Informational Breadth, Epistemic Depth, and Fragility

Swaminathan Balasubramaniam*

April 15, 2026

Abstract

We study coordination when information can spread widely without becoming mutually understood. In a two-state game, a higher probability of being informed is not unambiguously stabilizing: too little information induces panic by uninformed agents, while too much allows informed agents in fragile states to coordinate on attack. Greater epistemic depth relaxes this informed-agent constraint, enlarging the stable region and supporting stability under more fragile fundamentals. In a dynamic extension with fixed depth and rising information diffusion, early attack is driven by uninformed agents and later attack by informed agents. Greater depth delays both through strategic coupling.

JEL Classification: D82, D83

Keywords: Higher-order beliefs, Coordination games, Global games, Rationalizability

*NEOMA Business School, Rouen, France; s.balasubramaniam@neoma-bs.fr. This paper is based on the second chapter of my doctoral dissertation submitted to Washington University in St.Louis in 2021. I would like to thank Jonathan Weinstein for his invaluable advice throughout this project. I would especially like to thank two anonymous referees whose detailed suggestions contributed significantly to revising the paper in its current form. I would also like to thank Brian Rogers, Andres Perea, Elias Tsakas, Phil Dybvig and seminar participants at Epicenter, Research Center for Epistemic Game Theory at the Maastricht University, Midwest Economic Theory seminar, Washington University and the Research seminar at Olin Business School, Washington University for helpful comments. All errors are mine.

1 Introduction

Modern communication technologies spread information with remarkable speed, yet rapid dissemination does not by itself generate mutual knowledge: knowing that others know, and that others know that others know. Recent evidence suggests that digital platforms have fundamentally altered crisis dynamics by accelerating the spread of information,[1] but the theoretical mechanism by which broader communication can fail to stabilize coordination remains underexplored. We argue that the answer lies in a distinction between two margins of the communication environment: *informational breadth*, the probability that an agent receives direct exposure to the underlying information, and *epistemic depth*, the extent to which the communication environment sustains mutual knowledge through successive layers of confirmation. A communication regime may therefore generate broad exposure without generating much epistemic depth.

We study a two-player coordination game in the Morris-Shin tradition, in which a financial institution is either in a safe state, where *No Attack* is strictly dominant, or in a fragile but potentially solvent state, where outcomes depend on coordinated expectations. Under prior uncertainty, both actions remain rationalizable,[2] so the communication environment determines whether the economy lies in a *stable region*, where *No Attack* is uniquely rationalizable in both states, or in a *fragile region*, where multiplicity remains.

Our starting point is a benchmark in which each agent is directly informed of the realized state with probability p, and otherwise remains uninformed and knows only the prior. Our first result shows that this margin is not unambiguously stabilizing. When p is too low, uninformed agents remain insufficiently reassured and panic is supportable. When p is too high, agents who observe the fragile state assign too much probability to the opponent also being informed, so coordinated attack again becomes supportable. *No Attack* is therefore uniquely rationalizable only for an intermediate range of breadth. The basic static tension is that broader dissemination relaxes the panic problem of uninformed agents while simultaneously tightening the coordination problem of informed agents in the fragile state.

We then introduce epistemic depth as a second, distinct margin. After becoming informed, an agent may send confirmation messages of the form "I am informed," each of which fails with some

[1]See Cookson et al. [2026] for empirical evidence on the role of social media in modern bank runs. This phenomenon is also central to political mobilization and regime change games; see, for example, Chwe [2000], Edmond [2013], and Barberà and Jackson [2020].

[2]An action is rationalizable if it survives iterated elimination of actions that are never a best response to any belief consistent with the opponent's rationalizability [Bernheim, 1984, Pearce, 1984]. Unique rationalizability means that this process leaves a single surviving action. This solution concept is natural here because the paper's mechanism operates through infection-style logic, in which dominance propagates through overlapping information sets, rather than through an equilibrium selection rule imposed inside the fragile region; see Morris et al. [1995].

probability. The parameter k indexes how many rounds of such confirmation the communication environment can sustain before agents act. This construction yields a belief potential for an informed agent in the fragile state: the probability that the opponent remains on the dominance-anchored branch. Greater epistemic depth raises that belief potential. The key static asymmetry is that depth relaxes only the informed-agent constraint. It does not affect the uninformed panic condition, which depends only on the first-order probability that the opponent directly observed the dominant state. As a result, greater epistemic depth widens the stable interval in breadth only from above.

This asymmetry also has an immediate implication for fundamentals. Because depth relaxes the informed-agent constraint while leaving the panic condition unchanged, it enlarges the set of fragile environments for which the economy can still be placed in the stable region. The maximum level of fundamental fragility compatible with unique rationalizability rises strictly with depth and converges to a sharp upper bound. But this force cannot be arbitrarily strong. A separate general result in the paper shows that any admissible ex ante structure that robustly implements the safe action must confront a pooled bottleneck type at which the belief potential of the seed states cannot exceed their posterior mass. In the two-state environment, that logic yields the Rubinstein frontier as a hard upper bound on implementability. Finite depth therefore expands the stable region, but only up to a limit.

The paper then embeds this static logic in a dynamic environment in which epistemic depth is fixed as a primitive of the communication platform, while informational breadth rises deterministically over time as direct exposure spreads. This yields a timing problem rather than a static classification. Early in the diffusion process, attack is driven by panic among still-uninformed agents. Later in the diffusion process, attack is driven by directly informed agents coordinating on fragile-state information. Between these two forces lies a waiting phase. The dynamic model therefore provides a temporal analogue of the static intermediate-informativeness result.

The dynamic extension also delivers a coupling implication absent from the static benchmark. In the static model, depth affects only the informed-fragile margin directly. In the dynamic model, depth again delays attack directly for informed agents in the fragile state, but this delay also feeds back onto uninformed agents because their stopping problem depends on the future behavior of informed agents. Greater depth therefore raises the continuation value of waiting, shifts both attack thresholds outward, and widens the waiting region. This, in turn, yields a dynamic fragility frontier: the set of fragile environments for which an initial waiting phase can be sustained is increasing in epistemic depth.

The mechanism is especially relevant in digital environments, where broad exposure can spread rapidly even when the platform itself supports only limited higher-order mutual knowledge. More broadly, the paper highlights a distinction between the reach of communication and the depth of mutual awareness it can sustain. Communication regimes that broaden exposure without providing sufficient epistemic depth may intensify fragility, whereas environments that support deeper mutual knowledge can stabilize coordination over a wider range of fundamentals and over a longer interval of the diffusion process.

Our paper is closest in spirit to the literature on higher-order uncertainty in coordination games, beginning with Rubinstein [1989], Carlsson and Van Damme [1993], and Morris et al. [1995]. That literature shows that strategic outcomes can hinge on small departures from common knowledge and on the way iterated reasoning propagates through agents' beliefs about one another's beliefs.[3] Our paper uses this logic in a different way. Rather than studying perturbations around common knowledge or equilibrium selection under vanishing noise, we use higher-order reasoning to characterize when the safe action is uniquely rationalizable, how epistemic depth expands the feasible region, and where that expansion stops.

Our paper is also related to the coordination, bank-run, and transparency literatures on how information shapes strategic behavior [Morris and Shin, 2001, Goldstein and Pauzner, 2005, Corsetti et al., 2004, Angeletos and Werning, 2006, Goldstein and Sapra, 2014, Bouvard et al., 2015]. Our difference is that wider dissemination is non-monotone: it first reassures uninformed agents, but beyond a point it facilitates coordination on attack by informed agents unless it is supported by sufficient higher-order confirmation. In the dynamic model, the same logic appears as early panic-driven attack by uninformed agents and later information-driven attack by informed agents, with epistemic depth delaying both thresholds.

Finally, the paper is related to the higher-order information-design literature, in which information affects behavior through the full hierarchy of induced beliefs. Mathevet et al. [2020] is the main reference here: it shows how information design in games operates through first-order and higher-order beliefs. Related contributions study informed principals, multi-agent persuasion, manipulation through local obfuscation, and asymmetric information allocation in coordination environments.[4] Our contribution is not to solve the unrestricted design problem, but to provide an applied way to think about higher-order information design by isolating two economically transparent margins of the communication environment: breadth and epistemic depth. Within that

[3]See also Monderer and Samet [1989], Geanakoplos [1992], Kajii and Morris [1997], Dekel et al. [2005], and Weinstein and Yildiz [2007] for foundational treatments of approximate common knowledge, robustness, and the topology of higher-order beliefs.

[4]See Koessler and Skreta [2023] on informed information design, Hoshino [2022] on multi-agent persuasion, Li et al. [2023] on global manipulation through local obfuscation, and Moriya and Yamashita [2020] on asymmetric information allocation to avoid coordination failure. Friedenberg [2019] studies higher-order optimism in bargaining, and Morris et al. [2024] provides general implementation results for binary-action supermodular games.

restricted but interpretable class, we obtain explicit comparative statics, a sharp implementability frontier, and dynamic threshold results that are hard to derive in more general higher-order design environments.

The remainder of the paper proceeds as follows. Section 2 develops the static framework of coordination under strategic uncertainty. Section 3 studies the dynamic environment with deterministic diffusion and fixed epistemic depth. Section 4 concludes. All proofs are presented in the Appendix.

2 Coordination under Strategic Uncertainty

There are two payoff-relevant states, $\Omega = \{l, m\}$, with $\Pr(\omega = m) = q \in (0, 1)$. Two agents simultaneously choose between *Attack* and *No Attack*. Payoffs follow Morris and Shin [2001]:

	Attack	No Attack
Attack	θ	$\theta - 1$
No Attack	0	0

where $\theta \in \{\theta_l, \theta_m\}$. We maintain $\theta_l < 0 < \theta_m < 1$, so that *No Attack* is strictly dominant in state l, while both actions remain rationalizable in state m. We also assume

$$\mathbb{E}[\theta] := (1 - q)\theta_l + q\theta_m > 0,$$

so that without informational structure the economy lies in the fragile region even when fundamentals are favorable.

