum_{i=1}^k p_i \phi_i(x'_1, x_1, x_2) \tag{30}$$

or, the distributions are same for all $i \in [k]$.

		$X_2 = 0$		$X_2 = 1$	
		$U_2 = 0$	$U_2 = 1$	$U_2 = 0$	$U_2 = 1$
$X_1 = 0$	$U_1 = 0$	-2, 2	-2, -6	-2, -6	-2, 2
	$U_1 = 1$	2, -2	-4, 0	-4, 0	2, -2
$X_1 = 1$	$U_1 = 0$	2, -2	-4, 0	-4, 0	2, -2
	$U_1 = 1$	-2, 2	-2, -6	-2, -6	-2, 2

 Table 3: Y_1, Y_2 as a function of U_1, U_2, X_1, X_2 for SCM in Table 10

Next, we show some example distributions and conditions that satisfy the above assumption.

Example 11 (Deterministic Function) Consider the case when $P(\mathbf{Y}_{x_1, x_2} \mid x'_1, x'_2)$ has all its mass on a single point. In addition, assume that

$$E[\mathbf{Y}_{x_1, x_2} \mid x'_1, x'_{2i}] \neq E[\mathbf{Y}_{x_1, x_2} \mid x'_1, x'_{2j}]$$

for $i \neq j$. Then, for each (x'_1, x_1, x_2) we will get k distinct values of \mathbf{Y} , and we can map each $(\mathbf{Y}_i, x'_1, x_1, x_2)$ to a particular i and $p_i = P(\mathbf{Y}_i \mid x'_1, x_1, x_2)$ for $i \in [k]$.

Example 12 (Gaussian Mixtures) Yakowitz and Spragins [43] showed that mixture of multi-variate Gaussians are identifiable. Hence, we can get the mixing proportions and the mean of the Gaussians from the sufficient amount of data.

The next assumption ensures that Player 1 can correctly deduce the intuition of the other player from the observations.

Assumption 39 For all assignments x'_2, x''_2 to the natural intuition of the second player $P(x'_1, x'_2) \neq P(x'_1, x''_2)$.

Note that if $P(x'_1, x'_2)$ are sampled from a continuous distribution then the assumption is true almost surely.

E.4. Ctf-Nash Learning on Causal Prisoner's Dilemma

Table 4: Payoff Matrix learned by Player 1 in Causal Prisoner's Dilemma

	$X_2 = X'_2$	$do(X_2 = 0)$	$do(X_2 = 1)$	$X_2 = 1 - X'_2$
$X_1 = X'_1$	-2.443, -2.450	-1.218, -2.684	-8.892, 0.000	-7.668, -0.239
$do(X_1 = 0)$	-2.683, -1.235	-0.983, -0.983	-6.932, -0.490	-5.232, -0.239
$do(X_1 = 1)$	0.000, -8.848	-0.475, -6.951	-1.960, -1.897	-2.435, 0.000
$X_1 = 1 - X'_1$	-0.240, -7.637	-0.240, -5.250	0.000, -2.387	0.000, 0.000

This section shows the results of applying Ctf-Nash-Learning on Causal Prisoner's Dilemma. The experiment was carried out on 100K samples of $(x'_1, x_1, x_2, \mathbf{y})$ when both agents were playing

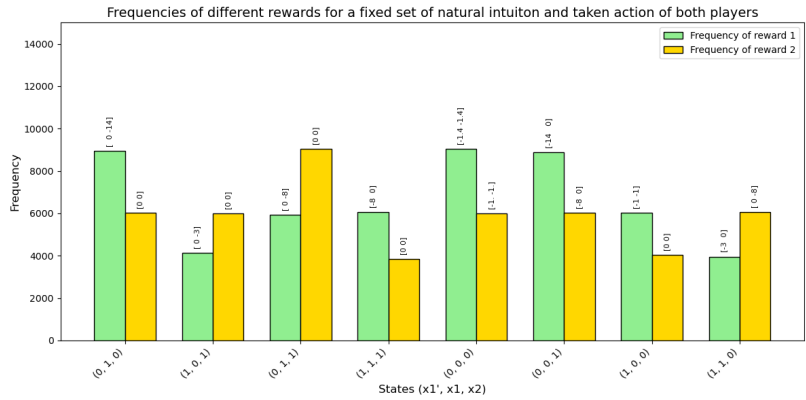


Figure 22: Frequencies of rewards observed for a particular tuple (x'_1, x_1, x'_2)




Figure 23: Probabilities of the rewards observed for a particular tuple (x'_1, x_1, x'_2)

Ctf-RCT. The rewards were assumed to be deterministic, that is, $P(\mathbf{y}_{x_1, x_2} \mid x'_1, x'_2)$ has a point mass. Now, for each tuple (x'_1, x_1, x_2) the frequencies of \mathbf{y} obtained are shown in Fig. 22. For example, when (x'_1, x_1, x'_2) is $(0, 1, 0)$, then the reward $(0, -14)$ was observed nearly 9000 times while $(0, 0)$ was observed 6000 times, and when it is $(0, 1, 1)$, then the reward $(0, -8)$ occurs nearly 6000 times and $(0, 0)$ occurs nearly 9000 times, and so on.

