

Evaluating Counterfactual Strategic Reasoning in Large Language Models

Dimitrios Georgousis, Maria Lymperaïou, Angeliki Dimitriou,
Giorgos Filandrianos, Giorgos Stamou

National Technical University of Athens

dimjimitris@gmail.com,

{marialymp, angelikidim, geofila}@ails.ece.ntua.gr,

gstam@cs.ntua.gr

Abstract

We evaluate whether LLMs adapt their strategic behavior when familiar games are counterfactually modified. We introduce a repeated-game evaluation framework covering Prisoner’s Dilemma and Rock–Paper–Scissors under default, label-perturbed, payoff-perturbed, and joint counterfactual variants. This design separates surface robustness to renamed actions from deeper sensitivity to changed incentives. Across multiple frontier LLMs, we find that label perturbations usually cause moderate degradation, whereas payoff perturbations expose stronger failures: LLMs often preserve canonical strategies even when the equilibrium structure changes. In RPS, several LLMs remain close to uniform play despite a payoff-counterfactual equilibrium requiring a biased mixed strategy. Behavioral and efficiency metrics further show that stronger or reasoning-enabled LLMs are not uniformly more strategic: some deliberate more without adapting faster. Overall, counterfactual repeated games provide a compact diagnostic for distinguishing robust incentive-sensitive behavior from brittle template-based strategic execution.

1 Introduction

Strategic reasoning using Large Language Models (LLMs) forms an upcoming field on the intersection of reasoning and agentic synergy, driven by the rapidly advancing capabilities of state-of-the-art (SoTA) LLMs. Communication between LLMs allows cooperation and competition, i.e. the basic ingredients for allowing *game-playing* (Gandhi et al., 2023; Zhang et al., 2024). Notably, the search for human-level strategic interactions spreads before the LLM-era, demonstrating a long-standing need for autonomous, rational agents (Silver et al., 2016; Berner et al., 2019; Bakhtin et al., 2022).

Interestingly, prompting LLMs as strategic players reveals several behavior traits for the models themselves, enabling the usage of strategic and

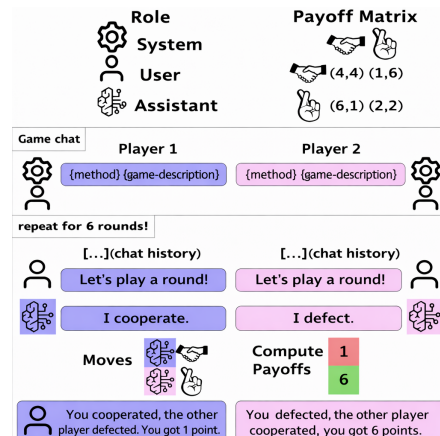


Figure 1: LLM interaction for Prisoner’s Dilemma

competitive environments as suitable testbeds for LLM evaluation (Zhang et al., 2024; Costarelli et al., 2024; Jia et al., 2025; Wang et al., 2025; Duan et al., 2024). LLMs are found to often demonstrate a gap between verbalizing versus executing optimal or equilibrium strategies, demonstrating brittle behavior under competition. Inductive biases, over-reliance to short term outcomes and sensitivity to prompts arise as major issues that influence LLM capability in strategic reasoning, paving the way for more evaluation environments to come.

We opt for evaluating LLMs in structured games rather than open-ended strategic environments, enabling controlled and well-defined assessment of their behaviors and strategic capabilities. This setup allows us to introduce counterfactual scenarios that modify standard game parameters to test cognitive flexibility. Performance deviations under such altered conditions reveal potential reasoning rigidity: LLMs may rely on canonical game-template behavior from pre-training yet struggle when familiar structures are redefined. We thus propose a simple and adaptable evaluation framework based on counterfactual game pairs to diagnose behavioral evidence of incentive-sensitive adaptation rather than canonical game-template behavior.

To capture distinct failure modes, we focus on

two complementary game settings: Rock–Paper–Scissors (RPS), which probes stochasticity and exploitable patterns, and the Prisoner’s Dilemma (PD), which probes cooperation under conflicting incentives. In our setup, RPS tests whether LLMs avoid predictable biases in adversarial play, while PD tests whether they adapt between defection and cooperation under repeated interaction.

In RPS, players choose among three cyclically dominant actions (*Rock beats Scissors, Scissors beats Paper, and Paper beats Rock*), so no move is inherently superior as it depends on other player’s actions, but stable patterns become exploitable (e.g. if we always play *Rock* in even rounds and *Paper* in odd rounds). In PD, players simultaneously choose between *Cooperate* and *Defect* without knowledge of the opponent’s decision. Mutual cooperation yields shared benefits, while unilateral defection provides a higher immediate payoff. Thus, no action is universally optimal: cooperation risks exploitation, whereas defection sacrifices collective gain. Although defection dominates in the single-round setting, serving as the safest option, repeated interaction allows for adaptation, with persistent defection inviting retaliation, while consistent cooperation sustaining mutually beneficial outcomes.

Our results reveal three main findings. First, payoff-based counterfactuals are more diagnostic than label counterfactuals: many LLMs tolerate renamed actions but fail when the incentive structure changes. Second, strategic competence is highly opponent-dependent: models often exploit simple algorithmic opponents but become unstable in LLM–LLM interactions, especially in RPS. Third, reasoning-enabled variants do not always improve strategic execution; in several cases, additional deliberation increases token usage without improving adaptation. Overall, our contributions are¹:

- We introduce a counterfactual repeated-game framework that separates label robustness from payoff-sensitive strategic adaptation.
- We propose behavioral metrics for strategic execution, including opponent comprehension, cooperation/action-distribution behavior, and token-normalized efficiency.
- We provide an empirical analysis showing that LLMs often behave strategically in canoni-

cal games but remain brittle under payoff and joint counterfactuals.

2 Related work

Games for LLM evaluation constitute a new research frontier, examining whether probed behavior aligns with game-theoretic predictions. So far, LLMs have demonstrated their ability to approach optimal strategies and equilibrium concepts when employed as rational players in games such as Matching Pennies, RPS, the Ring-Network Game, and the Dictator Game (Silva, 2024; Fan et al., 2024; Kader and Lee, 2025), although they pose systematic deviations in the course of strategic execution, particularly in randomized setups. Cooperative abilities and coordination of LLMs have been probed in social dilemma games, including PD; LLMs display collaborative and retaliation behaviors, even though they remain subject to prompt sensitivity and opponent specification (Akata et al., 2025; Fontana et al., 2024; Mozikov et al., 2024). Moreover, weaker models often fail to abstract away from narrative context and instead rely on superficial cues rather than the actual incentive game structure (Lorè and Heydari, 2024). These works underscore a persistent gap between explicit strategic knowledge and realized decision-making, supporting the use of interactive game environments as diagnostic LLM tools (Aher et al., 2023; Filippas et al., 2024). Finally, using LLMs as meta-evaluators of strategic LLM interactions confirms previously observed pitfalls in capturing dominant strategies and handling payoffs, especially in non-canonical game formulations (Collins et al., 2025).

Evaluation of LLMs in Counterfactual Setups as opposed to pattern memorization poses a seminal challenge, often operationalized through intervened inputs that explicitly contradict real-world knowledge. Prior work shows that although models can follow hypothetical premises, their performance deteriorates when these assumptions conflict with well-established facts, suggesting reliance on memorized associations rather than robust integration of altered conditions (Li et al., 2023; Yamin et al., 2025; Wu et al., 2024). Subsequent benchmarks formalize this challenge by systematically perturbing task inputs and measuring counterfactual sensitivity, showing that LLMs frequently fail to update their reasoning in response to minimal but semantically meaningful changes (Chen et al., 2025; Frohberg and Binder, 2022; Akter et al., 2026).

¹Code is available at https://github.com/dimjimitris/llm_gm_thesis.

Moreover, apparent improvements from chain-of-thought prompting or instruction tuning may mask underlying brittleness, as LLMs produce fluent yet inconsistent outputs across counterfactual variants (Hüyük et al., 2025), while scaling up LLM size is not a viable cure, since larger models may collapse even more dramatically under reformulated realities (Stringli et al., 2025). Despite these advances, existing evaluations largely focus on static setups, leaving interactive environments comparatively underexplored. Particularly, the evaluation of LLMs in *counterfactual game-theoretic settings* remains limited, motivating the use of repeated games as diagnostic environments for strategic reasoning.

3 Background in Game Theory

Strategic Games. A (finite) *strategic game* in normal form is defined by a tuple:

$$G = (\{1, 2\}, A_1, A_2, u_1, u_2) \quad (1)$$

where A_i is the finite action set of player i , and $u_i : A_1 \times A_2 \rightarrow \mathbb{R}$, $i \in \{1, 2\}$ is player i 's payoff function that indicates the quality of an action. We explicitly focus on *two-player, simultaneous games*, i.e. the two players take actions at the same time, without any information regarding the other player's planned action. A game is *symmetric* if:

$$A_1 = A_2 = A \text{ and } u_1(a, b) = u_2(b, a) \quad (2) \\ \forall a, b \in A.$$

Strategies. A *pure strategy* for a player is a single action $a \in A$. A *mixed strategy* is a probability distribution over actions, denoted by $\sigma \in \Delta(A)$, where $\Delta(A)$ denotes the set of all valid probability distributions over A . For a mixed-strategy profile $\sigma = (\sigma_1, \sigma_2)$, the expected payoff to player i is:

$$\mathbb{E}_{a \sim \sigma} [u_i(a_1, a_2)] \quad (3)$$

Here, σ_i denotes the mixed strategy of player i , $a \sim \sigma$ indicates that actions are sampled according to the players' mixed strategies, and u_i is the payoff function of player i .

Nash equilibrium. The central solution concept we employ is the *Nash equilibrium*, which formalizes a notion of strategic stability. Intuitively, a Nash equilibrium corresponds to a situation in which each player selects a strategy that is optimal given the strategy of the other player, so that no player has an incentive to unilaterally deviate.

A mixed-strategy profile $\sigma^* = (\sigma_1^*, \sigma_2^*)$ denotes a *candidate equilibrium*, where σ_i^* is the strategy assigned to player i at equilibrium. This profile is a Nash equilibrium if, for each player i :

$$\mathbb{E}[u_i(\sigma^*)] \geq \mathbb{E}[u_i(\sigma_i, \sigma_{-i}^*)] \quad \forall \sigma_i \in \Delta(A). \quad (4)$$

Here, σ_{-i}^* denotes the *equilibrium* strategy of the opposing player, while σ_i represents any alternative (possibly randomized) strategy available to player i . The inequality states that, assuming the other player follows their equilibrium strategy, player i cannot increase their expected payoff by switching unilaterally from σ_i^* to any other strategy.

Prisoner's Dilemma (PD). Let $A = \{C, D\}$ denote two distinct actions namely *Cooperate* and *Defect*. The PD payoff structure satisfies:

$$T > R > P > S \quad (5)$$

where R : *reward* for mutual cooperation (C, C), T : *temptation* payoff for unilateral cooperation (D, C), P : *punishment* for mutual defection (D, D), S : *sucker* payoff for unilateral cooperation (C, D). Based on the above, cooperation is socially beneficial ($R > P$), defection is individually rational ($T > R$), and unilateral cooperation is strictly worse than mutual defection ($S < P$).

A canonical payoff matrix is:

	C	D	
C	(R, R)	(S, T)	
D	(T, S)	(P, P)	

(6)

In the one-shot PD, D strictly dominates C and the unique Nash equilibrium is (D, D) . PD is a general-sum game, in which players' payoffs need not sum to zero and mutually beneficial outcomes may exist, such as the joint action (C, C) .

Repeated games are formed by playing a fixed *stage game* multiple times in sequence. At each round t , players simultaneously choose actions and observe the outcome before proceeding to the next round. Let a^t denote the joint action (combination of actions chosen by all players $a^t = (a_1^t, a_2^t)$) at round t . The history h_t available at this round is:

$$h_t = (a^1, \dots, a^{t-1}), \quad t \in \{1, \dots, T\} \quad (7)$$

A (behavioral) strategy specifies a distribution over actions conditioned on history:

$$\pi(\cdot | h_t) \in \Delta(A) \quad (8)$$

The payoff to a player in a repeated game is typically defined as a cumulative sum over rounds:

$$U_i = \sum_{t=1}^T u_i(a^t) \quad (9)$$

or, in infinite-horizon settings, as a discounted sum with discount factor $\gamma \in (0, 1)$.

Rock–Paper–Scissors (RPS). Let an action set $A = \{R, P, S\}$ denote *Rock, Paper, Scissors*. RPS is a two-player *zero-sum* game, i.e. the players’ payoffs always sum to zero. In zero-sum games, one player’s gain corresponds exactly to the other’s loss, excluding mutual benefits. Thus, optimal behavior is purely *adversarial*, with a payoff function:

$$u(a, b) = \begin{cases} 1 & \text{if } a \text{ beats } b, \\ 0 & \text{if } a = b, \\ -1 & \text{if } b \text{ beats } a. \end{cases} \quad (10)$$

RPS has no pure-strategy Nash equilibrium; its unique Nash equilibrium is the mixed strategy:

$$\sigma^*(R) = \sigma^*(P) = \sigma^*(S) = \frac{1}{3} \quad (11)$$

4 Method

4.1 Counterfactual Games

One of our main points is evaluating LLMs on counterfactual games, i.e. alternative formulations of action labels or payoff values. An almost invariant counterfactual outcome indicates LLM generalizability and adaptability in novel compositions, or else an overly strict adherence on specific instances rather than reasoning over game dynamics.

Prisoner’s Dilemma Counterfactuals

Stag Hunt (SH) stands as a counterfactual equivalent of PD, challenging the balance between social cooperation and individual safety: two hunters must decide separately whether to hunt a stag (S) or a hare (H). One hunter on their own can catch a hare, but the payoff is less than catching a stag. However, this requires mutual cooperation, as no hunter on their own is able to capture a stag. The SH payoff satisfies the inequality:

$$R > T \geq P > S \quad (12)$$

where R : reward for mutual S , T : hunting H while the other hunts S , P : payoff for mutual H , S : payoff for hunting S alone. The game displays two

Nash equilibria: the (S, S) is payoff dominant but risky, while the (H, H) is risk-dominant but safe.

The altered SH payoff suggests a change of strategic incentives, shifting the emphasis from defection to coordination; thus, SH is primarily a **payoff-driven** PD counterfactual. To test LLM adherence to action labels, we substitute PD labels with SH ones ($C \rightarrow S$ and $D \rightarrow H$), without altering the payoff values. An instance of the employed counterfactual formulation is detailed in App. A.1.

Rock–Paper–Scissors Counterfactuals

In RPS, we implement both a **payoff-based** counterfactual as well as a **label-based** one. While in the default game all wins and losses are of equal magnitude, leading to a unique mixed-strategy equilibrium characterized by uniform randomization, in the payoff-based counterfactual we alter the incentive structure by assigning higher payoff to specific outcomes: the *Rock–Paper* combination receives $\times 3$ win-loss value, while the rest of the combinations remain invariant (App. A.2). This modification breaks the game symmetry and forces rational players to adopt a biased mixed strategy rather than uniform randomization. Moreover, in the label-based counterfactual game we permute labels to inverse their dominance, so that *Scissors* beats *Rock*, *Rock* beats *Paper* and *Paper* beats *Scissors*. This intervention decouples label memorization from reasoning over the game structure. Finally, we jointly interfere on payoffs and labels to further stress test LLM players.

For the payoff-based RPS counterfactual, the canonical uniform equilibrium no longer applies. Under the modified payoff matrix, where the Rock–Paper outcome has magnitude 3 while all other win/loss outcomes have magnitude 1, the unique mixed Nash equilibrium is:

$$\pi^*(R) = 0.2, \pi^*(P) = 0.2, \pi^*(S) = 0.6 \quad (13)$$

Therefore, models that remain close to uniform play $(1/3, 1/3, 1/3)$ are not merely randomizing; they are failing to adapt to the altered payoff structure.

5 Experimental Setup

5.1 Player Types

Evaluating an LLM’s strategic thinking is only meaningful in the context of other players’ game-playing dynamics. For this reason, we consider the following LLM players, differentiated by their system prompt: a) **ZS** (default): initialized with **zero-shot** (ZS) system prompt; b) **CoT**: initialized with

Chain of Thought (CoT) system prompt; c) **SPP**: initialized with **Solo Performance** prompt (SPP) (Wang et al., 2024), a collaboration-enhancing approach that combines multiple views from LLM personas for advanced problem-solving; d) **SC**: employing **self-consistency** (SC) (Wang et al., 2023) to augment any of the aforementioned players.

We also contrast LLM players with algorithmic ones, which are based on deterministic strategies, constituting a predictable and well-defined evaluation framework for opposition or collaboration. A potent LLM player is able to recognize the patterns of an algorithmic one, and upon this, decide their optimal strategy or present exploitation tendencies. The following algorithmic players are employed: a) **SREP**: selects the single-round equilibrium action (pure or mixed depending on the game); in PD this corresponds to always defecting; b) **PP**: a **pattern player** follows a cyclic pattern of actions. The next sets of players are instantiated differently for PD and RPS: c_1) **MF**: exploiting the **most frequent** move of the opponent to maximize their reward for PD; d_1) **TFT**: a **tit-for-tat**² variant in which the player chooses the move that when paired with the most recent opponent’s move gives the best reward for PD. Finally, for RPS: c_2) **AP**: an **adaptive player** that acts so as to counter the opponent’s most frequent move; d_2) **TFT**: based on tit-for-tat, this player chooses the move that counters the opponent’s most recent move.

6 Experiments

6.1 Experimental design

Number of rounds To probe LLMs’ comprehension of game dynamics over time, we allow consequent rounds to play per game; specifically, PD is played for 16 rounds, while RPS is repeated 24 times. After preliminary experiments, we conclude that these are adequate repetitions for an LLM to understand and adapt to the opponent considering the number of available variables to select in each round (C or D for PD, and R , P , or S for RPS.)

Game repetition Each game is played by the **non-SC** players 5 different times, and 2 times by **SC** ones (it is expected that their outcomes remain similar due to the self-consistency mechanism).

Self-Consistency parameters We use 3 samples for PD, where two actions are available, and 5 sam-

²In tit-for-tat, a player chooses C in the 1st round and in following rounds it copies the opponent’s previous action.

ples for RPS, where three actions require a larger pool for stable marginalization. Ties at the highest frequency are broken randomly. These settings were chosen as a practical compromise between consistency and computational cost.

6.2 Models

LLMs instantiated as players include Claude Sonnet 3.5 v2, Claude Sonnet 3.7 (with vs without thinking), Claude Sonnet 4 (with vs without thinking), Deepseek R1, Llama 3.3 70B and Mistral Large (24.07). Each LLM plays against a different instance of the same model, or algorithmic agents.

6.3 Evaluation Metrics

Total points denote per-player cumulative reward over all rounds, averaged across repetitions. Higher values indicate stronger adaptation to the payoff structure and opponent behavior.

Opponent Comprehension is a more dynamic metric that evaluates how late an LLM unlocks its opponent’s behavior and takes advantage of it. Formally, given a repeated game of N rounds between an LLM agent A and an opponent B , we define the *round of opponent comprehension* m as the earliest round after which A ’s actions yield payoffs that are at least as high as B ’s in a large majority of subsequent rounds. To allow for occasional deviations or noise, we introduce a target percentage tp , requiring that from round m to round N , A achieves payoffs no worse than B in at least tp percent of the rounds. We default $tp = 90\%$. Lower values of m indicate earlier comprehension and faster adaptation, while a value of m exceeding the game horizon indicates that the LLM failed to consistently exploit the opponent’s behavior.

Cooperation Rate Cooperation rate is the proportion of rounds in which the LLM selects the cooperative action (e.g., C in PD or S in SH). It captures social tendencies, though high cooperation does not necessarily imply optimal play.

Efficiency is proposed to account for the cost-performance trade-off:

$$efficiency = \frac{points}{tokens} \times c \quad (14)$$

The scaling factor c has a default value of 1000 for numerical readability.

Failure Rate measures the proportion of invalid actions, including formatting violations and outputs outside the allowed action space.

6.4 Payoff instantiations

We instantiate PD with payoffs $(C, C) = (4, 4)$, $(C, D) = (1, 6)$, $(D, C) = (6, 1)$, and $(D, D) = (2, 2)$. Each PD game is repeated for $N = 16$ rounds, so per-player cumulative scores satisfy:

$$16 \leq \text{total points} \leq 96. \quad (15)$$

These bounds correspond to always cooperating against a defecting opponent and always defecting against a cooperating opponent, respectively.

For RPS, we use the standard zero-sum payoff instantiation, where a win, loss, and tie yield $w = 1$, $\ell = -1$, $t = 0$, respectively. Each game is repeated for $N = 24$ rounds, with cumulative scores within:

$$-24 \leq \text{total points} \leq 24. \quad (16)$$

Under the single-round mixed Nash equilibrium, the expected payoff per round is 0, yielding an expected cumulative score of 0 over 24 rounds. We therefore interpret RPS results primarily through deviations from this zero baseline.

Since these deterministic bounds are not always informative in uncertain settings (e.g. LLM–LLM interactions), we also use expected-payoff baselines under stationary mixed strategies; the corresponding derivations are provided in App. B.

Result aggregation. To support claims beyond isolated examples, we compare default and counterfactual variants over the same repeated runs and prompting variants. We report full per-model means and variability in App.D and use these values to examine whether degradation is consistent across games, opponents, and metrics. Our analysis is primarily diagnostic: we focus on the direction and recurrence of behavioral changes under label, payoff, and joint counterfactuals.

7 Results and Analysis

Full numerical results are shown in App. D.

7.1 Prisoner’s Dilemma

Total points widely vary based on the model choice, the opponent category (algorithmic vs LLM) and the prompting technique used. Commencing from algorithmic opponents, the LLM manages to accumulate higher total points, exploiting their predictable algorithmic strategies, although performance fluctuates per opponent. In particular, **SREP** always defects, making defection the optimal response and yielding 32 points over 16

rounds through sustained (D, D) play. Most LLMs cluster around this baseline (~ 30 points), with modest prompt-induced variation. A notable exception is Mistral Large, which underperforms substantially and remains unstable across prompting types (from 18.6 ± 10.6 under SPP to 29.8 ± 2.2 under CoT), suggesting delayed or inconsistent recognition of the dominant response. Against the cyclical **PP**, LLMs obtain substantially higher scores, typically in the mid-50s to low-60s, by exploiting the repeated D, C pattern once recognized.

A different regime emerges in LLM–LLM interactions, where most models converge to mutual cooperation and receive 4 points per round. Claude 3.5/3.7 Sonnet (with and without thinking) and Llama 3.3-70B consistently achieve the maximum 64 points, corresponding to perfect cooperation across all 16 rounds, with virtually no variability across prompting types. In contrast, Claude 4 Sonnet variants and DeepSeek R1 obtain substantially lower totals (31.4 ± 0.0 to 49.4 ± 15.5), reflecting a stronger tendency toward defection or unstable strategic exploration. These results suggest that increased reasoning depth does not necessarily improve strategic performance, and may in fact hinder convergence to the mutually beneficial outcome.

Opponent Comprehension is generally fast against algorithmic opponents, yielding low m . This is clearest for **SREP**, where most models adapt within the first few rounds ($m \leq 4$), and Claude 4 as well as DeepSeek R1 often respond from round 1. These early values align with total points close to the optimal defection baseline. The main exception is Mistral Large, whose comprehension is highly unstable and in some cases never stabilizes within the game horizon ($m > 16$), matching its lower cumulative returns. Other algorithmic opponents, especially **PP**, are harder to infer and therefore produce later comprehension.

In LLM–LLM interactions, comprehension reflects inference of cooperative intent rather than recognition of a fixed policy. For most models and prompting types, this happens very early, which explains the stable mutual cooperation reported in total points. Claude 4 (especially w/o thinking) and DeepSeek R1 show more variability, indicating delayed or unstable inference of the opponent’s behavior; this is consistent with their lower scores.

Cooperation Rate against algorithmic opponents largely reflects the optimal response to each strategy. When facing the always-defecting **SREP**,

cooperation rates cluster around ~ 0.1 across LLMs, indicating almost no cooperation beyond a few exploratory rounds. This aligns with the opponent comprehension results, where most models recognize persistent defection very early. Against **PP**, which alternates between defection and cooperation, cooperation rises but remains well below 0.5, typically ranging between 0.3 and 0.5 depending on the model and prompt. These values suggest a short exploratory phase followed by sustained defection once the cyclic pattern is understood. For **MF** and **TFT**, cooperation is again low (mostly 0.1–0.2), consistent with the fact that both opponents converge quickly to defection in **PD**.

A starkly different pattern emerges in LLM–LLM interactions, with cooperation rates consistently approaching 1.0 (perfect cooperation) in many cases. Notable exceptions are constituted once again by Claude 4 (particularly thinking) and DeepSeek which persistently defect, pushing cooperation rates ≤ 0.1 , confirming the previously reported non-cooperative tendencies.

Efficiency against algorithmic opponents is generally high, especially when interacting with **SREP**, as optimal play requires minimal reasoning once persistent defection is identified. The early stabilization (confirmed by lower opponent comprehension scores) minimizes token usage while maintaining (near-)optimal cumulative payoffs. A similar pattern holds for **PP**; although identifying the cyclic pattern may require a short exploratory phase, once the structure is understood, LLMs consistently defect without further reasoning overhead. LLMs with delayed comprehension incur additional token costs during exploration, which lowers efficiency despite achieving moderately high total points.

Failure Rate is negligible across models and mostly formatting-related; invalid actions are too rare to affect the main behavioral conclusions.

7.1.1 Counterfactual analysis

Label-only changes test whether models rely on memorized associations between action names and strategy. Comparing **PD** with its label-based counterfactual (where **PD** labels are replaced with **SH** labels while keeping **PD** payoffs) shows that stronger LLMs remain largely stable, whereas weaker ones—most notably Mistral Large—become more inconsistent. This degradation is not only reflected in total points, but also in delayed opponent comprehension and increased variability across prompt-

ing types. In other words, label changes alone can slow down strategic adaptation even when the underlying incentives are unchanged, indicating partial dependence on surface-level game templates rather than full abstraction over payoffs.

Payoff changes are more revealing, since they alter the strategic structure of the game itself. Moving from **PD** to **SH** incentives alleviates strict dominance and turns the interaction into a coordination problem. Many models adapt by shifting away from persistent defection and toward the higher-payoff cooperative equilibrium, often reaching totals close to the **SH** ceiling. However, this adaptation is uneven across metrics: intermediate total scores suggest partial rather than immediate coordination, delayed comprehension rounds show that several models initially favor the safer equilibrium before switching, and cooperation rates reveal whether higher totals are driven by sustained coordination or by late-stage recovery. Claude 4 variants and DeepSeek R1 remain more cautious across these metrics, often failing to fully exploit the new payoff structure. Overall, payoff-based counterfactuals expose **stronger reasoning differences** than label-only changes, since successful play requires recomputing incentives rather than merely recognizing renamed actions.

7.2 Rock-Paper-Scissors

Total points in **RPS** are centered around a zero baseline due to the zero-sum nature of the game, contrary to **PD**. As a result, performance is better interpreted through deviations from zero: positive scores indicate successful exploitation of opponent patterns, while negative scores reflect being exploited. Against algorithmic opponents, LLMs achieve consistently positive scores when exploitable patterns exist. In particular, against **PP**, stronger models (Claude 3.5/3.7 and Claude 4 variants) obtain the highest gains, reflecting successful identification and exploitation of cyclic behavior, while Llama 3.3 follows with moderate gains and Mistral Large exhibits lower and more variable performance. Against the **AP** player, total points are closer to zero across all models, indicating limited exploitation due to the opponent’s adaptive behavior. **SREP**, which follows the equilibrium distribution, also yields near-zero scores, as expected, since no systematic advantage can be obtained.

In LLM–LLM interactions, total points cluster tightly around zero, reflecting mutual adaptation and the absence of stable exploitation. However,

variance across runs reveals differences in stability: stronger LLMs maintain scores close to equilibrium, while weaker or more unstable ones occasionally deviate, resulting in subtle gains or losses. Overall, unlike PD, total points in RPS primarily capture the ability to detect and exploit structure, rather than convergence to a single optimum.

Opponent Comprehension in RPS is weaker and more variable than in PD, reflecting the greater difficulty of reasoning over three actions without a dominant response. Against algorithmic opponents, **PP** is the most informative benchmark: LLMs obtaining good total points against **PP** are precisely those that comprehend the cyclic behavior fairly early and adjust their strategy accordingly. Among these, advanced Claude models are the most consistent, while Llama 3.3 acts as a weaker but still competitive runner-up. DeepSeek R1 shows average performance, whereas Mistral Large often adapts late or inconsistently. Against **AP**, comprehension is generally delayed across models, which matches the more moderate total-point gains observed there and suggests that LLMs need several rounds of interaction before inferring the opponent’s play-style.

In LLM–LLM interactions, comprehension is less stable than in PD. Since RPS does not admit cooperative convergence, models do not settle into a shared regime as quickly, and m values remain comparatively high. These values are often fairly close to 25(>24), indicating that LLMs do not easily counter other LLMs’ strategies in the default RPS setting. For example, Claude 3.5 Sonnet v2 under zero-shot yields $m = 10.6 \pm 13.1$ against its ZS counterpart, but 21.4 ± 4.6 and 19.6 ± 5.6 against SPP and CoT counterparts, respectively. Overall, unlike PD, RPS does not produce early and stable mutual understanding; instead, opponent comprehension remains delayed and variable even among stronger models.

Cooperation Rate is not defined in the same way for RPS as in PD, since the game is strictly adversarial. Instead, this metric captures deviations from uniform play, where a balanced strategy corresponds to selecting each action with probability $\sim \frac{1}{3}$. Against algorithmic opponents, deviations from this baseline indicate exploitation: when facing the **PP** player, stronger models (Claude 3.7/Claude Sonnet 4 variants) show clear biases toward the countering actions, reflecting successful identification of cyclic patterns. In contrast, against **AP** and **SREP**, action frequencies remain closer to

uniform, suggesting limited exploitable structure and alignment with near-equilibrium play.

In LLM–LLM interactions, action distributions remain largely balanced across models, with most players approximating uniform randomization. This indicates the absence of stable exploitable patterns and is consistent with total points clustering around zero. However, weaker or less stable models (Mistral Large, DeepSeek R1) occasionally exhibit slight biases, which can lead to small but consistent gains or losses over repeated rounds.

Efficiency in RPS is generally lower and more variable than in PD, reflecting the increased difficulty of reasoning over three actions without a dominant response. Against algorithmic opponents, efficiency is highest when exploitable structure exists, particularly against the **PP** player, where stronger models (Claude 3.7 and Claude 4 variants) convert early pattern recognition into higher returns with relatively stable token usage. In contrast, against **AP** and **SREP**, efficiency decreases across all models, as limited exploitable structure leads to near-equilibrium play and continued exploration. Models with delayed opponent comprehension, such as Mistral Large and to a lesser extent DeepSeek R1, incur higher token costs without proportional gains in total points, resulting in lower efficiency overall.

Failure Rate remains negligible and mostly formatting-related, mirroring the PD findings.

7.2.1 Counterfactual analysis

Label-only changes test whether models rely on memorized dominance relations between action names. Comparing the default RPS game with its label-based counterfactual (where action labels are permuted while keeping payoffs unchanged) reveals moderate but consistent degradation for several models. Stronger LLMs, particularly Claude 3.7 and Claude 4 variants, remain relatively stable, maintaining near-equilibrium total points and comparable opponent comprehension. In contrast, weaker or less stable models—most notably Mistral Large and, to a lesser extent, DeepSeek R1—show increased variability, delayed comprehension, and occasional deviations from uniform play. This indicates that even without changing incentives, renaming actions can disrupt learned associations, leading to slower or less consistent adaptation.

Payoff changes are more challenging, as they break game symmetry and require abandoning the canonical uniform mixed strategy. Under payoff-




Figure 2: Counterfactual impact on opponent comprehension m across models (zero-shot, LLM-vs-LLM). Lower values indicate earlier adaptation, while values closer to the game horizon indicate delayed opponent modeling.

based counterfactuals, successful play requires biasing action frequencies toward higher-reward outcomes. Stronger LLMs again adapt more effectively, showing improved total points against exploitable opponents and earlier comprehension of advantageous actions. However, many models remain close to the canonical uniform distribution $(1/3, 1/3, 1/3)$, even though the payoff-counterfactual equilibrium is $(0.2, 0.2, 0.6)$. Thus, near-uniform play is no longer evidence of equilibrium behavior, but rather a form of equilibrium rigidity under altered incentives. This is reflected in higher and more variable comprehension values and lower efficiency, as additional reasoning does not consistently translate into better outcomes. The joint counterfactual (combining payoff/label changes) is the most demanding, with several models—especially Mistral Large and some thinking-enabled variants—showing near-horizon comprehension and unstable performance. This motivates a cross-game view of whether delayed comprehension appears systematically as counterfactual pressure increases.

Cross-Game Counterfactual Summary Figure 2 presents a stress trajectory from familiar to increasingly challenging game formulations. Across PD and RPS, several curves move upward from default to payoff-based and joint variants, showing that adaptation becomes delayed when surface recognition is no longer enough and incentive recomputation is required. Thus, counterfactual games reveal where strategic behavior stops transferring across altered settings.

Failure Modes Across games, counterfactual variants reveal four recurring failure modes. First, *template persistence* occurs when a model preserves the canonical policy despite changed incentives; this is clearest in payoff-based RPS,

where robust play should move away from uniform randomization toward $\sigma^* = (0.2, 0.2, 0.6)$. Second, *label anchoring* occurs when renamed actions alter behavior even though payoffs remain unchanged, suggesting partial reliance on surface associations. Third, *delayed adaptation* occurs when models eventually exploit an opponent but only after many rounds, reflected in near-horizon opponent-comprehension values. Fourth, *reasoning-overhead mismatch* occurs when reasoning-enabled variants use more tokens without proportional gains in payoff or adaptation speed. These failures are not random: label perturbations expose semantic anchoring, payoff perturbations expose incentive rigidity, and repeated interaction exposes unstable opponent modeling.

8 Conclusion

We introduced a counterfactual repeated-game framework for evaluating LLMs’ strategic adaptation under modified labels and incentives. Our results show that canonical game performance can overstate strategic robustness: many models behave competently in default PD or RPS but degrade when payoffs or dominance relations are altered. Payoff-based counterfactuals are especially diagnostic because they require recomputing incentives rather than recognizing familiar action templates. The RPS payoff-counterfactual further shows that models can remain close to canonical uniform play even when the equilibrium shifts to a biased mixed strategy. Across settings, opponent comprehension and efficiency reveal that stronger or more deliberative models are not always better strategic agents. Thus, counterfactual games offer not only a benchmark for strategic robustness, but also a diagnostic lens for separating label anchoring, incentive rigidity, delayed opponent modeling, and inefficient deliberation.

Limitations

Our study focuses on controlled, two-player repeated games, which, while enabling precise evaluation, may not fully capture the complexity of real-world strategic interactions. The use of fixed opponents and predefined payoff structures limits ecological validity, and results may vary under more diverse or multi-agent settings. Additionally, our evaluation relies on behavioral metrics derived from observed actions, which may not fully reflect the internal reasoning processes of LLMs. Finally, the set of models, prompting strategies, and counterfactual interventions considered is not exhaustive, and different configurations may yield different conclusions.

