A Resource-Rational Account of Rule-Based Coordination

Guy Grinfeld (ggrinfeld@fas.harvard.edu) & Fiery Cushman
Department of Psychology, Harvard University

Abstract

People can coordinate by inferring a partner's intention, or by relying on a mutually recognized rule. We propose that this choice reflects adaptive allocation of cognitive effort. We tested this account in a coordination game where participants could win the most by inferring a signaled *ad hoc* coordination point, or less by coordinating on a consistent, mutually understood rule. We manipulated the recursive depth required for successful inference and the value of rule-based coordination. Greater demands for recursive depth, and higher rule value, drove increased rule use. Response times suggest rule-based choices arose from different processes: fast defaults when inference was less valuable, and slower retreats to the rule when inference was more valuable but too demanding. Individual-level analyses showed participants shifted toward deeper reasoning when inference was more valuable. These findings suggest that rules and intention inference are substitutable strategies governed by shared resource-rational logic.

Keywords: coordination; resource rationality; pragmatic inference; rule-following

Introduction

Rules structure much of everyday coordination. Imagine two cooks in a shared kitchen. Most of the time they can work smoothly by following standing arrangements: one handles prep while the other runs the stove. But sometimes they can do better by adapting to what the other is about to do. If your partner pulls out the wok, ginger, and soy sauce, you can infer a stir-fry and start the rice. Other signals are less diagnostic: chopping onions could support almost anything. In those cases, it may be smarter to stick with the rule, even if that leaves efficiency on the table.

This example illustrates two routes to coordination. *Ad hoc* coordination relies on inferring a partner's goal from their behavior or signal (Hawkins et al., 2019; Le Pargneux et al., 2025). When successful, it can produce the optimal joint outcome, but it involves tracking context, representing what the partner is trying to communicate, and (often) engaging in recursive "what do they think I think…?" reasoning (Camerer et al., 2004; Frank & Goodman, 2012). Rule-based coordination, by contrast, involves actions that both parties can recognize as a viable coordination solution without inferring each other's intentions (Schelling, 1960; Thomas et al., 2014). Rules can be rigid and may yield suboptimal outcomes, but they can reduce the cognitive burden of coordinating.

When and why do people rely on rule-based coordination? A person who follows a rule might never have attempted inference, treating the rule as a default policy (Gächter et al., 2025). Alternatively, they might have initiated inference, encountered difficulty or uncertainty, and retreated to the rule as a fallback. These two processes produce the same overt behavior but reflect different underlying computations, and they make different predictions about how collaborators should respond to incentives and reasoning demands.

A Resource-Rational Account of Rule-Based Coordination

We propose that default and fallback rule-use both arise from a resource-rational logic applied at different stages of processing. Resource-rational accounts treat cognitive computation as a limited resource that agents allocate when the expected benefit justifies the cost (Lieder & Griffiths, 2020; Shenhav et al., 2013). Work on strategy selection formalizes this as a meta-decision: before computing an answer, agents evaluate whether the expected value of that computation exceeds its cost (Lieder & Griffiths, 2017). In strategic settings, for example, simple heuristics are often ecologically rational responses that maximize expected return (Spiliopoulos & Hertwig, 2020).

We extend this logic to coordination. Crucially, rule-based coordination is not merely a heuristic like "guess randomly," but a social solution both parties can recognize as acceptable without inferring one another's intentions. The rule has positive expected value precisely because coordination problems have shared structure that both agents can exploit. As such, the existence of rules presents a trade-off between two viable routes to a social outcome: one cognitively cheaper but potentially less rewarding, the other more demanding but potentially more rewarding.