Our solution concept is rationalizability [Bernheim, 1984, Pearce, 1984]. We say that an information environment lies in the *stable region* if *No Attack* is uniquely rationalizable in both states, and in the *fragile region* if multiplicity remains. The mechanism throughout is an infection argument [Rubinstein, 1989, Morris et al., 1995]: dominance in the safe state l propagates through overlapping information sets and may, depending on the information architecture, pin down *No Attack* even in the fragile state m.

We begin from a Fine/Coarse benchmark. Each agent i is *Fine* with probability p_i and *Coarse* with probability $1 - p_i$, independently across agents and independently of the state. A Fine agent observes the realized state; a Coarse agent knows only the prior. In the symmetric case, $p_1 = p_2 = p$. The parameter p is the *breadth* of the information environment: it captures how widely the initial signal is disseminated. This benchmark already contains the basic tension in the model. Raising p reassures Coarse agents by making it more likely that the opponent directly observes the dominant state l, but it also makes a Fine agent in state m more likely to expect the opponent to be informed as well.

2.1 Intermediate Informativeness

We begin with the benchmark information structure. In this case, each agent knows only its own type, and a Fine agent who observes state m assigns probability $1 - p$ to the opponent being Coarse and therefore pinned to *No Attack* by dominance. Breadth is therefore not unambiguously stabilizing. Greater breadth relaxes the Coarse panic constraint, but it simultaneously tightens the Fine-m coordination constraint. Stability requires intermediate informativeness.

It is convenient to represent the benchmark on the augmented state space $\{F, C\}^2 \times \{l, m\}$, where a generic element $T_1T_2\omega$ records the two information types and the payoff state. Since types and states are drawn independently,

$$\Pr(T_1T_2\omega) = \Pr(T_1)\Pr(T_2)\mu(\omega).$$

Player 1 observes its own type and, if Fine, the realized state. Its partition is therefore

$$\mathcal{P}_1 = \big\{\underbrace{\{FFl, FCl\}}_{\text{Fine-}l}, \underbrace{\{CFl, CCl, CCm, CFm\}}_{\text{Coarse}}, \underbrace{\{FCm, FFm\}}_{\text{Fine-}m}\big\},$$

and symmetrically for player 2. The Fine-l cell is the seed event: there *No Attack* is strictly dominant, and the infection argument asks how far that dominance anchor propagates.

Proposition 1 (Intermediate informativeness)**.** *In the benchmark Fine/Coarse architecture,* No Attack *is uniquely rationalizable in both states if and only if*

$$\frac{\mathbb{E}[\theta]}{1-q} \;\le\; p \;\le\; 1 - \theta_m. \tag{1}$$

This interval is nonempty if and only if $q \le \bar{q}_1 := (1 - \theta_m - \theta_l)/(1 - \theta_l)$.

The lower bound in (1) is the Coarse panic constraint. A Coarse agent refrains from attacking only if the probability that the opponent is already pinned to *No Attack* is large enough, and in the benchmark that event is precisely that the opponent is Fine in state l, which occurs with probability $(1 - q)p$. The upper bound is the Fine-m coordination constraint. A Fine-m agent refrains only if the opponent is sufficiently likely to be Coarse and hence safely anchored by dominance; in the

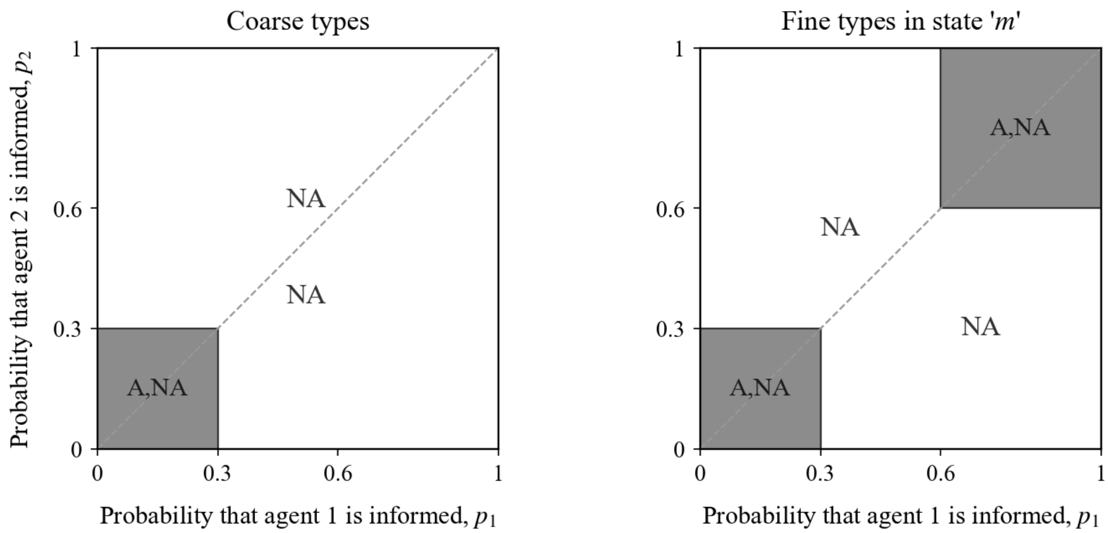

Figure 1: Rationalizable strategies as a function of (p_1, p_2). Parameters: $q = 0.5$, $\theta_m = 0.4$, $\theta_l = -0.1$. Left: Coarse types. Right: Fine types in state m. Dark: both actions rationalizable. Light: *No Attack* uniquely rationalizable. Dashed: symmetric locus $p_1 = p_2$.

benchmark, that probability is $1 - p$. Stability therefore requires that breadth be high enough to reassure Coarse agents but not so high that informed agents in the fragile state expect coordinated attack to be supportable.

Figure 1 illustrates these two margins in the (p_1, p_2) plane.

2.2 Epistemic Depth

The benchmark treats each agent's information as fixed at the moment of action. We now enrich the architecture by allowing a finite amount of higher-order confirmation before agents move. This is the source of *epistemic depth*. Let $k \in \mathbb{N}$ index how many layers of mutual knowledge the communication architecture can sustain before agents act. The benchmark analyzed above corresponds to the case in which no confirmation is exchanged beyond the initial Fine/Coarse assignment. Greater depth adds a finite confirmation process on top of that benchmark.

After becoming informed, a Fine agent may send a confirmation signal to the opponent saying, in effect, "I am informed."[5] Each confirmation fails to reach its recipient with probability $\varepsilon \in (0, 1)$, and a failed signal is observationally invisible. Signals are attempted sequentially for up to $k - 1$ rounds. Epistemic depth therefore measures how many layers of mutual knowledge the

[5]This differs from Rubinstein's (1989) email game, where messages communicate the realized state. Here the state is already known to Fine agents; the friction concerns whether agents know that others are informed.

communication architecture can sustain before agents act.

This finite confirmation process generates an augmented partition representation, developed formally in Appendix A. For the main text, what matters is the induced *belief potential*. At the deepest relevant information set under depth k, a Fine-m agent has received no confirmation after $k-1$ rounds. At that node, the opponent may be Coarse, in which case no confirmation was ever sent, or Fine, in which case all $k-1$ confirmations failed. Since the latter event occurs with probability ε^{k-1}, Bayes' rule implies that the Fine-m agent assigns probability

$$r_k(p,\varepsilon) = \frac{1-p}{(1-p)+p\,\varepsilon^{k-1}} \tag{2}$$

to the opponent being on the dominance-anchored branch. As ε^{k-1} shrinks with k, the belief potential $r_k(p,\varepsilon)$ rises.

The key static implication is asymmetric. Breadth affects both margins of the problem: it changes the probability that a Coarse agent assigns to meeting an opponent who is directly informed of the dominant state, and it changes the probability that a Fine-m agent assigns to the opponent being effectively anchored by dominance. Depth, by contrast, operates only through the informed-agent channel. It changes how a Fine-m agent reasons about whether the opponent remains on the dominance-anchored branch, but it does not change the first-order seed probability relevant for Coarse-agent panic.

Proposition 2 (Static effect of epistemic depth)**.** *Fix $k \in \mathbb{N}$ and $\varepsilon \in (0,1)$.* No Attack *is uniquely rationalizable in both states if and only if*

$$p \in \left[\frac{\mathbb{E}[\theta]}{1-q},\ \bar{p}_k(\varepsilon)\right],$$

where

$$\bar{p}_k(\varepsilon) := \frac{1-\theta_m}{(1-\theta_m)+\theta_m\,\varepsilon^{k-1}}. \tag{3}$$

The lower bound $\mathbb{E}[\theta]/(1-q)$ is independent of k, while $\bar{p}_k(\varepsilon)$ is strictly increasing in k.