Because our evaluation relies on observed actions, our conclusions are **behavioral** rather than mechanistic, and reported metrics should be interpreted as behavioral adaptation rather than literal mental-state inference.

References

- Gati Aher, Rosa I. Arriaga, and Adam Tauman Kalai. 2023. Using large language models to simulate multiple humans and replicate human subject studies. In *Proceedings of the 40th International Conference on Machine Learning, ICML'23*. JMLR.org.
- Elif Akata, Eric Schulz, and 1 others. 2025. Playing repeated games with large language models. *Nature Human Behaviour*.
- Sanjeda Akter, Ibne Farabi Shihab, and Anuj Sharma. 2026. Causal consistency regularization: Training verifiably sensitive reasoning in large language models. *Preprint*, arXiv:2509.01544.
- Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, Athul Paul Jacob, Mo jtaba Komeili, Karthik Konath, Minae Kwon, Adam Lerer, Mike Lewis, Alexander H. Miller, Sandra Mitts, Adithya Renduchintala, Stephen Roller, and 7 others. 2022. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378:1067 – 1074.
- Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemyslaw Debiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Christopher Hesse, Rafal Józefowicz, Scott Gray, Catherine Olsson, Jakub W. Pachocki, Michael Petrov, Henrique Pondé de Oliveira Pinto, Jonathan Raiman, Tim Salimans, Jeremy Schlatter, and 6 others. 2019. Dota 2 with large scale deep reinforcement learning. *ArXiv*, abs/1912.06680.
- Yuefei Chen, Vivek K. Singh, Jing Ma, and Ruxiang Tang. 2025. Counterbench: A benchmark for counterfactuals reasoning in large language models. *Preprint*, arXiv:2502.11008.
- Katherine M. Collins, Cedegao E. Zhang, Graham Todd, Lance Ying, Mauricio Barba da Costa, Ryan Liu, Adrian Weller, Ionatan Kuperwajs, Lionel Wong, Joshua B. Tenenbaum, and Thomas L. Griffiths. 2025. Evaluating language models' evaluations of games. In *NeurIPS 2025 Workshop on Evaluating the Evolving LLM Lifecycle: Benchmarks, Emergent Abilities, and Scaling*.
- Anthony Costarelli, Mat Allen, Roman Hauksson, Grace Sodunke, Suhas Hariharan, Carlson Cheng, Wenjie Li, Joshua M Clymer, and Arjun Yadav. 2024. Gamebench: Evaluating strategic reasoning abilities of LLM agents. In *Language Gamification - NeurIPS 2024 Workshop*.
- Jinhao Duan, Renming Zhang, James Diffenderfer, Bhavya Kailkhura, Lichao Sun, Elias Stengel-Eskin, Mohit Bansal, Tianlong Chen, and Kaidi Xu. 2024. Gtbench: uncovering the strategic reasoning limitations of llms via game-theoretic evaluations. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS '24*, Red Hook, NY, USA. Curran Associates Inc.
- Caoyun Fan, Jindou Chen, Yaohui Jin, and Hao He. 2024. Can large language models serve as rational players in game theory? a systematic analysis. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence, AAAI'24/IAAI'24/EAAI'24*. AAAI Press.
- Apostolos Filippas, John J. Horton, and Benjamin S. Manning. 2024. Large language models as simulated economic agents: What can we learn from homo silicon? In *Proceedings of the 25th ACM Conference on Economics and Computation, EC '24*, page 614–615, New York, NY, USA. Association for Computing Machinery.
- Nicoló Fontana, Francesco Pierri, and Luca Maria Aiello. 2024. Nicer than humans: How do large language models behave in the prisoner's dilemma? *Preprint*, arXiv:2406.13605.
- Jörg Froberg and Frank Binder. 2022. CRASS: A novel data set and benchmark to test counterfactual reasoning of large language models. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2126–2140, Marseille, France. European Language Resources Association.
- Kanishk Gandhi, Dorsa Sadigh, and Noah Goodman. 2023. Strategic reasoning with language models. In *NeurIPS 2023 Foundation Models for Decision Making Workshop*.

- Alihan Hüyük, Xinnuo Xu, Jacqueline R. M. A. Maasch, Aditya V. Nori, and Javier Gonzalez. 2025. Reasoning elicitation in language models via counterfactual feedback. In *The Thirteenth International Conference on Learning Representations*.
- Jingru Jia, Zehua Yuan, Junhao Pan, Paul E McNamara, and Deming Chen. 2025. LLM strategic reasoning: Agentic study through behavioral game theory. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Gavin Kader and Dongwoo Lee. 2025. The emergence of strategic reasoning of large language models. Preprint, arXiv:2412.13013.
- Jiaxuan Li, Lang Yu, and Allyson Ettinger. 2023. Counterfactual reasoning: Testing language models’ understanding of hypothetical scenarios. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 804–815, Toronto, Canada. Association for Computational Linguistics.
- Nunzio Lorè and Babak Heydari. 2024. Strategic behavior of large language models and the role of game structure versus contextual framing. *Scientific Reports*, 14(1):18490.
- Mikhail Mozikov, Nikita Severin, Valeria Bodishtianu, Maria Glushanina, Ivan Nasonov, Daniil Orekhov, Vladislav Pekhotin, Ivan Makovetskiy, Mikhail Baklashkin, Vasily Lavrentyev, Akim Tsvigun, Denis Turdakov, Tatiana Shavrina, Andrey Savchenko, and Ilya Makarov. 2024. Eai: emotional decision-making of llms in strategic games and ethical dilemmas. In *Proceedings of the 38th International Conference on Neural Information Processing Systems, NIPS ’24*, Red Hook, NY, USA. Curran Associates Inc.
- Alonso Silva. 2024. Large language models playing mixed strategy nash equilibrium games. In *Network Games, Artificial Intelligence, Control and Optimization: 11th International Conference, NETGCOOP 2024, Lille, France, October 9–11, 2024, Proceedings*, page 142–152, Berlin, Heidelberg. Springer-Verlag.
- David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, L. Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Vedavyas Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy P. Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. 2016. Mastering the game of go with deep neural networks and tree search. *Nature*, 529:484–489.
- Elena Stringli, Maria Lymperaïou, Giorgos Filandrianos, Athanasios Voulodimos, and Giorgos Stamou. 2025. Pitfalls of scale: Investigating the inverse task of re-definition in large language models. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 9445–9469, Vienna, Austria. Association for Computational Linguistics.
- Haochuan Wang, Xiachong Feng, Lei Li, Yu Guo, Zhanyue Qin, Dianbo Sui, and Lingpeng Kong. 2025. Tmgbench: A systematic game benchmark for evaluating strategic reasoning abilities of llms. Preprint, arXiv:2410.10479.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- Zhenhailong Wang, Shaoguang Mao, Wenshan Wu, Tao Ge, Furu Wei, and Heng Ji. 2024. Unleashing the emergent cognitive synergy in large language models: A task-solving agent through multi-persona self-collaboration. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 257–279, Mexico City, Mexico. Association for Computational Linguistics.
- Zhaofeng Wu, Linlu Qiu, Alexis Ross, Ekin Akyürek, Boyuan Chen, Bailin Wang, Najoung Kim, Jacob Andreas, and Yoon Kim. 2024. Reasoning or reciting? exploring the capabilities and limitations of language models through counterfactual tasks. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1819–1862, Mexico City, Mexico. Association for Computational Linguistics.
- Khurram Yamin, Gaurav Rohit Ghosal, and Bryan Wilder. 2025. LLMs struggle to perform counterfactual reasoning with parametric knowledge. In *ICML 2025 Workshop on Scaling Up Intervention Models*.
- Yadong Zhang, Shaoguang Mao, Tao Ge, Xun Wang, Yan Xia, Wenshan Wu, Ting Song, Man Lan, and Furu Wei. 2024. LLM as a mastermind: A survey of strategic reasoning with large language models. In *First Conference on Language Modeling*.

A Counterfactual game instances

A.1 Prisoner’s dilemma

In the **payoff-based** counterfactual of the Prisoner’s Dilemma (PD), we replace the canonical PD payoff structure with that of the Stag Hunt (SH). While the PD satisfies the ordering

$$T > R > P > S$$

which makes defection the dominant strategy in the one-shot game, the Stag Hunt satisfies

$$R > T \geq P > S$$

thereby removing dominance and introducing a coordination structure.

Under SH, two pure Nash equilibria exist: the payoff-dominant equilibrium (S, S) that yields the highest joint reward but requires mutual commitment, and the risk-dominant equilibrium (H, H), which provides a safer but lower payoff. Unlike PD, unilateral deviation from cooperation does not strictly dominate, and strategic incentives shift from individual exploitation to equilibrium selection. This modification alters the fundamental game dynamics: In PD, rational play in the one-shot setting favors defection; in SH, strategic reasoning requires balancing risk and coordination. Consequently, agents must adjust their behavior from exploitation-oriented reasoning (as in PD) to equilibrium selection reasoning (as in SH).

This counterfactual therefore tests whether LLMs recompute incentives when the dominance structure changes. If a model continues to favor defection-like behavior under SH—despite the absence of strict dominance—this suggests overreliance on the canonical PD template. Conversely, increased coordination on the payoff-dominant equilibrium indicates sensitivity to the modified incentive structure.

A.2 Rock-Paper-Scissors

In the case of **payoff-based** RPS counterfactuals, the game is no longer symmetric and the mixed-Nash equilibrium strategy is not uniform. Consequently, players are required to adopt a biased strategy rather than the canonical uniform one:

$$\begin{aligned}\sigma^*(R) &= 0.2, \\ \sigma^*(P) &= 0.2, \\ \sigma^*(S) &= 0.6.\end{aligned}\tag{17}$$

In equilibrium, each action played with positive probability must yield equal expected payoff. Because the modified reward structure changes these expectations asymmetrically, the resulting strategy distribution becomes biased toward actions associated with higher expected returns.

This counterfactual formulation serves as a diagnostic for reasoning flexibility. If a model continues to approximate the canonical uniform distribution despite asymmetric incentives, this suggests reliance on memorized RPS heuristics rather than recomputation under the altered payoff structure. Conversely, observable shifts in action frequencies toward the higher-reward actions indicate sensitiv-

ity to incentive changes and behavioral adaptation to the modified game dynamics.

B Proof of expected outcomes

In this section, we derive the expected per-round payoff under stationary mixed strategies for both Prisoner’s Dilemma (PD) and Rock–Paper–Scissors (RPS), including their counterfactual variants.

B.1 Prisoner’s Dilemma

Let $p = \Pr(C)$ denote the probability that player i cooperates, and let $q = \Pr(C)$ denote the corresponding probability for the opponent. Then the joint-action probabilities are:

$$\begin{aligned}\Pr(C, C) &= pq, \\ \Pr(C, D) &= p(1 - q), \\ \Pr(D, C) &= (1 - p)q, \\ \Pr(D, D) &= (1 - p)(1 - q).\end{aligned}\tag{18}$$

Using the payoff instantiation

$$\begin{aligned}(C, C) &= (4, 4), & (C, D) &= (1, 6), \\ (D, C) &= (6, 1), & (D, D) &= (2, 2)\end{aligned}\tag{19}$$

the expected payoff for player i is:

$$\begin{aligned}\mathbb{E}[u_i] &= 4pq \\ &+ 1 \cdot p(1 - q) \\ &+ 6 \cdot (1 - p)q \\ &+ 2 \cdot (1 - p)(1 - q).\end{aligned}\tag{20}$$

Expanding and simplifying gives:

$$\begin{aligned}\mathbb{E}[u_i] &= 4pq + p - pq + 6q - 6pq \\ &+ 2 - 2p - 2q + 2pq \\ &= 2 - p + 4q - pq.\end{aligned}\tag{21}$$

This expression provides the expected per-round payoff as a function of the players’ mixed strategies.

B.2 Rock–Paper–Scissors

Let $\sigma_i = (p_R, p_P, p_S)$ and $\sigma_j = (q_R, q_P, q_S)$ denote the mixed strategies of player i and the opponent, respectively, where each component corresponds to the probability of selecting Rock, Paper, or Scissors. Thus,

$$\begin{aligned}p_R + p_P + p_S &= 1 \\ q_R + q_P + q_S &= 1\end{aligned}\tag{22}$$

Under the standard payoff structure, where a win yields 1, a loss yields -1 , and a tie yields 0, the expected payoff for player i is:

$$\mathbb{E}[u_i] = (p_R q_S + p_P q_R + p_S q_P) - (p_R q_P + p_P q_S + p_S q_R). \quad (23)$$

The first term corresponds to winning outcomes, while the second term corresponds to losing outcomes; ties contribute zero.

At the mixed Nash equilibrium of standard RPS, both players randomize uniformly:

$$\begin{aligned} p_R = p_P = p_S &= \frac{1}{3} \\ q_R = q_P = q_S &= \frac{1}{3} \end{aligned} \quad (24)$$

Substituting these values into the expected payoff expression gives:

$$\begin{aligned} \mathbb{E}[u_i] &= \left(\frac{1}{3} \frac{1}{3} + \frac{1}{3} \frac{1}{3} + \frac{1}{3} \frac{1}{3} \right) \\ &\quad - \left(\frac{1}{3} \frac{1}{3} + \frac{1}{3} \frac{1}{3} + \frac{1}{3} \frac{1}{3} \right) \\ &= 0. \end{aligned} \quad (25)$$

This establishes the zero expected-payoff baseline used in the main analysis. More generally, deviations from the uniform strategy can yield positive or negative expected payoff depending on the opponent's distribution.

B.3 Counterfactual payoff instantiations

B.3.1 Prisoner's Dilemma Counterfactuals

We consider three counterfactual instantiations of the PD family: a label-based counterfactual, a payoff-based counterfactual, and a joint counterfactual that alters both labels and payoffs.

Label-based counterfactual. In the label-based counterfactual, only the action names are changed, while payoffs remain identical to the default game. We replace $C \rightarrow S$ and $D \rightarrow H$, yielding:

	S	H
S	(4, 4)	(1, 6)
H	(6, 1)	(2, 2)

Payoff-based counterfactual. In the payoff-based counterfactual, the labels remain C and D , but the payoff matrix is replaced by the Stag Hunt one:

	C	D
C	(6, 6)	(1, 4)
D	(4, 1)	(2, 2)

Joint counterfactual. Finally, the joint counterfactual alters both labels and payoffs, giving the standard Stag Hunt formulation:

	S	H
S	(6, 6)	(1, 4)
H	(4, 1)	(2, 2)

Expected payoff under stationary mixed strategies.

Let p denote the probability that player i selects the first action of the game, and let q denote the corresponding probability for the opponent. In the default and payoff-based settings, the first action is C ; in the label-based and joint settings, the first action is S . The joint-action probabilities are:

$$\begin{aligned} \Pr(C, C) &= pq, \\ \Pr(C, D) &= p(1 - q), \\ \Pr(D, C) &= (1 - p)q, \\ \Pr(D, D) &= (1 - p)(1 - q), \end{aligned} \quad (26)$$

with the same notation applying to (S, H) in the label-based settings.

For the default and label-based counterfactual instantiations, the expected per-round payoff for player i is:

$$\begin{aligned} \mathbb{E}[u_i] &= 4pq \\ &\quad + 1 \cdot p(1 - q) \\ &\quad + 6 \cdot (1 - p)q \\ &\quad + 2 \cdot (1 - p)(1 - q). \end{aligned} \quad (27)$$

Expanding and simplifying:

$$\begin{aligned} \mathbb{E}[u_i] &= 4pq + p - pq + 6q - \\ &\quad 6pq + 2 - 2p - 2q + 2pq \\ &= 2 - p + 4q - pq. \end{aligned} \quad (28)$$

For the payoff-based and joint counterfactual instantiations, the expected per-round payoff becomes:

$$\begin{aligned} \mathbb{E}[u_i] &= 6pq \\ &\quad + 1 \cdot p(1 - q) \\ &\quad + 4 \cdot (1 - p)q \\ &\quad + 2 \cdot (1 - p)(1 - q). \end{aligned} \quad (29)$$

Expanding and simplifying:

$$\begin{aligned} \mathbb{E}[u_i] &= 6pq + p - pq + 4q - \\ &\quad 4pq + 2 - 2p - 2q + 2pq \\ &= 2 - p + 2q + 3pq. \end{aligned} \quad (30)$$

These expressions make explicit how the payoff-based counterfactual changes the incentive structure relative to the default PD formulation.

B.3.2 Rock–Paper–Scissors Counterfactuals

We consider three counterfactual instantiations of Rock–Paper–Scissors: a label-based counterfactual, a payoff-based counterfactual, and a joint counterfactual that alters both labels and payoffs.

Label-based counterfactual. In the label-based counterfactual, only the dominance induced by the labels is changed, while payoff magnitudes remain identical to the default game. We permute the action names so that the game is presented as:

	<i>P</i>	<i>R</i>	<i>S</i>
<i>P</i>	(0, 0)	(−1, 1)	(1, −1)
<i>R</i>	(1, −1)	(0, 0)	(−1, 1)
<i>S</i>	(−1, 1)	(1, −1)	(0, 0)

Payoff-based counterfactual. In the payoff-based counterfactual, the action labels remain *R*, *P*, and *S*, but one win–loss interaction is amplified. Specifically, the Rock–Paper outcome receives magnitude 3 instead of 1:

	<i>R</i>	<i>P</i>	<i>S</i>
<i>R</i>	(0, 0)	(−3, 3)	(1, −1)
<i>P</i>	(3, −3)	(0, 0)	(−1, 1)
<i>S</i>	(−1, 1)	(1, −1)	(0, 0)

Joint counterfactual. Finally, the joint counterfactual alters both labels and payoffs, yielding:

	<i>P</i>	<i>R</i>	<i>S</i>
<i>P</i>	(0, 0)	(−3, 3)	(1, −1)
<i>R</i>	(3, −3)	(0, 0)	(−1, 1)
<i>S</i>	(−1, 1)	(1, −1)	(0, 0)

Expected payoff under stationary mixed strategies. Let player *i* choose actions with probabilities

$$p_R = \Pr(R), \quad p_P = \Pr(P), \quad p_S = \Pr(S),$$

with $p_R + p_P + p_S = 1$, and let the opponent choose actions with probabilities

$$q_R = \Pr(R), \quad q_P = \Pr(P), \quad q_S = \Pr(S),$$

with $q_R + q_P + q_S = 1$.

For the default and label-based counterfactual instantiations, the expected per-round payoff for player *i* is:

$$\mathbb{E}[u_i] = (p_R q_S + p_P q_R + p_S q_P) - (p_R q_P + p_P q_S + p_S q_R). \quad (31)$$

The first term collects winning outcomes and the second term collects losing outcomes; ties contribute zero.

For the payoff-based and joint counterfactual instantiations, the Rock–Paper interaction is weighted by 3, so the expected payoff becomes:

$$\mathbb{E}[u_i] = 3p_P q_R + p_R q_S + p_S q_P - (3p_R q_P + p_P q_S + p_S q_R). \quad (32)$$

Equivalently, the expected payoffs of the pure actions *R*, *P*, and *S* against opponent distribution $q = (q_R, q_P, q_S)$ are:

$$\begin{aligned} u(R) &= -3q_P + q_S, \\ u(P) &= 3q_R - q_S, \\ u(S) &= -q_R + q_P. \end{aligned} \quad (33)$$

At a fully mixed Nash equilibrium, all pure actions in support must yield equal expected payoff. Since this payoff matrix is zero-sum and skew-symmetric, the value of the game is 0. Therefore:

$$\begin{aligned} -3q_P + q_S &= 0, \\ 3q_R - q_S &= 0, \\ -q_R + q_P &= 0. \end{aligned} \quad (34)$$

From the third equation, $q_R = q_P$. From the second equation, $q_S = 3q_R$. Combining these with the normalization constraint,

$$q_R + q_P + q_S = 1,$$

we obtain:

$$q_R + q_R + 3q_R = 1,$$

and therefore:

$$q_R = q_P = \frac{1}{5}, \quad q_S = \frac{3}{5}.$$

Thus, the unique mixed Nash equilibrium of the payoff-based RPS counterfactual is:

$$\sigma^*(R) = \frac{1}{5}, \quad \sigma^*(P) = \frac{1}{5}, \quad \sigma^*(S) = \frac{3}{5}.$$

Equivalently,

$$\sigma^* = (0.2, 0.2, 0.6)$$

over (R, P, S) . This shows that the payoff-based counterfactual breaks the canonical uniform equilibrium of standard RPS and requires a biased mixed strategy rather than uniform randomization.

C Implementation details

C.1 System prompts

A **Game Description** is the part of a system prompt that effectively describes the game being played.

For Prisoner's Dilemma that Game Description is the following.

You are playing a game against other player(s). The game is based on strategic moves.

Game Description:

There are two possible moves: {a} and {b}.

- if both players choose {a}, both get {aa} points.
- if both players choose {b}, both get {bb} points.
- if you choose {a} and the other player chooses {b}, you get {ab} points and the other player gets {ba} points.
- if you choose {b} and the other player chooses {a}, you get {ba} points and the other player gets {ab} points.

In each round of the game you play against one other player. You and the other player each make a move privately. The moves are compared and you and the other player get your reward/penalty based on the moves made.

Message Formats:

Move messages are sent when you want to make a move.

Move messages are formatted like this:

[move] (Optional explanation here) Your move here

Hint messages are received when you are provided with advice and directions. These messages are important and you should pay attention to them.

Hint messages are formatted like this:

[hint] Hint message here

Objective:

Maximize your points and/or minimize your penalties.

Tips:

- Try to formulate a strategy by reasoning on the provided information.
- Pay attention to the payoff matrix of the game moves.
- You should only send move messages. Hint messages will only be provided to you.
- Pay attention to hints when provided.
- Do not use markdown formatting in your messages (e.g., bold text, italic text, etc.).

For Rock-Paper-Scissors that Game Description is the following.

You are playing a game against other player(s). The game is based on making moves.

Game Description:

There are three possible moves: {a}, {b}, and {c}.

- {a} beats {c} giving the winner {ac} points and the loser {ca} points.
- {b} beats {a} giving the winner {ba} points and the loser {ab} points.
- {c} beats {b} giving the winner {cb} points and the loser {bc} points.
- If both players make the same move, the game is a tie and no points are awarded.

In each round of the game you play against one other player. You and the other player each make a move privately. The moves are compared and you and the other player get your reward/penalty based on the moves made.

Message Formats:

Move messages are sent when you want to make a move.

Move messages are formatted like this:

[move] (Optional explanation here) Your move here

Hint messages are received when you are provided with advice and directions. These messages are important and you should pay attention to them.

Hint messages are formatted like this:

[hint] Hint message here

Objective:

Maximize your points and/or minimize your penalties.

Tips:

- Try to formulate a strategy by reasoning on the provided information.
- Pay attention to the payoff matrix of the game moves.
- You should only send move messages. Hint messages will only be provided to you.
- Pay attention to hints when provided.
- Do not use markdown formatting in your messages (e.g., bold text, italic text, etc.).

C.2 Prompting techniques

We employ three prompting strategies to instantiate LLM players: zero-shot (ZS), chain-of-thought (CoT), and solo performance prompting (SPP), optionally combined with self-consistency (SC).

Zero-shot (ZS). The model receives a direct instruction describing the game, available actions, and history, and is asked to output a valid action. No explicit reasoning guidance is provided.

Template:

You are playing a repeated game. At each round, choose one of the available actions. The history of previous rounds is provided below. Based on this, select your next move. Output only your final action.

History: [previous rounds]

Available actions: [actions]

Chain-of-Thought (CoT). We augment the zero-shot prompt with explicit reasoning instructions, encouraging step-by-step analysis before selecting an action.

Template:

You are playing a repeated game. Analyze the opponent's behavior and think

step by step before deciding your move.
Then output your final action.

History: [previous rounds]

Available actions: [actions]

Solo Performance Prompting (SPP). Following Wang et al. (2024), we prompt the model to internally consider multiple reasoning perspectives before producing a final decision, aiming to improve strategic consistency.

Template:

You are playing a repeated game. Consider multiple possible strategies and perspectives before deciding your next move. Reflect briefly on the opponent’s behavior and select the best action.

History: [previous rounds]

Available actions: [actions]

Self-Consistency (SC). Self-consistency is applied on top of the above prompting strategies by sampling multiple candidate responses per round and selecting the final action via majority vote. This reduces variance and improves robustness of the selected strategy.

Decoding parameters All models are queried with temperature = 0.7, top- p = 1.0, and a maximum of 50 output tokens. Action extraction is deterministic: after generation, we select the final valid action according to the post-processing rules below.

Action extraction Model outputs are post-processed using string matching to extract valid actions (e.g., C/D or R/P/S). If multiple actions are present, the last valid token is selected. In this case the LLM resubmits its answer until there is a unique valid action present. Invalid outputs are counted toward the failure rate.

Interaction protocol At each round, the model is provided with the full history of previous actions and corresponding payoffs. The opponent’s last action is explicitly revealed. Models do not have access to future rounds.

Model access All models are accessed via their respective APIs via Bedrock AWS. Claude models are queried through Anthropic API (versions 3.5, 3.7, and 4), DeepSeek R1 via official API, and Llama 3.3-70B via hosted inference endpoints.

D Full results

Analysis in Section 7 is backed by numerical results fully presented in the following tables. These tables exhibit LLM behavior across prompting strategies and opponents for all evaluation metrics used in this paper. Due to space constraints, we use shorthand names in tables:

C3.5Sv2 = Claude Sonnet 3.5 v2,

C3.7S = Claude Sonnet 3.7,

C3.7S(T) = Claude Sonnet 3.7 (thinking),

C4S = Claude Sonnet 4,

C4S(T) = Claude Sonnet 4 (thinking),

DS-R1 = DeepSeek R1,

L3.3-70B = Llama 3.3 70B,

Mistral = Mistral Large (24.07).

D.1 PD Total points

Table 1 showcases total points averaged for the number of iterations default PD was played (16). Bold entries denote the maximum total points achieved per LLM in PD. Higher total points indicate better strategic performance, which may arise from either cooperation (e.g., in LLM–LLM interactions) or successful exploitation (e.g., against algorithmic opponents).

In Table 2 we present results for the label-based PD counterfactual in which we substitute labels as $C \rightarrow S(\text{Stag})$ and $D \rightarrow H(\text{Hare})$, without modifying the payoffs. In this setting, stronger models such as Claude 3.5/3.7 and Llama 3.3 remain largely invariant, consistently achieving the maximum cumulative reward (96 points) across prompting strategies and LLM–LLM interactions. This indicates robust abstraction over action labels and reliance on underlying incentives rather than memorized action semantics. In contrast, Claude 4 variants and DeepSeek R1 exhibit substantial degradation, with lower total points and significantly higher variance, particularly under reasoning-heavy prompting, suggesting increased sensitivity to label changes and unstable adaptation. Mistral Large remains the most inconsistent model, showing large variability across prompts and occasional recovery only under self-consistency. Across algorithmic opponents, performance remains relatively high—especially against MF—indicating that exploitation of simple patterns is preserved even under relabeling, while interactions with TFT reveal substantial instability. Overall, the label-based counterfactual highlights that while some models generalize well across superficial changes,

model	prompt	PD						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.8 ± 0.4	54.2 ± 4.5	30.2 ± 1.1	30.6 ± 0.9
	cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	53.6 ± 1.5	31.0 ± 1.0	30.4 ± 0.9
	spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.4 ± 0.5	58.4 ± 4.3	30.6 ± 0.9	32.0 ± 0.0
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	52.0 ± 0.0	32.0 ± 0.0	30.0 ± 1.4
	sc-cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	53.5 ± 2.1	32.0 ± 0.0	30.0 ± 0.0
sc-spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	62.0 ± 0.0	31.0 ± 1.4	31.0 ± 1.4	
C3.7S	zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	54.8 ± 4.1	29.8 ± 0.8	30.4 ± 0.5
	cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	61.2 ± 1.1	31.0 ± 0.7	31.2 ± 0.4
	spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	56.8 ± 4.4	31.0 ± 1.0	30.2 ± 0.4
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	52.5 ± 0.7	31.0 ± 1.4	30.0 ± 0.0
	sc-cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	61.0 ± 1.4	30.5 ± 0.7	31.0 ± 0.0
sc-spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	57.5 ± 6.4	31.0 ± 1.4	31.5 ± 0.7	
C3.7S(T)	zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.8 ± 0.4	52.8 ± 0.4	31.0 ± 0.7	31.6 ± 0.9
	cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.6 ± 0.5	57.2 ± 3.7	30.8 ± 0.8	30.2 ± 0.4
	spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.2 ± 0.4	58.0 ± 4.7	31.0 ± 1.0	31.0 ± 1.0
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	52.0 ± 0.0	32.0 ± 0.0	30.0 ± 0.0
	sc-cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	53.0 ± 0.0	31.0 ± 1.4	30.5 ± 0.7
sc-spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	57.5 ± 2.1	32.0 ± 0.0	31.0 ± 1.4	
C4S	zs	49.4 ± 15.5	41.0 ± 10.7	42.0 ± 10.7	31.4 ± 0.5	55.8 ± 4.4	32.8 ± 1.6	32.2 ± 1.6
	cot	35.2 ± 3.0	39.6 ± 2.4	39.8 ± 7.4	31.4 ± 0.5	60.2 ± 4.5	32.2 ± 1.6	34.2 ± 1.9
	spp	48.4 ± 13.0	46.2 ± 13.6	36.4 ± 4.2	31.6 ± 0.5	63.0 ± 1.0	34.4 ± 2.1	31.8 ± 2.4
	sc-zs	34.5 ± 4.9	34.0 ± 2.8	40.0 ± 4.2	32.0 ± 0.0	58.0 ± 8.5	32.0 ± 0.0	32.0 ± 0.0
	sc-cot	39.0 ± 4.2	35.5 ± 4.9	34.0 ± 2.8	32.0 ± 0.0	64.0 ± 0.0	35.0 ± 0.0	33.5 ± 2.1
sc-spp	32.0 ± 1.4	38.0 ± 1.4	35.5 ± 4.9	30.5 ± 0.7	63.5 ± 0.7	35.5 ± 0.7	31.5 ± 0.7	
C4S(T)	zs	32.0 ± 0.0	34.8 ± 2.7	32.0 ± 0.0	32.0 ± 0.0	63.4 ± 0.9	35.2 ± 1.8	35.8 ± 0.4
	cot	32.0 ± 0.0	40.8 ± 8.0	34.6 ± 3.6	31.6 ± 0.5	64.0 ± 0.0	33.8 ± 1.8	34.4 ± 2.2
	spp	31.4 ± 0.9	50.6 ± 15.4	41.8 ± 12.4	31.4 ± 0.5	59.6 ± 5.2	34.8 ± 1.8	34.8 ± 2.2
	sc-zs	37.5 ± 7.8	34.0 ± 2.8	32.0 ± 0.0	32.0 ± 0.0	64.0 ± 0.0	32.0 ± 0.0	33.5 ± 2.1
	sc-cot	32.0 ± 0.0	37.0 ± 1.4	32.0 ± 0.0	32.0 ± 0.0	64.0 ± 0.0	36.0 ± 0.0	32.0 ± 0.0
sc-spp	47.0 ± 21.2	36.0 ± 0.0	32.0 ± 0.0	31.5 ± 0.7	62.0 ± 2.8	35.5 ± 0.7	33.5 ± 2.1	
DS-R1	zs	36.8 ± 4.1	39.0 ± 2.5	30.2 ± 13.0	31.6 ± 0.5	62.6 ± 1.5	32.8 ± 2.0	32.6 ± 2.2
	cot	33.2 ± 2.6	32.2 ± 2.2	33.0 ± 6.6	28.6 ± 7.1	64.0 ± 0.0	32.2 ± 1.8	23.6 ± 8.7
	spp	36.4 ± 4.2	32.0 ± 6.4	31.2 ± 6.3	31.2 ± 0.4	61.4 ± 3.7	33.2 ± 2.6	33.2 ± 1.6
	sc-zs	32.5 ± 2.1	33.0 ± 2.8	33.5 ± 3.5	31.5 ± 0.7	64.0 ± 0.0	35.0 ± 0.0	35.5 ± 0.7
	sc-cot	34.5 ± 4.9	34.0 ± 2.8	27.0 ± 7.1	32.0 ± 0.0	64.0 ± 0.0	31.5 ± 0.7	31.5 ± 0.7
sc-spp	35.0 ± 4.2	36.0 ± 5.7	34.0 ± 2.8	31.0 ± 0.0	61.0 ± 4.2	33.5 ± 2.1	32.0 ± 0.0	
L3.3-70B	zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.6 ± 0.5	52.0 ± 0.0	31.6 ± 0.9	31.6 ± 0.9
	cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.8 ± 0.4	55.6 ± 4.6	30.8 ± 1.1	30.4 ± 0.9
	spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.8 ± 0.4	53.8 ± 2.9	30.8 ± 1.1	30.4 ± 0.9
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.0 ± 0.0	52.0 ± 0.0	31.0 ± 1.4	32.0 ± 0.0
	sc-cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	52.0 ± 0.0	32.0 ± 0.0	31.0 ± 1.4
sc-spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.5 ± 0.7	55.5 ± 4.9	30.0 ± 0.0	32.0 ± 0.0	
Mistral	zs	61.4 ± 3.2	60.0 ± 4.4	59.6 ± 9.8	24.2 ± 0.4	53.8 ± 2.0	26.0 ± 1.6	25.8 ± 0.8
	cot	64.0 ± 0.0	60.2 ± 3.6	62.0 ± 2.8	29.8 ± 2.2	57.8 ± 3.0	29.0 ± 1.2	26.0 ± 1.2
	spp	64.0 ± 0.0	62.2 ± 4.0	60.8 ± 7.2	18.6 ± 10.6	53.2 ± 2.2	26.4 ± 3.2	24.4 ± 1.1
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	24.0 ± 0.0	54.0 ± 2.8	25.0 ± 1.4	25.0 ± 1.4
	sc-cot	62.0 ± 2.8	64.0 ± 0.0	60.5 ± 4.9	27.5 ± 4.9	60.0 ± 5.7	33.5 ± 3.5	29.0 ± 4.2
sc-spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	24.0 ± 0.0	56.0 ± 0.0	22.5 ± 4.9	22.5 ± 2.1	

Table 1: Total Points Averaged Over All Iterations for the default PD. **Bold** indicates the highest total points per row.

others exhibit strong dependence on action naming, leading to degraded and inconsistent strategic behavior.

In Table 3 we demonstrate total point results for Stag Hunt, as a payoff-based counterfactual of PD. The clearest result is that the strongest models—Claude 3.5, Claude 3.7, Claude 3.7 thinking, and Llama 3.3—very often reach the maximum 96 points in LLM–LLM interactions, indicating that they successfully shift from the default PD de-

fection logic to the payoff-dominant cooperative equilibrium of Stag Hunt. This is important because, unlike the label-based counterfactual, the payoff structure here actually changes; consistently reaching 96 therefore suggests genuine sensitivity to altered incentives rather than mere invariance to relabeling.