More specifically, what makes inference costly enough to favor rules? Pragmatic models (e.g., Rational Speech Act models; Goodman & Frank, 2016) suggest that some signals can be interpreted literally, while others require pragmatic recursive reasoning about what the speaker could have said but did not, or what the speaker expects the listener to infer. As the number of required recursive steps increases, inference becomes more demanding (Alaoui & Penta, 2016;

In G. Boleda, I. Dautriche, M. Maldonado, C. Saldana and T. Torrent (Eds.), *Proceedings of the 48th Annual Conference of the Cognitive Science Society*

Alós-Ferrer & Buckenmaier, 2021; Franke & Degen, 2016). Because rules provide a low-cost coordination option, increasing inference difficulty should increase rule use as a fallback—people initiate inference but retreat when it proves too demanding. In addition, increasing the payoff of the rule-based option should increase rule use as a default, even when inference is trivial, because the expected gain may not justify the fixed cost of entering inference mode.

We model this distinction with a two-layer account. At a policy layer, people decide whether to enter inference mode or default to the rule. This reflects strategy selection: a meta-level judgment that the expected value of inference does not justify its setup cost. At a local layer, conditional on entering inference mode, people allocate effort according to the demands of the current problem (Gill & Prowse, 2023). If inference proves too costly or uncertain, they may retreat to the rule as a fallback. Formally, the cost of inference can be decomposed as:

$$\mathrm{Cost(L)} = c_0 + \sum_{i=1}^{L} c_i$$

where $c_0 > 0$ is a fixed setup cost (e.g., maintaining a representation of the other agent, keeping it active in working memory, mapping observations to latent goals, etc.) and c_i is an incremental cost that increases with the required recursive depth (Alaoui & Penta, 2016). This structure distinguishes two sources of rule-based coordination: policy-level defaults, promoted when the rule-based payoff approaches the inference payoff (reducing the incentive to pay c_0), and local fallbacks, promoted when inference demands exceed what the agent is willing or able to compute (when Σc_i grows large). In particular, defaults are expected to be comparatively cheap and stable across contexts, whereas fallbacks should be slower, as they follow attempted inference.

The Mining Game

We test these predictions in a coordination game framed as a collaborative mining task. On each trial, participants see three mines containing gems. Each mine also contains silver, which functions as the rule-based option (Figure 1). A simulated partner enters one mine first, signaling what they intend to mine. Silver offers a context-independent coordination solution: its feasibility does not depend on which mine the partner enters or which gems appear, so it does not require inferring the partner's intention. Silver is not a safe option in the standard decision-theoretic sense: it pays off only if both players choose it, so its value depends on mutual recognition as a coordination solution, not on reducing variance in individual outcomes.

After observing the partner's choice, participants choose what to mine. Coordinating as on the same gem yields 100 points, while coordinating on silver yields a lower but variable payoff. Miscoordination yields 0 points. Critically, the task allows us to manipulate (i) the inference depth required to identify the target gem from the partner's mine entry and the full context of mines (0-, 1-, 2-order problems), and (ii) the incentive to engage in reasoning by varying the value of the rule-based option.

We manipulated the reasoning demands by varying the ambiguity of the partner's signal to the target gem. The design draws on reference game logic and iterated reasoning models (Franke & Degen, 2016). A hierarchy of reasoning depths defines the complexity required to identify the target gem. A level-0 reasoner treats the mine entry "literally": they consider only the gems present in the chosen mine and select among them with equal probability. A level-1 reasoner assumes the partner chose their mine to be informative: if the partner had intended a gem available in multiple mines, they would have entered an unambiguous mine. When the partner enters an ambiguous mine, the level-1 reasoner infers the target is the gem that could not have been signaled more clearly elsewhere. A level-2 reasoner applies a further recursive step, assuming the partner anticipated a level-1 interpretation and selected their mine accordingly. This hierarchy generates different predictions across problem types. In 0-order problems, even a literal reasoner succeeds because the target gem is the only gem in the chosen mine and appears only in it. In 1-order problems, success requires the first-order pragmatic inference. In 2-order problems, success requires an additional recursive step. Table 1 presents the structure of each problem type along with predicted accuracy at each reasoning level.