The asymmetry in Proposition 2 is the central static result. The lower bound does not move with k because the Coarse panic constraint depends only on the seed event that the opponent is Fine and observes state l, an event whose probability is $(1-q)p$ regardless of how many confirmation rounds the architecture permits. The upper bound rises with k because deeper epistemic structure raises the belief potential of Fine-m agents: additional layers of confirmation make it more plausible that the opponent remains on the dominance-anchored branch. Static depth therefore widens the

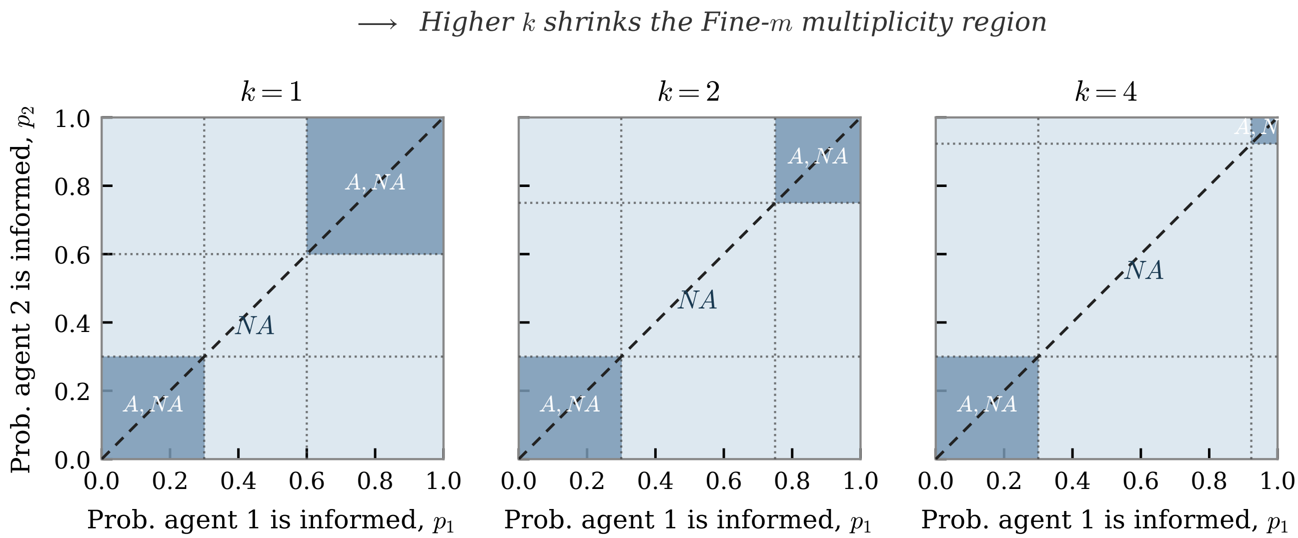

Figure 2: Greater epistemic depth shrinks the Fine-m multiplicity region (upper right) while leaving the Coarse multiplicity region (lower left) unchanged. Parameters: $\theta_l = -0.1$, $\theta_m = 0.4$, $q = 0.5$, $\varepsilon = 0.5$.

uniquely-rationalizable interval only from above. It allows greater breadth before informed agents in state m can sustain multiplicity, but it provides no relief on the Coarse margin.

Equivalently, greater epistemic depth expands the set of breadth values for which *No Attack* is uniquely rationalizable only through the Fine-m upper bound; it does not relax the Coarse panic constraint.

Figure 2 isolates the central static result. As epistemic depth rises, the Fine-m multiplicity region in the upper-right corner shrinks, while the Coarse multiplicity region in the lower-left corner is unchanged.

2.3 Epistemic Depth and the Fragility Frontier

Proposition 2 shows that deeper epistemic structure widens the stable interval in breadth by relaxing only the Fine-m constraint. An immediate implication is that it also enlarges the set of fundamentals for which the economy can be placed in the stable region. In that sense, greater epistemic depth allows *No Attack* to be sustained in more fragile environments, even though the Coarse panic constraint itself is unchanged.

Formally, the stable interval is nonempty if and only if

$$\bar{q}_k(\varepsilon) := \frac{\bar{p}_k(\varepsilon) - \theta_l}{\bar{p}_k(\varepsilon) - \theta_l + \theta_m}, \tag{4}$$

which is strictly increasing in k. The frontier $\bar{q}_k(\varepsilon)$ therefore summarizes, within the finite-depth Fine/Coarse family studied here, the maximum degree of fundamental fragility consistent with unique rationalizability at depth k.

The next proposition shows that this expansion is bounded. As depth rises, the Fine-m constraint is relaxed and the frontier shifts outward, but only up to a hard limit. That limiting logic is developed more generally in Section H, which isolates a pooling bottleneck common to admissible ex ante architectures. The present proposition is the specialization of that logic to the two-state finite-depth environment studied here.

Proposition 3 (Rubinstein frontier)**.** *The frontier $\bar{q}_k(\varepsilon)$ is strictly increasing in k and converges to*

$$\bar{q}_\infty := \frac{1-\theta_l}{(1-\theta_l)+\theta_m}. \tag{5}$$

For every $q < \bar{q}_\infty$, there exists a finite depth k at which No Attack *is uniquely rationalizable. For $q > \bar{q}_\infty$, no depth in the finite-depth Fine/Coarse family achieves the stable region.*

The intuition is straightforward. Since depth raises the upper bound $\bar{p}_k(\varepsilon)$ while leaving the lower Coarse bound $\mathbb{E}[\theta]/(1-q)$ unchanged, the interval of breadth values consistent with unique rationalizability remains nonempty for a wider set of fragile environments as k rises. This is exactly what the frontier $\bar{q}_k(\varepsilon)$ captures. But this force cannot be arbitrarily strong. Section H shows, at a more general level, that robust implementation must confront a pooled type at which the belief potential of the seed cannot exceed the posterior mass of the seed states themselves. In the two-state environment, that bottleneck yields the bound $1-q \geq \mathbb{E}[\theta]$, which is equivalent to $q \leq \bar{q}_\infty$. The proof of Proposition 3 is given in Appendix D. Figure 3 illustrates the convergence, with more reliable confirmation reaching the frontier faster.

Appendix E collects related derived results, including the minimum depth $k^*(q)$ required for a given fragility level, its analogue $k^*(q, p)$ at fixed breadth, the marginal value of one additional epistemic layer, and numerical threshold tables.

3 Dynamic Diffusion, Epistemic Depth, and the Timing of Attack

The static analysis shows that, for any fixed epistemic depth, informational breadth has a non-monotone effect on coordination. When breadth is too low, uninformed agents are insufficiently reassured and panic remains supportable. When breadth is too high, agents who observe the

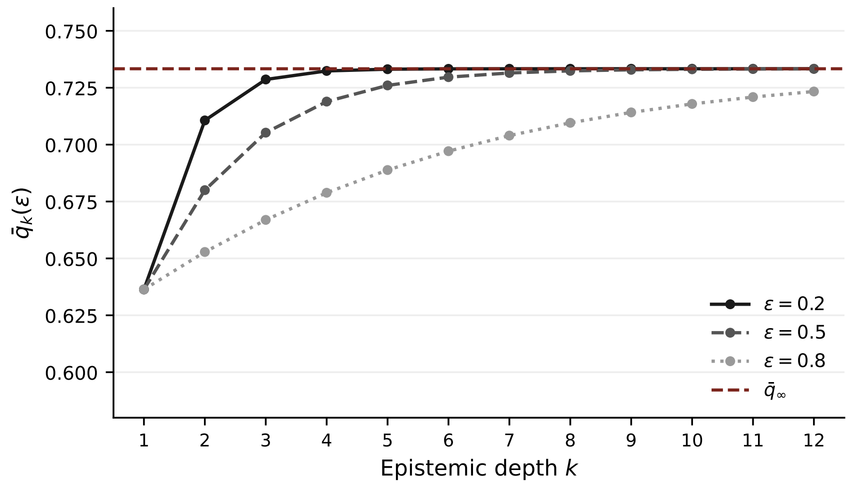

Figure 3: $\bar{q}_k(\varepsilon)$ rises with k and converges to $\bar{q}_\infty$ (dashed). Parameters: $\theta_l = -0.1$, $\theta_m = 0.4$.

fragile state expect the opponent to be similarly informed, and coordinated attack again becomes supportable. Only at intermediate breadth is *No Attack* uniquely rationalizable.

The dynamic analogue keeps epistemic depth fixed as a primitive of the communication platform. At each date, agents may exchange up to $k-1$ rounds of confirmation of the form "I am informed," with each round subject to the same failure probability ε as in the static model. We use a Markov representation in which the current date and the fixed platform depth k summarize the payoff-relevant consequences of communication, so that agents condition only on the current diffusion state and the platform's epistemic capacity when deciding whether to attack or wait.[6]

The primitives of Section 2 are unchanged. States, payoffs, and the finite-depth confirmation structure are exactly as before. The new element is that direct exposure spreads over time. The question is therefore no longer whether a given triplet (q, p, k) lies in the stable region, but when attack becomes optimal along a known path of growing breadth.

For each player i, let $\tau_i^I \in [0, \infty]$ denote the date at which the player becomes directly informed of the realized state. Conditional on the state, exposure dates are independent across players and identically distributed with cumulative distribution function

$$\Pr(\tau_i^I \le t) = p_t, \qquad t \ge 0, \tag{6}$$

where $p_t \in (0,1)$ is deterministic, continuous, and strictly increasing. Thus p_t is the dynamic

[6]This is in the spirit of the broader literature that studies dynamic interaction under finite-memory or Markov communication structures; see, for example, Bhaskar et al. [2013].

counterpart of the static breadth parameter. Write

$$h(t) := \frac{\dot{p}_t}{1 - p_t} \tag{7}$$

for the hazard at which a player who is still uninformed at date t becomes directly informed in the next instant.

As in the static model, once a player is directly informed and observes θ_l, *No Attack* is strictly dominant. The dynamically relevant information states are therefore the uninformed state, denoted U, and the directly informed fragile-state condition, denoted F. A player in state U knows only the prior, the deterministic diffusion path $(p_t)_{t\geq 0}$, and the fixed depth-k communication architecture. A player in state F observes θ_m and assigns belief potential

$$r_k(p_t, \varepsilon) = \frac{1 - p_t}{(1 - p_t) + p_t \varepsilon^{k-1}} \tag{8}$$

to the opponent being on the dominance-anchored branch.

The dynamic problem inherits the same two primitive attack margins as the static model. For an uninformed player at date t, the gain from attack relative to *No Attack* is

$$\Gamma^U(t; q) = \mathbb{E}[\theta] - (1 - q)p_t, \qquad \mathbb{E}[\theta] = (1 - q)\theta_l + q\theta_m. \tag{9}$$

This is the dynamic analogue of the Coarse-agent panic margin. For a directly informed player who observes θ_m, the gain from attack is

$$\Gamma^F(t; k) = \theta_m - r_k(p_t, \varepsilon), \tag{10}$$

which is the dynamic analogue of the Fine-m margin.

The first implication is immediate. Because p_t rises over time, the uninformed attack gain falls, while the informed fragile-state attack gain rises. Early in the diffusion process, attack is driven by panic among uninformed agents; later, attack is driven by informed agents coordinating on fragile-state information. Between these two forces lies a waiting phase in which exposure is high enough to relax panic but not yet so high as to sustain coordinated attack by informed agents.

