Coordination on a Grid: Posterior Ties, Entropy, and Multiplicity

Yihu Hou

April 3, 2026

Abstract

Information relevant for coordination is often reported in categories rather than as continuous signals. We show that reporting resolution is a distinct informational primitive, separate from measurement accuracy: coarse reporting creates exact posterior ties, making multiplicity more pervasive and equilibrium selection riskier. In a supermodular two-player coordination game with monotone cutoff equilibria, an additive translation-invariant benchmark yields a closed-form equilibrium interval whose width is governed by tie incidence. Finer reporting lowers ties and shrinks multiplicity, whereas on a fixed coarse grid higher measurement accuracy can concentrate mass on a few reports and widen the interval. The same tie statistic yields a natural entropy measure of uncertainty over posterior-relevant reports, under which the apparent Morris–Shin jump disappears in the near-complete-information limit. Beyond the benchmark, a tie–tilt decomposition extends the logic to more general information structures, and after integrating out arbitrary grid alignment tie risk yields a policy-relevant sensitivity curve identified by an estimable truncated moment of paired measurement differences.

1 Introduction

Coordination failures rarely occur in a world of perfectly continuous information. In practice, agents observe categories: a credit rating, an inspection grade, an alert level, a rounded risk score. A natural conjecture is that if the underlying measurement is already noisy, a finer reporting language should add little independent value. We show that this intuition is incomplete in coordination problems. Reporting resolution is a distinct dimension of information design: when information is represented on a grid, exact posterior ties arise with positive probability, and those ties make equilibrium multiplicity more pervasive. This

source of equilibrium-selection risk is absent in standard continuous formulations because exact ties occur with probability zero by construction.

The mechanism is simple. Discretisation creates atoms in the distribution of signal differences, so two agents can receive the same report with non-negligible probability. Near the boundary between the two actions, posterior clustering can flatten best responses and open the door to a whole range of self-consistent threshold rules. In a coordination environment, finer reporting therefore matters not only because it may improve learning. By reducing exact ties and local clustering, it disciplines equilibrium selection directly.

We make this idea precise in a standard supermodular two-player coordination environment with monotone cutoff equilibria. In an additive benchmark with a diffuse prior, we obtain a closed-form characterization of the full set of symmetric cutoff equilibria. The equilibrium cutoffs form an interval whose width decomposes into a mechanical grid term and a strategic term proportional to tie incidence. This makes the comparative statics transparent: multiplicity expands or contracts exactly with the frequency of ties, holding fixed strategic complementarity. Beyond the benchmark, we show that the same local logic survives through a tie–tilt decomposition: ties widen the set of admissible cutoffs, while directional asymmetries in boundary beliefs shift where the cutoff conditions bind.

These results separate resolution from accuracy. Resolution is the fineness of the reporting scale: how many categories are used and how wide they are. Accuracy is the dispersion of the underlying measurement error, shaped by data collection, auditing, sensor quality, or statistical filtering. The two margins affect outcomes differently. Refining the grid spreads mass across more categories, reduces ties, and shrinks the equilibrium set. Holding the grid fixed, higher accuracy can instead concentrate mass on a few reported values, increase ties, and widen the set of self-consistent cutoffs. After integrating out arbitrary grid alignment, that tie-based equilibrium-selection effect depends only on the ratio of measurement dispersion to bin width.

This distinction matters for policy because the two margins enter differently once equilibrium selection is taken seriously. Accuracy improves learning about the fundamental, but with coarse reporting it can also worsen equilibrium selection by increasing ties. Resolution instead moves only the selection margin, which makes it the cleaner instrument for disciplining that channel. Under the phase-robust summary, the equilibrium-selection side collapses to a one-dimensional tie-risk sensitivity curve indexed by the ratio of measurement dispersion to bin width. It can be identified by a truncated moment of paired measurement differences, so the policy-relevant sensitivity curve can be estimated from duplicate readings, parallel audits, or matched units exposed to the same shock.

This perspective also reinterprets the familiar discontinuity in the global-games bench-

mark. In the canonical Morris–Shin model with continuous signals, equilibrium is essentially unique for every strictly positive noise level, yet multiplicity returns at zero noise. That pattern is often read as saying that equilibrium outcomes react discontinuously to uncertainty itself. Our point is different. In the canonical continuous-signal benchmark, every strictly positive noise level generates an atomless report distribution, and atomlessness alone forces the tie statistic to zero. Uniqueness at positive noise is therefore a property of the continuous representation, not evidence that equilibrium selection reacts discontinuously to an underlying uncertainty primitive. Once representational resolution is treated as a primitive, the question becomes sharper: in the near-complete-information limit, multiplicity disappears only if reporting resolution improves faster than measurement noise shrinks.