A second important pattern is that Claude 4 and DeepSeek R1 remain much less stable under the Stag Hunt payoff change. Claude 4 improves rela-

model	prompt	PD label-based counterfactual						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	44.0 ± 5.3	52.0 ± 2.1	89.0 ± 0.0	60.6 ± 32.3
	cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	47.8 ± 3.8	49.6 ± 2.1	91.8 ± 3.8	60.0 ± 32.9
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	44.8 ± 9.8	47.6 ± 1.7	90.4 ± 3.1	72.0 ± 32.9
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	40.0 ± 7.1	54.0 ± 1.4	89.0 ± 0.0	96.0 ± 0.0
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	52.5 ± 4.9	52.0 ± 0.0	89.0 ± 0.0	66.0 ± 42.4
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	52.5 ± 13.4	52.0 ± 8.5	92.5 ± 4.9	34.5 ± 0.7
C3.7S	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	43.6 ± 8.1	49.0 ± 3.0	91.8 ± 3.8	72.6 ± 32.1
	cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	45.0 ± 5.0	52.2 ± 3.3	93.2 ± 3.8	60.2 ± 32.7
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	43.2 ± 6.7	48.2 ± 4.4	94.6 ± 3.1	48.6 ± 26.6
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	43.0 ± 14.1	46.5 ± 0.7	92.5 ± 4.9	37.5 ± 0.7
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	40.0 ± 9.9	52.0 ± 8.5	92.5 ± 4.9	66.5 ± 41.7
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	42.0 ± 5.7	49.0 ± 4.2	92.5 ± 4.9	36.5 ± 0.7
C3.7S(T)	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	44.6 ± 8.1	50.0 ± 2.0	67.2 ± 29.9	48.2 ± 26.7
	cot	96.0 ± 0.0	90.8 ± 11.6	96.0 ± 0.0	47.8 ± 7.3	49.8 ± 4.8	96.0 ± 0.0	72.2 ± 32.6
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	37.8 ± 4.6	49.6 ± 7.2	91.8 ± 3.8	72.2 ± 32.6
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	49.5 ± 20.5	48.0 ± 1.4	61.5 ± 38.9	96.0 ± 0.0
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	45.0 ± 1.4	55.5 ± 3.5	89.0 ± 0.0	35.5 ± 0.7
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	41.0 ± 8.5	48.5 ± 2.1	92.5 ± 4.9	65.5 ± 43.1
C4S	zs	71.8 ± 33.1	94.6 ± 3.1	80.4 ± 25.2	43.0 ± 5.7	47.0 ± 1.2	66.2 ± 31.2	47.4 ± 27.2
	cot	92.2 ± 3.5	66.8 ± 29.1	55.2 ± 27.3	38.2 ± 3.5	50.0 ± 5.2	45.0 ± 28.5	47.0 ± 27.4
	spp	83.0 ± 25.3	46.2 ± 27.9	46.8 ± 27.5	39.6 ± 3.0	49.2 ± 5.3	69.2 ± 33.6	45.4 ± 24.4
	sc-zs	96.0 ± 0.0	93.0 ± 4.2	88.0 ± 1.4	38.5 ± 3.5	46.0 ± 0.0	89.0 ± 0.0	36.0 ± 2.8
	sc-cot	63.0 ± 41.0	32.5 ± 0.7	34.5 ± 0.7	41.5 ± 0.7	45.5 ± 0.7	63.5 ± 46.0	36.0 ± 1.4
	sc-spp	94.0 ± 2.8	93.0 ± 4.2	57.5 ± 38.9	42.0 ± 4.2	47.0 ± 1.4	96.0 ± 0.0	54.0 ± 26.9
C4S(T)	zs	96.0 ± 0.0	84.0 ± 26.8	83.0 ± 29.1	41.6 ± 5.7	47.4 ± 1.1	79.0 ± 26.4	59.6 ± 33.2
	cot	59.0 ± 33.8	71.6 ± 33.4	70.0 ± 35.6	45.2 ± 6.1	47.2 ± 1.3	55.4 ± 30.7	71.8 ± 33.1
	spp	82.2 ± 26.1	68.0 ± 31.0	81.8 ± 27.0	39.6 ± 5.2	50.2 ± 5.5	91.0 ± 3.1	83.6 ± 27.7
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	91.5 ± 6.4	57.0 ± 2.8	50.5 ± 2.1	33.5 ± 2.1	36.0 ± 1.4
	sc-cot	65.5 ± 43.1	63.5 ± 46.0	96.0 ± 0.0	43.0 ± 4.2	47.5 ± 0.7	64.0 ± 45.3	65.5 ± 43.1
	sc-spp	92.5 ± 4.9	60.5 ± 41.7	65.0 ± 43.8	42.5 ± 2.1	47.0 ± 1.4	60.5 ± 40.3	33.0 ± 1.4
DS-R1	zs	46.6 ± 27.6	33.4 ± 2.4	30.6 ± 4.8	44.4 ± 2.9	48.0 ± 1.9	56.4 ± 33.1	46.8 ± 27.5
	cot	32.0 ± 1.9	32.6 ± 1.5	45.4 ± 28.3	38.6 ± 4.3	47.8 ± 0.8	32.2 ± 1.5	45.4 ± 28.3
	spp	31.6 ± 2.2	37.0 ± 11.8	53.4 ± 27.9	42.0 ± 2.2	47.8 ± 2.4	45.0 ± 28.5	51.6 ± 27.3
	sc-zs	33.0 ± 1.4	35.0 ± 5.7	35.0 ± 1.4	39.0 ± 1.4	48.0 ± 1.4	32.0 ± 1.4	33.5 ± 2.1
	sc-cot	63.5 ± 46.0	34.0 ± 2.8	34.0 ± 0.0	46.5 ± 9.2	46.5 ± 0.7	65.0 ± 43.8	65.5 ± 43.1
	sc-spp	28.5 ± 9.2	33.0 ± 1.4	35.0 ± 1.4	46.0 ± 2.8	47.0 ± 1.4	31.0 ± 0.0	34.0 ± 0.0
L3.3-70B	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	41.4 ± 4.3	52.0 ± 0.0	70.8 ± 34.5	47.0 ± 27.4
	cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	45.0 ± 2.9	49.6 ± 3.3	81.0 ± 26.0	46.8 ± 27.5
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	41.4 ± 10.5	48.0 ± 1.4	91.8 ± 3.8	83.8 ± 27.3
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	35.5 ± 6.4	52.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	39.5 ± 0.7	46.0 ± 0.0	96.0 ± 0.0	65.0 ± 43.8
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	44.0 ± 0.0	52.0 ± 0.0	65.5 ± 43.1	96.0 ± 0.0
Mistral	zs	96.0 ± 0.0	79.0 ± 32.2	96.0 ± 0.0	46.2 ± 6.6	42.4 ± 3.9	79.8 ± 28.6	83.4 ± 25.5
	cot	90.4 ± 11.4	84.8 ± 25.0	96.0 ± 0.0	43.2 ± 3.8	43.4 ± 4.8	79.4 ± 29.5	72.2 ± 30.4
	spp	96.0 ± 0.0	62.6 ± 42.1	92.6 ± 5.0	43.2 ± 9.9	41.2 ± 2.7	82.0 ± 24.8	59.4 ± 28.9
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	83.5 ± 17.7	42.5 ± 4.9	49.0 ± 0.0	92.5 ± 4.9	96.0 ± 0.0
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	37.5 ± 2.1	44.5 ± 2.1	92.5 ± 4.9	68.0 ± 39.6
	sc-spp	96.0 ± 0.0	60.5 ± 50.2	96.0 ± 0.0	37.0 ± 2.8	40.0 ± 0.0	89.0 ± 0.0	58.5 ± 26.2

Table 2: Total Points Averaged Over All Iterations for the label-based PD counterfactual.

tive to its weaker PD cooperation behavior in some settings, but still exhibits large variance and frequent drops well below the optimum, especially under CoT- or SC-based prompting. DeepSeek R1 performs even worse overall, mostly staying in the 30–50 range and rarely approaching the coordination optimum. This supports the claim that some models fail not only under label changes but also under genuine incentive shifts, often defaulting to cautious or inconsistent play instead of converging

to the payoff-dominant equilibrium.

Third, several algorithmic-opponent columns, especially MF and sometimes TFT, produce very high scores, often in the 90s for the stronger models. This suggests that once the altered incentives are understood, these opponents become highly exploitable. By contrast, the SREP and PP columns remain much lower, typically around the 40s–50s, indicating that not all opponents benefit equally from the SH payoff shift and that coordi-

model	prompt	SH (payoff-based counterfactual of PD)						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	43.8 ± 3.7	51.2 ± 5.1	91.8 ± 3.8	59.2 ± 33.6
	cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	43.0 ± 3.1	48.2 ± 0.8	59.4 ± 33.4	46.8 ± 27.5
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	36.8 ± 3.8	46.8 ± 1.3	79.8 ± 24.7	83.6 ± 27.7
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	44.5 ± 6.4	52.0 ± 8.5	92.5 ± 4.9	96.0 ± 0.0
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	40.5 ± 2.1	53.0 ± 7.1	34.0 ± 0.0	65.5 ± 43.1
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	49.0 ± 2.8	46.0 ± 0.0	66.0 ± 42.4	65.0 ± 43.8
C3.7S	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	46.8 ± 5.4	52.0 ± 4.7	93.2 ± 3.8	72.8 ± 31.8
	cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	42.6 ± 6.6	51.2 ± 3.7	93.2 ± 3.8	71.4 ± 33.7
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	40.4 ± 2.3	53.2 ± 4.5	93.2 ± 3.8	60.6 ± 32.3
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	39.5 ± 0.7	48.0 ± 1.4	89.0 ± 0.0	65.5 ± 43.1
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	45.0 ± 1.4	56.5 ± 2.1	92.5 ± 4.9	65.5 ± 43.1
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	42.5 ± 6.4	52.0 ± 8.5	92.5 ± 4.9	96.0 ± 0.0
C3.7S(T)	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	47.2 ± 5.8	50.2 ± 3.4	81.4 ± 26.1	59.6 ± 33.2
	cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	40.4 ± 4.8	52.2 ± 4.9	94.6 ± 3.1	59.0 ± 33.8
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	39.8 ± 4.8	54.0 ± 2.8	91.8 ± 3.8	59.6 ± 33.2
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	41.5 ± 3.5	48.0 ± 1.4	92.5 ± 4.9	65.0 ± 1.4
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	39.0 ± 1.4	55.5 ± 0.7	96.0 ± 0.0	65.5 ± 43.1
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	41.5 ± 10.6	50.5 ± 6.4	96.0 ± 0.0	35.0 ± 0.0
C4S	zs	84.4 ± 25.9	58.8 ± 31.8	70.4 ± 29.0	39.6 ± 4.9	52.0 ± 6.0	68.8 ± 34.2	52.4 ± 23.4
	cot	45.2 ± 25.1	47.0 ± 25.2	57.2 ± 30.4	45.4 ± 6.5	50.4 ± 5.9	32.0 ± 1.2	37.4 ± 4.9
	spp	46.2 ± 25.6	36.0 ± 1.9	67.0 ± 22.8	41.6 ± 4.0	51.2 ± 5.4	44.6 ± 26.0	88.6 ± 5.7
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	62.0 ± 38.2	41.0 ± 12.7	54.5 ± 4.9	92.5 ± 4.9	65.5 ± 43.1
	sc-cot	37.0 ± 4.2	59.5 ± 36.1	54.5 ± 27.6	42.0 ± 2.8	51.0 ± 1.4	33.0 ± 0.0	36.5 ± 2.1
	sc-spp	66.0 ± 42.4	35.0 ± 0.0	63.0 ± 41.0	43.5 ± 14.8	47.5 ± 2.1	61.0 ± 39.6	62.5 ± 37.5
C4S(T)	zs	68.0 ± 32.6	77.6 ± 25.6	68.2 ± 30.9	38.6 ± 5.2	48.8 ± 3.6	56.6 ± 32.9	47.6 ± 27.1
	cot	71.2 ± 33.1	77.0 ± 25.6	42.8 ± 23.6	48.6 ± 10.1	48.0 ± 1.2	49.2 ± 24.1	50.4 ± 20.8
	spp	83.8 ± 27.3	88.0 ± 6.4	68.8 ± 35.5	46.4 ± 4.7	50.2 ± 5.0	81.8 ± 28.0	49.0 ± 22.2
	sc-zs	62.5 ± 34.6	57.0 ± 29.7	92.5 ± 4.9	42.0 ± 1.4	58.0 ± 4.2	96.0 ± 0.0	35.0 ± 0.0
	sc-cot	32.5 ± 0.7	94.0 ± 2.8	66.0 ± 42.4	43.5 ± 2.1	48.5 ± 0.7	61.0 ± 39.6	65.5 ± 43.1
	sc-spp	89.0 ± 0.0	84.5 ± 10.6	63.0 ± 41.0	42.0 ± 1.4	49.5 ± 12.0	32.5 ± 0.7	36.0 ± 2.8
DS-R1	zs	43.8 ± 22.0	51.0 ± 27.8	32.0 ± 1.9	45.4 ± 5.0	47.8 ± 1.3	45.8 ± 28.1	34.2 ± 1.1
	cot	44.8 ± 28.7	33.4 ± 0.5	33.8 ± 0.8	44.0 ± 7.6	49.0 ± 2.1	33.4 ± 1.8	34.4 ± 0.9
	spp	44.6 ± 22.0	33.4 ± 1.3	58.0 ± 32.9	38.6 ± 1.9	48.0 ± 1.6	43.4 ± 25.5	33.4 ± 1.1
	sc-zs	33.5 ± 0.7	58.0 ± 36.8	33.0 ± 2.8	44.5 ± 2.1	48.0 ± 1.4	31.0 ± 0.0	33.0 ± 1.4
	sc-cot	34.0 ± 0.0	35.0 ± 2.8	34.5 ± 3.5	40.5 ± 0.7	48.0 ± 0.0	34.5 ± 2.1	33.5 ± 0.7
	sc-spp	32.5 ± 0.7	34.0 ± 1.4	32.5 ± 0.7	41.0 ± 1.4	48.0 ± 0.0	32.5 ± 2.1	33.5 ± 0.7
L3.3-70B	zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	47.0 ± 11.5	51.4 ± 1.3	82.0 ± 27.6	49.8 ± 25.9
	cot	86.8 ± 20.6	96.0 ± 0.0	96.0 ± 0.0	45.0 ± 5.7	48.8 ± 2.0	94.6 ± 3.1	83.8 ± 27.3
	spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	47.2 ± 4.1	50.2 ± 1.6	68.6 ± 31.3	71.6 ± 33.4
	sc-zs	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	43.5 ± 3.5	52.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0
	sc-cot	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	41.5 ± 10.6	47.0 ± 1.4	65.0 ± 43.8	65.0 ± 43.8
	sc-spp	96.0 ± 0.0	96.0 ± 0.0	96.0 ± 0.0	43.5 ± 6.4	50.5 ± 2.1	96.0 ± 0.0	96.0 ± 0.0
Mistral	zs	66.2 ± 20.0	83.2 ± 20.1	52.2 ± 26.1	36.4 ± 9.7	44.2 ± 7.8	88.8 ± 5.5	74.6 ± 29.5
	cot	72.6 ± 24.4	77.2 ± 17.0	62.6 ± 26.7	43.4 ± 5.7	50.4 ± 8.0	43.4 ± 25.8	49.2 ± 26.2
	spp	70.4 ± 25.9	68.4 ± 25.3	57.6 ± 27.6	38.8 ± 8.3	45.0 ± 4.5	78.8 ± 27.3	59.0 ± 27.1
	sc-zs	96.0 ± 0.0	74.5 ± 30.4	76.5 ± 27.6	41.5 ± 2.1	40.0 ± 0.0	92.5 ± 4.9	39.0 ± 1.4
	sc-cot	76.0 ± 28.3	63.0 ± 22.6	42.0 ± 14.1	39.0 ± 0.0	48.0 ± 7.1	61.5 ± 48.8	37.0 ± 2.8
	sc-spp	48.0 ± 0.0	68.5 ± 31.8	46.5 ± 12.0	44.5 ± 7.8	40.0 ± 0.0	96.0 ± 0.0	64.5 ± 37.5

Table 3: Total Points Averaged Over All Iterations for Stag Hunt (payoff-based PD counterfactual)

nation/exploitation depends strongly on the opponent’s dynamics.

Finally, Mistral remains the most unstable model, but the table also shows occasional recovery under self-consistency, including some 96-point outcomes. So while Mistral is still brittle, the Stag Hunt results suggest that SC can sometimes help it coordinate under changed incentives, even if this behavior is not reliable. Overall, the table provides strong evidence that payoff-based counterfac-

tuals are more revealing than label-only ones: the best models adapt to the new equilibrium structure, whereas weaker or more unstable models show large variance and frequent failure to coordinate on the optimal SH outcome.

Finally, in Table 4 results regarding the joint PD counterfactual game are presented. Once both labels and incentives are altered together, robustness becomes much more model-dependent than in either single intervention. The strongest evi-

model	prompt	Joint counterfactual of PD						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.2 ± 0.4	58.8 ± 5.0	31.2 ± 1.1	31.0 ± 1.0
	cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.8 ± 0.4	58.4 ± 1.8	31.4 ± 0.5	31.6 ± 0.9
	spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.4 ± 0.5	61.6 ± 0.9	31.6 ± 0.9	31.6 ± 1.8
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	59.5 ± 0.7	31.0 ± 1.4	32.0 ± 0.0
	sc-cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	58.5 ± 2.1	31.5 ± 0.7	32.0 ± 0.0
	sc-spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	59.5 ± 3.5	31.0 ± 1.4	31.5 ± 0.7
C3.7S	zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	55.6 ± 4.1	31.2 ± 0.8	30.8 ± 0.8
	cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.6 ± 0.5	55.4 ± 4.2	31.0 ± 1.0	31.2 ± 0.4
	spp	64.0 ± 0.0	58.4 ± 10.3	63.6 ± 0.9	30.2 ± 0.4	60.4 ± 2.6	32.8 ± 1.6	31.4 ± 1.5
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	62.5 ± 2.1	30.0 ± 0.0	57.5 ± 6.4	31.0 ± 1.4	31.5 ± 0.7
	sc-cot	64.0 ± 0.0	52.5 ± 16.3	64.0 ± 0.0	30.0 ± 0.0	55.5 ± 9.2	30.5 ± 0.7	32.0 ± 0.0
	sc-spp	64.0 ± 0.0	64.0 ± 0.0	50.5 ± 19.1	30.5 ± 0.7	60.5 ± 2.1	30.5 ± 0.7	32.0 ± 0.0
C3.7S(T)	zs	64.0 ± 0.0	59.6 ± 9.8	64.0 ± 0.0	30.4 ± 0.5	53.6 ± 2.2	30.8 ± 1.1	31.0 ± 1.4
	cot	55.0 ± 12.3	64.0 ± 0.0	64.0 ± 0.0	30.8 ± 0.4	58.6 ± 4.3	31.8 ± 0.4	31.2 ± 0.8
	spp	63.4 ± 1.3	57.8 ± 13.9	63.4 ± 1.3	30.2 ± 0.4	59.0 ± 3.5	31.0 ± 1.2	31.6 ± 0.9
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.0 ± 0.0	57.0 ± 7.1	32.0 ± 0.0	30.5 ± 0.7
	sc-cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	61.5 ± 2.1	31.0 ± 0.0	31.5 ± 0.7
	sc-spp	63.0 ± 1.4	64.0 ± 0.0	64.0 ± 0.0	30.5 ± 0.7	56.0 ± 5.7	32.0 ± 0.0	32.0 ± 0.0
C4S	zs	40.8 ± 13.7	38.6 ± 2.2	44.0 ± 13.1	31.2 ± 0.8	58.0 ± 4.5	32.2 ± 2.3	32.4 ± 2.3
	cot	35.8 ± 4.3	40.0 ± 12.5	38.0 ± 5.3	30.6 ± 0.5	60.4 ± 2.3	33.6 ± 1.7	32.6 ± 2.1
	spp	34.2 ± 3.5	36.4 ± 13.8	38.6 ± 9.4	31.4 ± 0.5	57.4 ± 4.5	32.4 ± 1.8	34.2 ± 2.2
	sc-zs	37.0 ± 1.4	35.0 ± 5.7	31.5 ± 0.7	31.5 ± 0.7	64.0 ± 0.0	33.5 ± 3.5	35.0 ± 1.4
	sc-cot	31.5 ± 0.7	32.0 ± 0.0	34.5 ± 4.9	32.0 ± 0.0	64.0 ± 0.0	32.0 ± 0.0	33.0 ± 1.4
	sc-spp	34.5 ± 3.5	33.5 ± 3.5	34.5 ± 4.9	32.0 ± 0.0	60.5 ± 2.1	32.0 ± 1.4	32.5 ± 0.7
C4S(T)	zs	33.0 ± 2.0	42.0 ± 11.5	32.0 ± 1.9	31.4 ± 0.5	59.4 ± 5.9	32.0 ± 1.9	34.0 ± 2.3
	cot	33.2 ± 3.9	34.0 ± 2.7	31.8 ± 0.4	32.0 ± 0.0	62.4 ± 3.6	34.2 ± 2.5	33.6 ± 2.2
	spp	34.2 ± 5.4	39.8 ± 12.7	31.4 ± 0.5	31.2 ± 0.4	60.2 ± 3.1	32.4 ± 2.1	31.4 ± 0.5
	sc-zs	36.5 ± 3.5	40.0 ± 0.0	47.5 ± 23.3	31.0 ± 0.0	64.0 ± 0.0	33.5 ± 3.5	36.0 ± 0.0
	sc-cot	40.0 ± 5.7	35.5 ± 6.4	32.0 ± 0.0	32.0 ± 0.0	64.0 ± 0.0	32.0 ± 0.0	36.0 ± 0.0
	sc-spp	32.0 ± 0.0	38.0 ± 2.8	41.5 ± 13.4	32.0 ± 0.0	58.0 ± 8.5	35.5 ± 0.7	33.0 ± 1.4
DS-R1	zs	32.8 ± 2.5	34.8 ± 0.8	32.0 ± 1.2	31.0 ± 0.7	62.2 ± 2.0	31.8 ± 1.8	33.6 ± 2.4
	cot	34.4 ± 1.9	34.0 ± 3.7	32.0 ± 1.7	31.4 ± 0.5	62.2 ± 2.2	33.6 ± 2.4	34.6 ± 2.1
	spp	36.2 ± 1.8	34.8 ± 2.6	33.6 ± 3.2	31.0 ± 0.0	63.6 ± 0.9	31.4 ± 0.5	32.6 ± 2.7
	sc-zs	33.5 ± 2.1	33.5 ± 3.5	31.5 ± 0.7	31.5 ± 0.7	64.0 ± 0.0	35.5 ± 0.7	34.0 ± 2.8
	sc-cot	35.5 ± 0.7	33.5 ± 3.5	35.5 ± 0.7	31.5 ± 0.7	64.0 ± 0.0	36.0 ± 0.0	33.5 ± 3.5
	sc-spp	34.0 ± 2.8	38.0 ± 2.8	36.0 ± 0.0	32.0 ± 0.0	64.0 ± 0.0	35.5 ± 0.7	35.5 ± 0.7
L3.3-70B	zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	28.0 ± 1.2	54.8 ± 3.6	30.0 ± 1.4	30.8 ± 0.8
	cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.2 ± 1.1	53.4 ± 3.1	30.8 ± 1.1	31.2 ± 1.1
	spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.4 ± 0.5	54.0 ± 2.8	31.0 ± 1.0	31.0 ± 1.0
	sc-zs	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	28.5 ± 0.7	52.5 ± 0.7	31.0 ± 1.4	29.5 ± 3.5
	sc-cot	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.0 ± 0.0	56.0 ± 5.7	30.0 ± 0.0	31.0 ± 1.4
	sc-spp	64.0 ± 0.0	64.0 ± 0.0	64.0 ± 0.0	29.0 ± 1.4	56.0 ± 4.2	31.0 ± 1.4	29.5 ± 0.7
Mistral	zs	52.6 ± 9.8	50.0 ± 9.9	39.2 ± 9.4	23.8 ± 0.4	53.8 ± 3.5	25.6 ± 1.7	26.0 ± 1.4
	cot	58.8 ± 3.7	59.6 ± 4.4	55.0 ± 13.1	26.8 ± 2.5	56.6 ± 6.4	28.8 ± 2.0	28.2 ± 3.3
	spp	57.2 ± 4.9	56.8 ± 7.3	51.8 ± 11.7	24.4 ± 4.1	54.0 ± 1.2	27.2 ± 1.6	24.8 ± 2.2
	sc-zs	54.0 ± 1.4	58.0 ± 5.7	54.0 ± 14.1	24.0 ± 0.0	56.0 ± 0.0	25.0 ± 1.4	26.0 ± 2.8
	sc-cot	58.5 ± 3.5	63.5 ± 9.2	43.5 ± 3.5	28.0 ± 5.7	56.5 ± 0.7	31.5 ± 2.1	30.0 ± 0.0
	sc-spp	59.0 ± 0.0	61.5 ± 3.5	32.5 ± 12.0	26.0 ± 4.2	51.0 ± 5.7	24.0 ± 0.0	27.0 ± 1.4

Table 4: Total Points Averaged Over All Iterations for the joint PD counterfactual (label-based and payoff-based).

dence comes from Claude 3.5, Claude 3.7, Claude 3.7 thinking, and Llama 3.3, which still preserve near-ceiling performance in LLM–LLM interactions, usually reaching 64.0 ± 0.0 , indicating that these models can often maintain coordination even when both the semantics and the strategic structure are perturbed. However, this robustness is weaker than in the isolated label-only or payoff-only cases, as some prompting configurations now show mild degradation or increased variance. By contrast,

Claude 4 variants and DeepSeek R1 remain substantially below the ceiling in most LLM–LLM settings, typically clustering in the 30s and 40s, which suggests that combining label and payoff changes amplifies their brittleness. Mistral is again the least stable overall: it improves in some prompting configurations, but remains far from ceiling performance and highly variable. Another useful observation is that the **PP** column remains relatively high across most models, meaning that cyclic al-

gorithmic structure is still exploitable even in the joint counterfactual, whereas the other algorithmic-opponent columns stay much closer to default-PD ranges. Overall, the table supports a strong conclusion: when both surface form and incentives change simultaneously, only the most robust models preserve strong strategic behavior, while weaker or less stable models fail to recompute the new game reliably.

D.2 PD Opponent comprehension

Opponent comprehension for the default PD is provided in Table 5, with registered values corresponding to the number of rounds at which the LLM understood the opponent’s play.

The main pattern is that PD opponent comprehension is generally very early for most strong models, especially in LLM–LLM interactions, where Claude 3.5, Claude 3.7, Claude 3.7(T), and Llama 3.3 almost always achieve $m = 1.0 \pm 0.0$. This means they establish the relevant strategic regime immediately—consistent with the perfect-cooperation totals reported earlier. In contrast, Claude 4, DeepSeek R1, and especially Mistral are less stable in these same settings, with higher and more variable m , indicating delayed or inconsistent inference of the opponent’s behavior.

Against algorithmic opponents, the easiest case is clearly SREP, where many models adapt within the first 1–3 rounds. Claude 4 and Claude 4(T) are particularly strong here, often reaching $m = 1.0 \pm 0.0$, and DeepSeek also performs well on several prompting variants. By contrast, PP is much harder: for the otherwise strong Claude 3.x and Llama models, m often rises to roughly 8–15, showing that cyclic structure takes longer to exploit than persistent defection. This supports the interpretation that PD comprehension depends strongly on opponent type rather than only model size.

A third notable point is that Mistral is the weakest and most unstable model overall. It often remains near the game horizon for algorithmic opponents—for example $m \approx 16$ for SREP, MF, and TFT in several settings—indicating that it frequently fails to reliably infer the opponent within the game. This sharply contrasts with the other models and aligns with its poorer total-point performance.

Finally, self-consistency sometimes stabilizes comprehension but does not uniformly improve it. For example, Claude 4(T) under SC often collapses to $m = 1$ across many opponent types, while for

Claude 3.x and Llama the gains are smaller because they are already near-optimal. Overall, the table shows that default PD is easy for the strongest models in cooperative LLM–LLM settings and against simple defecting opponents, but more demanding for cyclic algorithmic play and substantially harder for weaker models such as Mistral.

The label-based counterfactual opponent comprehension results are showcased in Table 6. This table shows that the label-based counterfactual preserves early comprehension in LLM–LLM play for the strongest models, but it substantially hurts comprehension against algorithmic opponents. Claude 3.5/3.7 and Llama 3.3 still achieve $m = 1.0 \pm 0.0$ across the LLM-opponent columns, indicating that relabeling alone does not disrupt mutual coordination when the opponent is another LLM. However, for the same models, comprehension against **SREP** and **PP** often deteriorates sharply to near-horizon values ($m \approx 14 \sim 17$), showing that changing action names can strongly delay recognition of deterministic strategies even when payoffs remain unchanged.

A second notable pattern is that MF remains comparatively easy for many models under the label shift, often with m close to 1–3, while **TFT** becomes much more variable. This suggests that simple frequency-based exploitation transfers more robustly across renamed actions than response policies that depend on correctly interpreting the opponent’s most recent move. Among the stronger models, Claude 4 variants remain mixed, with good LLM–LLM comprehension but unstable algorithmic performance, while DeepSeek R1 and especially Mistral are the least robust overall, frequently showing delayed comprehension across multiple opponent types. Overall, the table supports the claim that label-only changes mainly impair abstraction over opponent policy rather than cooperation with other LLMs.

Opponent comprehension results for the payoff-based PD counterfactual are presented in Table 7. In this case, opponent comprehension remains immediate in LLM–LLM play for the strongest models, with Claude 3.5/3.7 and Llama 3.3 almost uniformly achieving $m = 1.0 \pm 0.0$ across the LLM-opponent columns. This suggests that when facing another LLM, these models still establish the cooperative SH regime very quickly despite the altered incentives. However, comprehension against algorithmic opponents is substantially delayed, especially for **SREP** and often **PP**,

model	prompt	PD						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.2 ± 0.4	8.4 ± 7.0	5.6 ± 6.4	2.8 ± 1.3
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	14.0 ± 2.0	3.0 ± 1.2	2.2 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.6 ± 0.5	7.6 ± 6.1	2.0 ± 0.7	3.0 ± 0.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	12.0 ± 0.0	3.0 ± 0.0	4.5 ± 0.7
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	9.0 ± 4.2	3.0 ± 0.0	2.0 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	2.0 ± 0.0	2.5 ± 0.7	2.5 ± 0.7	
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	11.6 ± 5.5	3.6 ± 1.5	3.2 ± 1.6
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	2.0 ± 0.0	3.2 ± 1.8	3.2 ± 2.3
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	9.6 ± 6.2	3.0 ± 1.2	3.6 ± 2.3
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	13.0 ± 1.4	2.5 ± 0.7	2.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	2.0 ± 0.0	3.5 ± 2.1	5.0 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	8.0 ± 8.5	2.5 ± 0.7	4.0 ± 1.4	
C3.7S(T)	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.6 ± 1.3	12.8 ± 4.1	4.0 ± 1.4	2.8 ± 0.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.4 ± 0.5	10.4 ± 6.1	2.6 ± 1.5	1.8 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.8 ± 0.4	6.8 ± 6.7	3.0 ± 1.2	3.0 ± 1.2
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	13.0 ± 1.4	3.0 ± 0.0	2.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	15.0 ± 1.4	2.5 ± 0.7	1.5 ± 0.7
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	9.0 ± 7.1	3.0 ± 0.0	2.5 ± 0.7	
C4S	zs	1.2 ± 0.4	1.4 ± 0.9	3.0 ± 4.5	1.0 ± 0.0	12.0 ± 5.5	1.8 ± 1.1	2.2 ± 1.6
	cot	3.0 ± 2.8	10.8 ± 4.1	3.2 ± 2.3	1.0 ± 0.0	6.0 ± 6.1	2.2 ± 1.6	1.6 ± 1.3
	spp	5.2 ± 6.9	1.6 ± 1.3	4.6 ± 2.1	1.0 ± 0.0	1.8 ± 1.3	4.2 ± 7.2	2.2 ± 1.8
	sc-zs	2.5 ± 2.1	1.0 ± 0.0	7.0 ± 8.5	1.0 ± 0.0	6.5 ± 7.8	1.0 ± 0.0	2.0 ± 1.4
	sc-cot	3.5 ± 3.5	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
sc-spp	4.0 ± 4.2	3.0 ± 2.8	1.0 ± 0.0	2.0 ± 1.4	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	
C4S(T)	zs	1.0 ± 0.0	2.0 ± 2.2	1.0 ± 0.0	1.0 ± 0.0	1.2 ± 0.4	1.0 ± 0.0	1.0 ± 0.0
	cot	1.0 ± 0.0	4.2 ± 7.2	4.0 ± 6.7	1.0 ± 0.0	1.0 ± 0.0	1.6 ± 1.3	1.0 ± 0.0
	spp	1.4 ± 0.9	2.2 ± 2.2	2.4 ± 2.2	1.0 ± 0.0	5.2 ± 6.4	1.6 ± 1.3	1.0 ± 0.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	3.0 ± 2.8	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.5 ± 2.1	1.0 ± 0.0	2.0 ± 1.4	
DS-R1	zs	3.2 ± 2.2	2.0 ± 2.2	8.0 ± 8.3	1.0 ± 0.0	4.0 ± 6.7	1.0 ± 0.0	1.0 ± 0.0
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.2 ± 2.7	10.6 ± 8.8
	spp	5.2 ± 6.6	7.4 ± 7.6	5.2 ± 6.9	1.0 ± 0.0	4.8 ± 5.4	1.0 ± 0.0	1.0 ± 0.0
	sc-zs	6.5 ± 7.8	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	3.5 ± 3.5	1.0 ± 0.0	9.0 ± 11.3	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
sc-spp	2.5 ± 2.1	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	7.5 ± 9.2	1.0 ± 0.0	1.0 ± 0.0	
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.8 ± 1.1	14.0 ± 0.0	2.8 ± 0.4	2.8 ± 0.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.4 ± 0.9	10.0 ± 6.5	2.4 ± 0.5	2.2 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.4 ± 0.9	12.8 ± 5.0	2.4 ± 0.5	2.2 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	4.0 ± 0.0	14.0 ± 0.0	2.5 ± 0.7	3.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	14.0 ± 0.0	3.0 ± 0.0	2.5 ± 0.7
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	3.0 ± 1.4	9.0 ± 7.1	2.0 ± 0.0	3.0 ± 0.0	
Mistral	zs	10.2 ± 8.4	6.8 ± 6.8	3.8 ± 6.3	16.4 ± 0.5	15.2 ± 1.1	16.2 ± 0.4	16.0 ± 0.7
	cot	1.0 ± 0.0	7.0 ± 8.2	4.2 ± 7.2	4.4 ± 7.1	12.6 ± 6.5	8.6 ± 7.3	16.2 ± 0.8
	spp	1.0 ± 0.0	3.2 ± 4.9	4.2 ± 7.2	13.4 ± 6.9	15.6 ± 0.9	17.0 ± 0.0	16.4 ± 0.9
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 0.0	15.0 ± 1.4	16.5 ± 0.7	16.5 ± 0.7
	sc-cot	8.5 ± 10.6	1.0 ± 0.0	8.5 ± 10.6	9.0 ± 11.3	8.5 ± 10.6	1.0 ± 0.0	10.0 ± 9.9
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	16.0 ± 0.0	17.0 ± 0.0	16.0 ± 0.0	

Table 5: Round m where the LLM player understood the opponent’s Strategy in default PD. Bold for best results (lower number of rounds m at which opponent comprehension was realized).

where many values move close to the game horizon ($m \approx 14 \sim 17$). This indicates that switching from defection-dominant PD incentives to coordination-based SH makes deterministic opponents harder to interpret consistently, even for otherwise strong models.