Lemma 1 (Monotone attack margins)**.** *Suppose p_t is continuous and strictly increasing. Then $\Gamma^U(t; q)$ is continuous and strictly decreasing in t, while $\Gamma^F(t; k)$ is continuous and strictly increasing in t.*

The timing problem can be written recursively. Let $V^F(t)$ denote the continuation value of a player who is directly informed and observes θ_m at date t, and let $V^U(t)$ denote the continuation value of a player who is still uninformed at date t. A directly informed fragile-state player has no further information arrival, so the only trade-off is between attacking immediately and waiting. The value function solves

$$\max\left\{\Gamma^F(t;k) - V^F(t),\ \dot{V}^F(t) - \rho V^F(t)\right\} = 0. \tag{11}$$

An uninformed player, by contrast, may become directly informed while waiting. Conditional on still being uninformed at date t, direct exposure arrives at hazard $h(t)$. With probability $1-q$, exposure reveals the safe state and the continuation value is zero; with probability q, exposure reveals the fragile state and the player enters the informed-fragile state with value $V^F(t)$. The uninformed value function therefore solves

$$\max\left\{\Gamma^U(t;q) - V^U(t),\ \dot{V}^U(t) - (\rho + h(t))V^U(t) + q\,h(t)V^F(t)\right\} = 0. \tag{12}$$

The dynamic structure is thus asymmetric in exactly the way the static analysis suggests, but with an additional coupling. For the informed-fragile player, epistemic depth enters directly through $\Gamma^F(t;k)$. For the uninformed player, epistemic depth does not enter the primitive attack gain directly, but it enters the value of waiting through the prospect of future transition into the informed fragile state. This is the dynamic channel absent from the static classification.

To obtain threshold behavior, define

$$\Phi^F(t;k) := e^{-\rho t}\Gamma^F(t;k). \tag{13}$$

Assume that $\Phi^F(\cdot;k)$ is single-peaked and that the uninformed value-matching condition admits a unique crossing. Then the stopping problems are characterized by threshold dates.

Proposition 4 (Dynamic threshold characterization)**.** *Suppose p_t is continuous and strictly increasing, $\Phi^F(\cdot;k)$ is single-peaked, and the uninformed value-matching condition admits a unique crossing. Then there exist threshold dates $t_F^*(k)$ and $t_U^*(k)$ such that: a directly informed player who observes θ_l never attacks; a directly informed player who observes θ_m waits for $t < t_F^*(k)$ and attacks for $t \geq t_F^*(k)$; and an uninformed player waits for $t < t_U^*(k)$ and attacks for $t \geq t_U^*(k)$ if still uninformed.*

This threshold representation is the temporal analogue of the static intermediate-informativeness result. The static model partitions the breadth space into a low-breadth panic region, an intermediate no-attack region, and a high-breadth informed-attack region. The dynamic model maps the same forces into calendar time along the deterministic path $(p_t)_{t\geq 0}$. Early in time, the uninformed panic motive is strongest. Later in time, the informed-fragile coordination motive becomes strongest. Between them lies a waiting phase in which attack is delayed by both types.

3.1 Epistemic Depth and the Dynamic Fragility Frontier

The main dynamic comparative static concerns epistemic depth. Holding the diffusion path $(p_t)_{t\geq 0}$ fixed, deeper epistemic structure makes attack less attractive for directly informed agents in the fragile state at every date. This delays the informed-fragile stopping threshold. Because the uninformed stopping problem is coupled to the continuation value of becoming informed in the fragile state, the uninformed stopping threshold also shifts outward. Depth therefore lengthens the waiting phase from both sides.

Proposition 5 (Epistemic depth and waiting-time thresholds)**.** *Fix the deterministic diffusion path $(p_t)_{t\geq 0}$. As epistemic depth k rises, the informed-fragile waiting threshold $t^*_F(k)$ shifts weakly later. Under the single-crossing condition for the uninformed stopping problem, the uninformed waiting threshold $t^*_U(k)$ also shifts weakly later.*

Proposition 5 clarifies how the dynamic model differs from the static one. In the static classification, depth affects only the Fine-m constraint directly, because the Coarse discipline condition depends only on the seed event in which the opponent is directly informed of the dominant state l. In the dynamic problem, by contrast, depth affects the timing of both margins. It delays attack directly for informed agents in the fragile state, and it delays panic indirectly for uninformed agents through the value of waiting for future information arrival.

This same logic yields a dynamic counterpart to the static fragility frontier. For a fixed diffusion path and fixed depth, greater fragility makes waiting harder to sustain, because it raises the relative attractiveness of attack for uninformed agents and increases the weight placed on transition into the fragile informed state. But greater epistemic depth works in the opposite direction: it raises the continuation value of waiting by delaying informed attack, and through that channel enlarges the set of fragile environments for which an initial waiting phase can be sustained.

Define the dynamic fragility frontier by

$$\bar{q}_k^{dyn} := \sup\{q \in (0,1) : t^*_U(k;q) > 0\}, \tag{14}$$

that is, the largest fragility level for which an uninformed player strictly prefers not to attack immediately at date 0.

Corollary 1 (Dynamic fragility frontier)**.** *Fix the deterministic diffusion path* $(p_t)_{t\geq 0}$ *and suppose the threshold characterization of Proposition 4 holds. Then the dynamic fragility frontier* $\bar{q}_k^{dyn}$ *is weakly increasing in epistemic depth* k*.*

Thus deeper epistemic structure supports waiting under more fragile fundamentals in the dynamic model as well. The difference from the static benchmark is that the mechanism is stronger here. In the static model, depth enlarges the feasible region only by relaxing the Fine-m upper bound. In the dynamic model, the same direct effect on informed agents also feeds back onto uninformed behavior through the continuation value of waiting. In that sense, deeper epistemic structure not only widens the static stable region, but also makes dynamic coordination more resilient.

3.2 A Numerical Illustration

To illustrate the dynamic threshold structure, suppose diffusion follows the tractable deterministic path

$$p_t = 1 - e^{-\alpha t}, \qquad t \geq 0, \tag{15}$$

so that uninformed agents become directly informed at constant hazard

$$h(t) = \frac{\dot{p}_t}{1 - p_t} = \alpha. \tag{16}$$

This specification keeps the diffusion process monotone and transparent, while delivering a simple benchmark for the timing problem.

Under (15), the uninformed attack gain

$$\Gamma^U(t; q) = \mathbb{E}[\theta] - (1 - q)(1 - e^{-\alpha t})$$

is strictly decreasing in t, while the informed-fragile gain

$$\Gamma^F(t; k) = \theta_m - \frac{e^{-\alpha t}}{e^{-\alpha t} + (1 - e^{-\alpha t})\varepsilon^{k-1}}$$

is strictly increasing in t. Since $\Gamma^F(t; k)$ rises from a low initial value and remains bounded above by θ_m, the discounted objective

$$\Phi^F(t; k) = e^{-\rho t}\Gamma^F(t; k)$$

is single-peaked for the parameter values used below. The numerical exercise therefore delivers the

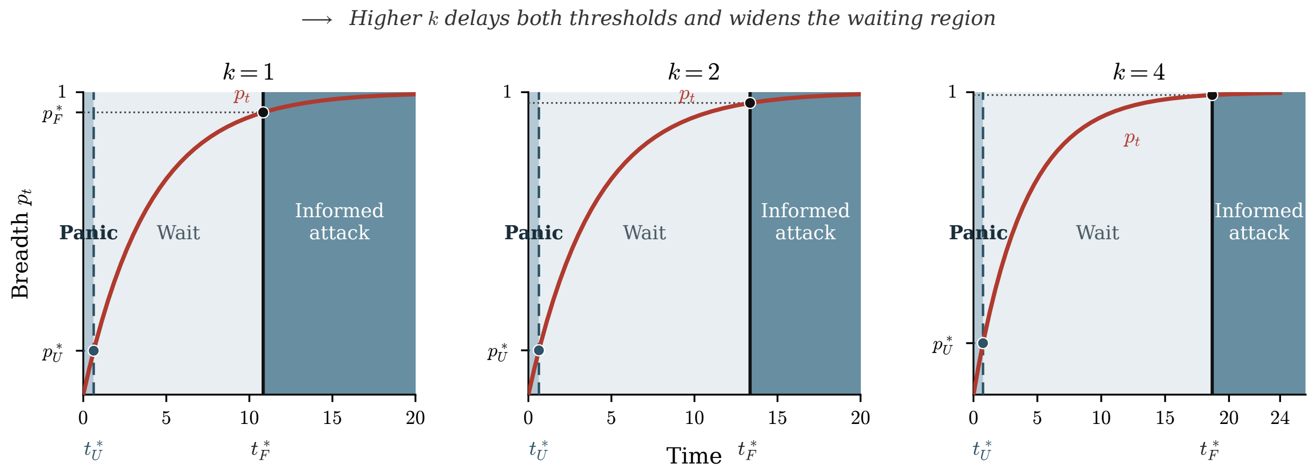

Figure 4: Deeper epistemic structure delays both the uninformed panic threshold t_U^* and the informed attack threshold t_F^*, widening the waiting region as breadth $p_t = 1 - e^{-\alpha t}$ diffuses upward.

threshold pattern characterized in Proposition 4.

In the numerical illustration we set

$$\theta_l = -0.1, \qquad \theta_m = 0.4, \qquad q = 0.5, \qquad \varepsilon = 0.5, \qquad \rho = 0.05, \qquad \alpha = 0.25.$$

Figure 4 plots the dynamic attack regions for $k \in \{1, 2, 4\}$. In each panel, the early region is one in which attack is driven by uninformed panic, the late region is one in which attack is driven by informed agents coordinating on fragile-state information, and the intermediate region is a waiting phase. As epistemic depth rises, the informed-fragile attack threshold shifts later directly, and the uninformed threshold also shifts later through the continuation value of waiting. Deeper epistemic structure therefore widens the dynamic waiting region from both sides.

4 Conclusion

This paper studies coordination under strategic uncertainty when the communication environment has two distinct margins: informational breadth and epistemic depth. Breadth captures how widely information is directly disseminated; depth captures how much mutual knowledge the communication environment can sustain before agents act. The central message is that these two margins play different economic roles, and that distinguishing between them helps organize both the static and dynamic determinants of fragility.