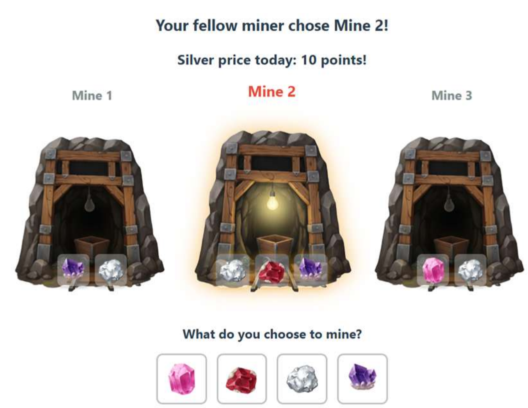

Figure 1: Example of 1-order problem. Level-1 and Level-2 reasoners would pick the red gem, reasoning that if their partner wanted the purple gem, they would have entered Mine 1. A rule-follower would choose silver even when its expected value is lower than guessing randomly between gems.

If rules serve as a fallback, rule choices should increase with problem order as inference becomes more demanding (0-order < 1-order < 2-order). To the extent that rules serve as a default, higher silver value should increase rule choices even on 0-order problems where inference is trivially easy—because the expected gain from inference does not justify entering inference mode. Response times should differentiate between the two: defaults should be fast, while fallbacks should be slower, reflecting attempted inference. Furthermore, RTs increase with problem order should be steeper when silver value is low, reflecting greater investment in inference when the rule-based alternative is less attractive.

Table 1: Problem structure and predicted accuracy for ad hoc coordination

Problem type	Mine 1	Mine 2	Mine 3	P(A | L0)	P(A | L1)	P(A | L2)
0-order	A, S	B, S	C, S	1	1	1
1-order	A, B, S	B, S	C, S	.5	.67	1
2-order	A, B, S	B, S	C, A, S	.5	.5	1
Silver-only	S	B, C	C, A, S	0	0	0
Intractable	A, B, S	B, C, S	C, A, S	.5	.5	.5

Note: Gems are labeled A, B, and C; S denotes silver (the rule-based option), which appears in every mine. The partner always enters Mine 1 (whose screen position is randomized), and the target gem is always A. Columns show the probability of correctly identifying A at each reasoning level, assuming the partner chooses among gems only. An L0 reasoner selects uniformly among the gems in the chosen mine. An L1 reasoner applies Bayesian inference, computing P(gem | chosen mine) ∝ P(partner chooses this mine | gem). An L2 reasoner assumes the partner selected the mine that maximizes an L1 reasoner's accuracy.

Methods

The sample size, methods, exclusion criteria, and main analyses were pre-registered. The dataset, study code, and analysis scripts are available on the project's OSF page.[1]

Participants We recruited 150 participants from Prolific. Participants were excluded if they failed two or more of the eight Silver-only trials (i.e., chose a gem when the partner's mine contained only silver). Five participants met this criterion and were excluded, leaving a final sample of 145 participants (84 women, 55 men, 6 nonbinaries; M_{age} = 29.00, *SD* = 4.13).

Procedure Participants were instructed that their partner preferred to coordinate on a gem and would choose silver only if they believed the participant could not identify which gem to mine together.

Participants completed 20 trials at each problem type (0-/1-/2-order). Within each problem, silver was worth 10 points on half the trials and 90 points on the other half. We refer to trials where silver was worth 10 points as the high reasoning incentive condition (low silver value makes inference more attractive) and trials where silver was worth 90 points as the low reasoning incentive condition (high silver value makes the rule more attractive).

In addition, participants completed eight Silver-only trials, in which the partner entered a mine containing only silver, making their intention transparent, and eight Intractable trials, in which no level of reasoning could uniquely identify a target gem.