From this frequency table, we can compute the probabilities as shown in Fig. 23. For instance, when (x'_1, x_1, x'_2) is $(0, 1, 0)$, the two values of x'_2 are taken with a probability of roughly 0.6 and 0.4. The same is observed for the tuple $(0, 1, 1)$. Hence, one value of x'_2 occurs with a probability of 0.6 and the other with 0.4. However, there is no way to know whether it is 0 or 1 that occurs with a probability 0.6. This results in a permutation of the actions. For example, if x'_2 is correctly identified, that is $\hat{x}'_2 = x'_2$ then $(\hat{x}'_2 = 0, x_2 = 0)$, $(\hat{x}'_2 = 1, x_2 = 1)$ would correspond to the natural instinct or action ($X_2 = X'_2$). However, if they are not correctly identified, that is $\hat{x}'_2 = 1 - x'_2$, then the same would correspond to acting against intuition ($X_2 = 1 - X'_2$).

The learned payoff matrix is shown in Table. 4. Even if the actions are permutations of the original payoff matrix, the equilibrium remains same. Hence, the algorithm will be able to find the equilibrium correctly.

The code is available at <https://anonymous.4open.science/r/CGT-2025/>.

E.5. Forgetting in Prisoner’s Dilemma

As noted in Appendix B, it is not always in the best interest of the agents to forget. Consider the classical prisoners dilemma: the choices of the action spaces are $\{C\}$, $\{D\}$ and $\{C, D\}$ and forget about whatever is not included in the sets. The metagame over the choice of the action spaces is shown below.

P1 \ P2	$\{C\}$	$\{D\}$	$\{C, D\}$
$\{C\}$	-1, -1	-7, -0.5	-7, -0.5
$\{D\}$	-0.5, -7	-1.9, -1.9	-1.9, -1.9
$\{C, D\}$	-0.5, -7	-1.9, -1.9	-1.9, -1.9

Table 5: Extended Prisoner’s Dilemma with a third action CD

Note, if both the agents forget about defecting D , and plays only C , then one of the agents can move to the action space $\{C, D\}$ and get a better payoff while the other is worse of. The resultant NE of this metagame thus also turns out to be $(-1.9, -1.9)$ with action spaces $\{D\}$ or $\{C, D\}$. Thus in classical prisoners dilemma, agents do not have advantage with forgetting.

Appendix F. PCH Layers and Level- k Rationality: Orthogonal Axes

PCH and Level- k rationality are orthogonal. Layers of PCH provides different ways of interacting with the environment, while Level- k rationality is a form of rationality assumption about the other agents. The former can be present even in single agent settings, while the later is only possible in multi-agent systems.

F.1. An Example:

Let's consider the following CMAS:

$$Y_1 = \begin{cases} 0 & X_1 = U_1, X_2 = U_2 \\ 2 & X_1 \neq U_1, X_2 = U_2 \\ -4 & X_1 \neq U_1, X_2 \neq U_2 \\ 2 & X_1 = U_1, X_2 \neq U_2 \end{cases} \quad (31)$$

and $Y_2 = Y_1$.

The result is a co-operative game. The payoff table induced by the L_1 and L_2 actions is shown below:

	$X_2 = X'_2$	$X_2 = 0$	$X_2 = 1$
$X_1 = X'_1$	0, 0	1, 1	1, 1
$X_1 = 0$	1, 1	0, 0	0, 0
$X_1 = 1$	1, 1	0, 0	0, 0

Now, say that the L_1 actions are the default actions in Level-0 rationality. Then, a Level-1 agent assumes that the other agent is performing Level-0 or L_1 action and finds the best response to that, which in this case is any of the L_2 actions $do(X = 0)$ or $do(X = 1)$. Now, given that the other agent is performing an L_2 action, the best response is just following the L_1 action again. And this cycle goes on and on in an alternate fashion, until we realize that the Nash Equilibrium can be best realized by an L_3 action space, where the payoff is (2, 2). The payoff table for the whole of L_1, L_2 and L_3 action space are given below.