In the static model, broader dissemination is not unambiguously stabilizing. When breadth is too low, uninformed agents remain insufficiently reassured and panic is supportable. When breadth is too high, informed agents in the fragile state can coordinate on attack. *No Attack* is therefore uniquely rationalizable only for an intermediate range of breadth. Greater epistemic depth relaxes this problem asymmetrically: it disciplines informed agents in the fragile state, but does not directly relax the panic motive of uninformed agents. As a result, deeper epistemic structure widens the stable region and allows coordination to be sustained in more fragile environments, but only up to a sharp upper bound.

The dynamic extension embeds the same logic in an environment where epistemic depth is fixed as a primitive of the communication platform, while informational breadth rises over time. This yields a timing problem in which early attacks are driven by uninformed panic and later attacks by informed agents coordinating on fragile-state information. Here the role of depth is stronger than in the static benchmark. Depth delays attack directly for informed agents and, through the continuation value of waiting, indirectly for uninformed agents as well. Deeper epistemic structure therefore widens the dynamic waiting region and supports waiting under more fragile fundamentals.

The broader implication is that the economically relevant design object is not only how far information spreads, but how communication is converted into higher-order beliefs and mutual knowledge. Environments that broaden exposure without generating sufficient epistemic depth may intensify fragility rather than reduce it.

References

George-Marios Angeletos and Ivan Werning. Crises and prices: Information aggregation, multiplicity, and volatility. *American Economic Review*, 96(5):1720–1736, 2006. 1

Salvador Barberà and Matthew O Jackson. A model of protests, revolution, and information. *Available at SSRN 2841440*, 2020. 1

B. Douglas Bernheim. Rationalizable strategic behavior. *Econometrica*, 52(4):1007–1028, 1984. 2, 2

V. Bhaskar, George J. Mailath, and Stephen Morris. A foundation for markov equilibria in sequential games with finite social memory. *The Review of Economic Studies*, 80(3):925–948, 2013. doi: 10.1093/restud/rds047. 6

Matthieu Bouvard, Pierre Chaigneau, and Adolfo De Motta. Transparency in the financial system: Rollover risk and crises. *Journal of Finance*, 70(4):1519–1553, 2015. 1

Hans Carlsson and Eric Van Damme. Global games and equilibrium selection. *Econometrica: Journal of the Econometric Society*, pages 989–1018, 1993. 1

Michael Suk-Young Chwe. Communication and coordination in social networks. *The Review of Economic Studies*, 67(1):1–16, 2000. 1

J Anthony Cookson, Corbin Huang, Gil Itzkowitz, and Jennifer Itzkowitz. Social media as a bank run catalyst. *Journal of Financial Economics*, 176:104218, 2026. 1

Giancarlo Corsetti, Amil Dasgupta, Stephen Morris, and Hyun Song Shin. Does one soros make a difference? a theory of currency crises with large and small traders. *The Review of Economic Studies*, 71(1):87–113, 2004. 1

Eddie Dekel, Drew Fudenberg, and Stephen Morris. Topologies on types. 2005. 3

Chris Edmond. Information and the limits to autocracy. *The Quarterly Journal of Economics*, 128(3):1421–1466, 2013. 1

Amanda Friedenberg. Bargaining under strategic uncertainty: The role of second-order optimism. *Econometrica*, 87(5):1835–1865, 2019. 4

John Geanakoplos. Common knowledge. *Journal of Economic Perspectives*, 6(4):53–82, 1992. 3

Itay Goldstein and Ady Pauzner. Demand–deposit contracts and the probability of bank runs. *the Journal of Finance*, 60(3):1293–1327, 2005. 1

Itay Goldstein and Haresh Sapra. Should banks' stress test results be disclosed? an analysis of the costs and benefits. *Foundations and Trends in Finance*, 8(1):1–54, 2014. 1

Takuya Hoshino. Multi-agent persuasion: Leveraging strategic uncertainty. *International Economic Review*, 63(2):755–776, 2022. 4

Atsushi Kajii and Stephen Morris. The robustness of equilibria to incomplete information. *Econometrica: Journal of the Econometric Society*, pages 1283–1309, 1997. 3

Frédéric Koessler and Vasiliki Skreta. Informed information design. *Journal of Political Economy*, 131(11):3135–3180, 2023. 4

Fei Li, Yangbo Song, and Mofei Zhao. Global manipulation by local obfuscation. *Journal of Economic Theory*, 207:105575, 2023. 4

Laurent Mathevet, Jacopo Perego, and Ina Taneva. On information design in games. *Journal of Political Economy*, 128(4):1370–1404, 2020. 1

Dov Monderer and Dov Samet. Approximating common knowledge with common beliefs. *Games and Economic Behavior*, 1(2):170–190, 1989. 3

Fumitoshi Moriya and Takuro Yamashita. Asymmetric-information allocation to avoid coordination failure. *Journal of Economics & Management Strategy*, 29(1):173–186, 2020. 4

Stephen Morris and Hyun Song Shin. Global games: Theory and applications. 2001. 1, 2

Stephen Morris, Rafael Rob, and Hyun Song Shin. p-dominance and belief potential. *Econometrica: Journal of the Econometric Society*, pages 145–157, 1995. 2, 1, 2

Stephen Morris, Daisuke Oyama, and Satoru Takahashi. Implementation via information design in binary-action supermodular games. *Econometrica*, 92(3):775–813, 2024. 4

David G. Pearce. Rationalizable strategic behavior and the problem of perfection. *Econometrica*, 52(4):1029–1050, 1984. 2, 2

Ariel Rubinstein. The electronic mail game: Strategic behavior under" almost common knowledge". *The American Economic Review*, pages 385–391, 1989. 1, 2

Jonathan Weinstein and Muhamet Yildiz. A structure theorem for rationalizability with application to robust predictions of refinements. *Econometrica*, 75(2):365–400, 2007. 3

A Partition Representation

This appendix formalizes the augmented partition representation induced by the Fine/Coarse benchmark and the finite confirmation process described in Section 2. Its purpose is to make the underlying information sets explicit and to justify the reduced-form belief potential $r_k(p, \varepsilon)$ used in the main analysis.

Benchmark augmented state space. The Fine/Coarse benchmark generates uncertainty jointly over opponent type and payoff state. We represent this on

$$\{F, C\} \times \{F, C\} \times \{l, m\},$$

where a generic element $T_1T_2\omega$ records player 1's information type, player 2's information type, and the realized state. Since types and states are drawn independently,

$$\Pr(T_1T_2\omega) = \Pr(T_1)\Pr(T_2)\mu(\omega).$$

Player 1 observes its own type and, if Fine, the realized state. Its partition is therefore

$$\mathcal{P}_1 = \Big\{\underbrace{\{FFl, FCl\}}_{\text{Fine-}l},\ \underbrace{\{CFl, CCl, CCm, CFm\}}_{\text{Coarse}},\ \underbrace{\{FCm, FFm\}}_{\text{Fine-}m}\Big\},$$

and symmetrically for player 2,

$$\mathcal{P}_2 = \Big\{\{FFl, CFl\},\ \{FCl, CCl, CCm, FCm\},\ \{CFm, FFm\}\Big\}.$$

The Fine-l cell is the seed event on which *No Attack* is strictly dominant.

Finite confirmation and epistemic depth. For depth $k \geq 2$, we augment the benchmark by allowing up to $k-1$ rounds of confirmation after agents become informed. A Fine player may send a confirmation signal to the opponent. A Coarse player never sends a signal. Each confirmation fails to reach its recipient with probability $\varepsilon \in (0, 1)$, and a failed signal is observationally invisible.

The relevant augmented histories can be indexed by

$$(n, T_2, \omega), \qquad n \in \{0, \ldots, k-1\},$$

where n counts completed confirmation rounds.

The key observational equivalence is the following. At the deepest relevant information set under depth k, a Fine-m agent has received no confirmation after $k-1$ rounds. At that node, the opponent may be Coarse, in which case no confirmation was ever sent, or Fine, in which case all $k-1$ confirmations failed. Since the latter event occurs with probability ε^{k-1}, Bayes' rule implies

$$r_k(p, \varepsilon) = \frac{1-p}{(1-p) + p\varepsilon^{k-1}}.$$

Thus the reduced-form statistic $r_k(p, \varepsilon)$ used in the main text is exactly the probability that a Fine-m agent assigns to the opponent being on the dominance-anchored branch.

Connection to Rubinstein. As $k \to \infty$, the finite confirmation process approaches Rubinstein's (1989) email-game logic. Even if both agents are Fine, no finite confirmation history delivers common knowledge because any finite sequence is compatible with the possibility that one more confirmation failed. In the present model, this means that deeper epistemic structure can keep raising the belief potential of the dominance-anchored branch, but only up to the hard upper bound developed in Section H.

B Proof of Proposition 1

Step 1: Types. Under the symmetric Fine/Coarse benchmark, player i has three possible interim types:

$$F_{il}, \qquad F_{im}, \qquad C_i,$$

with probabilities

$$p(1-q), \qquad pq, \qquad 1-p,$$

respectively.

Step 2: Seed event. The seed is

$$E = \{FFl, FCl\},$$

that is, the event that player 1 is Fine and observes state l. Since $\theta_l < 0$, *No Attack* is strictly dominant on E.