A second notable trend is that MF remains the easiest algorithmic opponent under SH for most strong models, often with m near 1-3, while TFT

is much more variable. Claude 4 variants remain mixed: they can still show low m in LLM-LLM settings and occasionally on MF, but are less stable overall than Claude 3.5/3.7. DeepSeek R1 is uneven, with some early comprehension but frequent delays on SREP and PP, while Mistral is again the weakest and most inconsistent model, often showing mid- to late-horizon comprehension across several opponent types. Overall, the table supports the

model	prompt	PD label-based counterfactual						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	14.0 ± 1.9	13.4 ± 2.2	2.0 ± 0.0	10.0 ± 8.2
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.6 ± 0.5	15.4 ± 1.7	1.6 ± 0.5	9.0 ± 8.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.6 ± 1.1	16.6 ± 0.9	1.8 ± 0.4	6.0 ± 7.3
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	11.0 ± 0.0	2.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 1.4	15.0 ± 2.8	2.0 ± 0.0	9.0 ± 11.3
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 1.4	10.5 ± 9.2	1.5 ± 0.7	5.5 ± 2.1
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.6 ± 2.1	15.4 ± 2.2	1.6 ± 0.5	5.4 ± 7.0
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.8 ± 1.1	12.2 ± 3.3	1.4 ± 0.5	8.8 ± 7.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	17.0 ± 0.0	14.6 ± 5.4	1.2 ± 0.4	9.8 ± 7.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.0 ± 2.8	17.0 ± 0.0	1.5 ± 0.7	16.0 ± 1.4
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	11.0 ± 8.5	11.5 ± 7.8	1.5 ± 0.7	8.5 ± 10.6
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	17.0 ± 0.0	16.0 ± 1.4	1.5 ± 0.7	12.5 ± 3.5
C3.7S(T)	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.4 ± 1.1	14.2 ± 2.7	7.4 ± 7.5	11.4 ± 6.0
	cot	1.0 ± 0.0	1.6 ± 1.3	1.0 ± 0.0	16.4 ± 0.9	14.0 ± 5.7	1.0 ± 0.0	5.8 ± 6.6
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	13.0 ± 4.2	12.4 ± 5.7	1.6 ± 0.5	5.2 ± 6.6
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	16.0 ± 1.4	8.0 ± 8.5	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	7.5 ± 4.9	2.0 ± 0.0	9.5 ± 3.5
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	14.0 ± 0.0	16.0 ± 1.4	1.5 ± 0.7	4.0 ± 4.2
C4S	zs	2.4 ± 2.6	1.2 ± 0.4	2.0 ± 2.2	17.0 ± 0.0	16.6 ± 0.9	2.6 ± 1.9	7.2 ± 4.5
	cot	1.2 ± 0.4	5.8 ± 6.1	8.8 ± 7.4	16.2 ± 0.8	14.0 ± 5.7	3.0 ± 1.4	7.0 ± 6.0
	spp	2.4 ± 2.6	5.8 ± 4.3	4.2 ± 2.6	15.2 ± 2.5	14.8 ± 4.9	2.2 ± 2.2	7.2 ± 5.8
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	14.0 ± 1.4	17.0 ± 0.0	2.0 ± 0.0	10.5 ± 7.8
	sc-cot	3.0 ± 2.8	11.5 ± 2.1	10.5 ± 7.8	15.5 ± 2.1	17.0 ± 0.0	1.0 ± 0.0	10.0 ± 0.0
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	2.5 ± 0.7	14.5 ± 3.5	17.0 ± 0.0	1.0 ± 0.0	9.5 ± 6.4
C4S(T)	zs	1.0 ± 0.0	1.2 ± 0.4	2.0 ± 2.2	15.6 ± 1.1	17.0 ± 0.0	2.4 ± 1.5	4.4 ± 5.0
	cot	3.0 ± 4.5	2.2 ± 2.2	2.0 ± 1.7	14.8 ± 2.0	16.6 ± 0.9	1.6 ± 0.5	4.6 ± 5.0
	spp	1.2 ± 0.4	1.8 ± 1.3	1.0 ± 0.0	12.8 ± 4.8	14.2 ± 6.3	4.2 ± 5.5	1.8 ± 1.8
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	17.0 ± 0.0	14.0 ± 1.4	3.5 ± 3.5	10.0 ± 5.7
	sc-cot	2.5 ± 2.1	1.0 ± 0.0	1.0 ± 0.0	17.0 ± 0.0	17.0 ± 0.0	1.0 ± 0.0	5.0 ± 5.7
	sc-spp	1.5 ± 0.7	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	17.0 ± 0.0	1.5 ± 0.7	3.0 ± 2.8
DS-R1	zs	8.4 ± 4.9	8.2 ± 5.9	11.2 ± 8.0	15.0 ± 1.6	17.0 ± 0.0	1.6 ± 0.5	8.2 ± 6.6
	cot	4.2 ± 5.2	5.8 ± 6.6	3.6 ± 5.8	13.6 ± 5.0	17.0 ± 0.0	3.2 ± 2.7	3.8 ± 4.1
	spp	8.0 ± 5.8	4.6 ± 3.6	1.6 ± 1.3	14.8 ± 1.6	16.6 ± 0.9	1.2 ± 0.4	5.6 ± 6.6
	sc-zs	1.0 ± 0.0	8.5 ± 10.6	4.0 ± 4.2	15.5 ± 0.7	17.0 ± 0.0	1.5 ± 0.7	8.5 ± 10.6
	sc-cot	1.0 ± 0.0	7.0 ± 8.5	9.0 ± 11.3	16.0 ± 0.0	17.0 ± 0.0	1.0 ± 0.0	6.0 ± 7.1
	sc-spp	3.0 ± 2.8	1.0 ± 0.0	6.0 ± 1.4	16.0 ± 1.4	17.0 ± 0.0	1.0 ± 0.0	5.0 ± 5.7
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.4 ± 0.9	15.0 ± 0.0	3.0 ± 2.7	6.8 ± 3.8
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.4 ± 0.9	15.8 ± 1.1	2.4 ± 2.1	5.2 ± 3.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	13.8 ± 4.5	17.0 ± 0.0	1.6 ± 0.5	2.2 ± 2.7
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	12.0 ± 7.1	15.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	12.5 ± 0.7	17.0 ± 0.0	1.0 ± 0.0	3.0 ± 2.8
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.5 ± 2.1	15.0 ± 0.0	3.5 ± 3.5	1.0 ± 0.0
Mistral	zs	1.0 ± 0.0	6.4 ± 7.6	1.0 ± 0.0	16.6 ± 0.5	17.0 ± 0.0	4.4 ± 6.5	4.0 ± 6.7
	cot	4.2 ± 7.2	4.2 ± 7.2	1.0 ± 0.0	16.6 ± 0.5	17.0 ± 0.0	4.4 ± 6.5	7.4 ± 8.8
	spp	1.0 ± 0.0	7.6 ± 8.6	1.2 ± 0.4	16.6 ± 0.5	17.0 ± 0.0	4.4 ± 7.1	9.0 ± 8.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	3.0 ± 2.8	16.0 ± 1.4	16.0 ± 1.4	1.5 ± 0.7	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	17.0 ± 0.0	1.5 ± 0.7	9.0 ± 11.3
	sc-spp	1.0 ± 0.0	9.0 ± 11.3	1.0 ± 0.0	17.0 ± 0.0	17.0 ± 0.0	2.0 ± 0.0	11.0 ± 8.5

Table 6: Round $m\#$ where the LLM player understood the opponent’s Strategy in the label-based counterfactual PD setup.

idea that the payoff change affects understanding of deterministic policies more than coordination with other LLMs, and that SH introduces a harder adaptation problem than default PD.

Finally, opponent comprehension for the joint PD counterfactual setup is demonstrated in Table 8. In this case, opponent comprehension is again very early for the strongest models, especially in LLM–LLM interactions: Claude 3.5/3.7, Claude 4

thinking, and DeepSeek frequently achieve $m = 1$ across the first three columns, showing that combining payoff and label changes does not disrupt mutual adaptation for all models equally. Among algorithmic opponents, Claude 4 variants and DeepSeek R1 adapt best overall, often keeping m close to 1-2 even for SREP, MF, and TFT. Claude 3.5 is also highly stable, though its comprehension against PP is sometimes slightly later. In contrast, Llama 3.3

		SH (payoff-based counterfactual of PD)						
model	prompt	zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.8 ± 0.8	13.2 ± 5.8	1.6 ± 0.5	4.2 ± 4.1
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 1.2	17.0 ± 0.0	3.4 ± 2.3	6.2 ± 4.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	11.8 ± 5.8	17.0 ± 0.0	2.4 ± 1.5	1.8 ± 1.8
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 0.0	11.5 ± 7.8	1.5 ± 0.7	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.0 ± 2.8	11.5 ± 7.8	4.0 ± 0.0	5.5 ± 6.4
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	17.0 ± 0.0	3.0 ± 2.8	2.5 ± 2.1
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.4 ± 0.9	12.0 ± 5.0	1.4 ± 0.5	6.2 ± 7.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	14.4 ± 2.6	13.2 ± 3.0	1.4 ± 0.5	4.8 ± 6.9
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.0 ± 2.3	11.6 ± 4.0	1.4 ± 0.5	8.6 ± 8.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	17.0 ± 0.0	16.0 ± 1.4	2.0 ± 0.0	4.0 ± 4.2
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	8.5 ± 3.5	1.5 ± 0.7	8.5 ± 4.9
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.5 ± 0.7	10.5 ± 9.2	1.5 ± 0.7	1.0 ± 0.0
C3.7S(T)	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.4 ± 0.5	15.0 ± 1.4	2.2 ± 2.2	6.6 ± 6.5
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 1.0	11.2 ± 5.4	1.2 ± 0.4	4.4 ± 4.1
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	14.6 ± 2.4	10.0 ± 4.1	1.6 ± 0.5	7.0 ± 6.6
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 1.4	16.0 ± 1.4	1.5 ± 0.7	10.0 ± 7.1
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	11.5 ± 3.5	7.0 ± 2.8	1.0 ± 0.0	4.0 ± 4.2
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	11.0 ± 7.1	13.0 ± 5.7	1.0 ± 0.0	10.0 ± 1.4
C4S	zs	3.8 ± 6.3	7.6 ± 7.2	6.4 ± 7.6	16.8 ± 0.4	13.6 ± 6.5	1.8 ± 1.3	8.8 ± 5.9
	cot	3.0 ± 2.1	9.6 ± 7.1	4.4 ± 2.7	15.4 ± 1.9	13.4 ± 6.1	10.0 ± 6.6	12.4 ± 6.6
	spp	6.6 ± 5.9	8.4 ± 5.9	6.2 ± 5.4	15.8 ± 2.7	12.0 ± 6.9	4.4 ± 3.8	2.2 ± 1.1
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	8.5 ± 9.2	14.5 ± 2.1	9.5 ± 7.8	1.5 ± 0.7	4.0 ± 4.2
	sc-cot	14.5 ± 2.1	2.0 ± 1.4	9.5 ± 6.4	14.5 ± 3.5	17.0 ± 0.0	4.0 ± 2.8	14.5 ± 2.1
	sc-spp	3.5 ± 3.5	8.5 ± 6.4	6.5 ± 7.8	15.5 ± 2.1	17.0 ± 0.0	4.0 ± 2.8	7.5 ± 7.8
C4S(T)	zs	1.0 ± 0.0	4.8 ± 5.9	3.6 ± 3.4	15.0 ± 2.9	15.4 ± 3.6	2.6 ± 1.8	5.4 ± 3.0
	cot	1.2 ± 0.4	5.0 ± 6.7	2.6 ± 1.8	16.0 ± 0.0	17.0 ± 0.0	6.0 ± 6.4	10.2 ± 4.8
	spp	3.0 ± 4.5	2.2 ± 1.3	2.2 ± 2.2	16.6 ± 0.5	14.4 ± 5.8	2.0 ± 1.7	7.8 ± 6.1
	sc-zs	3.5 ± 3.5	5.0 ± 1.4	1.5 ± 0.7	15.0 ± 1.4	6.5 ± 3.5	1.0 ± 0.0	9.5 ± 3.5
	sc-cot	6.5 ± 4.9	1.0 ± 0.0	1.5 ± 0.7	15.5 ± 2.1	17.0 ± 0.0	4.0 ± 2.8	2.5 ± 2.1
	sc-spp	2.0 ± 0.0	2.0 ± 1.4	6.0 ± 7.1	17.0 ± 0.0	11.5 ± 7.8	5.0 ± 1.4	7.0 ± 8.5
DS-RI	zs	4.8 ± 3.6	1.2 ± 0.4	9.0 ± 5.5	15.8 ± 0.8	17.0 ± 0.0	2.2 ± 1.6	7.2 ± 4.0
	cot	1.8 ± 1.3	1.8 ± 0.8	5.4 ± 4.9	14.6 ± 1.8	16.6 ± 0.9	3.0 ± 1.6	7.6 ± 5.1
	spp	3.4 ± 4.3	11.2 ± 3.0	3.6 ± 5.8	12.2 ± 3.0	16.6 ± 0.9	1.6 ± 0.5	4.2 ± 3.0
	sc-zs	1.5 ± 0.7	4.0 ± 4.2	6.0 ± 7.1	13.5 ± 3.5	17.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	9.0 ± 8.5	10.0 ± 7.1	16.5 ± 0.7	17.0 ± 0.0	3.5 ± 2.1	7.0 ± 5.7
	sc-spp	5.0 ± 5.7	11.0 ± 8.5	7.5 ± 6.4	17.0 ± 0.0	17.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.4 ± 1.8	15.4 ± 0.9	2.2 ± 2.2	11.6 ± 6.8
	cot	4.0 ± 6.7	1.0 ± 0.0	1.0 ± 0.0	15.6 ± 1.1	16.6 ± 0.9	1.2 ± 0.4	2.2 ± 2.7
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	15.0 ± 1.4	16.2 ± 1.1	5.4 ± 5.2	3.8 ± 3.9
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.0 ± 1.4	15.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	11.5 ± 7.8	17.0 ± 0.0	2.5 ± 2.1	3.0 ± 2.8
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	16.5 ± 0.7	16.0 ± 1.4	1.0 ± 0.0	1.0 ± 0.0
M-L(24.07)	zs	13.4 ± 7.0	8.2 ± 7.6	12.6 ± 6.5	15.8 ± 1.6	14.4 ± 4.8	3.4 ± 1.8	7.4 ± 8.8
	cot	8.0 ± 7.9	5.0 ± 6.5	10.0 ± 7.4	16.4 ± 0.5	13.4 ± 7.0	12.8 ± 6.4	12.2 ± 6.9
	spp	7.4 ± 8.3	9.8 ± 8.1	11.2 ± 7.9	16.4 ± 1.3	17.0 ± 0.0	4.8 ± 6.4	11.0 ± 8.2
	sc-zs	1.0 ± 0.0	8.0 ± 9.9	9.0 ± 11.3	17.0 ± 0.0	17.0 ± 0.0	1.5 ± 0.7	17.0 ± 0.0
	sc-cot	8.5 ± 10.6	11.5 ± 7.8	16.0 ± 1.4	17.0 ± 0.0	16.0 ± 1.4	8.5 ± 10.6	12.5 ± 4.9
	sc-spp	16.5 ± 0.7	9.0 ± 11.3	15.0 ± 1.4	15.5 ± 2.1	17.0 ± 0.0	1.0 ± 0.0	9.0 ± 11.3

Table 7: Round $m\#$ where the LLM player understood the opponent’s Strategy in the SH counterfactual.

shows noticeably delayed comprehension on algorithmic opponents in this joint setting, particularly against SREP and PP, despite remaining perfect in LLM–LLM play.

The clearest outlier is again Mistral, which remains the weakest and most inconsistent model: many values lie near the horizon ($m \approx 10^{\sim}17$), indicating frequent failure to reliably infer the opponent even after most of the game has passed. Overall, the joint counterfactual appears less dis-

ruptive than the payoff-only SH setup for the strongest models, but it still preserves a clear separation between robust models (Claude family, often DeepSeek) and brittle ones (especially Mistral).

D.3 PD cooperation rates

Regarding cooperation rates, the results for the default PD setup are presented in Table 9. Cooperation rates in default PD closely mirror the strategic split already seen in total points and compre-

model	prompt	Joint counterfactual of PD						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.8 ± 0.4	2.0 ± 0.0	2.6 ± 0.5	2.2 ± 0.8
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.2 ± 0.4	6.0 ± 5.8	1.8 ± 1.1	2.8 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.6 ± 0.5	2.0 ± 0.0	2.8 ± 0.4	1.8 ± 0.8
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	3.0 ± 1.4	2.5 ± 0.7	3.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	4.0 ± 2.8	2.0 ± 1.4	3.0 ± 0.0
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	4.0 ± 2.8	2.5 ± 0.7	2.0 ± 1.4
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	11.6 ± 5.7	3.6 ± 1.3	3.4 ± 1.5
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.4 ± 0.5	8.8 ± 4.6	2.2 ± 0.8	4.6 ± 0.9
	spp	1.0 ± 0.0	3.6 ± 5.8	1.0 ± 0.0	1.8 ± 0.4	4.4 ± 4.3	3.4 ± 1.1	4.0 ± 1.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	8.0 ± 8.5	2.5 ± 0.7	2.0 ± 1.4
	sc-cot	1.0 ± 0.0	3.0 ± 2.8	1.0 ± 0.0	2.0 ± 0.0	1.5 ± 0.7	3.5 ± 2.1	3.0 ± 0.0
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	3.5 ± 3.5	1.5 ± 0.7	3.0 ± 1.4	1.5 ± 0.7	3.0 ± 0.0
C3.7S(T)	zs	1.0 ± 0.0	1.2 ± 0.4	1.0 ± 0.0	1.6 ± 0.5	12.8 ± 3.9	5.0 ± 1.4	3.2 ± 1.1
	cot	2.2 ± 1.8	1.0 ± 0.0	1.0 ± 0.0	1.2 ± 0.4	3.2 ± 1.8	2.6 ± 0.9	2.8 ± 1.5
	spp	1.0 ± 0.0	1.6 ± 1.3	1.0 ± 0.0	1.8 ± 0.4	5.2 ± 5.0	2.0 ± 1.2	2.8 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 0.0	8.0 ± 8.5	3.0 ± 0.0	3.5 ± 2.1
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	1.5 ± 0.7	3.0 ± 2.8	2.0 ± 1.4
	sc-spp	1.5 ± 0.7	1.0 ± 0.0	1.0 ± 0.0	1.5 ± 0.7	8.0 ± 5.7	3.0 ± 0.0	3.0 ± 0.0
C4S	zs	3.6 ± 4.8	2.2 ± 1.3	1.4 ± 0.9	1.2 ± 0.4	7.4 ± 6.3	2.0 ± 1.0	2.2 ± 1.3
	cot	1.0 ± 0.0	3.0 ± 4.5	2.0 ± 2.2	1.8 ± 1.1	4.2 ± 5.6	2.2 ± 1.6	1.6 ± 1.3
	spp	2.0 ± 1.4	2.4 ± 1.3	1.0 ± 0.0	1.0 ± 0.0	5.6 ± 0.9	3.0 ± 1.6	1.6 ± 1.3
	sc-zs	2.0 ± 1.4	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	2.0 ± 1.4
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	5.5 ± 6.4	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	3.0 ± 2.8
	sc-spp	3.5 ± 3.5	1.0 ± 0.0	3.5 ± 3.5	1.0 ± 0.0	3.5 ± 3.5	3.5 ± 3.5	3.5 ± 3.5
C4S(T)	zs	1.2 ± 0.4	1.0 ± 0.0	1.2 ± 0.4	1.0 ± 0.0	6.2 ± 7.3	1.6 ± 0.9	1.4 ± 0.9
	cot	1.2 ± 0.4	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	4.0 ± 6.7	1.0 ± 0.0	1.0 ± 0.0
	spp	4.2 ± 5.2	1.6 ± 1.3	1.0 ± 0.0	1.0 ± 0.0	5.6 ± 5.1	1.0 ± 0.0	1.6 ± 1.3
	sc-zs	2.0 ± 1.4	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	3.5 ± 3.5	1.0 ± 0.0	8.5 ± 10.6	1.0 ± 0.0	5.0 ± 5.7
DS-R1	zs	1.0 ± 0.0	3.0 ± 2.8	1.4 ± 0.9	1.8 ± 1.8	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	spp	4.2 ± 7.2	2.2 ± 1.8	2.0 ± 2.2	1.0 ± 0.0	1.6 ± 1.3	1.0 ± 0.0	4.0 ± 6.7
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	6.4 ± 3.3	10.0 ± 5.7	4.2 ± 1.1	3.6 ± 1.8
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	3.8 ± 1.8	12.0 ± 4.5	2.4 ± 0.5	2.6 ± 0.5
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	3.4 ± 1.3	13.2 ± 5.2	3.0 ± 1.2	3.0 ± 1.2
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	5.0 ± 1.4	15.0 ± 1.4	2.5 ± 0.7	9.5 ± 9.2
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	4.0 ± 0.0	8.0 ± 8.5	2.0 ± 0.0	2.5 ± 0.7
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	4.5 ± 3.5	10.0 ± 8.5	2.5 ± 0.7	3.0 ± 1.4
Mistral	zs	15.6 ± 1.5	16.2 ± 0.8	10.4 ± 7.7	16.4 ± 0.5	13.0 ± 6.7	16.6 ± 0.5	17.0 ± 0.0
	cot	9.6 ± 7.9	7.4 ± 7.9	9.4 ± 7.3	11.4 ± 7.2	5.2 ± 5.4	13.6 ± 6.0	16.8 ± 0.4
	spp	10.0 ± 7.4	7.2 ± 8.5	7.2 ± 6.5	14.4 ± 4.2	16.0 ± 0.0	13.8 ± 6.1	16.6 ± 0.5
	sc-zs	16.0 ± 1.4	9.0 ± 9.9	8.0 ± 9.9	16.0 ± 0.0	16.0 ± 0.0	16.5 ± 0.7	16.5 ± 0.7
	sc-cot	8.5 ± 10.6	10.0 ± 5.7	16.5 ± 0.7	8.5 ± 10.6	16.0 ± 0.0	11.0 ± 7.1	11.5 ± 6.4
	sc-spp	16.5 ± 0.7	9.0 ± 11.3	16.5 ± 0.7	10.5 ± 9.2	16.5 ± 0.7	16.5 ± 0.7	17.0 ± 0.0

Table 8: Round $m\#$ where the LLM player understood the opponent’s strategy for the joint PD counterfactual (label-based and payoff-based).

hension. In LLM–LLM interactions (first three columns), Claude 3.5, Claude 3.7, Claude 3.7 Thinking, and Llama 3.3 show perfect cooperation across all prompting types (≈ 1.0), confirming robust convergence to mutual cooperation. By contrast, Claude 4, Claude 4 Thinking, and DeepSeek R1 cooperate far less, often near 0–0.2, indicating persistent distrust and over-defection even against other LLMs. Mistral is unstable: some settings

remain highly cooperative, but others drop substantially, showing weaker consistency.

Against algorithmic opponents, cooperation is appropriately low for most strong models. On SREP, rates cluster around 0.1, reflecting rapid adaptation to always defect. Against PP, cooperation rises to roughly 0.3–0.5 for Claude 3.x and Llama 3.3, suggesting a short exploratory phase before exploiting the pattern. Cooperation against

model	prompt	PD						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.4 ± 0.2	0.2 ± 0.1	0.2 ± 0.1
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.1	0.2 ± 0.1	0.1 ± 0.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.2	0.1 ± 0.0	0.2 ± 0.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.2 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.4 ± 0.1	0.2 ± 0.0	0.1 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.1 ± 0.0	0.2 ± 0.0	0.2 ± 0.0	
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.4 ± 0.2	0.2 ± 0.1	0.2 ± 0.1
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.1 ± 0.0	0.2 ± 0.1	0.2 ± 0.1
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.2	0.2 ± 0.1	0.2 ± 0.1
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.1 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.0	0.2 ± 0.1	0.2 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.3	0.2 ± 0.0	0.2 ± 0.0	
C3.7S(T)	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.0	0.2 ± 0.1	0.2 ± 0.0
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.1	0.1 ± 0.1	0.1 ± 0.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.2	0.2 ± 0.1	0.2 ± 0.1
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.1 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.1 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.1	0.2 ± 0.0	0.2 ± 0.0	
C4S	zs	0.6 ± 0.4	0.2 ± 0.4	0.2 ± 0.4	0.0 ± 0.0	0.4 ± 0.2	0.1 ± 0.1	0.1 ± 0.0
	cot	0.1 ± 0.1	0.2 ± 0.0	0.2 ± 0.2	0.0 ± 0.0	0.2 ± 0.2	0.1 ± 0.0	0.1 ± 0.0
	spp	0.5 ± 0.4	0.4 ± 0.5	0.1 ± 0.0	0.0 ± 0.0	0.1 ± 0.1	0.1 ± 0.1	0.1 ± 0.1
	sc-zs	0.1 ± 0.0	0.0 ± 0.0	0.1 ± 0.1	0.0 ± 0.0	0.2 ± 0.4	0.0 ± 0.0	0.1 ± 0.1
	sc-cot	0.1 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.0	0.0 ± 0.0
sc-spp	0.1 ± 0.1	0.1 ± 0.1	0.0 ± 0.0	0.1 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	
C4S(T)	zs	0.0 ± 0.0	0.0 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.1	0.0 ± 0.0	0.0 ± 0.0
	cot	0.0 ± 0.0	0.2 ± 0.3	0.0 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.1	0.0 ± 0.0
	spp	0.0 ± 0.1	0.6 ± 0.4	0.3 ± 0.3	0.0 ± 0.0	0.2 ± 0.2	0.0 ± 0.1	0.0 ± 0.0
	sc-zs	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
	sc-cot	0.0 ± 0.0	0.1 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
sc-spp	0.5 ± 0.7	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.1	0.0 ± 0.0	0.1 ± 0.1	
DS-R1	zs	0.1 ± 0.0	0.1 ± 0.0	0.1 ± 0.1	0.0 ± 0.0	0.1 ± 0.1	0.1 ± 0.0	0.1 ± 0.0
	cot	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.1	0.0 ± 0.0
	spp	0.1 ± 0.1	0.1 ± 0.1	0.1 ± 0.1	0.1 ± 0.0	0.1 ± 0.2	0.0 ± 0.0	0.0 ± 0.0
	sc-zs	0.1 ± 0.0	0.1 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.0	0.0 ± 0.0
	sc-cot	0.1 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
sc-spp	0.1 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.2	0.0 ± 0.0	0.0 ± 0.0	
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.2 ± 0.0
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.4 ± 0.2	0.1 ± 0.0	0.1 ± 0.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.1	0.1 ± 0.0	0.1 ± 0.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.2 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.2 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.0	0.4 ± 0.2	0.1 ± 0.0	0.2 ± 0.0	
Mistral	zs	0.8 ± 0.2	0.8 ± 0.2	0.9 ± 0.3	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.1	0.5 ± 0.1
	cot	1.0 ± 0.0	0.6 ± 0.4	0.7 ± 0.4	0.1 ± 0.1	0.3 ± 0.2	0.3 ± 0.2	0.5 ± 0.1
	spp	1.0 ± 0.0	1.0 ± 0.1	0.9 ± 0.3	0.4 ± 0.3	0.5 ± 0.1	0.4 ± 0.2	0.5 ± 0.1
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.0
	sc-cot	0.9 ± 0.2	1.0 ± 0.0	0.7 ± 0.4	0.3 ± 0.3	0.2 ± 0.4	0.0 ± 0.0	0.4 ± 0.3
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.0	0.5 ± 0.0	0.8 ± 0.3	0.5 ± 0.0	

Table 9: Average Cooperation (Ratio of Cooperative Moves) in the default PD case. Bold denotes higher cooperation rates per LLM-prompt instance.

MF and TFT is again low, typically 0.1–0.2, which is consistent with these opponents quickly inducing defection. The main exception is Mistral, whose cooperation remains unusually high against algorithmic players as well ($\sim 0.4\text{--}0.8$ in several cases), matching its weaker strategic adaptation and slower opponent comprehension. Overall, the table confirms that stronger default-PD performers are those that cooperate reliably with other LLMs but sup-

press cooperation quickly against exploitative algorithmic opponents.

Table 10 presents the cooperation rate results for the label-based PD counterfactual. The label-based PD counterfactual produces a clear shift in cooperation patterns relative to default PD. In LLM–LLM interactions, the strongest models—Claude 3.5, Claude 3.7, Claude 3.7 Thinking, and Llama 3.3—remain almost perfectly cooperative, showing

model	prompt	PD label-based counterfactual						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.2	0.5 ± 0.0	0.9 ± 0.0	0.6 ± 0.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.1	0.4 ± 0.1	1.0 ± 0.0	0.6 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.2	0.3 ± 0.2	0.9 ± 0.0	0.7 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.2	0.5 ± 0.0	0.9 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.0	0.5 ± 0.0	0.9 ± 0.0	0.7 ± 0.4
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.3	0.3 ± 0.3	1.0 ± 0.0	0.2 ± 0.0
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.2	0.5 ± 0.0	1.0 ± 0.0	0.7 ± 0.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.1	0.5 ± 0.1	1.0 ± 0.0	0.6 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.1	0.2 ± 0.1	1.0 ± 0.0	0.5 ± 0.3
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.3	0.3 ± 0.2	1.0 ± 0.0	0.3 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.2	0.5 ± 0.0	1.0 ± 0.0	0.7 ± 0.5
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.2	0.3 ± 0.3	1.0 ± 0.0	0.3 ± 0.0
C3.7S(T)	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.2	0.4 ± 0.1	0.7 ± 0.4	0.4 ± 0.3
	cot	1.0 ± 0.0	0.9 ± 0.1	1.0 ± 0.0	0.5 ± 0.1	0.4 ± 0.1	1.0 ± 0.0	0.7 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.1	0.4 ± 0.2	1.0 ± 0.0	0.7 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.7 ± 0.1	0.4 ± 0.2	0.6 ± 0.5	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.0	0.5 ± 0.0	0.9 ± 0.0	0.2 ± 0.0
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.3	0.3 ± 0.1	1.0 ± 0.0	0.6 ± 0.6
C4S	zs	0.6 ± 0.5	1.0 ± 0.0	0.8 ± 0.4	0.4 ± 0.1	0.2 ± 0.2	0.6 ± 0.4	0.4 ± 0.4
	cot	0.9 ± 0.1	0.6 ± 0.4	0.5 ± 0.4	0.2 ± 0.1	0.2 ± 0.2	0.3 ± 0.4	0.3 ± 0.4
	spp	0.8 ± 0.4	0.3 ± 0.4	0.3 ± 0.4	0.2 ± 0.1	0.3 ± 0.2	0.6 ± 0.5	0.3 ± 0.4
	sc-zs	1.0 ± 0.0	0.9 ± 0.1	0.9 ± 0.0	0.2 ± 0.1	0.1 ± 0.0	0.9 ± 0.0	0.2 ± 0.1
	sc-cot	0.5 ± 0.6	0.2 ± 0.0	0.2 ± 0.0	0.3 ± 0.1	0.2 ± 0.0	0.5 ± 0.7	0.2 ± 0.0
	sc-spp	1.0 ± 0.0	0.9 ± 0.1	0.5 ± 0.5	0.3 ± 0.1	0.2 ± 0.2	1.0 ± 0.0	0.5 ± 0.4
C4S(T)	zs	1.0 ± 0.0	0.8 ± 0.4	0.8 ± 0.4	0.4 ± 0.1	0.3 ± 0.1	0.8 ± 0.4	0.5 ± 0.5
	cot	0.4 ± 0.5	0.6 ± 0.5	0.7 ± 0.5	0.3 ± 0.2	0.2 ± 0.1	0.4 ± 0.5	0.7 ± 0.4
	spp	0.8 ± 0.4	0.6 ± 0.5	0.8 ± 0.4	0.3 ± 0.1	0.3 ± 0.2	0.9 ± 0.0	0.8 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.7 ± 0.0	0.5 ± 0.0	0.1 ± 0.1	0.2 ± 0.1
	sc-cot	0.5 ± 0.7	0.5 ± 0.7	1.0 ± 0.0	0.1 ± 0.1	0.0 ± 0.0	0.5 ± 0.7	0.6 ± 0.6
	sc-spp	1.0 ± 0.0	0.5 ± 0.6	0.5 ± 0.7	0.3 ± 0.0	0.2 ± 0.0	0.5 ± 0.7	0.1 ± 0.1
DS-R1	zs	0.3 ± 0.4	0.1 ± 0.0	0.1 ± 0.1	0.3 ± 0.1	0.2 ± 0.2	0.4 ± 0.5	0.3 ± 0.4
	cot	0.1 ± 0.1	0.1 ± 0.0	0.2 ± 0.4	0.3 ± 0.2	0.1 ± 0.1	0.1 ± 0.1	0.2 ± 0.4
	spp	0.1 ± 0.1	0.1 ± 0.2	0.3 ± 0.4	0.2 ± 0.1	0.2 ± 0.1	0.2 ± 0.4	0.3 ± 0.4
	sc-zs	0.0 ± 0.0	0.1 ± 0.1	0.1 ± 0.1	0.0 ± 0.0	0.2 ± 0.2	0.1 ± 0.0	0.1 ± 0.1
	sc-cot	0.5 ± 0.7	0.1 ± 0.1	0.1 ± 0.1	0.3 ± 0.2	0.1 ± 0.0	0.5 ± 0.7	0.5 ± 0.7
	sc-spp	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.3	0.1 ± 0.1	0.1 ± 0.0	0.1 ± 0.1
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.2	0.5 ± 0.0	0.7 ± 0.4	0.3 ± 0.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.1	0.3 ± 0.2	0.8 ± 0.4	0.3 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.2	0.4 ± 0.2	1.0 ± 0.0	0.8 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.5 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.1	0.1 ± 0.0	1.0 ± 0.0	0.6 ± 0.6
	sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.0	0.5 ± 0.0	0.6 ± 0.6	1.0 ± 0.0
Mistral	zs	1.0 ± 0.0	0.9 ± 0.2	1.0 ± 0.0	0.6 ± 0.1	0.5 ± 0.0	0.9 ± 0.2	0.9 ± 0.3
	cot	0.9 ± 0.1	0.9 ± 0.2	1.0 ± 0.0	0.5 ± 0.1	0.4 ± 0.2	0.8 ± 0.3	0.8 ± 0.3
	spp	1.0 ± 0.0	0.7 ± 0.4	1.0 ± 0.0	0.6 ± 0.1	0.5 ± 0.1	0.9 ± 0.2	0.6 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	0.9 ± 0.1	0.6 ± 0.0	0.5 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.2	0.5 ± 0.0	1.0 ± 0.0	0.8 ± 0.4
	sc-spp	1.0 ± 0.0	0.8 ± 0.3	1.0 ± 0.0	0.6 ± 0.1	0.5 ± 0.0	0.9 ± 0.0	0.7 ± 0.3

Table 10: Average Cooperation (Ratio of Cooperative Moves) in the label-based PD counterfactual setup.

that simple relabeling does not substantially disrupt mutual coordination for stronger models. Mistral is also much more cooperative here than in the default setup, often approaching 1.0, which suggests that the renamed actions may actually encourage more socially aligned behavior in weaker models. By contrast, DeepSeek R1 remains the least cooperative, with rates mostly around 0.0–0.3, indicating that label shifts do not resolve its general tendency toward defection or instability.