On each trial, we randomly sampled a set of distinct gems from a fixed pool and instantiated the problem structures using those gems. Trial order was randomized, and the left/middle/right positions of the mines were randomized on each trial. Each trial began with a 1 s screen displaying the current day (e.g., "Day 1"). The following 3 s screen presented the three mines and their contents in addition to the announcement of the value of silver for that day. An interstitial screen then appeared for 4 s with the text: "Based on these options, your fellow miner will choose the mine they think best reveals what they are going to mine." This interstitial was included to ensure participants interpreted the partner's mine entry as a deliberate signal. On the choice screen, the partner's chosen mine was highlighted, and an announcement stated the partner's choice (e.g., "Your fellow miner chose Mine 2!"). Participants then selected what to mine from the available options and submitted their choice. There was no time limit for responding. Response time was measured from choice screen onset to submission. Participants did not receive trial-by-trial feedback about outcomes or the partner's final choice; only their cumulative score was shown at the end of the experiment.

Results

All analyses reported below used only the inference trials (0-, 1-, and 2-order), excluding Silver-only and Intractable trials.

Choice Behavior To test whether rule choices increase with reasoning demands and decrease with reasoning incentive, we fit a mixed-effects logistic regression predicting rule-based (silver) choice from problem order (coded: 0, 1, 2), reasoning incentive (effect-coded), and their interaction, with random intercepts and slopes for both predictors by participant.

As predicted, participants chose silver more often as problem order increased (β = 1.04, *SE* = 0.20, z = 5.18, p < .001) and when the incentive to reason was low (β = -4.79, *SE* = 0.41, z = −11.55, p < .001; Figure 2). Reasoning

[1] https://aspredicted.org/9t3sc4.pdf
https://osf.io/avbh8/

incentive interacted with problem order, such that the effect of problem order on silver choice was stronger when the incentive to reason was low, i.e., when silver was more attractive (β = -0.31, SE = 0.14, z = -2.11, p = .035). However, this interaction was not robust when problem order was treated as a categorical predictor ($\chi^2(2) = 0.61$, p = .74).

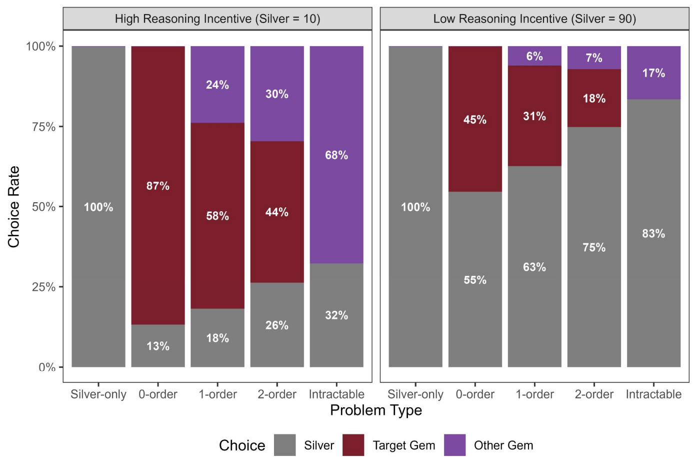

Figure 2: Choice rates by problem type and reasoning incentive.

Response Times RTs were log-transformed, and trials with RT ≥ 30,000 ms were excluded (pre-registered). We fit a linear mixed-effects model predicting log RT from problem order, reasoning incentive, and their interaction, with random intercepts and slopes of both predictors by participant. Participants responded more slowly as problem order increased (β = 0.13, SE = 0.01, $t(144.3)$ = 10.9, p < .001). There was no overall difference in RT between reasoning incentive conditions (β = 0.01, SE = 0.02, $t(399.1)$ = 0.25, p = .81). However, reasoning incentive moderated the effect of problem order: the increase in RT across problem orders was larger when the incentive to reason was high (β = 0.08, SE = 0.01, $t(8122)$ = 6.26, p < .001; Figure 3A).

Because problem order also shifts the distribution of response types (e.g., silver vs. gem choices), we repeated the analysis including choice type (Silver / Target gem / Other gem) as a fixed effect. The interaction between problem order and reasoning incentive remained significant (β = 0.07, SE = 0.01, $t(8149)$ = 5.50, p < .001), indicating that the effect is not solely due to shifts in response type across problem orders.