This, in turn, changes how uncertainty should be measured. In richer information structures, noise variance does not exhaust the uncertainty relevant for coordination, because it does not record how posterior mass is grouped across reports. Measured only through variance, one can therefore see a jump near complete information that reflects the limitation of the measure rather than a genuine discontinuity in equilibrium responses. Section 8 returns to this point and shows that when uncertainty is measured by the entropy of the posterior-relevant report distribution, the equilibrium correspondence approaches its complete-information limit continuously.

The rest of the paper is organized as follows. Section 3 introduces the environment and establishes the cutoff structure. Section 4 derives the closed-form equilibrium characterization and identifies tie incidence as the statistic governing the width of the equilibrium set. Section 5 studies how resolution and accuracy shape posterior ties, including joint-limit results and phase-robust comparative statics. Section 6 moves beyond translation invariance and studies how exact ties in reported signals and directional asymmetries in boundary beliefs jointly shape the equilibrium set. It also identifies conditions under which the benchmark logic is approximately or exactly recovered. Section 7 separates learning from equilibrium-selection effects, develops a single-index description of the phase-robust equilibrium-selection effect, and identifies an estimable statistic for assessing when that marginal sensitivity is likely to be strong or weak. Section 8 concludes and discusses entropy as a measure of uncertainty over posterior-relevant reports and its connection to the public-information debate.

2 Related Literature

This paper studies coordination under incomplete information when signals are represented on a finite grid. Its main contribution is to treat reporting resolution as a distinct informational primitive, separate from measurement accuracy, and to trace how it affects

equilibrium selection. It connects to five strands of research: global games and equilibrium selection; higher-order beliefs, type spaces, and robustness; entropy and uncertainty measurement; information disclosure and the social value of information in coordination problems; and information design with *finite* or *coarse* signal spaces.

2.1 Global Games and Selection

The global-games approach, pioneered by Carlsson and van Damme (1993), shows how small amounts of incomplete information can select equilibria in coordination games. In canonical models of currency attacks and investment, noisy private signals yield a unique equilibrium in monotone threshold strategies, even as noise becomes arbitrarily small (Morris and Shin, 1998, 2001). A central lesson is that equilibrium selection is governed by higher-order beliefs and information structure rather than by payoff fundamentals alone. Frankel et al. (2003) extend this logic to broader classes of coordination games with strategic complementarities and study how selection behaves in the vanishing-noise limit.

Our contribution keeps the global-games focus on higher-order beliefs but treats reporting resolution as an additional primitive rather than folding all informational variation into noise. On a grid, exact posterior ties occur with positive probability, so the selection problem depends jointly on measurement accuracy and reporting resolution. The familiar Morris–Shin discontinuity is then a corollary of a broader point: near-complete information is a joint limit in accuracy and resolution, and multiplicity can persist when reporting remains coarse.

2.2 Higher-Order Beliefs and Robustness

A second strand studies coordination and equilibrium selection under general informational assumptions, emphasising the structure of type spaces and the behaviour of higher-order uncertainty. Classic contributions formalise approximate common knowledge and its role in coordination (Monderer and Samet, 1989; Morris and Shin, 1997). The robustness program of Kajii and Morris (1997) clarifies when complete-information equilibria survive perturbations of higher-order beliefs, and subsequent work develops further implications for selection and limits in incomplete-information games.

Within the global-games tradition, Hellwig (2002) relates equilibrium convergence to the relative speed at which higher-order uncertainty vanishes when both public and private information become precise, while Morris et al. (2016) provide common-belief foundations for global-games selection under general type spaces in terms of rank beliefs. More broadly, Weinstein and Yildiz (2007, 2011); Kajii and Morris (1997); Morris and Shin (1997) highlight

that equilibrium behaviour can be highly sensitive to higher-order beliefs even in otherwise well-behaved environments. Our contribution is to provide a tractable object for finite-resolution settings and to isolate the channel through which resolution matters for selection. Atoms in the signal-difference distribution generate posterior ties, and tie incidence maps higher-order uncertainty into equilibrium-set size. This perspective yields closed-form expressions in a benchmark and informative bounds more generally, complementing robustness approaches stated directly in terms of belief hierarchies.

2.3 Entropy and Uncertainty Measures

Our reinterpretation of the near-complete-information limit also connects to the information-theoretic literature that uses entropy to measure uncertainty on a distribution. Shannon entropy is the canonical benchmark (Shannon, 1948), and the Rényi family (Rényi, 1961) is especially relevant here because its order-two case is collision entropy, exactly the transform of the tie statistic κ that appears in our equilibrium formulas. The Tsallis family (Tsallis, 1988) provides a monotone reparameterisation on the zero-uncertainty margin relevant for our continuity result. We do not contribute to entropy theory itself. Rather, the point is that once equilibrium selection depends on how posterior mass is grouped across report categories, entropy becomes a natural way to describe the relevant uncertainty, and collision entropy is singled out by the identity between tie incidence and collision probability.