	$X_2 = X'_2$	$X_2 = 0$	$X_2 = 1$	$X_2 = 1 - X'_2$
$X_1 = X'_1$	0, 0	1, 1	1, 1	2, 2
$X_1 = 0$	1, 1	0, 0	0, 0	-1, -1
$X_1 = 1$	1, 1	0, 0	0, 0	-1, -1
$X_1 = 1 - X'_1$	2, 2	-1, -1	-1, -1	-4, -4

Now, we present an example, where the Level-k rationality oscillates between L_1 and L_2 with default action being L_1 . However, a better equilibrium for both agents lie in L_3 , that is (3, 3).

	$X_2 = X'_2$	$X_2 = 0$	$X_2 = 1$	$X_2 = 1 - X'_2$
$X_1 = X'_1$	0, 0	1, 1	-3, -3	-2, -2
$X_1 = 0$	1, 1	0, 0	0, 0	-1, -1
$X_1 = 1$	-3, -3	0, 0	-1, -1	2, 2
$X_1 = 1 - X'_1$	-2, -2	-1, -1	2, 2	3, 3

F.2. Level-0 Default action vs L_1 action

We provide the following two examples as an illustration:

Consider the classical Prisoner’s Dilemma where the default action for both agents is: play $X = 1$ with probability 0.6 and $X = 0$ otherwise. The payoff is given by $(-2.46, -2.46)$. This is can be represented as the Markovian PD from 6, where the L_1 action is defined as $X = U_X$ and $P(U_X = 1) = 0.6$. We will get the L_1 action payoff as $(-2.46, -2.46)$.

Now, consider the classical prisoner’s dilemma, and we collect the data $P(X_1, X_2, Y_1, Y_2)$. We will see that $P(X_1 = 1) = P(X_2 = 1) = 0.6$. The payoff is $(0, 0)$. However, if we model it as a default action $X = 1$ with probability 0.6 and $X = 0$ otherwise, then we will get the payoff of $(-2.46, -2.46)$ as shown earlier.

Hence, L_1 can represent a larger class of irrationalities than default actions. In particular, all default actions can be represented by some L_1 action; however not all L_1 actions can be modelled as a default action.

Appendix G. FAQ

1. What are L_1 , L_2 , and L_3 actions, and why are they essential in modeling decision-making?

A: Pearl [31] and Bareinboim et al. [4] introduced a framework for studying real-world systems – ranging from experiments in medicine to analyses of climate models – using causal models such as Structural Causal Models (SCMs). An SCM induces three levels of distributions: observational (L_1), interventional (L_2), and counterfactual (L_3). One of the key results in this literature is the *Causal Hierarchy Theorem* (CHT), which states that these three levels of distributions form a strict hierarchy and do not collapse. That is, given an L_1 distribution, there may be multiple SCMs that induce the same L_1 distribution but yield different L_2 distributions. Likewise, given both L_1 and L_2 distributions, multiple models may still differ in their induced L_3 distributions [4, Thm. 1].

In decision-making systems, an agent interacts with the environment (and possibly with other agents) through its actions. If the agent does nothing and simply observes, this behavior corresponds to an L_1 action. If it disregards its intuition and performs an intervention (either hard or soft) it engages in an L_2 action. If the agent’s realized action depends on what it would have done under natural circumstances (i.e., its L_1 action), then the behavior corresponds to a counterfactual, or L_3 , action. These distinctions in decision-making have been studied extensively over the past decade, including in Bareinboim et al. [2, 3].

A particularly relevant work that forms the basis for our discussion in the single-agent setting is Bareinboim et al. [2], which introduces a scenario known as the Greedy Casino (reviewed in Appendix A). In this setting, there are machines that may blink and patrons who may be drunk. If a machine is blinking and a patron is intoxicated, the patron is more likely to play that machine instinctively (i.e., subconsciously). Conversely, if the patron is sober and the machine is not blinking, they tend to prefer a different option. This behavior reflects natural predispositions, biases, and inclinations, and is modeled as an L_1 action.

Of course, this is a stylized example, but human agents often behave in similar ways – acting without full awareness of the underlying causes of their decisions. This phenomenon has been extensively studied in behavioral decision-making and cognitive psychology, most notably in the work of Daniel Kahneman and collaborators [19]. In contrast, an L_2 action in a single-agent setting may involve flipping a coin and allowing the coin toss to determine which machine to play. This process effectively averages over the agent’s internal biases and corresponds to what is formally called the causal effect.