Step 3: Coarse constraint. Let π denote the probability the opponent plays *No Attack*. A Coarse agent's expected payoff from *Attack* is

$$\pi(\mathbb{E}[\theta] - 1) + (1 - \pi)\mathbb{E}[\theta] = \mathbb{E}[\theta] - \pi.$$

The payoff from *No Attack* is 0. Hence *No Attack* is strictly optimal if and only if

$$\pi > \mathbb{E}[\theta].$$

In the benchmark architecture, the only branch pinned to *No Attack* is that the opponent is Fine and observes state l, which occurs with probability $(1 - q)p$. The Coarse constraint is therefore

$$(1 - q)p \geq \mathbb{E}[\theta].$$

Step 4: Fine-m constraint. A Fine-m agent refrains from attacking if and only if the probability the opponent is Coarse and hence anchored by dominance is at least θ_m. In the benchmark architecture that probability is $1 - p$, so the Fine-m constraint is

$$1 - p \geq \theta_m.$$

Step 5: Combining. Both constraints hold if and only if

$$\frac{\mathbb{E}[\theta]}{1 - q} \leq p \leq 1 - \theta_m.$$

The interval is nonempty if and only if

$$\frac{\mathbb{E}[\theta]}{1 - q} \leq 1 - \theta_m,$$

which rearranges to

$$q \leq \frac{1 - \theta_m - \theta_l}{1 - \theta_l} = \bar{q}_1.$$

□

C Proof of Proposition 2

Coarse constraint. The Coarse constraint is identical to Step 3 of Appendix B. It is

$$(1-q)p \geq \mathbb{E}[\theta],$$

and is independent of epistemic depth.

Deriving $r_k(p,\varepsilon)$. At the deepest information set under depth k, a Fine-m agent has received no confirmation after $k-1$ rounds. A Coarse opponent reaches that node with probability one, because it never sends a signal. A Fine opponent reaches that same node only if all $k-1$ confirmations fail, which occurs with probability ε^{k-1}. Weighting by type probabilities therefore gives

$$r_k(p,\varepsilon) = \frac{1-p}{(1-p)+p\varepsilon^{k-1}}.$$

Fine-m constraint. The Fine-m agent refrains from attacking if and only if

$$r_k(p,\varepsilon) \geq \theta_m.$$

Solving this inequality yields

$$p \leq \bar{p}_k(\varepsilon) := \frac{1-\theta_m}{(1-\theta_m)+\theta_m\varepsilon^{k-1}}.$$

Since ε^{k-1} is strictly decreasing in k, $\bar{p}_k(\varepsilon)$ is strictly increasing in k.

Combining. Feasibility requires

$$\frac{\mathbb{E}[\theta]}{1-q} \leq p \leq \bar{p}_k(\varepsilon).$$

The lower bound is independent of k because it depends only on the seed event that the opponent is Fine in state l. The upper bound is strictly increasing in k because deeper confirmation raises the belief potential of Fine-m agents. The stable region therefore expands only from above. □

D Proof of Proposition 3

Part (i). By construction,

$$\bar{p}_k(\varepsilon) = \frac{1-\theta_m}{(1-\theta_m)+\theta_m\varepsilon^{k-1}}$$

is strictly increasing in k, and

$$\bar{p}_k(\varepsilon) \to 1 \qquad \text{as } k \to \infty.$$

Substituting into

$$\bar{q}_k(\varepsilon) = \frac{\bar{p}_k(\varepsilon) - \theta_l}{\bar{p}_k(\varepsilon) - \theta_l + \theta_m}$$

shows that $\bar{q}_k(\varepsilon)$ is strictly increasing in k and converges to

$$\bar{q}_\infty = \frac{1 - \theta_l}{(1 - \theta_l) + \theta_m}.$$

Part (ii). For any $q < \bar{q}_\infty$, monotone convergence implies that there exists a finite k such that $q \le \bar{q}_k(\varepsilon)$, so the stable region is nonempty at that depth. For $q > \bar{q}_\infty$, sharpness follows from Corollary 5 in Section H: no admissible ex ante architecture can robustly implement *No Attack* in both states once the posterior seed mass bottleneck is violated. The finite-depth Fine/Coarse family considered here is therefore bounded above by the same limit. □

E Static Extensions

This appendix collects several consequences of the main static asymmetry that epistemic depth relaxes only the Fine-m margin while leaving the Coarse panic constraint unchanged.

E.1 Minimum Epistemic Depth

Proposition 6 (Minimum epistemic depth)**.** *Define*

$$k^*(q) := \min\{k \in \mathbb{N} : q \le \bar{q}_k(\varepsilon)\}.$$

Then:

(i) $k^*(q)$ *is well defined and finite if and only if* $q < \bar{q}_\infty$.

(ii) $k^*(q)$ *is weakly increasing in* q.

(iii) $k^*(q) = 1$ *for* $q \le \bar{q}_1$. *For* $q \in (\bar{q}_1, \bar{q}_\infty)$, *define*

$$\delta(q) := \frac{(1 - \theta_m)\big[(1 - q) - \mathbb{E}[\theta]\big]}{\theta_m\, \mathbb{E}[\theta]}.$$

Then $\delta(q) \in (0,1)$ *and*

$$k^*(q) = 1 + \left\lceil \frac{\log \delta(q)}{\log \varepsilon} \right\rceil,$$

with $k^*(q) \to \infty$ *as* $q \nearrow \bar{q}_\infty$.

Proof. Part *(i)* follows from the strict monotonicity of $\bar{q}_k(\varepsilon)$ in k and its limit $\bar{q}_\infty$.

Part *(ii)* follows because raising q weakly shrinks the set

$$\{k \in \mathbb{N} : q \leq \bar{q}_k(\varepsilon)\}.$$

For part *(iii)*, if $q \leq \bar{q}_1$, the benchmark interval is already nonempty, so $k^*(q) = 1$. For $q \in (\bar{q}_1, \bar{q}_\infty)$, the condition

$$q \leq \bar{q}_k(\varepsilon)$$

is equivalent to

$$\frac{\mathbb{E}[\theta]}{1-q} \leq \bar{p}_k(\varepsilon),$$

which rearranges to

$$\varepsilon^{k-1} \leq \delta(q).$$

Since $q < \bar{q}_\infty$, one has $(1-q) > \mathbb{E}[\theta]$, so $\delta(q) > 0$. Since $q > \bar{q}_1$, one has

$$\frac{\mathbb{E}[\theta]}{1-q} > 1 - \theta_m,$$

which implies $\delta(q) < 1$. Taking logarithms with $\log \varepsilon < 0$ yields the stated formula. As $q \nearrow \bar{q}_\infty$, $\delta(q) \to 0^+$ and hence $k^*(q) \to \infty$. □

Corollary 2 (Depth thresholds)**.** $k^*(q) = 1$ *if and only if* $q \leq \bar{q}_1$. *For* $k \geq 2$,

$$k^*(q) = k \iff \bar{q}_{k-1}(\varepsilon) < q \leq \bar{q}_k(\varepsilon),$$

or equivalently,

$$k^*(q) = k \iff \bar{p}_{k-1}(\varepsilon) < \frac{\mathbb{E}[\theta]}{1-q} \leq \bar{p}_k(\varepsilon).$$

Proof. Immediate from the strict monotonicity of $\bar{q}_k(\varepsilon)$ and the equivalence

$$q \leq \bar{q}_k(\varepsilon) \iff \frac{\mathbb{E}[\theta]}{1-q} \leq \bar{p}_k(\varepsilon).$$

□

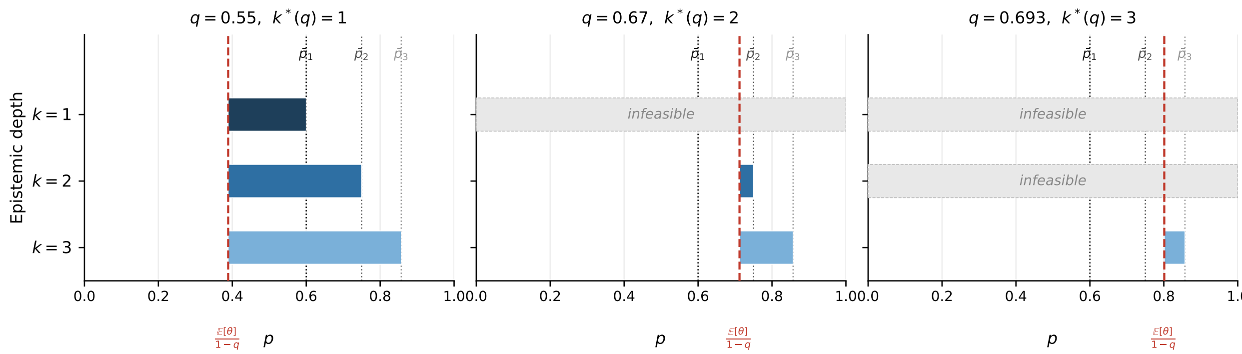

Figure 5: Intervals of p for which *No Attack* is uniquely rationalizable at each depth $k \in \{1, 2, 3\}$, for three values of q. The lower bound is fixed by the Coarse panic constraint, while the upper bound rises with depth through the Fine-m margin. Parameters: $\theta_l = -0.1$, $\theta_m = 0.4$, $\varepsilon = 0.5$.

Table 1: Depth thresholds. Parameters: $\theta_l = -0.1$, $\theta_m = 0.4$, $\varepsilon = 0.5$.

k	$\bar{p}_k(\varepsilon)$	$\bar{q}_k(\varepsilon)$	Region: $k^*(q) = k$
1	0.600	0.636	$q \le 0.636$
2	0.750	0.680	$0.636 < q \le 0.680$
3	0.857	0.705	$0.680 < q \le 0.705$
4	0.923	0.719	$0.705 < q \le 0.719$
5	0.960	0.726	$0.719 < q \le 0.726$
∞	1.000	0.733	$\bar{q}_\infty = 0.733$

The corollary makes the mechanism transparent. Each transition in $k^*(q)$ occurs when the fixed Coarse lower bound

$$\frac{\mathbb{E}[\theta]}{1-q}$$

overtakes the Fine-m upper bound available at the previous depth. One additional confirmation layer then restores a nonempty stable interval by pushing the Fine-m bound outward.

E.2 Minimum Depth as a Function of Breadth

Corollary 3 (Minimum depth at fixed breadth)**.** *Fix q and $\varepsilon \in (0,1)$. Define*

$$k^*(q,p) := \min\{k \in \mathbb{N} : (1-q)p \ge \mathbb{E}[\theta] \text{ and } r_k(p,\varepsilon) \ge \theta_m\},$$

with the convention that $k^*(q,p) = \infty$ if no such k exists. Then:

(i) *If $p < \mathbb{E}[\theta]/(1-q)$, then $k^*(q,p) = \infty$.*

(ii) *If* $\mathbb{E}[\theta]/(1-q) \le p \le 1-\theta_m$*, then* $k^*(q,p) = 1$*.*

(iii) *If* $p > 1-\theta_m$*, then*

$$k^*(q,p) = 1 + \left\lceil \frac{\log\big((1-p)(1-\theta_m)/(p\theta_m)\big)}{\log \varepsilon} \right\rceil,$$

which is strictly increasing in p*, with* $k^*(q,p) \to \infty$ *as* $p \to 1$*.*

Proof. If $p < \mathbb{E}[\theta]/(1-q)$, the Coarse constraint fails for all k, proving part *(i)*.