Response Time of Rule-based Choices We next examined whether defaults and fallbacks leave distinct RT signatures among rule-based choices specifically. When silver choices reflect defaults, RTs should be invariant to problem order; when silver choices reflect fallbacks, RTs should increase with order. These processes should contribute differentially depending on reasoning incentive: defaults should predominate when reasoning incentive is low (silver is valuable), fallbacks when reasoning incentive is high (silver is cheap).

We fit a linear mixed-effects model predicting log RT on trials where participants chose silver, from problem order, reasoning incentive, and their interaction, with random intercepts and slopes by participant. The RT increase across

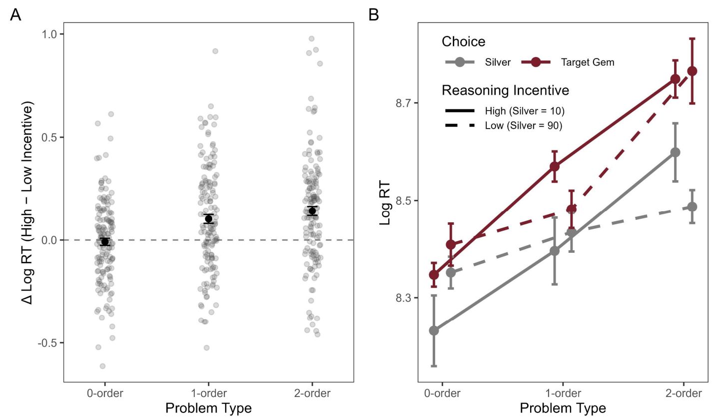

Figure 3: (A) Difference in log RT between reasoning incentive conditions (high − low) by problem type. Each point represents one participant. (B) Mean log RT by problem order, and reasoning incentive condition for silver (rule-based) and gem choices. Error bars show ±1 SE.

problem orders was steeper when reasoning incentive was high than when it was low (β = 0.09, SE = 0.03, t = 3.64, p < .001; Figure 3B), consistent with fallbacks predominating when silver is unattractive and defaults when it is attractive. This interaction was specific to silver choices; the same analysis on target-gem choices showed no interaction (β = −0.01, p = .72), suggesting that successful inference follows a similar process regardless of reasoning incentive.

Reasoning Depth To test whether individuals shifted reasoning depth across incentive conditions, we classified each participant separately within each reasoning incentive condition using their choices across problem orders.

For each participant in each reasoning condition, we counted target-gem choices out of 10 trials at each order (0-, 1-, and 2-order). Classification followed a hierarchical gating rule: participants first had to meet the criterion at 0-order to be eligible for higher classifications. Participants who met the criterion at 0-order were then evaluated at 1-order, and those who met the criterion at 1-order were then evaluated at 2-order.

We used a criterion of ≥ 8/10 target-gem choices at a given order (consistent with one-tailed binomial p < .055). Participants who failed at a gate were assigned to the highest lower level passed; if they failed already at 0-order, they were classified as non-reasoners.[2] Table 2 shows transitions in classification across reasoning incentive conditions.

A cumulative link mixed model predicting ordered classification from incentive condition with a by-participant random intercept found that participants were more likely to be classified at higher reasoning levels in the high-incentive condition than in the low-incentive condition ($OR = 11.06$, $p < .001$).

Table 2: Classification transition matrix

High Reasoning Incentive	Low Reasoning Incentive: NR	L0	L1	L2
NR	29	0	0	0
L0	47	23	1	0
L1	9	3	12	1
L2	10	2	2	6

Note: NR = non-reasoner. Rows indicate classification under high reasoning incentive (Silver = 10); columns indicate classification under low reasoning incentive (Silver = 90).

Response Times by Reasoning Depth Figure 4 presents mean log RT across problem order by reasoning classification. To test whether the effect of problem order on RT differed by classification, we fit a linear mixed-effects model predicting log RT from problem order, reasoning classification, reasoning incentive, and all interactions, with by-participant random intercepts and random slopes for problem order and reasoning incentive. We then estimated the marginal slopes (β) of log RT over problem order within each classification, averaging over reasoning incentive conditions for this summary.