These L_1 and L_2 behaviors are fundamentally distinct from a counterfactual scenario, in which an agent intends to play machine $X = x$ but ends up playing $X = x'$. Formally, this is represented by the counterfactual quantity $P(Y_{X=x} \mid X = x')$. Note that if the agent naturally plays $X = x'$, this implies they were not initially inclined to play $X = x$ (see Fig. 3 for an illustration). This type of counterfactual evaluation was made possible by the introduction of counterfactual randomization [3], precisely to decouple the agent from its natural flow and enable the estimation of Q .¹ The theoretical limits of which counterfactuals can be physically inferred from the world have been recently characterized [34]. In positive cases, once the counterfactual randomization step is performed, the agent’s intuition becomes a new type of information, which can then be conditioned upon.

2. Why are Causal Models essential if they can be converted into a matrix game with counterfactual actions?

A: Structural causal models (SCMs) are essential for constructing the payoff matrix (e.g., Fig. 9). In particular, determining the payoffs corresponding to L_1 and L_3 actions inherently relies on the underlying causal model. Once this matrix is computed, it may be viewed as a normal-form game. However, there are four critical phases where causal modeling plays a vital role:

- **Representation:** Traditional game theory lacks the concept of natural actions; it primarily deals with interventions, corresponding to L_2 actions in our causal framework. Once the existence of natural actions is acknowledged, the action space expands to include L_1 , L_2 , and L_3 , enabling the construction of a richer payoff matrix.

Example 1 illustrates this distinction more explicitly: two SCMs may yield identical payoffs and equilibria in the L_2 action space, yet diverge significantly when L_1 and L_3 actions are considered. This expanded matrix – and its associated equilibria – is difficult to recover without causal assumptions or an underlying structural model. For instance, suppose we observe repeated instances of the scenario in Ex. 5 and attempt to infer the payoffs for the actions (C, C) , (C, D) , (D, C) , and (D, D) from observational data. We might obtain $(-1.4, -1.4)$, $(-8, 0)$, $(0, -8)$, and $(0, 0)$, respectively. However, if agents follow a randomized controlled trial (RCT) protocol, the corresponding payoffs could be $(-1, -1)$, $(-7, -0.5)$, $(-0.5, -7)$, and $(-1.9, -1.9)$. An underlying causal model with unobserved confounders can explain this discrepancy. The three layers of the causal hierarchy – observational, interventional, and counterfactual – were formally introduced in Bareinboim et al. [4].

- **Agency and Execution:** Causal modeling also addresses the practical question of how agents can implement these actions. While the payoff matrix in Fig. 5 encodes the outcomes, it does not specify the mechanisms by which those actions are executed. From prior results in causal decision theory, we know that L_1 actions correspond to passively observing the system, L_2 actions to standard interventions (e.g., RCTs), and L_3 actions require more advanced techniques, such as counterfactual randomization (e.g., ctf-RCT). Thus, causal modeling provides a bridge between abstract game-theoretic strategies and their realizability in practice. (In fact, the notion of counterfactual realizability – achieved through a specific type of randomization – appears to exhaust what is physically implementable in the real world; for a more refined discussion, see Raghavan and Bareinboim [33].)

1. There are intriguing connections here to the notion of free will (see Pearl’s 2013 discussion, “The Curse of Free Will and the Paradox of Inevitable Regret”).

- **Solution Concept:** Our solution concept goes beyond merely computing a Nash equilibrium over an expanded action space. Causal modeling introduces a hierarchy of action spaces, allowing agents to commit to specific layers (e.g., L_1 , L_2 , $L_1 + L_2$, L_3) and ignore the others. This gives rise to a new “metagame,” where the strategic choice involves selecting a layer of the PCH, and the equilibrium is computed within the corresponding subspace. This layered structure adds a new dimension to strategic reasoning and highlights the importance of causal structure in shaping the space of strategic possibilities.
- **Learning:** Normal-form representations overlook the structure linking agents’ intuitions to their executed actions, whereas CNFGs can capture this relationship directly. In practice, agents may not observe the other agent’s intuition when learning the payoff matrix. In such cases, this structural distinction becomes essential, and is explicitly exploited in Alg. 3.

3. Is the causal graph necessary? How does it help?

A: Causal graphs contain information about the structural dependencies between variables. For example, when the SCM is Markovian in a CNFG, the optimal actions or equilibrium strategies will always lie in L_2 . However, when Markovianity does not hold, the optimal strategies may fall anywhere between L_1 and L_3 . Without further knowledge of the parameters – either through prior knowledge or interaction with the system—it is impossible to determine which strategy is better.

4. How does this work relate to prior works?

A: Our goal in this paper is to develop a model that captures both instinctive and deliberate decision-making processes of the human mind. Structural Causal Models provide a principled framework for representing reality as closely as possible. Some prior works have attempted to reconcile causal models and game theory. However, due to differing goals, their models can be extremely restrictive, as discussed in detail in Appendix A.7.