If

$$\frac{\mathbb{E}[\theta]}{1-q} \le p \le 1-\theta_m,$$

then the Coarse constraint holds and

$$r_1(p) = 1-p \ge \theta_m,$$

so part *(ii)* follows.

If $p > 1-\theta_m$, then $r_1(p) = 1-p < \theta_m$, so the benchmark depth fails. Solving

$$r_k(p,\varepsilon) \ge \theta_m$$

yields

$$\varepsilon^{k-1} \le \frac{(1-p)(1-\theta_m)}{p\theta_m}.$$

Taking logarithms gives the stated formula. The right-hand side is strictly decreasing in p and tends to 0^+ as $p \to 1$, implying the stated monotonicity and divergence. □

Figure 6 plots $k^*(q,p)$ on the (p,q) plane.

E.3 Marginal Value of Epistemic Depth

Corollary 4 (Non-monotone marginal value of depth)**.** *Fix* $k \in \mathbb{N}$ *and* $\varepsilon \in (0,1)$*. The marginal gain*

$$\Delta_k(p,\varepsilon) := r_{k+1}(p,\varepsilon) - r_k(p,\varepsilon)$$

is strictly positive for $p \in (0,1)$*, strictly single-peaked in* p*, and vanishes as* $p \to 0$ *or* $p \to 1$*. Its unique maximizer is*

$$p_k^* = \frac{1}{1+\varepsilon^{(2k-1)/2}}.$$

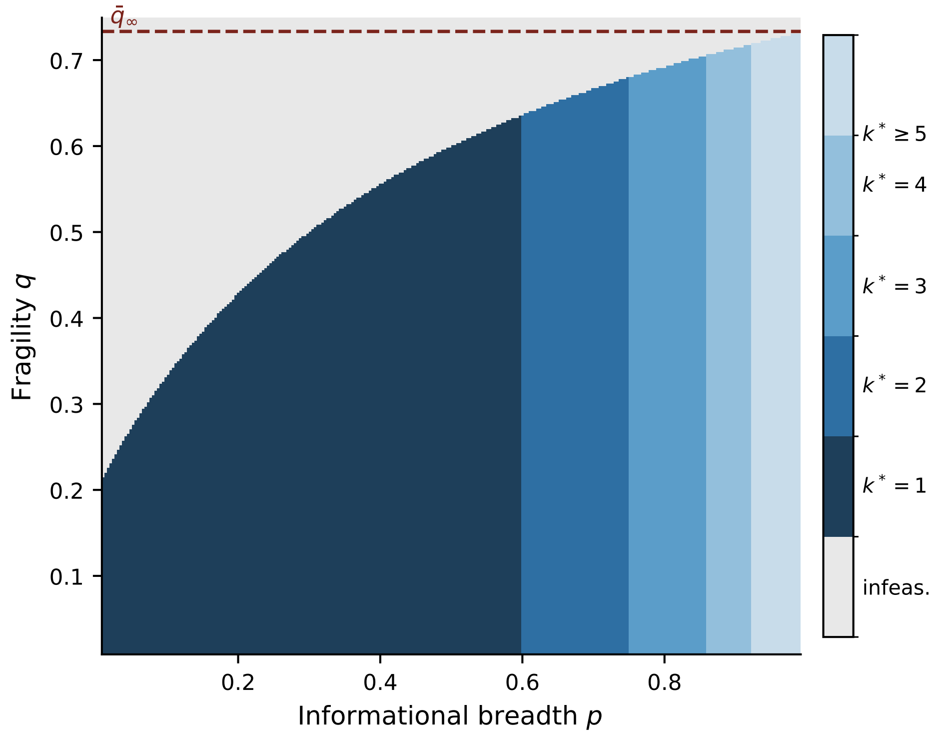

Figure 6: Minimum epistemic depth $k^*(q, p)$ on the (p, q) plane. Left: the Coarse constraint fails for all depths. Middle: the benchmark depth suffices. Right: additional depth is required as the Fine-m constraint tightens. Dashed: $\bar{q}_\infty$.

Proof. Write

$$\alpha = \varepsilon^{k-1}, \qquad \beta = \varepsilon^k,$$

so that $0 < \beta < \alpha < 1$. Then

$$\Delta_k(p, \varepsilon) = \frac{(1-p)p(\alpha - \beta)}{[(1-p) + p\beta][(1-p) + p\alpha]}.$$

Positivity is immediate, and the factor $(1-p)p$ implies $\Delta_k \to 0$ as $p \to 0$ or $p \to 1$.

Setting the derivative equal to zero and simplifying yields

$$\frac{1-p}{p} = \sqrt{\alpha\beta} = \varepsilon^{(2k-1)/2},$$

which gives

$$p_k^* = \frac{1}{1 + \varepsilon^{(2k-1)/2}}.$$

Since Δ_k is positive on $(0, 1)$ and vanishes at the endpoints, this critical point is the unique global maximum. □

The marginal value of one additional epistemic layer is therefore largest at intermediate breadth, where uncertainty about the opponent's information is most consequential, and small at either extreme, where the relevant uncertainty is already limited.

F Robustness: Dominant Risky Action

This appendix verifies that the main static asymmetry survives when the risky action, rather than the safe action, is dominant in the anchor state.

Environment. States are $\Omega = \{h, m\}$ with $\Pr(m) = q$. Payoffs are as in the main text, except that now $\theta_h > 1$, so *Attack* is strictly dominant in state h, while $\theta_m \in (0, 1)$, so both actions are rationalizable in state m. The Fine/Coarse benchmark and the finite confirmation process are otherwise unchanged.

Constraints. The seed event is that the opponent is Fine and observes state h, which occurs with probability $(1-q)p$. The Coarse constraint is therefore

$$(1-q)p \geq 1 - \mathbb{E}[\theta],$$

and is independent of depth. A Fine-m agent attacks if and only if

$$r_k(p, \varepsilon) \geq 1 - \theta_m,$$

which yields the upper bound

$$\tilde{p}_k(\varepsilon) := \frac{\theta_m}{\theta_m + (1-\theta_m)\varepsilon^{k-1}},$$

strictly increasing in k.

Proposition 7 (Frontier with a dominant risky action)**.** Attack *is uniquely rationalizable in both states if and only if*

$$p \in \left[\frac{1-\mathbb{E}[\theta]}{1-q},\ \tilde{p}_k(\varepsilon)\right],$$

which is nonempty if and only if

$$q \leq \tilde{q}_k(\varepsilon) := \frac{\theta_h + \tilde{p}_k(\varepsilon) - 1}{\theta_h + \tilde{p}_k(\varepsilon) - \theta_m}.$$

The lower bound is independent of k, while $\tilde{p}_k$ and $\tilde{q}_k$ are strictly increasing in k.

The structure is exactly parallel to Proposition 2. Depth again operates through the Fine-m channel and leaves the Coarse constraint unchanged. The only substantive difference is the payoff threshold relevant for Fine-m types: here it is $1-\theta_m$ rather than θ_m, reflecting that a Fine-m agent must now assign sufficiently high probability to coordinated *Attack*, rather than coordinated restraint, in order for attack to be uniquely optimal.

G Proofs for Section 3

Proof of Lemma 1. By (9),

$$\Gamma^U(t;q) = \mathbb{E}[\theta] - (1-q)p_t.$$

Since p_t is continuous and strictly increasing in t, $\Gamma^U(t;q)$ is continuous and strictly decreasing in t.

For the informed-fragile attack gain,

$$\Gamma^F(t;k) = \theta_m - r_k(p_t,\varepsilon), \qquad r_k(p_t,\varepsilon) = \frac{1-p_t}{(1-p_t)+p_t\varepsilon^{k-1}}.$$

Fix (k,ε). The function $r_k(p,\varepsilon)$ is strictly decreasing in p, while p_t is strictly increasing in t. Hence $r_k(p_t,\varepsilon)$ is strictly decreasing in t, so $\Gamma^F(t;k)$ is continuous and strictly increasing in t. □

Derivation of (11). Fix a date t and consider a player who is directly informed and observes $\theta=\theta_m$.

If the player attacks immediately, the gain from attack relative to *No Attack* is $\Gamma^F(t;k)$.

If instead the player waits over a short interval $[t,t+dt]$, then no further information arrives to this player. The continuation value at date $t+dt$ is $V^F(t+dt)$, discounted by $e^{-\rho dt}=1-\rho\,dt+o(dt)$. Hence the value of waiting is

$$(1-\rho\,dt)V^F(t+dt)+o(dt).$$

Using the expansion

$$V^F(t+dt) = V^F(t) + \dot{V}^F(t)\,dt + o(dt),$$

this becomes

$$V^F(t) + \big(\dot{V}^F(t) - \rho V^F(t)\big)dt + o(dt).$$

Optimality requires that $V^F(t)$ equal the maximum of the stopping and waiting payoffs. Dividing by dt and letting $dt \downarrow 0$ yields

$$\max\left\{\Gamma^F(t;k) - V^F(t),\ \dot{V}^F(t) - \rho V^F(t)\right\} = 0,$$

which is (11). □

Derivation of (12). Fix a date t and consider a player who is still uninformed.

If the player attacks immediately, the gain from attack relative to *No Attack* is $\Gamma^U(t;q)$.

If instead the player waits over a short interval $[t, t+dt]$, then with probability $1-h(t)\,dt+o(dt)$ no direct exposure arrives and the player remains uninformed, while with probability $h(t)\,dt+o(dt)$ direct exposure arrives. Conditional on direct exposure, state l is revealed with probability $1-q$, in which case the continuation value is zero because *No Attack* is strictly dominant, and state m is revealed with probability q, in which case the continuation value is $V^F(t)$.