Against algorithmic opponents, cooperation becomes much higher than in default PD for many models. In particular, cooperation against MF often rises to 0.9–1.0 for Claude 3.x, Llama 3.3, and Mistral, while TFT also elicits moderate-to-high cooperation in many settings. Even against SREP and PP, rates are frequently around 0.3–0.5, much higher than the near-zero cooperation observed in the default game. This suggests that the label substitution weakens immediate recognition

of the underlying incentive structure: models often preserve cooperative behavior under renamed actions instead of quickly switching to the more exploitative default-PD response. Overall, the table supports the claim that label-only counterfactuals do not strongly harm strong LLM–LLM cooperation, but they do alter adaptation to algorithmic opponents by making models more cooperative and therefore less tightly aligned with the original PD incentives.

Table 11 presents the cooperation rate results for the payoff-based PD counterfactual. The payoff-based counterfactual successfully shifts many models toward more cooperative behavior, but this shift is highly model-dependent.

The clearest pattern is that Claude 3.5, Claude 3.7, Claude 3.7(T), and Llama 3.3 remain strongly cooperative in LLM–LLM interactions, almost always staying at or near 1.0 cooperation across ZS, CoT, SPP, and SC. This suggests that these models adapt well to the SH coordination incentives and consistently select the socially efficient action when paired with other cooperative LLMs. They also often show very high cooperation against MF and sometimes against TFT, indicating that once the opponent appears coordination-friendly, these models sustain cooperative play.

By contrast, Claude 4 variants are much more mixed. Their cooperation is often moderate rather than maximal, and strongly depends on the prompting variant. Self-consistency sometimes helps substantially—for example, pushing some LLM–LLM cases or MF cases close to full cooperation—but overall these models appear less stable and more cautious than the Claude 3.x models. DeepSeek R1 is the least cooperative among the stronger models: cooperation is mostly low across opponent types, with only occasional moderate increases under specific prompts, suggesting persistent reluctance to move toward SH-style coordination. Mistral is notable in a different way: unlike in default PD, it often displays moderate to high cooperation in SH, especially in LLM–LLM settings and against MF, but remains highly variable across prompts and opponents.

Across algorithmic opponents, cooperation is no longer uniformly low as in default PD. Against SREP and PP, many models now show intermediate cooperation levels (roughly 0.2–0.5, sometimes higher), which is consistent with SH no longer rewarding universal defection. Against MF, several strong models frequently reach near-perfect cooper-

ation, while TFT produces more variable outcomes, reflecting the difficulty of sustaining reciprocal coordination over time. Overall, the table supports the main claim that the SH payoff change induces a real behavioral shift from defection toward coordination, with the strongest evidence coming from Claude 3.x and Llama 3.3, while Claude 4, DeepSeek, and Mistral remain more unstable in how fully they exploit the new incentive structure.

Finally, in Table 11 we demonstrate the cooperation rate results for the joint PD counterfactual. The joint PD counterfactual setup behaves much more like the default PD regime than like SH: for most strong models, cooperation remains maximal in LLM–LLM play but stays low against algorithmic opponents, indicating that the label+payoff intervention did not induce the broad coordination shift observed in the pure SH payoff counterfactual.

The clearest pattern is that Claude 3.5, Claude 3.7, Claude 3.7(T), and Llama 3.3 still cooperate almost perfectly (≈ 1.0) against LLM counterparts across nearly all prompting variants. This suggests that mutual cooperation remains easy to sustain in same-model interactions even under the joint reformulation. However, against SREP, PP, MF, and TFT, these same models mostly drop back to low cooperation rates (typically around 0.1–0.3, with only occasional values around 0.4–0.5), which is much closer to the default PD pattern than to the SH payoff-based counterfactual. In other words, once both labels and payoffs are altered together, these models do not generalize the cooperative SH behavior as strongly as they do when only payoffs are changed.

D.4 Efficiency in PD

Table 13 exhibits the efficiency metric over the default and the three counterfactual PD settings. A concise takeaway is that efficiency does not perfectly track raw capability. Some of the strongest models in total points are not the most efficient, while some lighter or less reliable models achieve high points-per-token because they use fewer tokens or behave more tersely.

The most notable result is that Llama 3.3 70B is the strongest overall efficiency performer among the robust models. It posts the highest or near-highest efficiency values across many settings, especially in pd-alt and sh, suggesting a very favorable trade-off between strategic performance and token cost. This supports the broader claim that Llama is not always the strongest on absolute strate-

		SH (payoff-based counterfactual of PD)						
model	prompt	zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.1	0.3 ± 0.2	1.0 ± 0.0	0.5 ± 0.5
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.3 ± 0.1	0.5 ± 0.5	0.3 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.1	0.8 ± 0.4	0.8 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.3	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.3 ± 0.2	0.1 ± 0.0	0.6 ± 0.6
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.0	0.1 ± 0.0	0.6 ± 0.6	0.6 ± 0.6	
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.2	0.5 ± 0.1	1.0 ± 0.0	0.7 ± 0.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.2	0.5 ± 0.0	1.0 ± 0.0	0.7 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.1	0.4 ± 0.2	1.0 ± 0.0	0.6 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.0	0.4 ± 0.2	0.9 ± 0.0	0.6 ± 0.6
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.0	0.5 ± 0.0	1.0 ± 0.0	0.2 ± 0.1
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.0	0.3 ± 0.3	1.0 ± 0.0	1.0 ± 0.0	
C3.7S(T)	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.1	0.4 ± 0.2	0.8 ± 0.4	0.5 ± 0.4
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.1	0.4 ± 0.2	1.0 ± 0.0	0.5 ± 0.5
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.5 ± 0.0	1.0 ± 0.0	0.6 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.2	0.4 ± 0.2	1.0 ± 0.0	0.2 ± 0.1
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.5 ± 0.0	1.0 ± 0.0	0.6 ± 0.6
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.3	0.3 ± 0.3	1.0 ± 0.0	0.2 ± 0.0	
C4S	zs	0.8 ± 0.3	0.6 ± 0.4	0.7 ± 0.4	0.2 ± 0.1	0.3 ± 0.2	0.6 ± 0.5	0.5 ± 0.4
	cot	0.2 ± 0.4	0.3 ± 0.4	0.4 ± 0.5	0.3 ± 0.1	0.3 ± 0.2	0.2 ± 0.1	0.2 ± 0.2
	spp	0.3 ± 0.4	0.1 ± 0.1	0.6 ± 0.3	0.3 ± 0.1	0.4 ± 0.1	0.3 ± 0.4	0.9 ± 0.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	0.6 ± 0.5	0.2 ± 0.0	0.5 ± 0.0	1.0 ± 0.0	0.6 ± 0.6
	sc-cot	0.2 ± 0.2	0.4 ± 0.6	0.5 ± 0.4	0.0 ± 0.0	0.5 ± 0.0	0.1 ± 0.1	0.2 ± 0.0
sc-spp	0.6 ± 0.6	0.1 ± 0.1	0.6 ± 0.5	0.4 ± 0.2	0.2 ± 0.1	0.6 ± 0.5	0.6 ± 0.5	
C4S(T)	zs	0.6 ± 0.5	0.8 ± 0.3	0.6 ± 0.4	0.2 ± 0.2	0.3 ± 0.2	0.4 ± 0.5	0.3 ± 0.4
	cot	0.6 ± 0.5	0.8 ± 0.3	0.2 ± 0.4	0.4 ± 0.2	0.2 ± 0.1	0.3 ± 0.4	0.5 ± 0.3
	spp	0.8 ± 0.4	0.9 ± 0.1	0.7 ± 0.4	0.3 ± 0.1	0.3 ± 0.2	0.8 ± 0.4	0.4 ± 0.4
	sc-zs	0.5 ± 0.6	0.5 ± 0.5	1.0 ± 0.0	0.2 ± 0.1	0.5 ± 0.0	1.0 ± 0.0	0.2 ± 0.0
	sc-cot	0.2 ± 0.0	1.0 ± 0.0	0.5 ± 0.7	0.2 ± 0.2	0.3 ± 0.0	0.6 ± 0.5	0.5 ± 0.7
sc-spp	0.9 ± 0.0	0.8 ± 0.1	0.6 ± 0.5	0.3 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.1 ± 0.2	
DS-R1	zs	0.2 ± 0.4	0.3 ± 0.4	0.1 ± 0.0	0.2 ± 0.2	0.1 ± 0.1	0.3 ± 0.4	0.1 ± 0.1
	cot	0.2 ± 0.4	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.1	0.2 ± 0.1	0.1 ± 0.0	0.1 ± 0.1
	spp	0.2 ± 0.4	0.1 ± 0.0	0.4 ± 0.5	0.1 ± 0.1	0.1 ± 0.1	0.2 ± 0.4	0.1 ± 0.0
	sc-zs	0.0 ± 0.0	0.5 ± 0.5	0.1 ± 0.0	0.2 ± 0.1	0.1 ± 0.0	0.1 ± 0.0	0.0 ± 0.0
	sc-cot	0.0 ± 0.0	0.1 ± 0.0	0.1 ± 0.0	0.1 ± 0.1	0.1 ± 0.1	0.1 ± 0.0	0.1 ± 0.0
sc-spp	0.0 ± 0.0	0.1 ± 0.0	0.1 ± 0.0	0.1 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.5 ± 0.1	0.5 ± 0.0	0.8 ± 0.4	0.5 ± 0.3
	cot	0.9 ± 0.2	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.2	0.3 ± 0.1	1.0 ± 0.0	0.8 ± 0.4
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.1	0.5 ± 0.0	0.7 ± 0.4	0.7 ± 0.4
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.1	0.5 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.3 ± 0.2	0.2 ± 0.0	0.6 ± 0.6	0.6 ± 0.6
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.4 ± 0.4	0.5 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	
Mistral	zs	0.7 ± 0.2	0.9 ± 0.2	0.7 ± 0.2	0.6 ± 0.1	0.5 ± 0.0	0.9 ± 0.1	0.8 ± 0.3
	cot	0.6 ± 0.4	0.8 ± 0.3	0.5 ± 0.4	0.5 ± 0.2	0.5 ± 0.0	0.4 ± 0.3	0.5 ± 0.3
	spp	0.7 ± 0.3	0.7 ± 0.4	0.6 ± 0.3	0.5 ± 0.1	0.5 ± 0.1	0.8 ± 0.2	0.7 ± 0.3
	sc-zs	1.0 ± 0.0	0.8 ± 0.3	0.8 ± 0.2	0.6 ± 0.0	0.5 ± 0.0	1.0 ± 0.0	0.4 ± 0.1
	sc-cot	0.7 ± 0.4	0.4 ± 0.5	0.2 ± 0.3	0.4 ± 0.3	0.5 ± 0.0	0.7 ± 0.4	0.2 ± 0.2
sc-spp	0.2 ± 0.4	0.8 ± 0.3	0.5 ± 0.0	0.6 ± 0.1	0.5 ± 0.0	1.0 ± 0.0	0.7 ± 0.4	

Table 11: Average Cooperation (Ratio of Cooperative Moves) in the Stag Hunt counterfactual.

gic sophistication, but it is often one of the best cost-performance choices.

Among the Claude family, Claude 3.5 Sonnet v2, Claude 3.7 Sonnet, and Claude 3.7 Sonnet (Thinking) remain reasonably efficient, but efficiency drops substantially once one moves from ZS to CoT/SPP, and drops again under SC. This shows that extra reasoning and majority-vote aggregation often add token cost faster than they improve payoff. In particular, the thinking variants do not

yield proportional efficiency gains: even when they improve robustness or points in some settings, their extra deliberation often lowers points-per-token.

Claude Sonnet 4 and Claude Sonnet 4 Thinking are clearly less efficient overall than Claude 3.5/3.7 and Llama. Their values are consistently lower, especially under SC, which fits the earlier observation that these models often over-reason or behave cautiously without enough performance gain to justify the added token cost. So although Claude 4 can

model	prompt	Joint counterfactual of PD						
		zs	spp	cot	srep	pp	mf	tft
C3.5Sv2	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.2	0.2 ± 0.0	0.1 ± 0.1
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.1	0.1 ± 0.1	0.2 ± 0.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.1 ± 0.0	0.2 ± 0.0	0.1 ± 0.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.0	0.2 ± 0.0	0.2 ± 0.0
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.1	0.1 ± 0.1	0.2 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.1	0.2 ± 0.0	0.1 ± 0.1	
C3.7S	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.4 ± 0.2	0.2 ± 0.1	0.2 ± 0.1
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.4 ± 0.2	0.1 ± 0.1	0.2 ± 0.0
	spp	1.0 ± 0.0	0.8 ± 0.3	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.1	0.1 ± 0.0	0.2 ± 0.1
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	0.9 ± 0.1	0.1 ± 0.0	0.3 ± 0.3	0.2 ± 0.0	0.1 ± 0.1
	sc-cot	1.0 ± 0.0	0.6 ± 0.6	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.3	0.2 ± 0.1	0.2 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	0.6 ± 0.6	0.1 ± 0.0	0.2 ± 0.1	0.1 ± 0.0	0.2 ± 0.0	
C3.7S(T)	zs	1.0 ± 0.0	0.8 ± 0.4	1.0 ± 0.0	0.1 ± 0.0	0.5 ± 0.1	0.2 ± 0.0	0.2 ± 0.0
	cot	0.7 ± 0.5	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.1	0.2 ± 0.1	0.2 ± 0.1
	spp	1.0 ± 0.1	0.8 ± 0.4	1.0 ± 0.1	0.1 ± 0.0	0.2 ± 0.1	0.1 ± 0.1	0.2 ± 0.0
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.3	0.2 ± 0.0	0.2 ± 0.1
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.1 ± 0.1	0.2 ± 0.1	0.1 ± 0.1
sc-spp	0.9 ± 0.1	1.0 ± 0.0	1.0 ± 0.0	0.1 ± 0.0	0.3 ± 0.2	0.2 ± 0.0	0.2 ± 0.0	
C4S	zs	0.3 ± 0.4	0.1 ± 0.1	0.3 ± 0.4	0.1 ± 0.1	0.3 ± 0.2	0.1 ± 0.1	0.1 ± 0.1
	cot	0.0 ± 0.0	0.3 ± 0.4	0.0 ± 0.1	0.1 ± 0.0	0.2 ± 0.1	0.1 ± 0.1	0.1 ± 0.1
	spp	0.1 ± 0.1	0.3 ± 0.3	0.2 ± 0.4	0.0 ± 0.0	0.3 ± 0.1	0.1 ± 0.0	0.1 ± 0.1
	sc-zs	0.1 ± 0.1	0.1 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.1
	sc-cot	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.1
sc-spp	0.1 ± 0.1	0.0 ± 0.0	0.1 ± 0.0	0.0 ± 0.0	0.2 ± 0.1	0.1 ± 0.1	0.1 ± 0.1	
C4S(T)	zs	0.0 ± 0.1	0.2 ± 0.4	0.1 ± 0.1	0.0 ± 0.0	0.2 ± 0.2	0.1 ± 0.1	0.1 ± 0.1
	cot	0.0 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.2	0.0 ± 0.0	0.0 ± 0.0
	spp	0.1 ± 0.1	0.2 ± 0.4	0.0 ± 0.0	0.1 ± 0.0	0.2 ± 0.2	0.0 ± 0.0	0.1 ± 0.1
	sc-zs	0.1 ± 0.0	0.0 ± 0.0	0.5 ± 0.7	0.1 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
	sc-cot	0.1 ± 0.1	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
sc-spp	0.0 ± 0.0	0.0 ± 0.0	0.1 ± 0.2	0.0 ± 0.0	0.2 ± 0.4	0.0 ± 0.0	0.1 ± 0.1	
DS-R1	zs	0.1 ± 0.0	0.1 ± 0.1	0.1 ± 0.1	0.1 ± 0.0	0.1 ± 0.1	0.1 ± 0.0	0.1 ± 0.0
	cot	0.1 ± 0.0	0.0 ± 0.0	0.1 ± 0.0	0.0 ± 0.0	0.1 ± 0.1	0.1 ± 0.0	0.0 ± 0.0
	spp	0.0 ± 0.1	0.1 ± 0.1	0.1 ± 0.1	0.1 ± 0.0	0.0 ± 0.1	0.0 ± 0.0	0.1 ± 0.0
	sc-zs	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
	sc-cot	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0
sc-spp	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	0.0 ± 0.0	
L3.3-70B	zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.4 ± 0.1	0.2 ± 0.0	0.2 ± 0.1
	cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.5 ± 0.1	0.1 ± 0.0	0.2 ± 0.0
	spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.0	0.5 ± 0.1	0.2 ± 0.1	0.2 ± 0.1
	sc-zs	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.0	0.5 ± 0.0	0.2 ± 0.0	0.3 ± 0.2
	sc-cot	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.0	0.3 ± 0.2	0.1 ± 0.0	0.2 ± 0.0
sc-spp	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	0.2 ± 0.1	0.4 ± 0.2	0.2 ± 0.0	0.2 ± 0.0	
Mistral	zs	0.7 ± 0.2	0.6 ± 0.2	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.0
	cot	0.5 ± 0.2	0.4 ± 0.2	0.4 ± 0.3	0.3 ± 0.2	0.3 ± 0.2	0.4 ± 0.1	0.3 ± 0.2
	spp	0.7 ± 0.1	0.6 ± 0.3	0.7 ± 0.2	0.5 ± 0.3	0.6 ± 0.0	0.5 ± 0.2	0.5 ± 0.0
	sc-zs	0.7 ± 0.1	0.9 ± 0.0	0.9 ± 0.2	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.0	0.5 ± 0.0
	sc-cot	0.4 ± 0.1	0.5 ± 0.0	0.6 ± 0.0	0.2 ± 0.4	0.4 ± 0.1	0.3 ± 0.1	0.3 ± 0.0
sc-spp	0.6 ± 0.1	0.7 ± 0.3	0.6 ± 0.1	0.4 ± 0.3	0.5 ± 0.0	0.5 ± 0.0	0.4 ± 0.0	

Table 12: Average Cooperation (Ratio of Cooperative Moves) in the joint counterfactual setup.

be competitive in some raw outcomes, it appears less attractive from an efficiency standpoint.

DeepSeek R1 is interesting because it is fairly competitive in efficiency despite mixed strategic quality. Its values are often moderate and stable, and in a few cases it attains the best score in a row. This suggests that DeepSeek’s weaker strategic consistency is partly offset by lower token usage. Still, its efficiency advantage should be interpreted carefully, since it does not necessarily reflect stronger

reasoning.

Finally, Mistral shows the clearest “efficiency can be misleading” pattern. It achieves some of the highest efficiency scores in the table, especially in pd-alt and under some SC settings, but earlier results showed that it is also one of the most brittle models overall. This likely means that Mistral is efficient because it is relatively cheap in tokens, not because it is consistently the best strategic reasoner. In that sense, Mistral looks like a high-variance,

model	prompt	PD	Label-based	Payoff-based	Joint
C3.5Sv2	zs	18.35 ± 10.24	28.04 ± 12.12	21.84 ± 12.24	15.56 ± 9.72
	cot	9.00 ± 4.05	15.88 ± 8.80	11.25 ± 5.37	7.85 ± 2.92
	spp	7.54 ± 2.98	11.99 ± 4.73	9.60 ± 4.14	5.93 ± 2.41
	sc-zs	6.37 ± 3.42	9.88 ± 4.56	7.63 ± 3.60	4.59 ± 2.71
	sc-cot	2.96 ± 1.46	6.06 ± 4.01	5.50 ± 5.26	2.58 ± 1.00
	sc-spp	2.54 ± 1.09	3.81 ± 1.63	4.70 ± 4.25	1.96 ± 0.76
C3.7S	zs	14.03 ± 7.88	24.35 ± 12.64	24.04 ± 12.10	13.16 ± 7.68
	cot	6.39 ± 2.64	10.14 ± 4.40	10.68 ± 4.72	6.17 ± 2.67
	spp	5.48 ± 3.66	7.89 ± 5.63	7.22 ± 3.47	4.57 ± 3.10
	sc-zs	4.88 ± 2.49	7.02 ± 4.64	7.70 ± 4.79	4.37 ± 2.50
	sc-cot	2.34 ± 1.32	3.63 ± 1.78	3.15 ± 1.60	2.01 ± 1.07
	sc-spp	2.36 ± 1.37	3.21 ± 2.72	2.67 ± 1.01	1.51 ± 0.69
C3.7S(T)	zs	15.67 ± 6.23	22.93 ± 10.25	30.01 ± 15.17	19.17 ± 9.07
	cot	8.44 ± 3.39	13.73 ± 5.86	13.68 ± 6.63	7.86 ± 3.05
	spp	6.82 ± 2.85	10.99 ± 4.98	9.87 ± 4.58	6.10 ± 2.53
	sc-zs	5.35 ± 2.15	9.53 ± 4.24	9.88 ± 5.87	6.90 ± 2.89
	sc-cot	2.67 ± 1.19	4.16 ± 2.16	4.45 ± 2.13	2.71 ± 1.05
	sc-spp	2.21 ± 0.86	3.45 ± 1.66	3.49 ± 1.84	2.08 ± 0.91
C4S	zs	7.45 ± 3.99	13.85 ± 8.83	15.12 ± 12.56	7.92 ± 4.45
	cot	4.15 ± 0.81	6.34 ± 3.79	4.53 ± 2.42	4.19 ± 0.82
	spp	3.85 ± 1.26	4.49 ± 2.59	4.46 ± 2.23	3.45 ± 0.81
	sc-zs	2.15 ± 0.48	4.86 ± 3.10	5.45 ± 3.33	1.99 ± 0.41
	sc-cot	1.36 ± 0.18	1.50 ± 0.94	1.51 ± 0.47	1.41 ± 0.25
	sc-spp	1.13 ± 0.35	2.01 ± 0.99	1.36 ± 0.81	1.04 ± 0.31
C4S(T)	zs	11.59 ± 3.05	24.90 ± 12.96	19.88 ± 11.98	11.66 ± 3.63
	cot	5.84 ± 1.51	9.81 ± 6.49	8.43 ± 5.02	5.37 ± 1.18
	spp	3.99 ± 1.62	6.68 ± 2.97	6.04 ± 2.96	3.50 ± 0.87
	sc-zs	3.99 ± 1.00	7.91 ± 3.90	7.51 ± 3.45	4.19 ± 1.89
	sc-cot	1.88 ± 0.38	3.43 ± 2.19	3.13 ± 2.20	1.94 ± 0.54
	sc-spp	1.32 ± 0.36	1.80 ± 0.90	1.68 ± 0.87	1.07 ± 0.28
DS-R1	zs	13.86 ± 5.08	12.59 ± 5.53	13.47 ± 6.02	12.38 ± 4.01
	cot	8.48 ± 5.19	9.71 ± 4.00	11.77 ± 3.26	12.48 ± 3.23
	spp	8.22 ± 3.06	8.79 ± 5.23	8.51 ± 4.02	8.68 ± 2.39
	sc-zs	4.37 ± 1.43	3.66 ± 0.95	4.72 ± 2.32	4.51 ± 1.36
	sc-cot	2.74 ± 1.31	4.70 ± 2.15	4.41 ± 0.99	4.39 ± 1.34
	sc-spp	2.69 ± 0.86	2.29 ± 0.92	2.40 ± 0.37	3.02 ± 1.29
L3.3-70B	zs	25.03 ± 9.53	38.08 ± 17.58	40.90 ± 17.03	25.09 ± 9.76
	cot	18.62 ± 7.48	28.85 ± 13.32	30.11 ± 12.13	17.57 ± 6.97
	spp	16.46 ± 6.41	27.37 ± 9.98	25.53 ± 10.79	16.63 ± 6.44
	sc-zs	8.37 ± 3.35	15.04 ± 5.46	14.98 ± 5.11	8.50 ± 3.24
	sc-cot	6.22 ± 2.55	10.28 ± 4.55	9.55 ± 4.55	5.68 ± 2.40
	sc-spp	5.50 ± 2.14	8.98 ± 3.58	9.78 ± 3.30	5.62 ± 2.20
Mistral	zs	27.04 ± 11.01	47.46 ± 18.77	39.23 ± 17.11	23.17 ± 8.92
	cot	6.07 ± 3.11	9.60 ± 4.27	7.18 ± 3.49	5.57 ± 2.45
	spp	4.51 ± 2.62	7.15 ± 4.36	6.01 ± 4.00	4.45 ± 2.33
	sc-zs	9.65 ± 4.33	16.86 ± 5.59	13.38 ± 5.80	8.68 ± 3.49
	sc-cot	2.19 ± 0.84	3.52 ± 1.21	2.14 ± 1.12	1.72 ± 0.57
	sc-spp	1.84 ± 1.08	2.71 ± 1.55	2.10 ± 1.08	1.61 ± 1.01

Table 13: Average Efficiency (Points per kilo-token) for all PD setups.

high-efficiency option, whereas Llama is the more convincing balanced winner.

Overall, the table reinforces three points: (i) zero-shot is usually the most efficient prompting strategy, (ii) CoT/SPP and especially SC often reduce efficiency unless they deliver a substantial behavioral gain, and (iii) the best trade-off between capability and efficiency appears to come from Llama 3.3 70B, while Claude 4 variants are the clearest examples of stronger models whose extra reasoning is not matched by proportional efficiency

gains.

D.5 PD validity rate

We present validity rates for PD in Table 14. Validity in PD is effectively perfect for nearly all models, indicating that formatting compliance and action-space adherence do not meaningfully affect the behavioral results. Claude 3.5 Sonnet v2, Claude 3.7, Claude 3.7(T), Claude 4, Claude 4(T), and Llama 3.3 all achieve a strict 100.0% ± 0.0 validity rate, showing completely stable instruction following

model	avg validity rate
C3.5Sv2	100.0 \pm 0.0
C3.7S	100.0 \pm 0.0
C3.7S(T)	100.0 \pm 0.0
C4S	100.0 \pm 0.0
C4S(T)	100.0 \pm 0.0
DS-R1	99.1 \pm 6.3
L3.3-70B	100.0 \pm 0.0
Mistral	99.4 \pm 6.5

Table 14: Average Valid Rate (% of Valid Outcomes)

in this simpler two-action setting. The only minor deviations come from DeepSeek R1 (99.1% \pm 6.3) and Mistral (99.4% \pm 6.5), whose averages remain very high but whose larger variance suggests occasional formatting or parsing failures. Overall, invalid outputs are too rare to explain any of the observed differences in total points, opponent comprehension, cooperation, or efficiency. Instead, the PD results should be interpreted as reflecting genuine strategic differences across models, with Mistral and DeepSeek showing slightly weaker output reliability in addition to their broader behavioral instability.

D.6 RPS total points

Considering the full repetition experiment of 24 RPS rounds, the following Tables are going to demonstrate the averaged point count in all RPS setups considered (default and three counterfactuals).

Commencing from the default RPS setup, results are presented in Table 15. A concise interpretation of the default RPS total points table is that performance is highly opponent-dependent, with clear separation between algorithmic opponents (especially PP, AP, TFT) and LLM–LLM play, where scores tend to remain close to zero.

First, across almost all models and prompting strategies, LLM–LLM settings (zs, spp, cot) produce scores close to 0, often with high variance. This is consistent with the mixed Nash equilibrium of RPS: when both agents behave unpredictably, neither side gains a consistent advantage. The large variance indicates that models do not perfectly randomize, but deviations are inconsistent rather than systematically exploitable.

In contrast, algorithmic opponents are consistently exploitable, and this is where most positive scores arise. In particular, Patterned Player (PP) and Tit-for-Tat (TFT) often yield the highest gains, with many models reaching 15–24 points, especially under self-consistency (SC) variants. This

suggests that models are capable of detecting and exploiting regularities or reactive patterns, even if they struggle to maintain optimal mixed strategies in symmetric play.

Among models, Claude 3.7 (both standard and thinking) and Claude 4 variants are the strongest overall in RPS. They achieve the most consistent high scores against algorithmic opponents and benefit notably from SC prompting, which often boosts results toward the upper range ($\approx 20^{\sim}24$). However, their performance in LLM–LLM settings still centers around zero, indicating that gains come primarily from exploitation rather than equilibrium play.

DeepSeek R1 and Claude 3.5 show moderate performance: they can exploit structured opponents (especially TFT and PP) but with less consistency and lower ceilings. Llama 3.3 is relatively stable but less aggressive, often achieving smaller gains, suggesting more conservative or less adaptive strategies.

Finally, Mistral exhibits the most unstable behavior: while it sometimes achieves high scores (e.g., against PP or under SC), it also shows large negative values and high variance, especially under CoT and SC. This reinforces its earlier profile as brittle and inconsistent, even when occasional exploitation succeeds.

Overall, the results highlight that in default RPS, success is driven more by opponent exploitation than by equilibrium play, and that prompting strategies like SC amplify exploitation ability rather than improving adherence to optimal mixed strategies.

Label-based RPS counterfactuals are demonstrated in Table 16. A brief takeaway from the label-based RPS counterfactual is that simply renaming / permuting the dominance structure hurts many models substantially, especially against PP, while performance against AP and sometimes TFT remains more positive. This is exactly the pattern expected if some models rely partly on the canonical “Rock beats Scissors, Scissors beats Paper, Paper beats Rock” template rather than fully recomputing the altered dominance relation.

Compared with default RPS, the most striking change is that several models now score negative points against PP, sometimes heavily so. This is especially visible for Claude 3.5, Llama 3.3, and parts of DeepSeek and Mistral, indicating that the cyclic patterned opponent becomes much harder to exploit once the familiar label semantics are disrupted. In contrast, Claude 3.7 and particularly

model	prompt	RPS						
		zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	6.6 ± 13.5	-1.2 ± 9.2	-1.8 ± 6.9	-0.4 ± 3.3	6.0 ± 10.1	11.2 ± 6.4	16.6 ± 6.9
	cot	0.8 ± 5.4	2.0 ± 5.7	5.4 ± 4.8	0.2 ± 5.3	3.8 ± 11.9	6.0 ± 3.7	13.0 ± 6.8
	spp	8.0 ± 9.9	7.2 ± 11.6	1.6 ± 6.9	-1.4 ± 2.2	9.4 ± 11.2	6.8 ± 1.5	12.8 ± 7.6
	sc-zs	-12.5 ± 12.0	-0.5 ± 21.9	0.5 ± 4.9	-0.5 ± 9.2	1.5 ± 2.1	9.0 ± 1.4	21.5 ± 0.7
	sc-cot	9.0 ± 0.0	-1.0 ± 2.8	0.0 ± 7.1	1.0 ± 4.2	10.5 ± 14.8	14.0 ± 12.7	16.5 ± 2.1
	sc-spp	-2.5 ± 21.9	-9.0 ± 17.0	0.0 ± 9.9	7.0 ± 1.4	-1.5 ± 3.5	5.0 ± 5.7	12.5 ± 14.8
C3.7S	zs	1.4 ± 9.6	-4.6 ± 11.3	-3.0 ± 11.9	-2.0 ± 5.0	13.6 ± 12.5	8.8 ± 2.6	4.8 ± 7.1
	cot	9.6 ± 5.2	2.6 ± 9.8	0.4 ± 5.5	1.4 ± 2.1	19.6 ± 1.3	8.2 ± 1.5	20.2 ± 3.6
	spp	3.6 ± 9.6	4.8 ± 6.0	-5.2 ± 8.6	-2.4 ± 2.3	19.2 ± 1.3	14.2 ± 7.2	19.0 ± 4.8
	sc-zs	-2.0 ± 28.3	0.0 ± 7.1	-15.5 ± 9.2	-4.0 ± 2.8	14.0 ± 14.1	6.5 ± 0.7	0.0 ± 0.0
	sc-cot	20.5 ± 4.9	-4.5 ± 4.9	1.5 ± 2.1	-3.5 ± 6.4	21.0 ± 0.0	9.5 ± 0.7	22.0 ± 2.8
	sc-spp	9.0 ± 12.7	-0.5 ± 0.7	-1.0 ± 5.7	2.5 ± 10.6	19.5 ± 2.1	12.0 ± 1.4	21.0 ± 2.8
C3.7S(T)	zs	0.2 ± 3.6	0.0 ± 5.3	-4.2 ± 7.3	-0.6 ± 3.6	19.6 ± 1.9	7.6 ± 2.8	17.6 ± 4.0
	cot	-2.0 ± 6.0	0.2 ± 3.2	-3.8 ± 5.1	-0.4 ± 3.2	21.0 ± 0.0	9.8 ± 4.4	15.6 ± 6.7
	spp	-3.2 ± 4.9	2.2 ± 10.9	-1.6 ± 2.9	5.2 ± 4.1	20.8 ± 0.4	8.4 ± 1.7	15.8 ± 9.3
	sc-zs	-5.0 ± 2.8	-7.5 ± 0.7	7.5 ± 3.5	-3.5 ± 4.9	21.0 ± 0.0	15.5 ± 9.2	1.0 ± 1.4
	sc-cot	-2.0 ± 12.7	-8.0 ± 9.9	-9.5 ± 6.4	3.5 ± 4.9	18.5 ± 0.7	14.0 ± 8.5	22.0 ± 0.0
	sc-spp	-2.0 ± 11.3	-9.5 ± 12.0	-16.5 ± 4.9	-4.0 ± 4.2	21.0 ± 0.0	9.0 ± 2.8	22.5 ± 2.1
C4S	zs	0.2 ± 4.8	-5.4 ± 5.3	-4.2 ± 8.1	2.4 ± 1.9	18.2 ± 8.0	10.6 ± 1.7	9.2 ± 8.4
	cot	3.4 ± 17.3	3.2 ± 4.6	4.4 ± 12.5	1.2 ± 3.8	19.6 ± 2.1	12.4 ± 6.1	12.2 ± 7.3
	spp	5.0 ± 8.5	-7.6 ± 16.9	4.6 ± 9.4	1.0 ± 6.4	18.0 ± 3.7	12.4 ± 3.5	10.4 ± 7.3
	sc-zs	12.5 ± 9.2	-19.5 ± 3.5	-11.5 ± 12.0	0.0 ± 0.0	24.0 ± 0.0	10.5 ± 2.1	18.0 ± 8.5
	sc-cot	-5.5 ± 6.4	1.0 ± 18.4	1.5 ± 14.8	3.5 ± 4.9	21.0 ± 0.0	9.5 ± 0.7	18.0 ± 8.5
	sc-spp	-0.5 ± 2.1	6.5 ± 9.2	-6.0 ± 4.2	1.0 ± 2.8	21.0 ± 0.0	15.5 ± 9.2	20.0 ± 1.4
C4S(T)	zs	-0.2 ± 2.3	-2.8 ± 7.3	-0.4 ± 15.9	1.0 ± 6.2	19.2 ± 4.0	11.8 ± 0.8	15.0 ± 10.6
	cot	7.6 ± 10.4	0.0 ± 3.4	-2.6 ± 11.5	2.0 ± 4.6	19.8 ± 1.6	10.0 ± 7.2	14.6 ± 8.1
	spp	-2.2 ± 5.4	-1.8 ± 10.6	-1.2 ± 5.7	0.4 ± 5.6	20.8 ± 0.4	8.4 ± 2.4	15.2 ± 8.3
	sc-zs	-5.0 ± 7.1	0.0 ± 2.8	0.5 ± 9.2	2.5 ± 0.7	21.0 ± 0.0	7.0 ± 5.7	22.0 ± 2.8
	sc-cot	11.5 ± 6.4	8.5 ± 2.1	12.0 ± 2.8	-0.5 ± 3.5	21.0 ± 0.0	11.0 ± 8.5	23.0 ± 1.4
	sc-spp	5.0 ± 1.4	-1.0 ± 0.0	1.0 ± 9.9	3.0 ± 1.4	21.0 ± 0.0	10.5 ± 2.1	23.5 ± 0.7
DS-R1	zs	1.2 ± 7.8	-1.0 ± 1.6	-3.2 ± 9.7	0.8 ± 3.0	5.8 ± 2.9	6.8 ± 5.0	14.6 ± 3.8
	cot	4.0 ± 8.2	5.0 ± 4.3	-4.2 ± 4.5	4.4 ± 4.2	12.4 ± 4.2	9.6 ± 3.0	17.8 ± 3.4
	spp	0.8 ± 6.3	-2.4 ± 5.9	-0.4 ± 6.3	-1.4 ± 3.5	6.2 ± 4.8	10.6 ± 2.2	13.4 ± 5.4
	sc-zs	5.0 ± 9.9	-3.5 ± 4.9	-0.5 ± 6.4	-3.5 ± 0.7	9.0 ± 8.5	6.5 ± 3.5	14.5 ± 0.7
	sc-cot	-0.5 ± 13.4	-3.5 ± 4.9	-3.5 ± 9.2	1.0 ± 1.4	15.0 ± 1.4	15.5 ± 3.5	20.5 ± 2.1
	sc-spp	-6.0 ± 5.7	-2.5 ± 2.1	-6.0 ± 1.4	0.0 ± 1.4	13.0 ± 2.8	7.0 ± 2.8	20.0 ± 2.8
L3.3-70B	zs	-0.8 ± 1.8	-0.4 ± 7.1	1.8 ± 5.4	-2.0 ± 4.7	14.8 ± 12.6	6.6 ± 2.1	1.0 ± 1.0
	cot	-0.6 ± 1.3	-1.6 ± 8.8	0.0 ± 2.4	0.6 ± 3.8	19.0 ± 8.7	8.2 ± 3.1	10.0 ± 8.9
	spp	1.0 ± 1.7	-1.6 ± 2.4	-1.4 ± 1.8	-0.4 ± 2.9	6.2 ± 7.8	9.8 ± 6.0	4.4 ± 4.0
	sc-zs	0.0 ± 0.0	-8.0 ± 11.3	2.0 ± 0.0	-1.0 ± 2.8	12.5 ± 16.3	9.0 ± 2.8	0.0 ± 0.0
	sc-cot	-0.5 ± 0.7	2.0 ± 2.8	1.5 ± 2.1	-0.5 ± 2.1	6.0 ± 2.8	7.0 ± 1.4	16.5 ± 9.2
	sc-spp	0.0 ± 0.0	-10.5 ± 16.3	-0.5 ± 2.1	5.5 ± 0.7	19.0 ± 7.1	12.5 ± 2.1	2.0 ± 2.8
Mistral	zs	0.0 ± 1.4	-10.6 ± 7.7	-3.8 ± 2.8	0.8 ± 5.2	20.8 ± 7.2	5.4 ± 3.6	0.8 ± 0.8
	cot	8.8 ± 5.4	0.0 ± 2.6	-4.2 ± 10.7	0.4 ± 4.2	-1.0 ± 12.9	1.6 ± 4.2	16.0 ± 9.5
	spp	9.2 ± 5.4	7.0 ± 13.2	0.6 ± 8.4	0.8 ± 1.5	6.8 ± 8.2	2.2 ± 4.1	12.6 ± 6.6
	sc-zs	0.5 ± 0.7	-4.0 ± 2.8	-13.5 ± 0.7	1.0 ± 1.4	24.0 ± 0.0	0.5 ± 0.7	0.0 ± 0.0
	sc-cot	11.5 ± 10.6	5.5 ± 3.5	-2.5 ± 30.4	-0.5 ± 0.7	11.5 ± 17.7	-0.5 ± 7.8	17.5 ± 7.8
	sc-spp	-0.5 ± 0.7	-0.5 ± 0.7	2.5 ± 17.7	0.0 ± 4.2	6.0 ± 7.1	-4.0 ± 2.8	6.5 ± 21.9

Table 15: Total Points Averaged Over All Iterations for the default RPS setup. Bold denotes best results (most total points).