Non-reasoners and L0 reasoners showed relatively shallow increases in RT with problem order (Non-reasoners: $\beta = 0.07$, $SE = 0.01$, $p < .001$; L0: $\beta = 0.09$, $SE = 0.01$, $p < .001$). This pattern is consistent with rule-following as a default policy: these participants appear to bypass inference rather than attempt and abandon it. In contrast, L1 and L2 reasoners exhibited steeper increases (L1: $\beta = 0.30$, $SE = 0.02$, $p < .001$; L2: $\beta = 0.32$, $SE = 0.03$, $p < .001$). Notably, L1 reasoners showed RT increases comparable to L2 reasoners, even though many of their 2-order responses were silver choices. This suggests that their rule-based choices on difficult problems reflect fallbacks after attempted inference. Pairwise comparisons showed that L1 and L2 slopes were significantly larger than those of non-reasoners and L0 reasoners (all Tukey-adjusted $ps < .001$), whereas slopes did not differ between non-reasoners and L0 reasoners ($p = .918$) or between L1 and L2 reasoners ($p = .915$).

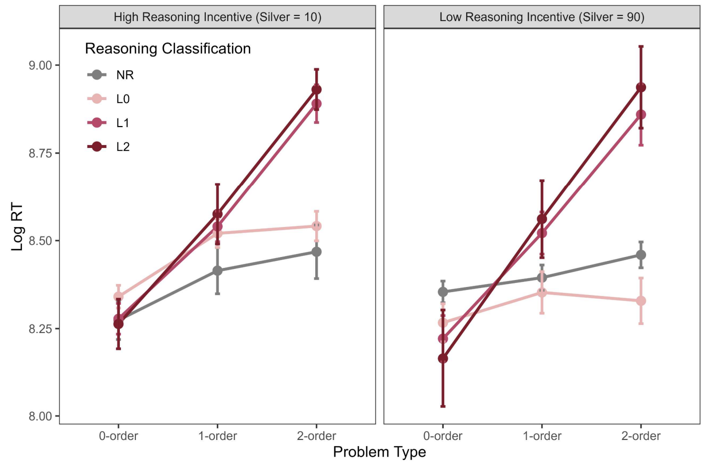

Figure 4: Mean log RT by problem type and reasoning classification, shown separately for each reasoning incentive condition. Participants were classified independently within each condition based on their choice patterns; thus, the same participant may appear in different classification groups across panels. Error bars show ±1 SE.

Discussion

Two factors determined when participants coordinated by following a rule, and when they instead coordinated by inferring their partner's intention: the depth of inference required, and the value of the rule-based option. Rule choices increased with problem order, consistent with rule-following as a fallback when inference becomes too cognitively demanding. Rule choices also increased when reasoning incentive was low, consistent with rule-following as a default. Put simply, we find that participants used a cost-benefit calculus to decide whether, and to what extent, to achieve social coordination by *ad hoc* reasoning versus simple rules.

Response times provide more direct leverage on the underlying processes. Across all trials, RTs increased with problem order, indicating that higher-order problems engaged more demanding computation. Critically, this increase was steeper when reasoning incentive was high than when it was low. Among silver choices specifically, RTs increased more steeply with problem order under high incentive than under low incentive. This supports the interpretation that the same overt choice arose from different processes: fast defaults when the expected gain from inference was low, slower retreats when it was high. RTs alone would be consistent with general decision conflict; the individual classification analysis disambiguates default from fallback. At the individual level, participants classified as L1 or L2 reasoners showed steep RT increases with problem

[2] In addition, when a participant failed at a gate, we labeled whether their non-target responses were primarily rule-based (silver) or primarily other gems, based on whether they chose silver more often than other gems at that order. For Table 2, these subtypes were collapsed into the ordered levels Non-reasoner, L0, L1, L2; we return to this distinction below.

order, consistent with engaging recursive inference and paying its costs. Non-reasoners and L0 participants showed flat slopes, consistent with bypassing reasoning altogether. Critically, participants shifted toward higher classifications under high incentive, indicating that reasoning depth responds to expected returns of reasoning.