Hence the expected continuation payoff from waiting is

$$(1-\rho\,dt)\Big[(1-h(t)\,dt)V^U(t+dt) + h(t)\,dt \cdot q\,V^F(t)\Big] + o(dt).$$

Using

$$V^U(t+dt) = V^U(t) + \dot{V}^U(t)\,dt + o(dt),$$

this becomes

$$V^U(t) + \Big(\dot{V}^U(t) - (\rho + h(t))V^U(t) + q\,h(t)V^F(t)\Big)dt + o(dt).$$

Optimality therefore implies

$$\max\left\{\Gamma^U(t;q) - V^U(t),\ \dot{V}^U(t) - (\rho + h(t))V^U(t) + q\,h(t)V^F(t)\right\} = 0,$$

which is (12). □

Characterization of the informed-fragile threshold. Consider the informed-fragile stopping problem (11). On any interval on which waiting is strictly optimal, one has $\Gamma^F(t;k) < V^F(t)$, so the continuation term must satisfy

$$\dot{V}^F(t) - \rho V^F(t) = 0.$$

The solution is

$$V^F(t) = C\,e^{\rho t}$$

for some constant C. If the stopping date is $t_F^*(k)$, value matching implies

$$V^F(t_F^*(k)) = \Gamma^F(t_F^*(k); k),$$

so

$$C = e^{-\rho t_F^*(k)} \Gamma^F(t_F^*(k); k).$$

Substituting back gives

$$V^F(t) = e^{-\rho(t_F^*(k)-t)} \Gamma^F(t_F^*(k); k), \qquad t < t_F^*(k).$$

The date $t_F^*(k)$ maximizes the discounted attack gain

$$\Phi^F(t; k) = e^{-\rho t} \Gamma^F(t; k).$$

Under the assumed single-peakedness of $\Phi^F(\cdot; k)$, the solution is threshold-based. If the optimum is interior, the first order condition is

$$\frac{d}{dt} \Phi^F(t; k) \Big|_{t=t_F^*(k)} = 0,$$

equivalently,

$$\dot{\Gamma}^F(t_F^*(k); k) = \rho \, \Gamma^F(t_F^*(k); k).$$

□

Derivation of the uninformed waiting value. On the uninformed waiting region, (12) reduces to

$$\dot{V}^U(t) - (\rho + h(t)) V^U(t) + q \, h(t) V^F(t) = 0.$$

This is a linear first-order ODE. Let

$$M(t) := \exp\Big(- \int_0^t (\rho + h(u)) \, du \Big)$$

be the integrating factor. Multiplying through yields

$$\frac{d}{dt} \big(M(t) V^U(t) \big) = - M(t) \, q \, h(t) V^F(t).$$

Integrating from t to t_U^* gives

$$V^U(t) = \int_t^{t_U^*} \exp\left(-\int_t^s (\rho + h(u))\, du\right) q\, h(s) V^F(s)\, ds + \exp\left(-\int_t^{t_U^*} (\rho + h(u))\, du\right) \Gamma^U(t_U^*; q),$$

where value matching at t_U^* has been used. □

Proof of Proposition 4. A directly informed player who observes θ_l never attacks, because in state l the stage-game structure implies that *No Attack* strictly dominates *Attack*.

For a directly informed player who observes θ_m, the dynamic problem is (11). By Lemma 1, $\Gamma^F(t;k)$ is strictly increasing in t. Under the assumption that $\Phi^F(\cdot;k)$ is single-peaked, the optimal stopping rule is characterized by a threshold date $t_F^*(k)$, with waiting for $t < t_F^*(k)$ and attack for $t \geq t_F^*(k)$.

For an uninformed player, the dynamic problem is (12). By Lemma 1, $\Gamma^U(t;q)$ is strictly decreasing in t, while waiting creates the option of transition into the informed-fragile state. Under the assumed unique crossing of the value-matching condition, the stopping rule is characterized by a threshold date $t_U^*(k)$, with waiting for $t < t_U^*(k)$ and attack for $t \geq t_U^*(k)$ if still uninformed.

Together these arguments yield the threshold characterization stated in the proposition. □

Proof of Proposition 5. Fix the deterministic path $(p_t)_{t\geq 0}$.

For any date t, the belief potential

$$r_k(p_t, \varepsilon) = \frac{1 - p_t}{(1 - p_t) + p_t \varepsilon^{k-1}}$$

is strictly increasing in k, because ε^{k-1} is strictly decreasing in k for $\varepsilon \in (0, 1)$. Hence

$$\Gamma^F(t;k) = \theta_m - r_k(p_t, \varepsilon)$$

is strictly decreasing in k at every date t. Under the single-peakedness assumption, this delays the informed-fragile optimal stopping date, so $t_F^*(k)$ shifts weakly later as k rises.

Now consider the uninformed threshold. The primitive gain $\Gamma^U(t;q)$ does not depend directly on k, but the waiting value depends on k through $V^F(\cdot;k)$. Because deeper epistemic structure delays the informed-fragile attack date, it raises the value to an uninformed player of waiting for possible future information arrival. Under the maintained single-crossing condition, the date at which the declining function $\Gamma^U(t;q)$ intersects the waiting value therefore shifts weakly later as well. Hence $t_U^*(k)$ also shifts weakly later with depth. □

Proof of Corollary 1. Fix the deterministic diffusion path and depth k. The date-0 uninformed attack gain is

$$\Gamma^U(0;q) = \mathbb{E}[\theta] - (1-q)p_0,$$

which is increasing in q because

$$\frac{\partial \mathbb{E}[\theta]}{\partial q} = \theta_m - \theta_l > 0$$

and the coefficient on p_0 also rises with q. More generally, higher q raises the relative attractiveness of attack for an uninformed player and lowers the value of waiting by placing greater weight on future transition into the fragile informed state.

By Proposition 5, increasing k raises the uninformed value of waiting and shifts $t^*_U(k;q)$ weakly later for every fixed q. Therefore the set

$$\{q \in (0,1) : t^*_U(k;q) > 0\}$$

is weakly larger at higher depths. Taking suprema yields that $\bar{q}_k^{dyn}$ is weakly increasing in k. □

H A General Pooling Bound

Greater epistemic depth relaxes the informed agent's coordination constraint, but this force cannot be arbitrarily strong. Any admissible ex ante architecture that robustly implements the safe action must confront a pooled type at which the belief potential of the seed states is bounded by their posterior mass. This section states that logic in a general form and explains why the frontier in the main model rises with epistemic depth yet stops at a hard upper bound.

Consider a two-player binary-action coordination problem. Let the preferred action be denoted by a^*. Let Ω be a finite state space with common prior μ. A subset $S \subseteq \Omega$ is a *seed region* if a^* is strictly dominant for both players at every state in S. In the main text, $S = \{l\}$.

An admissible ex ante architecture specifies a distribution over players' information structures before the state is realized. The realized structure induces an augmented state space and, for each player, a collection of interim information sets. For an information set T_i, let $\mu(S \mid T_i)$ denote the posterior probability that the true state lies in the seed region, and let $t_i(T_i) \in [0,1]$ denote the minimal probability with which player i must assign the opponent to the safe branch in order for a^* to be uniquely optimal at T_i.

Proposition 8 (Pooling bottleneck)**.** *Consider any admissible ex ante architecture in a binary-action coordination environment with seed region S. If the preferred action a^* is uniquely ratio-*

nalizable at every information set, then there exists an information set T_i that pools some state in S with some state in $\Omega \setminus S$. At that information set, the belief potential supporting a^* is bounded above by $\mu(S \mid T_i)$. Hence a necessary condition for robust implementation is

$$\mu(S \mid T_i) \geq t_i(T_i).$$

Proof. Suppose instead that every information set compatible with a non-seed state reveals that the realized state lies outside S. Then at any such information set the player faces a residual coordination problem in which a^* is not strictly dominant. In a binary-action coordination game, the non-preferred action then survives iterated elimination of strictly dominated actions, contradicting unique rationalizability of a^* everywhere. Therefore some information set T_i must pool at least one state in S with at least one state in $\Omega \setminus S$.

Now fix such a T_i. Any event on which a^* is strictly dominant must be contained in S. Hence the maximal belief potential available from seed states at T_i cannot exceed

$$\mu(S \mid T_i).$$

If unique optimality of a^* at T_i requires belief potential at least $t_i(T_i)$, then necessarily

$$\mu(S \mid T_i) \geq t_i(T_i).$$

□

Proposition 8 identifies the bottleneck behind the upper bound in the main model: robust implementation must pass through a pooled type at which the architecture cannot assign more probability to the safe branch than to the posterior mass of the seed region itself.

The next corollary specializes this logic to the two-state benchmark.

Corollary 5 (Two-state upper bound)**.** *Consider the two-state coordination environment of Section 2, with seed region $S = \{l\}$ and $\mu(m) = q$. Under any admissible ex ante architecture, robust implementation of* No Attack *in both states requires*

$$1 - q \geq (1 - q)\theta_l + q\theta_m.$$

Equivalently,

$$q \leq \bar{q}_\infty := \frac{1 - \theta_l}{(1 - \theta_l) + \theta_m}.$$

Proof. By Proposition 8, robust implementation requires an information set T_i that pools l with m. In the two-state environment, the maximal belief potential of the seed at such a type cannot exceed

$$\mu(\{l\} \mid T_i).$$

Under the admissible ex ante class considered in the paper, this upper bound is $1-q$. The relevant threshold at the pooled type is the expected-payoff cutoff for *No Attack*,

$$(1-q)\theta_l + q\theta_m.$$

Therefore robust implementation requires

$$1-q \geq (1-q)\theta_l + q\theta_m,$$

which rearranges to the stated bound. □

Corollary 5 recovers the limiting frontier in the main model. Increasing epistemic depth raises belief potential at informed types and relaxes the Fine-m constraint, but it does not remove the pooled bottleneck. This is why the feasible region expands with depth and yet converges to a hard upper bound.