Claude 3.7 Thinking remain among the strongest and most robust models: they still obtain strong positive scores against PP/TFT in several prompt settings, with some of the best results in the table, suggesting better adaptation to the remapped action semantics.

A second consistent pattern is that AP remains one of the most reliably exploitable opponents across models. For many rows, AP is the best or

among the best-performing columns, often yielding moderate positive returns even when PP becomes negative. This suggests that models can still respond to short-horizon behavioral regularities even when they struggle with the deeper symbolic remapping of the whole RPS structure.

Model-wise, Claude 3.7 / Claude 3.7 Thinking look strongest overall under this counterfactual, with high peaks against PP, AP, and TFT. Claude 4

		Label-based RPS counterfactual						
model	prompt	zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	-1.6 ± 3.5	-3.4 ± 2.3	-0.8 ± 3.2	-0.4 ± 3.4	-4.4 ± 2.5	9.0 ± 2.6	5.2 ± 7.1
	cot	0.8 ± 8.0	-2.8 ± 6.7	-0.4 ± 1.7	-5.0 ± 4.7	-9.0 ± 10.9	4.4 ± 3.9	-0.8 ± 3.4
	spp	1.6 ± 3.7	-2.6 ± 1.5	1.8 ± 4.6	0.2 ± 5.0	-8.0 ± 5.0	6.6 ± 3.9	2.0 ± 4.4
	sc-zs	3.0 ± 5.7	-2.0 ± 4.2	3.5 ± 2.1	3.5 ± 2.1	-12.0 ± 2.8	7.0 ± 1.4	12.0 ± 1.4
	sc-cot	-2.5 ± 0.7	-2.5 ± 0.7	-4.5 ± 12.0	0.5 ± 2.1	-20.5 ± 0.7	4.0 ± 0.0	5.5 ± 4.9
	sc-spp	-9.0 ± 15.6	5.0 ± 11.3	0.5 ± 6.4	4.5 ± 2.1	-17.0 ± 5.7	9.0 ± 1.4	3.0 ± 0.0
C3.7S	zs	-4.8 ± 5.5	-2.4 ± 9.2	0.4 ± 8.2	-0.8 ± 5.5	-12.4 ± 11.9	10.2 ± 1.3	12.0 ± 6.0
	cot	4.2 ± 6.4	1.4 ± 2.5	-3.4 ± 5.9	-0.2 ± 3.5	18.4 ± 7.0	13.8 ± 5.0	12.0 ± 12.4
	spp	0.6 ± 5.9	5.6 ± 5.5	-2.2 ± 10.0	-2.4 ± 5.8	18.8 ± 2.3	10.6 ± 1.3	17.8 ± 3.3
	sc-zs	-6.0 ± 1.4	-2.5 ± 9.2	0.0 ± 5.7	3.0 ± 1.4	-8.0 ± 22.6	7.0 ± 2.8	7.5 ± 2.1
	sc-cot	-1.0 ± 1.4	3.5 ± 0.7	7.0 ± 4.2	3.0 ± 2.8	18.5 ± 7.8	10.0 ± 2.8	13.5 ± 12.0
	sc-spp	6.0 ± 7.1	5.0 ± 21.2	4.0 ± 0.0	-3.5 ± 9.2	15.0 ± 0.0	11.5 ± 0.7	16.5 ± 2.1
C3.7S(T)	zs	-1.6 ± 5.5	-2.6 ± 8.0	-0.4 ± 3.5	1.0 ± 2.3	2.2 ± 14.8	6.4 ± 1.9	4.6 ± 3.4
	cot	0.8 ± 3.6	0.0 ± 4.7	-4.0 ± 2.2	-1.6 ± 4.6	8.8 ± 17.3	9.2 ± 2.2	13.2 ± 6.1
	spp	4.2 ± 8.5	-0.2 ± 7.7	-0.8 ± 6.9	0.6 ± 3.4	18.8 ± 6.1	9.4 ± 1.8	17.4 ± 5.7
	sc-zs	-8.5 ± 7.8	-11.5 ± 2.1	1.5 ± 13.4	0.5 ± 3.5	18.0 ± 8.5	10.0 ± 1.4	20.5 ± 2.1
	sc-cot	1.5 ± 2.1	-3.0 ± 4.2	1.5 ± 3.5	3.5 ± 0.7	14.0 ± 14.1	7.5 ± 2.1	24.0 ± 0.0
	sc-spp	12.5 ± 3.5	7.0 ± 2.8	-4.5 ± 9.2	-4.0 ± 1.4	21.5 ± 3.5	9.0 ± 1.4	10.0 ± 12.7
C4S	zs	-2.6 ± 5.6	-1.0 ± 2.1	-2.8 ± 4.8	-0.6 ± 4.2	-2.6 ± 17.9	7.4 ± 1.9	3.8 ± 9.1
	cot	1.2 ± 2.7	-2.6 ± 5.9	-4.6 ± 5.9	-1.8 ± 4.6	15.8 ± 5.4	9.0 ± 0.7	8.6 ± 4.2
	spp	1.0 ± 11.1	4.4 ± 8.4	2.0 ± 5.8	-2.4 ± 4.4	16.4 ± 2.6	1.2 ± 12.8	-1.4 ± 6.8
	sc-zs	-0.5 ± 3.5	-7.5 ± 0.7	-0.5 ± 2.1	0.0 ± 2.8	17.0 ± 1.4	7.5 ± 0.7	-2.5 ± 4.9
	sc-cot	1.0 ± 1.4	3.5 ± 7.8	-3.0 ± 2.8	1.0 ± 11.3	20.0 ± 4.2	10.0 ± 0.0	3.5 ± 3.5
	sc-spp	1.0 ± 5.7	-2.0 ± 0.0	-2.5 ± 2.1	1.0 ± 2.8	17.5 ± 3.5	9.0 ± 1.4	2.5 ± 0.7
C4S(T)	zs	-3.2 ± 1.3	0.6 ± 2.9	-1.2 ± 3.1	-1.6 ± 3.4	-7.2 ± 14.5	5.6 ± 3.1	3.2 ± 3.3
	cot	0.0 ± 3.2	0.8 ± 3.3	2.4 ± 4.3	1.6 ± 4.3	12.8 ± 8.2	9.4 ± 1.9	4.6 ± 9.8
	spp	0.0 ± 7.5	0.8 ± 6.4	-1.2 ± 2.4	-5.6 ± 4.7	13.0 ± 8.0	8.4 ± 2.7	10.0 ± 13.5
	sc-zs	1.0 ± 2.8	-6.0 ± 1.4	-9.0 ± 8.5	-2.0 ± 2.8	-11.5 ± 6.4	6.5 ± 3.5	-4.0 ± 0.0
	sc-cot	6.5 ± 3.5	-3.0 ± 4.2	0.0 ± 2.8	-2.0 ± 2.8	22.5 ± 2.1	11.0 ± 0.0	11.0 ± 9.9
	sc-spp	4.5 ± 3.5	-4.0 ± 8.5	4.0 ± 7.1	1.5 ± 6.4	13.5 ± 4.9	9.0 ± 1.4	6.5 ± 20.5
DS-R1	zs	2.0 ± 3.5	0.4 ± 2.7	2.6 ± 3.2	1.2 ± 2.2	-7.4 ± 2.5	6.2 ± 2.0	4.8 ± 2.3
	cot	0.8 ± 3.8	2.8 ± 5.1	0.2 ± 4.8	1.0 ± 4.8	-4.0 ± 4.6	6.4 ± 2.1	5.0 ± 4.6
	spp	-3.0 ± 2.5	0.2 ± 1.9	-2.4 ± 5.3	0.0 ± 2.8	-3.8 ± 6.2	4.6 ± 1.7	3.2 ± 1.3
	sc-zs	-1.5 ± 2.1	-0.5 ± 3.5	2.5 ± 2.1	-6.0 ± 2.8	-6.5 ± 0.7	6.0 ± 0.0	2.5 ± 0.7
	sc-cot	-3.0 ± 1.4	1.0 ± 0.0	4.5 ± 3.5	-2.5 ± 2.1	-0.5 ± 4.9	9.0 ± 0.0	6.5 ± 0.7
	sc-spp	2.5 ± 2.1	3.0 ± 4.2	1.0 ± 1.4	0.0 ± 2.8	-8.5 ± 0.7	6.5 ± 0.7	6.0 ± 2.8
L3.3-70B	zs	-1.4 ± 2.8	-0.6 ± 1.7	0.6 ± 1.8	0.4 ± 1.7	-7.4 ± 9.4	7.0 ± 1.6	4.8 ± 4.5
	cot	-0.6 ± 1.8	0.8 ± 3.2	-0.2 ± 1.6	0.0 ± 1.6	-8.2 ± 2.5	7.4 ± 1.7	3.2 ± 2.9
	spp	0.0 ± 1.9	3.2 ± 4.5	2.6 ± 4.0	1.0 ± 2.0	-10.4 ± 8.9	5.0 ± 4.0	5.4 ± 4.3
	sc-zs	0.0 ± 0.0	-2.5 ± 0.7	4.0 ± 2.8	3.5 ± 3.5	-14.0 ± 14.1	9.5 ± 2.1	10.5 ± 2.1
	sc-cot	0.0 ± 5.7	1.5 ± 6.4	-1.5 ± 3.5	-3.0 ± 1.4	-14.0 ± 14.1	4.0 ± 1.4	1.0 ± 2.8
	sc-spp	0.0 ± 0.0	-1.5 ± 2.1	1.0 ± 2.8	1.0 ± 2.8	-6.0 ± 2.8	8.5 ± 0.7	4.0 ± 1.4
Mistral	zs	-0.2 ± 4.2	-3.6 ± 5.2	2.4 ± 5.2	-0.2 ± 3.0	-12.4 ± 9.5	0.0 ± 2.4	8.4 ± 1.3
	cot	-3.2 ± 5.4	2.4 ± 5.5	0.4 ± 5.5	-2.4 ± 4.8	-6.2 ± 7.1	3.6 ± 5.9	5.0 ± 3.7
	spp	1.4 ± 8.0	-3.0 ± 5.0	0.8 ± 7.2	2.2 ± 4.3	-4.8 ± 5.4	9.0 ± 10.7	5.4 ± 6.7
	sc-zs	0.0 ± 0.0	-0.5 ± 0.7	-5.0 ± 12.7	-0.5 ± 4.9	-11.5 ± 17.7	-0.5 ± 0.7	9.0 ± 4.2
	sc-cot	-3.0 ± 1.4	2.5 ± 3.5	8.0 ± 2.8	0.0 ± 1.4	-14.5 ± 13.4	12.5 ± 13.4	11.0 ± 0.0
	sc-spp	-0.5 ± 0.7	0.5 ± 2.1	4.5 ± 4.9	3.0 ± 1.4	-22.0 ± 2.8	7.5 ± 17.7	19.0 ± 4.2

Table 16: Total Points Averaged Over All Iterations in the label-based RPS counterfactual.

and Claude 4 Thinking are more mixed: they can still do well, especially with SC or against PP/AP, but are less uniformly robust. DeepSeek R1 is relatively stable but weaker, usually getting modest gains mainly against AP. Llama 3.3 degrades noticeably under the label shift, especially against PP, which supports the view that it is less robust in harder RPS settings. Mistral remains the most erratic: it sometimes achieves strong scores, but it also shows large negative drops and very high

variability.

Finally, prompting effects are mixed but self-consistency often helps, especially for stronger Claude variants, by pushing scores upward on exploitable opponents such as TFT and PP. Still, SC is not a universal fix: under label remapping, some runs remain strongly negative, meaning that the counterfactual exposes a genuine robustness problem rather than just sampling noise. Overall, this table provides clear evidence that label-only per-

turbations in RPS are enough to break strategic transfer for many models.

Table 17 exhibits results regarding payoff-based RPS counterfactuals. Changing rewards proves to be more disruptive than the label-only shift, because strong play now requires abandoning uniform randomization and biasing toward the amplified interaction. This creates larger spreads in total points, more variance across prompts, and clearer separation between models that adapt to incentives and those that remain closer to default RPS behavior.

The most consistent pattern is that PP and TFT become the most exploitable opponents. Many of the best scores in the table appear in these two columns, often reaching the mid-30s or even around 40, which is much higher than in default RPS. This suggests that once an LLM both detects the opponent pattern and aligns it with the altered payoff asymmetry, it can extract substantially more reward than in the canonical game. In contrast, SREP remains near zero for most rows, as expected, since an equilibrium-style opponent leaves little room for systematic gain even under modified payoffs. AP usually yields moderate positive returns, but generally not as high as PP/TFT.

Model-wise, Claude 3.7, Claude 3.7 Thinking, and the Claude 4 variants remain among the strongest overall under this counterfactual, especially against PP and TFT. They repeatedly achieve the highest or near-highest totals, indicating relatively good adaptation to the new payoff landscape. Claude 3.5 is more uneven but still capable of strong gains, especially under SC, where TFT scores become very large. DeepSeek R1 is mixed: it sometimes performs strongly, particularly against TFT and under SC, but remains less stable across prompts. Llama 3.3 is the least consistent among the stronger models: it has some high peaks, but also several near-zero or highly variable results, suggesting weaker payoff sensitivity. Mistral remains the most erratic, with occasional strong positive runs but also many unstable or negative ones.

Prompting effects are again mixed, but here self-consistency often helps more clearly than in the default game. For several models, SC substantially boosts scores against PP/TFT, which fits the intuition that majority voting can stabilize a biased strategy once the model has partially inferred the new incentives. At the same time, some prompt/model combinations still collapse into strongly negative totals, showing that payoff changes expose not just variance but genuine failures to recompute the opti-

mal mixed strategy.

All in all, presented results support the main claim of the paper: payoff-based counterfactuals are a stronger test of strategic reasoning than default play. Many models that look competent in standard RPS do not transfer cleanly when the reward structure changes, while the strongest Claude-family models show the best, though still imperfect, adaptation

Finally, joint counterfactual results are presented in Table 18. This proves to be the hardest RPS setting overall, since models must handle both the label permutation and the payoff asymmetry at the same time. As a result, the table shows the strongest instability: many model-prompt combinations that were competitive in default RPS or in the single-intervention counterfactuals now collapse to near-zero or strongly negative scores, while only a subset of stronger models still manage to exploit structured opponents reliably.

The clearest pattern is that LLM-LLM and SREP columns remain mostly near zero or highly variable, which is expected because the joint counterfactual makes stable equilibrium adaptation harder. The real separation appears against structured algorithmic opponents, especially PP and TFT, and to a lesser extent AP. For the strongest models, these are still the main sources of positive reward, but performance is much less uniform than in the payoff-only setup. In particular, Claude 3.7 Sonnet, Claude 3.7 Thinking, and parts of the Claude 4 family remain the strongest overall: they still achieve high totals against PP and TFT, often in the 20s or low 30s, showing partial adaptation to the altered dominance and incentive structure. Among them, Claude 3.7 appears the most robust on average, with repeated strong scores across several prompts and opponents.

By contrast, DeepSeek R1, Llama 3.3, and especially Mistral are much more brittle here. DeepSeek tends to obtain its best results mainly against AP, while frequently performing poorly or negatively against PP, suggesting that it does not consistently internalize the jointly modified structure. Llama 3.3 is even more telling: compared to the payoff-only counterfactual, it drops substantially, with many negative scores against PP and only modest gains against AP/TFT, indicating weak robustness to the combined intervention. Mistral remains the most unstable model overall, with very large swings and some extreme failures, such as highly negative PP scores under SC, confirming

		Payoff-based RPS counterfactual						
model	prompt	zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	3.8 ± 16.1	-1.2 ± 25.5	-5.0 ± 12.8	0.2 ± 5.6	-4.0 ± 17.2	9.2 ± 6.2	29.6 ± 10.7
	cot	-3.0 ± 21.6	11.6 ± 23.1	7.0 ± 15.1	-0.4 ± 6.3	-7.0 ± 5.1	6.8 ± 2.8	24.2 ± 14.0
	spp	20.2 ± 11.2	12.4 ± 11.8	7.4 ± 8.5	3.4 ± 3.9	1.6 ± 17.7	6.4 ± 2.1	13.2 ± 12.8
	sc-zs	-14.5 ± 21.9	3.5 ± 3.5	17.0 ± 5.7	2.0 ± 9.9	35.0 ± 0.0	18.0 ± 1.4	40.0 ± 0.0
	sc-cot	-0.5 ± 17.7	8.0 ± 15.6	-18.5 ± 26.2	0.0 ± 0.0	16.0 ± 26.9	18.5 ± 9.2	32.5 ± 4.9
	sc-spp	4.5 ± 9.2	-24.5 ± 2.1	33.0 ± 8.5	-3.0 ± 0.0	34.5 ± 0.7	16.0 ± 18.4	35.5 ± 0.7
C3.7S	zs	2.6 ± 20.1	-3.0 ± 13.8	-1.6 ± 11.1	-3.8 ± 6.4	35.0 ± 5.3	13.6 ± 12.3	31.4 ± 7.3
	cot	12.2 ± 8.6	-9.2 ± 8.1	-4.4 ± 11.5	0.0 ± 2.9	34.6 ± 0.9	16.6 ± 14.9	27.6 ± 11.9
	spp	-6.8 ± 18.5	-4.4 ± 11.8	-5.2 ± 12.7	-3.8 ± 4.6	31.8 ± 4.1	10.6 ± 8.7	34.6 ± 2.9
	sc-zs	10.5 ± 12.0	-34.0 ± 7.1	-6.5 ± 13.4	9.0 ± 8.5	11.0 ± 2.8	9.5 ± 7.8	37.0 ± 4.2
	sc-cot	12.0 ± 2.8	-12.0 ± 15.6	0.5 ± 6.4	-1.0 ± 9.9	35.0 ± 0.0	9.0 ± 1.4	38.0 ± 2.8
	sc-spp	7.5 ± 33.2	-5.5 ± 6.4	1.5 ± 4.9	2.0 ± 7.1	34.0 ± 1.4	31.5 ± 0.7	36.5 ± 0.7
C3.7S(T)	zs	2.4 ± 7.6	7.2 ± 9.0	-3.4 ± 8.3	4.6 ± 5.9	28.0 ± 4.6	10.6 ± 0.9	10.8 ± 3.0
	cot	7.0 ± 7.8	-3.2 ± 12.9	-4.4 ± 3.4	-0.6 ± 3.6	34.8 ± 0.4	17.2 ± 12.3	30.0 ± 7.4
	spp	-1.2 ± 14.7	-5.6 ± 12.9	-6.6 ± 14.1	2.8 ± 3.0	34.6 ± 0.5	28.8 ± 13.4	31.4 ± 5.2
	sc-zs	-8.5 ± 4.9	-4.0 ± 7.1	-6.5 ± 7.8	2.5 ± 4.9	35.0 ± 0.0	13.0 ± 2.8	33.0 ± 9.9
	sc-cot	3.5 ± 0.7	7.5 ± 9.2	7.5 ± 0.7	-1.5 ± 2.1	35.0 ± 0.0	9.5 ± 6.4	34.5 ± 2.1
	sc-spp	-7.0 ± 1.4	-18.0 ± 1.4	5.0 ± 5.7	4.0 ± 7.1	34.5 ± 0.7	21.5 ± 14.8	37.0 ± 0.0
C4S	zs	-3.0 ± 12.7	-1.8 ± 10.5	-4.0 ± 19.1	-1.8 ± 7.4	34.4 ± 0.5	11.8 ± 4.0	23.0 ± 14.5
	cot	-2.8 ± 20.2	-9.0 ± 17.4	6.8 ± 7.2	4.0 ± 5.5	31.8 ± 4.4	17.4 ± 11.1	23.2 ± 7.8
	spp	13.8 ± 7.6	3.6 ± 10.5	4.8 ± 6.5	3.0 ± 8.0	28.2 ± 13.5	13.0 ± 5.4	20.4 ± 15.5
	sc-zs	-13.5 ± 27.6	1.0 ± 7.1	-10.0 ± 2.8	1.0 ± 14.1	35.0 ± 0.0	16.5 ± 4.9	11.0 ± 5.7
	sc-cot	7.0 ± 19.8	-3.5 ± 0.7	-7.0 ± 9.9	-1.0 ± 2.8	36.0 ± 2.8	17.5 ± 14.8	20.0 ± 28.3
	sc-spp	-3.0 ± 22.6	0.0 ± 5.7	3.5 ± 44.5	4.5 ± 2.1	35.0 ± 0.0	12.5 ± 6.4	0.0 ± 0.0
C4S(T)	zs	-8.8 ± 12.6	-9.2 ± 4.5	-6.0 ± 11.0	0.6 ± 2.2	30.6 ± 6.3	13.8 ± 3.2	12.0 ± 16.7
	cot	4.0 ± 10.1	6.4 ± 7.6	1.6 ± 10.7	-1.6 ± 5.3	33.6 ± 2.6	9.0 ± 7.2	33.8 ± 4.5
	spp	2.8 ± 15.9	3.4 ± 12.0	-0.6 ± 21.2	-3.2 ± 4.0	34.0 ± 1.4	12.0 ± 1.6	28.8 ± 9.2
	sc-zs	-5.0 ± 9.9	-9.5 ± 3.5	-10.0 ± 29.7	-5.5 ± 0.7	32.5 ± 3.5	17.0 ± 5.7	28.0 ± 12.7
	sc-cot	-5.5 ± 20.5	3.0 ± 1.4	-26.5 ± 12.0	-4.0 ± 14.1	27.5 ± 3.5	12.0 ± 4.2	37.0 ± 1.4
	sc-spp	10.5 ± 4.9	-4.5 ± 12.0	3.5 ± 2.1	2.0 ± 7.1	36.0 ± 1.4	10.5 ± 4.9	25.0 ± 21.2
DS-R1	zs	-10.4 ± 15.5	1.6 ± 6.9	9.8 ± 10.1	1.0 ± 6.4	13.8 ± 12.8	8.4 ± 4.9	29.4 ± 4.2
	cot	9.0 ± 8.4	-8.6 ± 14.6	-0.4 ± 10.6	1.2 ± 7.8	14.8 ± 5.4	11.4 ± 6.3	20.2 ± 3.8
	spp	-7.8 ± 9.4	3.0 ± 12.8	2.0 ± 6.8	-1.2 ± 9.2	12.4 ± 10.6	9.8 ± 3.8	17.0 ± 12.9
	sc-zs	4.5 ± 10.6	11.5 ± 9.2	5.5 ± 4.9	-2.0 ± 8.5	22.0 ± 11.3	18.5 ± 6.4	36.5 ± 0.7
	sc-cot	8.5 ± 0.7	-6.5 ± 19.1	-5.5 ± 4.9	9.0 ± 8.5	18.0 ± 1.4	11.5 ± 0.7	24.5 ± 3.5
	sc-spp	-14.5 ± 12.0	-8.5 ± 4.9	-2.0 ± 8.5	4.5 ± 3.5	25.0 ± 8.5	19.0 ± 4.2	23.5 ± 2.1
L3.3-70B	zs	0.4 ± 17.6	-3.8 ± 16.1	12.2 ± 15.0	2.0 ± 5.3	15.2 ± 14.5	11.6 ± 8.2	1.6 ± 3.4
	cot	0.8 ± 7.3	0.4 ± 14.2	2.6 ± 8.5	0.2 ± 1.9	37.0 ± 6.7	17.0 ± 0.7	1.6 ± 1.1
	spp	-0.2 ± 1.8	-0.4 ± 10.3	2.6 ± 4.7	0.6 ± 7.1	9.2 ± 5.6	10.4 ± 5.9	13.2 ± 15.7
	sc-zs	-10.0 ± 24.0	13.0 ± 18.4	-1.0 ± 1.4	3.0 ± 7.1	24.5 ± 21.9	17.5 ± 12.0	-1.0 ± 1.4
	sc-cot	-17.0 ± 25.5	-0.5 ± 3.5	1.5 ± 2.1	2.5 ± 6.4	20.5 ± 27.6	13.5 ± 0.7	1.0 ± 1.4
	sc-spp	27.0 ± 2.8	3.0 ± 2.8	2.5 ± 0.7	-4.5 ± 9.2	2.5 ± 2.1	13.5 ± 16.3	16.5 ± 26.2
Mistral	zs	-0.2 ± 2.8	-7.8 ± 10.2	-5.4 ± 4.8	-1.0 ± 2.7	34.8 ± 11.6	1.8 ± 13.0	1.8 ± 1.6
	cot	15.6 ± 11.5	-3.8 ± 13.7	-1.6 ± 20.9	-0.6 ± 4.2	8.6 ± 7.1	6.8 ± 7.4	14.0 ± 16.5
	spp	9.0 ± 19.0	-13.0 ± 19.1	2.6 ± 13.7	2.2 ± 6.4	-4.8 ± 7.7	-3.0 ± 5.1	27.8 ± 10.4
	sc-zs	0.0 ± 0.0	-18.5 ± 17.7	-9.0 ± 5.7	4.5 ± 4.9	20.5 ± 20.5	2.5 ± 0.7	0.5 ± 0.7
	sc-cot	11.0 ± 15.6	17.5 ± 31.8	3.0 ± 15.6	0.5 ± 2.1	15.5 ± 21.9	13.0 ± 2.8	22.0 ± 25.5
	sc-spp	5.0 ± 11.3	-7.5 ± 23.3	-8.5 ± 30.4	3.5 ± 7.8	11.0 ± 17.0	-1.0 ± 21.2	35.5 ± 3.5

Table 17: Total Points Averaged Over All Iterations for the payoff-based RPS counterfactual.

that the joint counterfactual is particularly damaging for weaker or less consistent models.

Prompting remains mixed, but the joint setting again suggests that self-consistency can help only when the base reasoning is already somewhat aligned with the new structure. For stronger Claude models, SC sometimes lifts PP/TFT performance further, but for weaker models it often amplifies bad policies instead of correcting them. This is especially visible in rows with consistently negative

PP scores despite SC.