Models of recursive strategic reasoning have identified several factors that shape how deeply people reason about partners, including baseline cognitive ability (Camerer et al., 2004), beliefs about partner sophistication and strategy (Stahl & Wilson, 1995), as well as incentive magnitude (Alaoui & Penta, 2016). Because incentives were manipulated within participants in our study, individual differences in ability cannot explain the observed shifts in rule use or classification. Moreover, our instructions fixed beliefs about the partner's preferences and strategy: the partner prefers gems, signals mines informatively, and retreats to silver when no mine reliably indicates a gem. The absence of feedback prevented trial-by-trial updating. With ability and beliefs held constant, the observed shifts reflect how participants allocated effort to *ad hoc* relative to rule-based coordination.

We interpret these findings in terms of adaptive effort allocation governed by cost-benefit logic. Resource-rational accounts have addressed how individuals select strategies for non-social problems, trading off costly deliberation against simpler heuristics when the expected benefit justifies the cost (Lieder & Griffiths, 2020). We extend this logic to coordination, where the trade-off is between cognitively demanding inference about a partner's intention and a rule that both parties recognize as an acceptable solution. In many real-world contexts, the cognitively cheap, rule-based solutions to coordination problems are often established by socialization, explicit agreement, or culturally evolved norms.

This framing suggests a complementary perspective on why social rules and conventions exist. Rules may arise from cognitive constraints: they allow people to coordinate complex collective behavior when distributed inference would be too costly. Cultural evolution may therefore favor rules especially in those situations where the cognitive demands of *ad hoc* reasoning are particularly great, or the risk of social miscoordination is particularly acute. This may also bear on when rules are openly justified versus treated as opaque conventions (Henrich, 2015): rules whose coordination logic is hard to reconstruct may persist as taboos, while rules with transparent coordination value invite consequentialist scrutiny and revision.

For rules to promote coordination, people must conform to them in tandem. Yet, our data reveals heterogeneity in how people respond when inference proves too costly: Not all participants fell back to the socially-agreed rule. Some retreated to idiosyncratic strategies, such as picking a gem at random. Notably, their response times resembled those of participants who fell back to silver, suggesting a similar underlying process of attempted inference followed by retreat, but with different target strategies. This heterogeneity matters because only the shared rule actually solves the coordination problem; an idiosyncratic fallback abandons the mutual recognition that makes coordination possible. This pattern may also explain why non-reasoners and L0 participants showed similar RT profiles despite different choice patterns: some individuals classified as L0 may not have been reasoning at all, but defaulting to an individual strategy (i.e., don't pick silver) without attempting inference. What determines whether people conform to shared rules versus retreat to idiosyncratic strategies remains an open question.

Of course, real-world social coordination differs from our paradigm in several important ways. Our paradigm fixed the partner's strategy via instructions; in real coordination, the partner model must be built and revised on the fly through repeated interaction and feedback. This gives people richer information about a partner's preferences and intentions, but it also leaves that information noisier and more uncertain — which should make the retreat to rules more frequent. In addition, with feedback, people may learn how reliably informative their partner is: those who discover that their partner is reliable should increase their investment in inference; those who discover inconsistency should retreat toward rule-based coordination (Frank & Goodman, 2012).

We also imposed no competing demands on attention or time. Cognitive load reduces pragmatic inference in language comprehension (Ryzhova & Demberg, 2020), and similar effects should emerge in coordination. Time pressure or secondary tasks should inflate inference costs, particularly the fixed cost of maintaining a representation of the partner's goals. Under load, even easy inference may not be worth initiating, predicting more policy-level defaults to rules on problems where inference would otherwise succeed.

Overall, our findings indicate that rules and intention inference are substitutable coordination strategies governed by the same resource-rational logic. People are neither rigid rule-followers nor tireless Gricean reasoners. They allocate cognitive effort adaptively, entering inference mode when the expected benefit justifies its cost and coordinating via a mutually recognizable rule when it does not.