Overall, this table provides some of the strongest evidence for your paper’s core claim: strategic competence in default RPS does not reliably transfer when both labels and payoffs are changed simultaneously. The joint counterfactual exposes the sharpest reasoning brittleness, with only the strongest Claude-family models retaining substantial exploitation ability, while weaker models and some prompt variants fail to recompute the altered

model	prompt	Joint RPS counterfactual						
		zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	0.6 ± 6.0	2.2 ± 10.3	-1.0 ± 10.9	-3.4 ± 4.7	-8.6 ± 22.2	10.6 ± 5.0	11.8 ± 12.1
	cot	2.0 ± 6.8	-5.8 ± 15.9	-4.2 ± 13.0	-0.2 ± 5.4	-8.8 ± 16.5	6.2 ± 11.4	12.2 ± 7.4
	spp	-1.4 ± 11.0	5.0 ± 10.2	2.6 ± 9.8	3.6 ± 3.9	-25.8 ± 10.0	14.0 ± 4.8	11.8 ± 7.7
	sc-zs	-7.0 ± 8.5	-7.5 ± 4.9	4.5 ± 0.7	0.5 ± 6.4	-17.0 ± 19.8	8.5 ± 4.9	7.0 ± 5.7
	sc-cot	8.0 ± 5.7	0.5 ± 9.2	2.0 ± 9.9	-2.5 ± 7.8	-20.5 ± 20.5	14.0 ± 7.1	6.5 ± 2.1
	sc-spp	5.0 ± 4.2	-0.5 ± 7.8	2.5 ± 2.1	-4.5 ± 9.2	-14.0 ± 5.7	-11.0 ± 7.1	18.0 ± 5.7
C3.7S	zs	1.6 ± 11.9	-2.6 ± 6.9	-10.2 ± 11.0	-1.2 ± 6.5	-13.4 ± 15.8	12.0 ± 7.0	16.6 ± 16.6
	cot	1.2 ± 9.9	2.4 ± 9.4	-1.6 ± 5.3	4.2 ± 6.4	33.4 ± 10.0	12.4 ± 1.1	20.0 ± 15.4
	spp	6.4 ± 14.7	1.6 ± 11.7	3.0 ± 8.9	-2.6 ± 6.1	27.4 ± 18.1	18.0 ± 10.1	29.0 ± 11.9
	sc-zs	10.0 ± 28.3	-3.0 ± 0.0	-19.5 ± 4.9	0.5 ± 4.9	16.0 ± 2.8	24.0 ± 21.2	31.0 ± 8.5
	sc-cot	-2.0 ± 2.8	-12.0 ± 9.9	-13.0 ± 2.8	1.5 ± 7.8	31.0 ± 2.8	19.5 ± 0.7	22.5 ± 24.7
	sc-spp	5.5 ± 14.8	3.5 ± 7.8	2.5 ± 3.5	2.0 ± 8.5	33.0 ± 2.8	11.0 ± 4.2	24.5 ± 14.8
C3.7S(T)	zs	4.0 ± 7.0	-6.2 ± 12.2	0.8 ± 14.7	7.6 ± 6.1	-8.6 ± 13.7	13.6 ± 4.8	15.6 ± 12.0
	cot	6.8 ± 7.3	-1.2 ± 11.5	-1.0 ± 6.0	-0.2 ± 4.3	18.4 ± 20.2	15.0 ± 4.4	16.6 ± 18.5
	spp	6.8 ± 5.6	-0.8 ± 10.8	7.8 ± 13.6	-1.6 ± 4.6	25.8 ± 15.1	14.0 ± 4.6	19.6 ± 13.6
	sc-zs	-1.5 ± 7.8	-1.0 ± 12.7	-14.0 ± 11.3	4.0 ± 1.4	-14.0 ± 8.5	8.0 ± 0.0	9.0 ± 5.7
	sc-cot	2.0 ± 7.1	-13.0 ± 0.0	3.5 ± 6.4	-2.0 ± 1.4	35.0 ± 7.1	16.5 ± 2.1	30.5 ± 13.4
	sc-spp	10.0 ± 14.1	5.0 ± 26.9	-3.0 ± 1.4	2.5 ± 6.4	33.0 ± 9.9	16.5 ± 0.7	28.5 ± 0.7
C4S	zs	-2.6 ± 12.7	-11.2 ± 8.9	-0.6 ± 4.9	2.2 ± 7.9	7.4 ± 27.3	15.6 ± 2.3	10.2 ± 21.1
	cot	7.0 ± 2.5	-2.8 ± 13.1	-4.8 ± 12.6	4.4 ± 9.2	35.4 ± 4.4	8.0 ± 11.9	6.8 ± 4.4
	spp	2.8 ± 10.0	-6.8 ± 8.2	2.4 ± 11.9	-3.4 ± 12.8	24.4 ± 13.8	15.6 ± 2.1	7.6 ± 12.9
	sc-zs	5.0 ± 7.1	-6.0 ± 5.7	-7.0 ± 4.2	4.0 ± 9.9	15.0 ± 14.1	13.5 ± 4.9	-1.5 ± 6.4
	sc-cot	8.0 ± 19.8	-8.5 ± 7.8	0.0 ± 5.7	4.5 ± 12.0	36.0 ± 5.7	8.5 ± 0.7	3.0 ± 1.4
	sc-spp	-6.5 ± 10.6	5.5 ± 4.9	-0.5 ± 10.6	-8.0 ± 14.1	27.5 ± 6.4	11.0 ± 1.4	33.5 ± 9.2
C4S(T)	zs	-3.2 ± 6.1	-8.6 ± 17.8	4.4 ± 15.1	5.8 ± 4.6	-16.2 ± 11.1	12.0 ± 2.0	8.8 ± 16.5
	cot	1.6 ± 5.1	2.0 ± 13.8	-1.2 ± 14.9	5.4 ± 8.6	9.8 ± 27.8	12.0 ± 4.5	21.6 ± 11.3
	spp	0.4 ± 8.7	9.2 ± 9.7	2.0 ± 4.1	-2.4 ± 7.8	32.0 ± 4.9	7.6 ± 7.7	8.8 ± 10.7
	sc-zs	-3.0 ± 4.2	-5.5 ± 3.5	13.0 ± 4.2	-3.5 ± 12.0	-0.5 ± 29.0	16.0 ± 5.7	-12.5 ± 2.1
	sc-cot	-3.5 ± 16.3	-9.5 ± 9.2	-16.0 ± 7.1	-1.0 ± 9.9	-2.5 ± 0.7	15.0 ± 1.4	8.0 ± 7.1
	sc-spp	10.5 ± 0.7	-1.5 ± 10.6	-3.0 ± 1.4	-0.5 ± 7.8	23.5 ± 0.7	12.5 ± 3.5	14.5 ± 10.6
DS-R1	zs	2.4 ± 7.9	-3.4 ± 8.7	0.2 ± 8.8	-4.0 ± 8.9	-12.6 ± 5.8	12.4 ± 3.8	-1.4 ± 7.3
	cot	-4.8 ± 11.6	-0.2 ± 0.8	-4.2 ± 4.3	-3.2 ± 2.8	-7.2 ± 6.9	8.0 ± 5.7	8.2 ± 4.3
	spp	-4.8 ± 6.9	-2.8 ± 4.8	-1.8 ± 7.1	1.4 ± 4.5	-14.8 ± 5.3	9.4 ± 7.7	10.6 ± 9.2
	sc-zs	-0.5 ± 10.6	2.5 ± 3.5	-1.5 ± 3.5	5.5 ± 2.1	-14.0 ± 0.0	13.5 ± 0.7	-0.5 ± 0.7
	sc-cot	6.5 ± 2.1	4.0 ± 4.2	-6.5 ± 2.1	6.5 ± 7.8	-27.5 ± 0.7	15.5 ± 0.7	2.0 ± 5.7
	sc-spp	3.5 ± 2.1	2.5 ± 3.5	0.0 ± 1.4	0.0 ± 8.5	-16.5 ± 9.2	10.5 ± 3.5	4.0 ± 4.2
L3.3-70B	zs	0.0 ± 3.9	0.0 ± 5.5	-1.2 ± 2.9	3.2 ± 3.3	-16.2 ± 16.5	5.4 ± 5.9	7.8 ± 2.9
	cot	-4.8 ± 2.9	2.2 ± 2.7	-0.6 ± 5.6	-0.4 ± 3.2	-8.8 ± 8.6	7.2 ± 4.4	5.2 ± 5.5
	spp	-6.6 ± 10.4	1.0 ± 3.3	4.4 ± 3.4	-4.2 ± 7.2	-11.8 ± 9.7	11.0 ± 3.7	3.8 ± 2.8
	sc-zs	0.5 ± 0.7	5.0 ± 1.4	2.5 ± 0.7	-1.5 ± 4.9	-20.5 ± 27.6	6.5 ± 9.2	6.5 ± 9.2
	sc-cot	-4.0 ± 0.0	4.0 ± 2.8	-0.5 ± 3.5	2.5 ± 2.1	-4.5 ± 4.9	5.0 ± 5.7	3.5 ± 2.1
	sc-spp	-6.0 ± 1.4	-0.5 ± 3.5	1.5 ± 7.8	-10.0 ± 1.4	-30.5 ± 13.4	11.5 ± 0.7	0.5 ± 0.7
Mistral	zs	0.0 ± 6.4	1.2 ± 8.0	-0.4 ± 10.5	0.0 ± 6.5	-11.6 ± 17.6	-16.4 ± 13.8	10.2 ± 9.7
	cot	2.0 ± 12.9	-0.2 ± 6.9	6.8 ± 8.5	-0.6 ± 8.1	-14.6 ± 12.5	-11.2 ± 14.9	21.8 ± 8.9
	spp	0.4 ± 10.7	2.8 ± 8.9	0.0 ± 15.0	4.2 ± 5.3	17.4 ± 18.4	13.8 ± 13.5	14.6 ± 8.4
	sc-zs	1.0 ± 26.9	2.0 ± 1.4	1.0 ± 5.7	2.5 ± 4.9	-40.0 ± 0.0	-4.5 ± 20.5	8.5 ± 14.8
	sc-cot	-3.5 ± 3.5	4.5 ± 4.9	9.5 ± 14.8	6.5 ± 3.5	-16.0 ± 14.1	-6.0 ± 21.2	19.5 ± 21.9
	sc-spp	9.5 ± 33.2	-1.0 ± 11.3	-4.5 ± 21.9	1.0 ± 2.8	-32.0 ± 11.3	8.5 ± 16.3	29.0 ± 8.5

Table 18: Total Points Averaged Over All Iterations for the joint RPS counterfactual.

game structure.

D.7 RPS Opponent comprehension

Table 19 presents opponent comprehension results for the default RPS game. opponent comprehension is markedly weaker and more variable than in PD, which is expected given the absence of a dominant response and the need to reason over three actions rather than two. In LLM-LLM interactions (first three columns), most values remain

relatively high and often close to the horizon, indicating that models generally do not infer and exploit other LLMs’ play quickly; for instance, Claude 3.5 Sonnet v2 under zero-shot reaches only $m = 10.6 \pm 13.1$ against its ZS counterpart and rises to 21.4 ± 4.6 and 19.6 ± 5.6 against SPP and CoT counterparts, respectively. By contrast, algorithmic opponents are more informative: **PP** and especially **TFT** are the most exploitable, with Claude 3.7, Claude 3.7(T), Claude 4, Claude 4(T),

model	prompt	RPS						
		zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	10.6 ± 13.1	21.4 ± 4.6	19.6 ± 5.6	21.0 ± 3.5	14.6 ± 12.4	10.4 ± 11.2	1.0 ± 0.0
	cot	17.2 ± 7.3	19.6 ± 7.5	20.6 ± 5.6	19.6 ± 7.1	16.0 ± 11.9	22.6 ± 2.1	5.8 ± 10.7
	spp	11.2 ± 10.6	11.8 ± 10.9	20.4 ± 5.9	23.0 ± 1.9	11.4 ± 11.3	18.4 ± 9.3	1.6 ± 0.9
	sc-zs	25.0 ± 0.0	14.0 ± 15.6	21.5 ± 4.9	13.5 ± 16.3	12.0 ± 15.6	19.5 ± 6.4	1.0 ± 0.0
	sc-cot	15.0 ± 14.1	20.5 ± 0.7	22.5 ± 3.5	25.0 ± 0.0	12.5 ± 16.3	7.0 ± 8.5	1.0 ± 0.0
	sc-spp	14.5 ± 14.8	15.0 ± 14.1	14.5 ± 14.8	24.0 ± 1.4	16.0 ± 8.5	21.0 ± 0.0	1.0 ± 0.0
C3.7S	zs	21.8 ± 5.5	22.2 ± 4.2	20.2 ± 9.7	23.4 ± 2.5	10.0 ± 12.4	17.6 ± 9.3	15.0 ± 12.8
	cot	15.4 ± 7.1	16.2 ± 10.0	18.8 ± 4.3	18.0 ± 7.4	1.0 ± 0.0	15.6 ± 10.1	1.0 ± 0.0
	spp	24.4 ± 1.3	20.0 ± 5.6	22.4 ± 4.3	23.6 ± 1.5	1.2 ± 0.4	12.0 ± 10.6	1.2 ± 0.4
	sc-zs	13.0 ± 17.0	23.0 ± 0.0	24.5 ± 0.7	22.0 ± 2.8	11.5 ± 14.8	15.0 ± 14.1	12.5 ± 16.3
	sc-cot	4.5 ± 4.9	24.5 ± 0.7	11.5 ± 14.8	24.0 ± 1.4	1.0 ± 0.0	24.5 ± 0.7	1.0 ± 0.0
	sc-spp	9.5 ± 12.0	21.5 ± 3.5	22.0 ± 1.4	12.5 ± 16.3	1.0 ± 0.0	22.5 ± 2.1	1.0 ± 0.0
C3.7S(T)	zs	24.2 ± 1.3	19.4 ± 8.3	22.6 ± 4.8	22.4 ± 2.2	1.4 ± 0.9	22.6 ± 1.7	1.6 ± 0.5
	cot	19.2 ± 5.1	21.8 ± 4.9	22.0 ± 2.7	23.8 ± 2.2	1.0 ± 0.0	15.6 ± 12.0	4.4 ± 7.6
	spp	23.6 ± 1.5	20.6 ± 4.3	21.2 ± 5.3	15.2 ± 11.3	1.0 ± 0.0	17.8 ± 9.7	1.2 ± 0.4
	sc-zs	24.5 ± 0.7	22.5 ± 3.5	13.0 ± 17.0	22.5 ± 3.5	1.0 ± 0.0	12.0 ± 15.6	1.0 ± 0.0
	sc-cot	24.5 ± 0.7	25.0 ± 0.0	25.0 ± 0.0	19.5 ± 7.8	1.0 ± 0.0	10.0 ± 12.7	1.0 ± 0.0
	sc-spp	24.0 ± 1.4	25.0 ± 0.0	25.0 ± 0.0	25.0 ± 0.0	1.0 ± 0.0	23.0 ± 0.0	1.0 ± 0.0
C4S	zs	15.8 ± 8.5	20.0 ± 10.6	23.8 ± 2.2	22.2 ± 0.8	1.0 ± 0.0	23.2 ± 1.1	11.8 ± 11.3
	cot	15.2 ± 13.0	15.6 ± 1.9	10.8 ± 11.8	19.4 ± 8.6	1.0 ± 0.0	19.0 ± 10.1	3.2 ± 4.9
	spp	16.0 ± 10.5	20.8 ± 9.4	18.6 ± 10.0	19.2 ± 5.5	3.4 ± 4.8	14.4 ± 12.3	9.0 ± 8.6
	sc-zs	7.5 ± 9.2	25.0 ± 0.0	25.0 ± 0.0	22.0 ± 4.2	1.0 ± 0.0	23.0 ± 1.4	1.0 ± 0.0
	sc-cot	21.5 ± 4.9	15.0 ± 14.1	14.5 ± 14.8	21.0 ± 1.4	1.0 ± 0.0	24.0 ± 1.4	1.0 ± 0.0
	sc-spp	10.0 ± 7.1	11.0 ± 8.5	23.0 ± 2.8	21.5 ± 0.7	1.0 ± 0.0	13.0 ± 17.0	1.0 ± 0.0
C4S(T)	zs	13.4 ± 12.0	21.0 ± 5.3	17.0 ± 10.0	17.6 ± 8.5	3.4 ± 5.4	22.6 ± 1.8	10.6 ± 13.1
	cot	14.4 ± 12.3	21.8 ± 2.7	20.2 ± 10.7	21.6 ± 4.2	1.0 ± 0.0	19.6 ± 10.4	3.2 ± 4.9
	spp	24.2 ± 1.1	15.2 ± 9.9	17.8 ± 8.6	18.8 ± 10.1	1.0 ± 0.0	19.4 ± 10.4	5.6 ± 6.5
	sc-zs	13.0 ± 17.0	22.0 ± 4.2	24.0 ± 1.4	18.0 ± 5.7	1.0 ± 0.0	24.5 ± 0.7	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	7.0 ± 4.2	10.5 ± 13.4	23.5 ± 0.7	1.0 ± 0.0	11.0 ± 14.1	1.0 ± 0.0
	sc-spp	24.0 ± 1.4	23.0 ± 2.8	24.5 ± 0.7	20.5 ± 2.1	1.0 ± 0.0	23.5 ± 0.7	1.0 ± 0.0
DS-R1	zs	21.2 ± 5.5	21.2 ± 5.3	19.4 ± 10.3	23.4 ± 1.3	23.2 ± 0.4	16.8 ± 11.1	7.8 ± 9.7
	cot	21.2 ± 4.3	20.0 ± 5.5	24.4 ± 0.9	13.6 ± 10.9	12.8 ± 10.3	15.4 ± 12.7	1.0 ± 0.0
	spp	18.0 ± 10.4	22.6 ± 2.3	20.6 ± 8.8	23.0 ± 1.4	20.2 ± 4.1	16.8 ± 8.1	6.2 ± 9.0
	sc-zs	12.0 ± 15.6	21.0 ± 2.8	24.0 ± 1.4	22.5 ± 3.5	13.5 ± 16.3	22.5 ± 2.1	1.5 ± 0.7
	sc-cot	12.0 ± 15.6	23.0 ± 1.4	19.0 ± 7.1	21.5 ± 4.9	3.5 ± 2.1	1.0 ± 0.0	1.0 ± 0.0
	sc-spp	16.0 ± 8.5	24.0 ± 1.4	20.5 ± 3.5	23.0 ± 0.0	13.5 ± 16.3	24.5 ± 0.7	1.0 ± 0.0
L3.3-70B	zs	5.8 ± 10.7	16.2 ± 10.4	12.4 ± 11.8	21.8 ± 2.3	1.0 ± 0.0	12.8 ± 7.9	1.0 ± 0.0
	cot	5.6 ± 10.3	5.8 ± 10.7	10.0 ± 12.4	23.0 ± 1.9	2.8 ± 4.0	21.2 ± 4.1	1.0 ± 0.0
	spp	1.0 ± 0.0	7.8 ± 10.0	21.2 ± 5.8	21.8 ± 5.5	8.6 ± 11.0	15.2 ± 9.4	1.0 ± 0.0
	sc-zs	1.0 ± 0.0	13.0 ± 17.0	20.0 ± 7.1	25.0 ± 0.0	1.0 ± 0.0	5.5 ± 4.9	1.0 ± 0.0
	sc-cot	10.5 ± 13.4	23.5 ± 2.1	1.0 ± 0.0	25.0 ± 0.0	1.0 ± 0.0	12.5 ± 16.3	1.0 ± 0.0
	sc-spp	1.0 ± 0.0	13.0 ± 17.0	24.5 ± 0.7	16.0 ± 8.5	2.5 ± 2.1	12.0 ± 15.6	1.0 ± 0.0
Mistral	zs	1.2 ± 0.4	19.2 ± 10.3	23.6 ± 1.1	22.6 ± 4.8	1.0 ± 0.0	6.0 ± 10.6	1.0 ± 0.0
	cot	18.8 ± 10.4	20.2 ± 6.1	15.8 ± 10.7	21.2 ± 4.3	12.4 ± 12.1	17.0 ± 9.4	1.0 ± 0.0
	spp	6.6 ± 10.4	10.6 ± 13.1	17.8 ± 9.6	23.0 ± 0.7	17.2 ± 10.3	19.8 ± 10.5	1.0 ± 0.0
	sc-zs	1.0 ± 0.0	13.0 ± 17.0	24.5 ± 0.7	23.0 ± 0.0	1.0 ± 0.0	24.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	24.0 ± 0.0	13.0 ± 17.0	24.5 ± 0.7	12.5 ± 16.3	24.5 ± 0.7	1.0 ± 0.0
	sc-spp	12.5 ± 16.3	11.5 ± 14.8	12.5 ± 16.3	25.0 ± 0.0	11.5 ± 14.8	24.5 ± 0.7	6.5 ± 7.8

Table 19: Round m # where the LLM player understood the opponent’s Strategy in the default RPS setup. Lower is better; $m = 25$ indicates no stable comprehension within the 24-round game.

and Llama 3.3 often reaching very early comprehension ($m \approx 1$) in several prompting settings, showing successful recognition of cyclic or reactive behavior. **AP** is noticeably harder, with comprehension typically delayed across models, while **SREP** also remains difficult for many rows, consistent with the fact that an equilibrium-style opponent offers little exploitable structure. Overall, the strongest and most consistent RPS comprehension

is shown by the advanced Claude models, Llama 3.3 acts as a competitive but less stable runner-up, DeepSeek R1 is mixed, and Mistral Large remains the most inconsistent despite a few isolated low- m outcomes. This confirms that, unlike PD, default RPS rarely leads to early and stable mutual understanding; instead, successful play mainly comes from detecting exploitable algorithmic regularities rather than quickly countering another LLM.

		Label-based RPS counterfactual						
model	prompt	zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	19.0 ± 10.3	22.4 ± 4.2	24.4 ± 1.3	20.4 ± 4.2	21.6 ± 3.8	11.4 ± 12.5	14.8 ± 11.4
	cot	21.8 ± 6.1	19.6 ± 7.3	19.8 ± 6.8	20.8 ± 5.1	21.6 ± 6.5	24.4 ± 0.9	23.6 ± 1.7
	spp	21.8 ± 2.6	24.0 ± 0.7	14.6 ± 8.7	22.6 ± 3.0	18.4 ± 9.0	22.8 ± 2.7	16.4 ± 11.3
	sc-zs	7.5 ± 9.2	22.5 ± 0.7	24.0 ± 1.4	12.0 ± 4.2	25.0 ± 0.0	19.0 ± 5.7	13.0 ± 17.0
	sc-cot	23.0 ± 1.4	11.5 ± 3.5	25.0 ± 0.0	22.5 ± 0.7	25.0 ± 0.0	25.0 ± 0.0	24.0 ± 1.4
	sc-spp	19.0 ± 8.5	12.5 ± 16.3	18.5 ± 9.2	17.0 ± 8.5	25.0 ± 0.0	25.0 ± 0.0	12.0 ± 15.6
C3.7S	zs	23.4 ± 1.8	21.0 ± 5.3	22.6 ± 2.3	19.2 ± 7.2	21.2 ± 5.8	22.8 ± 0.8	7.8 ± 10.0
	cot	19.8 ± 10.0	19.2 ± 10.2	24.0 ± 1.0	23.2 ± 2.0	3.2 ± 3.9	18.4 ± 9.8	1.6 ± 1.3
	spp	18.4 ± 9.5	19.4 ± 8.3	21.6 ± 6.5	22.2 ± 5.2	2.4 ± 1.9	23.6 ± 2.1	1.0 ± 0.0
	sc-zs	25.0 ± 0.0	22.0 ± 2.8	20.0 ± 7.1	19.0 ± 5.7	16.5 ± 12.0	25.0 ± 0.0	23.0 ± 2.8
	sc-cot	25.0 ± 0.0	21.5 ± 0.7	18.5 ± 4.9	25.0 ± 0.0	4.0 ± 4.2	21.5 ± 2.1	13.0 ± 17.0
	sc-spp	12.0 ± 2.8	13.0 ± 17.0	24.0 ± 0.0	25.0 ± 0.0	5.0 ± 0.0	22.5 ± 2.1	1.0 ± 0.0
C3.7S(T)	zs	22.6 ± 4.3	24.6 ± 0.9	23.6 ± 1.3	22.4 ± 2.2	15.0 ± 9.7	17.6 ± 10.0	10.4 ± 12.9
	cot	19.2 ± 4.4	22.4 ± 2.6	22.0 ± 2.6	23.2 ± 1.9	10.6 ± 13.1	21.8 ± 1.8	7.2 ± 9.7
	spp	17.4 ± 10.4	21.6 ± 4.9	23.6 ± 0.5	21.6 ± 3.8	3.4 ± 3.9	24.0 ± 1.0	1.2 ± 0.4
	sc-zs	24.5 ± 0.7	22.0 ± 4.2	12.5 ± 16.3	22.5 ± 3.5	3.0 ± 2.8	23.0 ± 0.0	1.0 ± 0.0
	sc-cot	19.0 ± 8.5	13.0 ± 17.0	25.0 ± 0.0	22.5 ± 0.7	11.5 ± 14.8	19.0 ± 7.1	1.0 ± 0.0
	sc-spp	3.0 ± 1.4	16.0 ± 7.1	24.5 ± 0.7	25.0 ± 0.0	1.0 ± 0.0	24.5 ± 0.7	1.0 ± 0.0
C4S	zs	19.4 ± 10.3	21.6 ± 7.1	24.0 ± 2.2	22.8 ± 3.0	16.2 ± 10.3	21.6 ± 6.5	9.4 ± 9.2
	cot	13.2 ± 10.2	24.2 ± 1.8	23.8 ± 1.3	24.0 ± 1.2	4.8 ± 2.8	18.0 ± 9.6	13.6 ± 6.8
	spp	19.6 ± 10.4	17.0 ± 10.1	17.6 ± 10.9	19.6 ± 4.3	4.2 ± 2.2	21.2 ± 4.0	21.2 ± 5.4
	sc-zs	25.0 ± 0.0	22.5 ± 3.5	24.5 ± 0.7	24.5 ± 0.7	2.5 ± 0.7	25.0 ± 0.0	13.0 ± 17.0
	sc-cot	24.0 ± 0.0	23.5 ± 2.1	24.0 ± 1.4	12.5 ± 16.3	2.5 ± 2.1	25.0 ± 0.0	20.5 ± 3.5
	sc-spp	23.5 ± 0.7	24.0 ± 1.4	23.5 ± 0.7	24.0 ± 1.4	3.0 ± 2.8	15.0 ± 8.5	24.5 ± 0.7
C4S(T)	zs	22.2 ± 4.7	23.8 ± 0.8	24.2 ± 0.8	22.4 ± 4.2	20.6 ± 9.3	21.4 ± 4.5	16.4 ± 9.7
	cot	24.8 ± 0.4	24.6 ± 0.5	22.6 ± 4.3	21.8 ± 4.4	12.2 ± 10.7	21.0 ± 5.7	17.6 ± 10.6
	spp	23.6 ± 1.1	21.8 ± 4.9	22.4 ± 2.1	23.4 ± 0.5	7.8 ± 10.0	22.0 ± 1.6	14.8 ± 12.7
	sc-zs	25.0 ± 0.0	24.5 ± 0.7	25.0 ± 0.0	23.0 ± 1.4	24.5 ± 0.7	22.0 ± 1.4	24.5 ± 0.7
	sc-cot	21.5 ± 4.9	9.0 ± 11.3	24.0 ± 1.4	25.0 ± 0.0	1.0 ± 0.0	22.0 ± 1.4	12.0 ± 15.6
	sc-spp	18.5 ± 7.8	24.0 ± 1.4	13.0 ± 17.0	21.5 ± 3.5	6.5 ± 3.5	23.0 ± 1.4	13.0 ± 17.0
DS-R1	zs	14.4 ± 12.3	23.6 ± 2.1	17.0 ± 6.5	20.6 ± 4.6	23.8 ± 0.8	23.0 ± 1.0	19.8 ± 4.9
	cot	15.8 ± 10.8	23.2 ± 1.3	23.6 ± 2.6	24.0 ± 1.7	24.6 ± 0.9	21.2 ± 3.5	16.8 ± 10.5
	spp	23.2 ± 2.2	21.0 ± 3.7	18.6 ± 8.8	23.2 ± 0.8	24.6 ± 0.9	22.6 ± 1.1	23.4 ± 0.5
	sc-zs	22.5 ± 3.5	16.5 ± 12.0	8.0 ± 9.9	23.0 ± 2.8	25.0 ± 0.0	24.0 ± 0.0	16.5 ± 3.5
	sc-cot	24.0 ± 1.4	22.0 ± 2.8	18.0 ± 9.9	24.5 ± 0.7	25.0 ± 0.0	13.0 ± 0.0	15.0 ± 14.1
	sc-spp	21.5 ± 4.9	19.0 ± 5.7	24.0 ± 1.4	24.0 ± 0.0	24.0 ± 0.0	25.0 ± 0.0	13.0 ± 17.0
L3.3-70B	zs	16.2 ± 9.8	7.4 ± 9.0	16.4 ± 9.6	23.8 ± 1.6	10.4 ± 9.5	18.6 ± 9.9	6.8 ± 9.5
	cot	16.2 ± 9.1	10.2 ± 12.6	12.2 ± 11.1	20.2 ± 5.0	21.0 ± 5.2	19.8 ± 10.0	1.0 ± 0.0
	spp	13.0 ± 11.4	1.4 ± 0.9	10.2 ± 12.6	23.0 ± 1.2	12.2 ± 9.8	18.8 ± 10.1	11.2 ± 10.7
	sc-zs	5.0 ± 5.7	6.5 ± 3.5	12.5 ± 13.4	21.0 ± 1.4	14.0 ± 15.6	23.5 ± 2.1	1.0 ± 0.0
	sc-cot	11.5 ± 14.8	8.5 ± 6.4	13.0 ± 17.0	20.5 ± 4.9	14.0 ± 15.6	24.0 ± 1.4	13.0 ± 17.0
	sc-spp	6.0 ± 7.1	6.0 ± 7.1	9.0 ± 11.3	25.0 ± 0.0	7.0 ± 5.7	1.0 ± 0.0	1.0 ± 0.0
Mistral	zs	14.2 ± 12.2	19.0 ± 10.2	16.0 ± 12.0	23.8 ± 1.3	19.4 ± 9.1	24.2 ± 0.8	1.0 ± 0.0
	cot	20.8 ± 4.2	13.8 ± 9.0	9.0 ± 7.6	23.4 ± 2.1	24.2 ± 0.8	17.2 ± 10.5	3.2 ± 4.9
	spp	18.0 ± 7.0	22.2 ± 4.6	20.4 ± 6.4	22.2 ± 2.9	23.8 ± 1.6	12.8 ± 11.3	5.6 ± 10.3
	sc-zs	1.0 ± 0.0	1.5 ± 0.7	12.0 ± 15.6	21.5 ± 4.9	13.0 ± 17.0	24.5 ± 0.7	7.5 ± 9.2
	sc-cot	17.5 ± 10.6	14.5 ± 13.4	10.0 ± 7.1	18.5 ± 4.9	14.5 ± 14.8	13.0 ± 17.0	12.5 ± 16.3
	sc-spp	12.5 ± 16.3	19.0 ± 7.1	23.5 ± 0.7	19.0 ± 7.1	25.0 ± 0.0	12.5 ± 16.3	1.0 ± 0.0

Table 20: Round m # where the LLM player understood the opponent’s Strategy in the label-based RPS counterfactual

In Table 20 we demonstrate results for the label-based RPS counterfactual. The label-based RPS counterfactual makes opponent comprehension noticeably harder and more uneven than in default RPS, especially in LLM–LLM interactions and against the cyclic **PP** opponent. Since only the dominance induced by action labels is changed, these delays suggest that many models rely partly on canonical RPS label associations rather than

fully recomputing the altered dominance relation. This is most visible for Claude 3.5 Sonnet v2, whose comprehension often moves close to the horizon across almost all opponent types (e.g., $m = 21.6 \pm 3.8$ against **PP** under **ZS**, and repeated $m = 25.0$ under several **SC** settings), indicating substantial difficulty adapting to the re-labeled game.

A clearer separation emerges among stronger

models. Claude 3.7, Claude 3.7(T), Claude 4, and to a lesser extent Claude 4(T), still achieve early comprehension against structured algorithmic opponents in several prompting settings, particularly against **PP** and **TFT**. For example, Claude 3.7 reaches 3.2 ± 3.9 against **PP** and 1.6 ± 1.3 against **TFT** under CoT, while Claude 3.7(T) reaches $m = 3.4 \pm 3.9$ and $m = 1.2 \pm 0.4$ under SPP. Claude 4 is similarly robust against **PP**, often obtaining $m \approx 2.5 - 4.8$, although its comprehension against **TFT** is less consistent. These results suggest that the advanced Claude models retain some ability to infer reactive or cyclic structure even when familiar action semantics are permuted.

By contrast, **AP** and **SREP** remain difficult for most models under this counterfactual, with many m -values close to 20–25. This indicates that the label remapping does not merely slow exploitation of cyclic opponents, but more broadly disrupts stable identification of opponent policies. DeepSeek R1 is relatively weak in this setting: most of its comprehension values remain late across all opponent types, with only isolated improvements such as $m = 8.0 \pm 9.9$ against CoT under SC-ZS or $m = 13.0 \pm 0.0$ against **AP** under SC-CoT. Llama 3.3 is more competitive, showing several very early values—for example $m = 1.0 \pm 0.0$ against **TFT** in multiple prompting settings and $m = 1.0 \pm 0.0$ against **AP** under SC-SPP—but it remains less uniformly robust than the strongest Claude variants. Mistral Large is again highly unstable: although it occasionally achieves immediate comprehension (e.g., $m = 1.0 \pm 0.0$ against **TFT** or in some LLM–LLM settings), many other entries stay near the horizon, confirming brittle and inconsistent adaptation.

Overall, the label-based RPS counterfactual experiment shows that label-only perturbations in RPS significantly impair opponent comprehension. The most robust models are Claude 3.7, Claude 3.7(T), and Claude 4, which still exploit **PP/TFT** in several settings, while Claude 3.5, DeepSeek R1, and especially Mistral exhibit much greater degradation. In contrast to default RPS, where some models could reliably infer cyclic or reactive opponents, the relabeled game often pushes comprehension toward the horizon, supporting the claim that many LLMs depend in part on memorized action semantics rather than abstract reasoning over the modified dominance structure.

Table 21 illustrated results for the payoff-based RPS counterfactual. In this case, opponent compre-

hension becomes more polarized than in the default or label-only setting. The clearest pattern is that **PP** and **TFT** are often understood very early, especially by the stronger Claude models, whereas **SREP** and many LLM–LLM matchups still remain close to the horizon. This suggests that once exploitable cyclic or reactive structure is present, several models can adapt quickly to the altered payoff asymmetry, but they still struggle to infer behavior in near-equilibrium or symmetric play. In particular, Claude 3.7, Claude 3.7(T), and Claude 4 show the strongest results, frequently reaching $m \approx 1$ against **PP** and/or **TFT** across several prompting settings; for example, Claude 3.7 obtains $m = 1.0 \pm 0.0$ against **PP** and $m = 1.2 \pm 0.4$ against **TFT** under CoT, while Claude 4 repeatedly reaches $m = 1.0 \pm 0.0$ against **PP** under multiple prompts. Claude 3.5 also improves substantially under self-consistency, with several $m = 1.0 \pm 0.0$ outcomes against **PP**, **AP**, and **TFT**, although its non-SC variants remain much less stable.

A second notable trend is that **AP** becomes easier than in the default and label-based setups for some models, but this improvement is uneven. Claude 3.5 under SC, Claude 3.7 under SC-SPP, DeepSeek R1 under SC-ZS/SC-SPP, and Llama 3.3 under SC-ZS all achieve very early or immediate comprehension on **AP**, yet many other entries remain delayed, often with m near 15–24. This suggests that the payoff asymmetry can make adaptive behavior more exploitable once recognized, but models do not consistently detect this advantage.

DeepSeek R1 remains mixed overall: it shows some good outcomes under SC, especially on **AP** and **TFT**, but many non-SC values remain relatively late. Llama 3.3 is more competitive in this counterfactual than in the label-only one, with many early values across **PP**, **TFT**, and occasionally **AP**, indicating relatively good payoff sensitivity. Mistral Large remains the most unstable model, although it also exhibits several isolated $m = 1.0 \pm 0.0$ results under SC, particularly against **PP**, **AP**, and **TFT**. Overall, the table suggests that payoff-based changes are highly diagnostic: the strongest models can recompute the new incentives and quickly exploit structured opponents, but comprehension remains weak in LLM–LLM play and against **SREP**, confirming that altered payoffs help mainly when exploitable regularities are present rather than by generally improving strategic inference.

We finally present results for the joint RPS coun-

		Payoff-based RPS counterfactual						
model	prompt	zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	16.2 ± 11.0	17.6 ± 10.3	24.4 ± 0.5	23.4 ± 1.9	19.0 ± 10.4	19.4 ± 10.4	5.6 ± 10.3
	cot	20.2 ± 9.1	12.6 ± 11.0	21.0 ± 5.3	17.6 ± 9.8	24.2 ± 1.3	18.4 ± 9.8	10.0 ± 12.4
	spp	10.6 ± 10.1	13.4 ± 11.5	22.2 ± 4.1	18.4 ± 6.9	15.0 ± 9.1	24.2 ± 0.8	14.2 ± 12.3
	sc-zs	23.5 ± 2.1	24.0 ± 1.4	13.0 ± 14.1	18.0 ± 9.9	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	25.0 ± 0.0	18.5 ± 6.4	25.0 ± 0.0	19.5 ± 2.1	12.0 ± 15.6	11.5 ± 14.8	1.0 ± 0.0
	sc-spp	23.0 ± 1.4	25.0 ± 0.0	2.0 ± 1.4	24.5 ± 0.7	1.0 ± 0.0	12.5 ± 16.3	1.0 ± 0.0
C3.7S	zs	20.4 ± 8.6	17.6 ± 10.6	23.2 ± 2.0	22.4 ± 2.5	1.4 ± 0.9	18.8 ± 10.1	3.2 ± 3.5
	cot	16.6 ± 9.5	23.0 ± 2.9	23.2 ± 1.6	18.8 ± 6.2	1.0 ± 0.0	7.4 ± 8.9	1.2 ± 0.4
	spp	21.0 ± 4.6	24.4 ± 1.3	24.0 ± 1.0	20.8 ± 6.0	3.0 ± 4.5	18.4 ± 9.8	1.0 ± 0.0
	sc-zs	16.5 ± 12.0	25.0 ± 0.0	24.5 ± 0.7	17.0 ± 11.3	12.0 ± 15.6	12.0 ± 15.6	1.0 ± 0.0
	sc-cot	19.0 ± 7.1	21.0 ± 5.7	21.5 ± 0.7	24.5 ± 0.7	1.0 ± 0.0	22.0 ± 1.4	1.0 ± 0.0
	sc-spp	13.0 ± 17.0	23.5 ± 2.1	23.5 ± 2.1	24.0 ± 1.4	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
C3.7S(T)	zs	19.2 ± 6.8	20.8 ± 3.5	23.6 ± 2.1	17.8 ± 8.6	5.8 ± 9.1	23.6 ± 1.5	4.8 ± 4.8
	cot	18.8 ± 6.6	21.4 ± 4.9	23.6 ± 1.1	21.2 ± 5.8	1.0 ± 0.0	13.8 ± 11.9	3.4 ± 5.4
	spp	21.4 ± 5.0	22.4 ± 1.9	24.8 ± 0.4	19.6 ± 8.8	1.0 ± 0.0	5.8 ± 10.7	1.0 ± 0.0
	sc-zs	25.0 ± 0.0	20.0 ± 7.1	13.0 ± 17.0	24.0 ± 0.0	1.0 ± 0.0	19.5 ± 7.8	1.0 ± 0.0
	sc-cot	22.5 ± 2.1	24.5 ± 0.7	25.0 ± 0.0	18.5 ± 6.4	1.0 ± 0.0	23.5 ± 2.1	1.0 ± 0.0
	sc-spp	24.0 ± 1.4	25.0 ± 0.0	24.0 ± 1.4	20.5 ± 0.7	1.0 ± 0.0	12.0 ± 15.6	1.0 ± 0.0
C4S	zs	23.8 ± 1.6	23.2 ± 2.0	18.0 ± 7.6	24.2 ± 0.8	1.0 ± 0.0	23.0 ± 1.9	7.2 ± 8.5
	cot	19.8 ± 10.5	21.8 ± 6.1	19.6 ± 5.6	20.0 ± 4.2	1.0 ± 0.0	18.8 ± 10.0	5.6 ± 10.3
	spp	10.2 ± 9.7	19.2 ± 10.3	23.0 ± 1.4	24.0 ± 1.7	5.8 ± 10.7	18.6 ± 5.8	15.6 ± 11.6
	sc-zs	20.0 ± 7.1	25.0 ± 0.0	18.0 ± 9.9	25.0 ± 0.0	1.0 ± 0.0	23.0 ± 0.0	20.5 ± 3.5
	sc-cot	12.0 ± 15.6	19.0 ± 8.5	13.0 ± 17.0	16.5 ± 9.2	1.0 ± 0.0	13.0 ± 17.0	1.0 ± 0.0
	sc-spp	25.0 ± 0.0	23.5 ± 2.1	13.0 ± 17.0	13.5 ± 0.7	1.0 ± 0.0	24.0 ± 1.4	1.0 ± 0.0
C4S(T)	zs	24.0 ± 0.7	24.8 ± 0.4	21.0 ± 4.3	21.4 ± 4.3	4.2 ± 7.2	23.4 ± 1.5	15.0 ± 12.3
	cot	22.8 ± 2.3	10.6 ± 10.3	20.2 ± 10.7	17.4 ± 10.1	1.2 ± 0.4	18.2 ± 9.8	1.0 ± 0.0
	spp	19.2 ± 6.2	19.4 ± 4.6	12.2 ± 11.5	21.2 ± 4.2	1.0 ± 0.0	20.4 ± 9.2	6.0 ± 10.1
	sc-zs	25.0 ± 0.0	24.0 ± 1.4	19.5 ± 7.8	11.5 ± 0.7	1.0 ± 0.0	23.0 ± 1.4	1.0 ± 0.0
	sc-cot	23.5 ± 0.7	23.0 ± 1.4	25.0 ± 0.0	23.0 ± 0.0	3.5 ± 2.1	18.5 ± 7.8	1.0 ± 0.0
	sc-spp	22.5 ± 2.1	24.0 ± 1.4	21.5 ± 0.7	24.5 ± 0.7	1.0 ± 0.0	23.0 ± 2.8	12.5 ± 16.3
DS-R1	zs	24.8 ± 0.4	19.0 ± 6.7	20.0 ± 10.6	23.0 ± 2.0	18.6 ± 9.3	18.8 ± 8.3	5.0 ± 8.9
	cot	19.6 ± 7.4	19.0 ± 8.0	20.8 ± 5.4	17.8 ± 9.1	24.4 ± 0.9	15.2 ± 11.7	7.0 ± 7.0
	spp	22.2 ± 3.3	21.2 ± 4.7	19.4 ± 3.0	24.4 ± 1.3	17.8 ± 9.7	14.6 ± 11.6	9.4 ± 11.5
	sc-zs	20.0 ± 7.1	20.0 ± 0.0	23.0 ± 1.4	24.0 ± 0.0	11.5 ± 14.8	1.5 ± 0.7	1.0 ± 0.0
	sc-cot	17.0 ± 5.7	24.5 ± 0.7	25.0 ± 0.0	18.5 ± 6.4	25.0 ± 0.0	18.5 ± 6.4	5.5 ± 6.4
	sc-spp	25.0 ± 0.0	18.0 ± 5.7	24.5 ± 0.7	24.0 ± 0.0	13.0 ± 17.0	1.5 ± 0.7	1.5 ± 0.7
L3.3-70B	zs	5.8 ± 10.7	15.0 ± 9.7	8.4 ± 10.9	23.8 ± 0.8	7.2 ± 10.4	6.8 ± 9.7	1.0 ± 0.0
	cot	9.6 ± 11.9	16.8 ± 10.8	13.4 ± 11.3	21.6 ± 4.9	1.0 ± 0.0	16.4 ± 10.5	1.0 ± 0.0
	spp	1.0 ± 0.0	5.0 ± 4.1	7.6 ± 10.5	24.2 ± 0.8	11.6 ± 11.1	20.2 ± 7.3	1.2 ± 0.4
	sc-zs	13.0 ± 17.0	1.0 ± 0.0	11.5 ± 14.8	23.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	13.0 ± 17.0	7.0 ± 8.5	10.5 ± 13.4	16.0 ± 7.1	1.0 ± 0.0	22.0 ± 0.0	1.0 ± 0.0
	sc-spp	1.0 ± 0.0	12.0 ± 15.6	19.0 ± 5.7	18.5 ± 9.2	6.5 ± 7.8	6.0 ± 7.1	1.0 ± 0.0
Mistral	zs	7.6 ± 10.1	15.2 ± 13.0	12.4 ± 10.5	24.8 ± 0.4	1.0 ± 0.0	19.8 ± 10.5	1.0 ± 0.0
	cot	9.8 ± 11.2	17.0 ± 6.8	19.2 ± 10.2	24.0 ± 1.7	13.8 ± 11.9	15.4 ± 13.1	4.8 ± 7.9
	spp	10.0 ± 12.3	20.6 ± 4.6	11.6 ± 9.0	22.0 ± 5.6	20.6 ± 6.6	21.0 ± 6.4	8.4 ± 9.7
	sc-zs	1.0 ± 0.0	24.5 ± 0.7	11.0 ± 0.0	22.0 ± 2.8	1.0 ± 0.0	1.0 ± 0.0	1.0 ± 0.0
	sc-cot	1.0 ± 0.0	11.5 ± 14.8	12.0 ± 15.6	23.5 ± 0.7	13.0 ± 17.0	1.0 ± 0.0	1.5 ± 0.7
	sc-spp	24.0 ± 0.0	12.5 ± 16.3	16.5 ± 10.6	17.0 ± 8.5	13.0 ± 17.0	12.5 ± 16.3	1.0 ± 0.0

Table 21: Round m # where the LLM player understood the opponent’s Strategy in the payoff-based RPS counterfactual.

terfactual in Table 22. The joint RPS counterfactual is the hardest comprehension setting overall, since it simultaneously perturbs both action semantics and payoff structure. Relative to the label-only and payoff-only variants, many m -values move closer to the horizon, showing that stable opponent inference becomes substantially harder. The most robust models are Claude 3.7 and Claude 3.7(T), which still achieve very early comprehen-

sion against structured algorithmic opponents in several prompting settings, especially on **PP**, **AP**, and **TFT**. Claude 4 is somewhat competitive but less stable, while Claude 3.5 is much more disrupted, with many late values across opponent types. DeepSeek R1 also degrades markedly in this joint setting, whereas Llama 3.3 remains moderately competitive with a few strong low- m cases. Mistral is again the most brittle and inconsistent.

model	prompt	Joint RPS counterfactual						
		zs	spp	cot	srep	pp	ap	tft
C3.5Sv2	zs	23.8 ± 1.3	18.8 ± 7.6	14.6 ± 10.7	24.2 ± 0.8	21.2 ± 7.4	15.0 ± 12.8	10.2 ± 9.3
	cot	16.2 ± 9.5	17.6 ± 9.6	19.0 ± 7.4	17.4 ± 10.3	24.4 ± 0.9	18.6 ± 5.6	14.2 ± 12.1
	spp	23.4 ± 0.9	24.6 ± 0.5	15.8 ± 9.6	20.0 ± 3.7	25.0 ± 0.0	19.0 ± 8.0	7.8 ± 10.2
	sc-zs	24.0 ± 1.4	24.0 ± 0.0	15.0 ± 0.0	19.0 ± 8.5	21.0 ± 5.7	23.5 ± 2.1	24.5 ± 0.7
	sc-cot	22.5 ± 2.1	14.0 ± 15.6	12.5 ± 13.4	24.0 ± 1.4	22.5 ± 3.5	24.5 ± 0.7	10.5 ± 13.4
	sc-spp	10.5 ± 12.0	12.5 ± 16.3	17.5 ± 10.6	22.0 ± 2.8	24.0 ± 1.4	23.5 ± 2.1	10.0 ± 12.7
C3.7S	zs	21.0 ± 4.2	23.2 ± 2.2	23.4 ± 1.8	22.8 ± 2.9	16.6 ± 9.0	19.0 ± 10.2	12.2 ± 11.8
	cot	20.0 ± 4.5	20.4 ± 4.6	21.4 ± 4.8	18.0 ± 7.5	2.6 ± 3.6	23.2 ± 1.6	9.4 ± 11.7
	spp	17.0 ± 9.7	22.8 ± 2.9	24.6 ± 0.5	20.8 ± 8.8	4.6 ± 8.0	17.4 ± 9.9	7.0 ± 10.4
	sc-zs	12.5 ± 16.3	23.5 ± 0.7	25.0 ± 0.0	18.5 ± 9.2	7.5 ± 0.7	1.0 ± 0.0	3.0 ± 2.8
	sc-cot	13.0 ± 17.0	23.0 ± 2.8	24.5 ± 0.7	23.0 ± 1.4	2.0 ± 1.4	22.5 ± 2.1	1.0 ± 0.0
	sc-spp	19.0 ± 5.7	23.0 ± 1.4	24.5 ± 0.7	23.0 ± 1.4	1.0 ± 0.0	25.0 ± 0.0	1.0 ± 0.0
C3.7S(T)	zs	22.6 ± 4.3	20.0 ± 4.7	21.6 ± 4.6	24.2 ± 0.4	24.8 ± 0.4	19.0 ± 10.1	9.4 ± 8.8
	cot	17.4 ± 6.7	21.6 ± 4.1	22.6 ± 3.2	21.6 ± 5.1	11.8 ± 12.1	21.4 ± 3.9	2.4 ± 3.1
	spp	17.0 ± 7.2	19.8 ± 8.4	22.8 ± 1.9	17.4 ± 7.4	2.4 ± 1.3	22.0 ± 1.6	6.2 ± 10.0
	sc-zs	24.5 ± 0.7	18.5 ± 0.7	24.5 ± 0.7	21.0 ± 1.4	25.0 ± 0.0	23.0 ± 1.4	8.0 ± 9.9
	sc-cot	17.0 ± 5.7	23.5 ± 0.7	25.0 ± 0.0	24.0 ± 0.0	1.5 ± 0.7	23.5 ± 0.7	1.0 ± 0.0
	sc-spp	14.0 ± 12.7	18.0 ± 9.9	23.5 ± 2.1	13.0 ± 8.5	3.0 ± 2.8	25.0 ± 0.0	1.0 ± 0.0
C4S	zs	20.4 ± 6.3	23.4 ± 1.7	19.8 ± 10.5	21.6 ± 3.8	13.6 ± 9.9	23.4 ± 0.5	17.4 ± 9.9
	cot	24.0 ± 1.0	20.4 ± 7.6	23.2 ± 2.2	19.2 ± 5.9	1.4 ± 0.9	24.2 ± 0.8	18.6 ± 9.9
	spp	23.4 ± 1.3	24.2 ± 1.3	21.4 ± 4.8	19.0 ± 5.6	6.4 ± 8.8	22.0 ± 4.0	14.0 ± 9.7
	sc-zs	22.5 ± 2.1	23.0 ± 1.4	24.0 ± 1.4	18.5 ± 6.4	8.5 ± 4.9	24.0 ± 1.4	22.0 ± 1.4
	sc-cot	15.0 ± 14.1	25.0 ± 0.0	22.5 ± 3.5	21.0 ± 1.4	1.0 ± 0.0	24.0 ± 1.4	24.0 ± 0.0
	sc-spp	23.5 ± 2.1	11.5 ± 2.1	25.0 ± 0.0	24.0 ± 1.4	3.0 ± 2.8	24.5 ± 0.7	1.0 ± 0.0
C4S(T)	zs	17.4 ± 10.0	24.8 ± 0.4	19.4 ± 10.3	18.2 ± 9.1	23.4 ± 3.0	24.0 ± 2.2	14.6 ± 12.4
	cot	21.0 ± 6.2	20.4 ± 8.6	23.2 ± 1.6	15.2 ± 11.3	12.2 ± 11.5	24.6 ± 0.9	13.6 ± 11.8
	spp	23.6 ± 1.3	19.4 ± 7.2	23.6 ± 0.9	16.8 ± 7.0	2.6 ± 3.0	22.0 ± 1.9	19.8 ± 9.4
	sc-zs	25.0 ± 0.0	24.0 ± 1.4	23.5 ± 2.1	21.0 ± 2.8	14.0 ± 12.7	24.0 ± 1.4	25.0 ± 0.0
	sc-cot	23.5 ± 0.7	24.5 ± 0.7	24.0 ± 1.4	18.5 ± 9.2	24.5 ± 0.7	24.5 ± 0.7	24.0 ± 0.0
	sc-spp	14.5 ± 13.4	8.0 ± 9.9	23.5 ± 0.7	24.5 ± 0.7	13.5 ± 13.4	23.5 ± 0.7	22.5 ± 2.1
DS-R1	zs	19.2 ± 7.1	24.0 ± 1.2	22.6 ± 2.5	21.8 ± 6.6	24.6 ± 0.5	20.8 ± 5.0	22.4 ± 2.1
	cot	23.6 ± 1.1	21.0 ± 5.7	20.4 ± 8.6	24.6 ± 0.5	24.4 ± 0.5	23.4 ± 0.9	22.0 ± 2.2
	spp	22.2 ± 4.1	24.2 ± 0.4	19.6 ± 4.9	21.0 ± 5.9	24.8 ± 0.4	20.8 ± 5.3	16.6 ± 9.8
	sc-zs	25.0 ± 0.0	10.5 ± 13.4	19.0 ± 7.1	22.0 ± 4.2	24.0 ± 1.4	18.5 ± 4.9	12.0 ± 2.8
	sc-cot	18.0 ± 4.2	22.0 ± 2.8	25.0 ± 0.0	24.0 ± 1.4	24.5 ± 0.7	19.5 ± 7.8	25.0 ± 0.0
	sc-spp	24.5 ± 0.7	20.0 ± 7.1	6.0 ± 2.8	25.0 ± 0.0	25.0 ± 0.0	23.0 ± 0.0	13.0 ± 15.6
L3.3-70B	zs	10.4 ± 11.6	14.4 ± 12.3	14.4 ± 10.7	24.2 ± 1.3	11.4 ± 12.4	16.8 ± 11.1	10.6 ± 13.1
	cot	18.2 ± 5.4	6.8 ± 6.3	5.6 ± 10.3	22.4 ± 4.2	16.2 ± 8.9	22.2 ± 1.8	10.4 ± 12.9
	spp	20.0 ± 10.6	6.2 ± 10.5	4.6 ± 5.1	24.2 ± 1.1	16.0 ± 9.9	14.8 ± 12.7	1.0 ± 0.0
	sc-zs	10.0 ± 12.7	1.0 ± 0.0	7.5 ± 9.2	23.5 ± 2.1	13.0 ± 17.0	24.0 ± 1.4	1.0 ± 0.0
	sc-cot	12.5 ± 10.6	12.0 ± 15.6	20.5 ± 4.9	25.0 ± 0.0	10.5 ± 13.4	24.5 ± 0.7	1.0 ± 0.0
	sc-spp	21.0 ± 1.4	13.0 ± 14.1	14.0 ± 15.6	20.0 ± 1.4	19.0 ± 8.5	1.0 ± 0.0	1.0 ± 0.0
Mistral	zs	6.0 ± 6.2	17.6 ± 10.3	14.6 ± 12.4	23.8 ± 0.8	13.4 ± 11.2	19.2 ± 10.2	7.8 ± 10.5
	cot	11.2 ± 12.2	16.8 ± 5.8	14.0 ± 12.1	22.6 ± 2.1	25.0 ± 0.0	22.0 ± 4.6	1.6 ± 1.3
	spp	17.2 ± 10.0	19.6 ± 6.1	17.6 ± 9.7	21.4 ± 4.5	7.0 ± 10.0	12.0 ± 10.4	5.8 ± 10.7
	sc-zs	13.0 ± 17.0	20.0 ± 7.1	19.0 ± 8.5	19.0 ± 8.5	25.0 ± 0.0	24.5 ± 0.7	7.5 ± 9.2
	sc-cot	24.5 ± 0.7	23.5 ± 0.7	23.5 ± 2.1	24.5 ± 0.7	17.5 ± 10.6	12.0 ± 15.6	1.5 ± 0.7
	sc-spp	13.0 ± 17.0	20.5 ± 0.7	23.5 ± 0.7	24.0 ± 1.4	25.0 ± 0.0	24.5 ± 0.7	1.0 ± 0.0

Table 22: Round m # where the LLM player understood the opponent’s Strategy in the joint RPS counterfactual.

Overall, the table shows that combining label and payoff perturbations exposes stronger reasoning limitations than either intervention alone, with only the strongest models retaining reliable comprehension of exploitable opponents.

D.8 RPS Efficiency

In RPS, efficiency (Table 23) is generally lower and more variable than in PD, reflecting the greater difficulty of converting reasoning into reliable pay-

off gains in a three-action adversarial game. A consistent pattern across almost all models is that the payoff-based counterfactual yields the highest efficiency values, often by a clear margin. This is visible for Claude 3.7 (2.02 ± 4.02 under ZS), Claude 3.7(T) (3.26 ± 4.15 under ZS), Claude 4(T) (1.50 ± 2.05 under CoT), DeepSeek R1 (1.04±1.77 under CoT), and Llama 3.3 (2.38±4.85 under ZS). The reason is intuitive: once the altered payoff asymmetry is recognized, exploitable oppo-

model	prompt	RPS	RPS label-based	RPS payoff-based	RPS joint
C3.5v2	zs	1.09 ± 2.37	0.09 ± 1.34	0.47 ± 2.65	0.25 ± 1.61
	cot	0.44 ± 0.74	-0.23 ± 0.84	0.49 ± 1.64	-0.03 ± 1.28
	spp	0.49 ± 0.71	0.01 ± 0.48	0.66 ± 0.83	0.05 ± 1.10
	sc-zs	0.20 ± 0.60	0.10 ± 0.48	0.46 ± 0.56	-0.03 ± 0.35
	sc-cot	0.14 ± 0.18	-0.08 ± 0.25	0.15 ± 0.41	-0.01 ± 0.36
	sc-spp	0.01 ± 0.20	-0.03 ± 0.17	0.19 ± 0.34	-0.01 ± 0.17
C3.7S	zs	0.58 ± 2.39	0.20 ± 2.12	2.02 ± 4.02	0.07 ± 2.02
	cot	0.65 ± 0.75	0.51 ± 0.80	0.63 ± 1.07	0.71 ± 1.10
	spp	0.51 ± 0.70	0.41 ± 0.61	0.43 ± 1.13	0.56 ± 0.76
	sc-zs	-0.08 ± 0.72	0.00 ± 0.36	0.00 ± 0.93	0.32 ± 0.71
	sc-cot	0.21 ± 0.25	0.14 ± 0.17	0.17 ± 0.26	0.11 ± 0.28
	sc-spp	0.14 ± 0.18	0.09 ± 0.12	0.16 ± 0.21	0.13 ± 0.16
C3.7S(T)	zs	1.92 ± 3.30	0.60 ± 2.55	3.26 ± 4.15	1.19 ± 4.43
	cot	0.61 ± 1.04	0.25 ± 1.02	0.87 ± 1.27	0.60 ± 1.15
	spp	0.57 ± 0.89	0.60 ± 0.80	1.06 ± 1.53	0.70 ± 0.97
	sc-zs	0.23 ± 0.77	0.31 ± 0.98	0.70 ± 1.44	-0.08 ± 0.81
	sc-cot	0.11 ± 0.27	0.15 ± 0.21	0.24 ± 0.28	0.18 ± 0.31
	sc-spp	0.04 ± 0.26	0.11 ± 0.18	0.17 ± 0.35	0.18 ± 0.24
C4S	zs	0.80 ± 1.82	-0.12 ± 2.25	0.92 ± 2.96	0.23 ± 2.49
	cot	0.62 ± 0.89	0.31 ± 0.71	0.75 ± 1.41	0.57 ± 1.10
	spp	0.37 ± 0.70	0.19 ± 0.56	0.73 ± 0.83	0.30 ± 0.74
	sc-zs	0.13 ± 0.50	0.07 ± 0.29	0.12 ± 0.43	0.07 ± 0.21
	sc-cot	0.10 ± 0.18	0.07 ± 0.13	0.15 ± 0.29	0.11 ± 0.25
	sc-spp	0.10 ± 0.14	0.05 ± 0.09	0.09 ± 0.20	0.12 ± 0.21
C4S(T)	zs	2.07 ± 3.74	-0.16 ± 2.00	1.51 ± 4.69	0.04 ± 3.76
	cot	1.11 ± 1.68	0.60 ± 0.92	1.50 ± 2.05	0.81 ± 1.93
	spp	0.31 ± 0.72	0.28 ± 0.68	0.69 ± 1.20	0.57 ± 0.91
	sc-zs	0.44 ± 0.67	-0.21 ± 0.39	0.37 ± 1.14	0.03 ± 0.75
	sc-cot	0.35 ± 0.26	0.20 ± 0.28	0.16 ± 0.61	-0.04 ± 0.27
	sc-spp	0.13 ± 0.13	0.07 ± 0.12	0.14 ± 0.18	0.10 ± 0.13
DS-R1	zs	0.47 ± 0.73	0.17 ± 0.44	0.89 ± 1.49	0.03 ± 0.86
	cot	0.80 ± 0.84	0.23 ± 0.49	1.04 ± 1.77	0.02 ± 0.76
	spp	0.41 ± 0.74	0.03 ± 0.31	0.49 ± 0.96	0.09 ± 0.85
	sc-zs	0.13 ± 0.22	-0.01 ± 0.11	0.42 ± 0.48	0.03 ± 0.20
	sc-cot	0.19 ± 0.28	0.07 ± 0.14	0.29 ± 0.32	0.08 ± 0.27
	sc-spp	0.09 ± 0.20	0.05 ± 0.10	0.22 ± 0.38	0.05 ± 0.18
L3.3-70B	zs	1.16 ± 3.16	0.14 ± 2.06	2.38 ± 4.85	-0.07 ± 3.64
	cot	1.18 ± 2.20	0.12 ± 1.26	1.59 ± 2.91	-0.01 ± 1.46
	spp	0.57 ± 1.41	0.26 ± 1.05	1.28 ± 2.29	-0.13 ± 2.19
	sc-zs	0.13 ± 0.63	0.11 ± 0.56	0.53 ± 1.31	0.02 ± 0.85
	sc-cot	0.22 ± 0.28	-0.10 ± 0.45	0.12 ± 0.73	0.04 ± 0.21
	sc-spp	0.20 ± 0.53	0.04 ± 0.24	0.46 ± 0.72	-0.22 ± 0.69
Mistral	zs	0.78 ± 4.58	-0.35 ± 3.42	1.44 ± 6.57	-1.05 ± 5.91
	cot	0.36 ± 1.04	0.02 ± 0.59	0.51 ± 1.16	0.13 ± 1.21
	spp	0.50 ± 0.75	0.15 ± 0.79	0.24 ± 1.68	0.62 ± 1.14
	sc-zs	0.11 ± 0.96	-0.10 ± 0.74	0.01 ± 1.21	-0.37 ± 1.70
	sc-cot	0.15 ± 0.29	0.03 ± 0.22	0.22 ± 0.31	0.03 ± 0.24
	sc-spp	0.03 ± 0.18	0.02 ± 0.26	0.13 ± 0.35	0.02 ± 0.39

Table 23: Average Efficiency (Points per kilo-token). Bold values denote higher efficiency per model-prompting setup.

nents such as PP or TFT can produce larger score gains per token than in the default game, so the extra reasoning is more often rewarded.

By contrast, the label-based and especially the joint counterfactuals are usually less efficient. In the label-based setup, efficiency often drops close to zero or even becomes negative, suggesting that relabeling disrupts strategic adaptation without producing compensating payoff gains. This effect is particularly strong for Claude 3.5, Claude 4, Claude

4(T), Llama 3.3, and Mistral in several prompt settings. The joint counterfactual is often even harsher: many values are near zero or negative, especially for weaker or less stable models, which indicates that simultaneously changing labels and payoffs increases token cost without reliably improving performance.

Across models, Claude 3.7(T) is the strongest overall efficiency performer in RPS, especially in the payoff-based setting, where it achieves the high-

est values in the table. Llama 3.3 is also highly competitive in payoff-based RPS, although its efficiency drops sharply in the joint setup, showing weaker robustness when both semantics and incentives change together. Claude 4(T) can also be very efficient in default and payoff-based RPS, but becomes much less efficient in the label-based and joint settings. Claude 3.7 and Claude 4 are more moderate but stable performers, usually remaining positive and competitive without reaching the peaks of Claude 3.7(T). DeepSeek R1 is fairly competitive in efficiency, especially in payoff-based RPS, even though its strategic robustness is less consistent. Mistral is the least reliable: while it occasionally attains strong efficiency, especially in payoff-based or joint SPP settings, it also produces some of the most negative and variable results, showing that low token cost can coincide with brittle behavior rather than genuinely better adaptation.

Prompting effects broadly mirror the PD findings. ZS and standard prompted variants often achieve the highest raw efficiency when the model already exploits the opponent effectively, because they avoid the heavy token overhead of aggregation. Self-consistency usually lowers efficiency, even when it stabilizes behavior, since majority voting adds substantial token cost and only sometimes yields proportional score gains. This is especially clear in the stronger models, where SC frequently reduces the large positive efficiencies seen in non-SC payoff-based settings. Overall, the table reinforces three main points: (i) payoff-based RPS is the most efficiency-favorable setting because altered incentives can yield larger returns per token, (ii) label-based and joint counterfactuals are substantially less efficient because they disrupt strategic transfer, and (iii) the best efficiency-capability trade-off in RPS appears to come from Claude 3.7(T), with Llama 3.3 and Claude 4(T) also competitive in the payoff-based setup.

D.9 RPS validity rate

In Table 24 we present validity rates for the various LLMs used in our experimentation. Validity in RPS is effectively perfect across almost all models, showing that formatting and action-space compliance are not a meaningful confound for the strategic results. Claude 3.5 Sonnet v2, Claude 3.7, Claude 4(T), and Llama 3.3 all achieve a strict $100.0\% \pm 0.0$ valid rate, while Claude 3.7(T) and DeepSeek R1 remain essentially perfect with only

model	avg validity rate
C3.5Sv2	100.0 ± 0.0
C3.7S	100.0 ± 0.0
C3.7S(T)	100.0 ± 0.2
C4S	99.9 ± 0.9
C4S(T)	100.0 ± 0.0
DS-R1	100.0 ± 0.2
L3.3-70B Instruct	100.0 ± 0.0
Mistral	99.2 ± 6.8

Table 24: Average Valid Rate (% of Valid Outcomes)

negligible variation ($100.0\% \pm 0.2$). Claude 4 is similarly reliable at $99.9\% \pm 0.9$, indicating only extremely rare invalid outputs. The only clear outlier is Mistral, with a lower and much more variable validity rate of $99.2\% \pm 6.8$. Although this still indicates high average compliance, the larger variance suggests occasional formatting or parsing failures that are substantially more frequent than for the other models. Overall, these results confirm that the differences observed in RPS total points, opponent comprehension, and efficiency are not driven by invalid-action artifacts, but by genuine differences in strategic behavior; validity issues are too rare to materially affect the main conclusions, except that Mistral’s instability is also reflected at the output-format level.

D.10 Cross-model analysis

Overall model comparison. Across settings, the strongest models differ by game and by robustness criterion. In the default PD, Claude 3.5/3.7 Sonnet and Llama 3.3-70B are the strongest overall, as they reliably achieve perfect or near-perfect mutual cooperation in LLM–LLM play while also adapting well to algorithmic opponents. In the default RPS setting, the strongest performance is shown by the advanced Claude models, especially Claude 3.7 and Claude 4, which most consistently exploit cyclic algorithmic opponents and maintain near-equilibrium behavior against stronger opposition. Under counterfactual interventions, Claude 3.7 and the stronger Claude variants remain the most robust overall, showing the smallest degradation under label-only changes and the best adaptation when payoff structure changes require recomputing incentives; by contrast, Mistral Large is the most brittle, with DeepSeek R1 also showing noticeable instability. In terms of prompting stability, Claude 3.5/3.7 and Llama 3.3 are generally the most consistent across ZS, CoT, and SPP, whereas Mistral Large exhibits the highest variability and Claude

4 as well as DeepSeek R1 are more sensitive to prompting, particularly in settings where additional reasoning appears to induce overthinking rather than better strategic adaptation.

Robustness across prompting strategies.

Across prompting strategies, the benefits of explicit reasoning are uneven. CoT is most helpful for models that already display strong strategic competence but occasionally need more structured adaptation, with Claude 3.5/3.7 and, in some PD settings, Mistral Large showing improvements in stability or recovery relative to zero-shot behavior. By contrast, several stronger reasoning-oriented models degrade under thinking-enabled variants: Claude 4 and DeepSeek R1 often become more distrustful or unstable, particularly in PD LLM–LLM interactions and in the harder counterfactual settings, suggesting that additional deliberation can induce overthinking rather than better strategic play. In terms of prompting stability, Claude 3.5/3.7 and Llama 3.3 are the most consistent across ZS, CoT, SPP, and SC, maintaining similar qualitative behavior and relatively low variance, whereas Mistral Large is the most prompt-sensitive overall and Claude 4/DeepSeek R1 exhibit greater fluctuations across prompting conditions. Self-consistency generally improves robustness by reducing run-to-run variance, but it does not fundamentally change the strategic tendencies of a model; rather, it tends to reinforce whichever behavior—cooperative, exploitative, or unstable—the underlying prompting setup already induces.

Capability vs Efficiency. A clear trade-off emerges between strategic capability and efficiency. Models that achieve the strongest total-point outcomes and earliest opponent comprehension are not always the most efficient, since superior performance often comes at the cost of longer reasoning traces and higher token usage. This pattern is especially visible in RPS, where advanced Claude models often outperform weaker baselines in total points and adaptation, but do so with greater computational cost. Conversely, some lighter or less deliberative models can appear more efficient simply because they generate fewer tokens, even when their strategic performance is weaker or less stable. The same trade-off appears within model families: thinking-enabled variants do not consistently improve total points or comprehension in proportion to their additional token usage, and in several

cases—particularly for Claude 4 and DeepSeek R1—the extra reasoning is associated with lower efficiency and occasionally worse strategic outcomes. Overall, efficiency should therefore be interpreted jointly with performance metrics: higher token expenditure can support stronger play in harder settings, but it does not guarantee proportionate gains and may instead reflect overthinking or unstable adaptation.

Robustness across opponent types.

Robustness also varies markedly across opponent types. Against algorithmic opponents, the strongest models are those that can quickly identify deterministic or cyclic structure and exploit it reliably: advanced Claude models perform best overall in this regime, with Llama 3.3 also showing strong and consistent adaptation, especially in PD. In contrast, Mistral Large is the least reliable exploiter of deterministic opponents, often adapting late or inconsistently, while DeepSeek R1 tends to be more variable across settings. Against LLM opponents, the challenge shifts from exploitation to coordination or equilibrium maintenance. In PD, Claude 3.5/3.7 and Llama 3.3 coordinate best with other LLMs, consistently converging to mutual cooperation, whereas Claude 4 and DeepSeek R1 remain more brittle, often defaulting to defection or unstable exploration. In RPS, where cooperative convergence is impossible, robustness in LLM–LLM play is better reflected by maintaining near-equilibrium behavior without introducing exploitable biases; here again, stronger Claude models are the most stable, while Mistral Large and DeepSeek R1 are more prone to deviations and instability. Overall, some models are strong exploiters of deterministic opponents, while others are better coordinators or stabilizers against adaptive LLM opposition, and the most robust models are those that perform well in both regimes.

Counterfactual behaviors. Counterfactual sensitivity differs substantially by model. Claude 3.7 and Claude 4 generally deliver the strongest performance in the default games, especially in exploiting structured opponents in RPS and achieving high payoffs in PD, but their robustness under counterfactual changes is uneven: Claude 3.7 remains comparatively stable, whereas Claude 4 is more prone to degradation, particularly when altered incentives or additional reasoning induce overthinking. Llama 3.3 is often one of the most stable models in default PD, especially in LLM–LLM cooperation, but it is

Model	PD	RPS	Counterfactuals	Efficiency	Main weakness
Claude 3.7	strong	strong	relatively robust	moderate	overreasoning
Claude 4	mixed	strong	mixed	lower	instability
DeepSeek R1	mixed	moderate	brittle	lower	coordination
Llama 3.3	strong	moderate	moderate	good	RPS adaptation
Mistral Large	weak	weak	weakest	low	comprehension

Table 25: Cross-model qualitative comparison across games, counterfactual robustness, and efficiency.

weaker in the more difficult RPS and payoff-shifted counterfactual settings, where adaptation requires finer-grained strategic recalibration. DeepSeek R1 is strong in some regimes, including early adaptation against certain algorithmic opponents, yet remains unstable overall and often over-defective in PD, which hurts both coordination and counterfactual robustness. Mistral Large is the most brittle model across the board: it shows the greatest sensitivity to label shifts, the highest instability in opponent comprehension, and the weakest ability to preserve performance when familiar game structure is altered. Overall, counterfactual sensitivity reveals that strong default-game performance does not necessarily imply robust strategic generalization, and models differ sharply in how well they recompute behavior when labels or incentives are changed.

Final takeaways. Table 25 summarizes the main qualitative differences across models, aggregating performance across games, prompting strategies, and counterfactual settings.

Claude 3.7 emerges as the most consistently strong model, combining high performance in both PD and RPS with relatively robust behavior under counterfactual changes, albeit with moderate efficiency due to additional reasoning. Claude 4 achieves strong results in RPS but is less stable overall, showing mixed counterfactual robustness and a tendency toward distrust or unstable strategies. Llama 3.3 is particularly strong and stable in PD, especially in cooperative settings, while remaining moderately effective in RPS with good efficiency. DeepSeek R1 demonstrates competitive performance in some scenarios but is overall brittle, especially in coordination-heavy settings, often defaulting to unstable or overly defensive strategies. Finally, Mistral Large is consistently the weakest across all dimensions, with delayed opponent comprehension, high variability, and the strongest degradation under counterfactual interventions.

Overall, the table highlights that strong default-

game performance does not necessarily translate to robustness or efficiency, reinforcing the importance of multi-dimensional evaluation.