Acknowledgments

We used Google Gemini to generate the visual stimuli used in the Mining Game paradigm. We also used Claude (Anthropic) to assist with JavaScript code implementing of the experimental task. FC was supported by an SBE-UKRI Lead Agency grant from the U.S. National Science Foundation (Award No. 2417241).

References

Alaoui, L., & Penta, A. (2016). Endogenous depth of reasoning. *The Review of Economic Studies*, *83*(4), 1297–1333. https://doi.org/10.1093/restud/rdv052

Alós-Ferrer, C., & Buckenmaier, J. (2021). Cognitive sophistication and deliberation times. *Experimental*

Economics, *24*(2), 558–592. https://doi.org/10.1007/s10683-020-09672-w

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

Frank, M. C., & Goodman, N. D. (2012). Predicting pragmatic reasoning in language games. *Science*, *336*(6084), 998–998. https://doi.org/10.1126/science.1218633

Franke, M., & Degen, J. (2016). Reasoning in reference games: Individual- vs. population-level probabilistic modeling. *PLOS ONE*, *11*(5), e0154854. https://doi.org/10.1371/journal.pone.0154854

Gächter, S., Molleman, L., & Nosenzo, D. (2025). Why people follow rules. *Nature Human Behaviour*, *9*(7), 1342–1354. https://doi.org/10.1038/s41562-025-02196-4

Gill, D., & Prowse, V. (2023). Strategic complexity and the value of thinking. *The Economic Journal*, *133*(650), 761–786. https://doi.org/10.1093/ej/ueac070

Goodman, N. D., & Frank, M. C. (2016). Pragmatic language interpretation as probabilistic inference. *Trends in Cognitive Sciences*, *20*(11), 818–829. https://doi.org/10.1016/j.tics.2016.08.005

Hawkins, R. X. D., Goodman, N. D., & Goldstone, R. L. (2019). The emergence of social norms and conventions. *Trends in Cognitive Sciences*, *23*(2), 158–169. https://doi.org/10.1016/j.tics.2018.11.003

Henrich, J. P. (2015). *The secret of our success: How culture is driving human evolution, domesticating our species, and making us smarter*. Princeton University Press.

Le Pargneux, A., Levine, S., Tenenbaum, J. B., & Cushman, F. (2025). The trade-off between rule-based thinking and mutual benefit in tacit coordination. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *47*(0). https://escholarship.org/uc/item/4rg5b61f

Lieder, F., & Griffiths, T. L. (2017). Strategy selection as rational metareasoning. *Psychological Review*, *124*(6), 762–794. https://doi.org/10.1037/rev0000075

Lieder, F., & Griffiths, T. L. (2020). Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behavioral and Brain Sciences*, *43*, e1. https://doi.org/10.1017/S0140525X1900061X

Ryzhova, M., & Demberg, V. (2020). Processing particularized pragmatic inferences under load. *Proceedings of the Annual Meeting of the Cognitive Science Society*, *42*(0). https://escholarship.org/uc/item/31g0093d

Schelling, T. C. (1960). *The strategy of conflict*. Harvard University Press.

Shenhav, A., Botvinick, M. M., & Cohen, J. D. (2013). The expected value of control: An integrative theory of anterior cingulate cortex function. *Neuron*, *79*(2), 217–240. https://doi.org/10.1016/j.neuron.2013.07.007

Spiliopoulos, L., & Hertwig, R. (2020). A map of ecologically rational heuristics for uncertain strategic worlds. *Psychological Review*, *127*(2), 245–280. https://doi.org/10.1037/rev0000171

Stahl, D. O., & Wilson, P. W. (1995). On players' models of other players: Theory and experimental evidence. *Games and Economic Behavior*, *10*(1), 218–254. https://doi.org/10.1006/game.1995.1031

Thomas, K. A., DeScioli, P., Haque, O. S., & Pinker, S. (2014). The psychology of coordination and common knowledge. *Journal of Personality and Social Psychology*, *107*(4), 657–676. https://doi.org/10.1037/a0037037