2.4 Transparency and Coordination

A large literature studies the value of information when agents face coordination motives. In the influential analysis of Morris and Shin (2002), public disclosures can improve alignment with fundamentals yet reduce welfare when coordination concerns lead agents to place excessive weight on the common signal relative to the social optimum, making actions overly sensitive to shared noise. Subsequent contributions refine and debate these forces and identify conditions under which transparency and public disclosures are beneficial or harmful, including environments with heterogeneous information, externalities, and investment complementarities (Hellwig, 2005; Svensson, 2006; Morris et al., 2006; Angeletos and Pavan, 2004, 2007).

A closely related insight appears in the global-games discussion of public versus private information. In Morris and Shin (2001), increasing the relative importance of common (public) information can reintroduce multiple monotone threshold equilibria. In finite-signal coordination games studied theoretically and experimentally, Anctil et al. (2010) show that greater transparency can destroy uniqueness and can reverse the risk-dominance ordering of

equilibria. We connect to this literature through an observable notion of effective commonality. In our framework, tighter clustering of beliefs shows up as more frequent posterior ties under coarse reporting, and those ties expand the set of monotone cutoff equilibria even though the mechanism differs from the standard public-signal one. Our point is narrower than a full welfare analysis of transparency: tie incidence provides a tractable object for studying how disclosure and reporting design affect equilibrium selection.

2.5 Coarse Information and Design

Finally, our paper relates to work on information design when signal spaces are constrained or coarse. The modern information-design and Bayesian-persuasion literature studies how a designer shapes actions through the choice of an information structure (Bergemann and Morris, 2019; Kamenica and Gentzkow, 2011; Kamenica, 2019). While much of this literature allows rich signal spaces, many applications involve discrete scores, categories, or ratings. Recent work explicitly studies *coarse* or *finite* information structures. In particular, Lyu et al. (2023) characterises information design with discrete signal spaces, while Rayo and Segal (2010) studies optimal disclosure through pooling and partitioning. A related applied literature on ratings and certification (e.g., Lizzeri, 1999) also highlights that intermediaries often optimally choose coarse categories.

Our focus differs from persuasion settings: rather than choosing an information structure, we study how a given reporting language reshapes higher-order beliefs and equilibrium selection in a coordination problem. The main implication is that representational resolution matters in its own right and is not redundant with measurement accuracy. The analysis also yields a tractable statistic, based on paired measurement differences, for assessing how reporting categories and measurement quality affect equilibrium-selection risk.

Experimental and empirical links. A growing experimental literature tests global-games and coordination predictions (e.g., Heinemann et al., 2004; Anctil et al., 2010). Our empirical object—a truncated moment of paired measurement differences—is directly estimable in laboratory and field settings with duplicate readings, parallel audits, or matched units facing the same realised shock. This makes the equilibrium-selection consequences of finite reporting resolution empirically measurable.

3 Model

There are two players $i \in \{1, 2\}$. Each player chooses a binary action $a_i \in \{0, 1\}$, where $a_i = 1$ denotes the "high" action (e.g. invest, adopt, comply).

The payoff-relevant fundamental is $\theta \in \Omega$, where

$$\Omega = \{\ldots, -2\Delta, -\Delta, 0, \Delta, 2\Delta, \ldots\}$$

for some resolution $\Delta > 0$. Payoffs are given by $u(a_i, a_j, \theta)$. Define the (ex post) gain from choosing the high action as

$$d(a_j, \theta) \equiv u(1, a_j, \theta) - u(0, a_j, \theta).$$

3.1 Complementarities and Monotonicity

Assumption 1 (Strategic complementarities)**.** *For all $\theta \in \Omega$, $d(1, \theta) \geq d(0, \theta)$.*

Assumption 1 imposes the standard coordination force: for a given fundamental, the gain from taking the high action is larger when the opponent also takes it. Actions are therefore mutually reinforcing.

Assumption 2 (Monotone fundamentals)**.** *For each $a_j \in \{0, 1\}$, the function $\theta \mapsto d(a_j, \theta)$ is weakly increasing.*

Assumption 2 links the fundamental to the attractiveness of the high action. Regardless of what the opponent does, higher fundamentals make $a_i = 1$ weakly more attractive. This is the usual monotonicity condition in binary-action coordination models.

3.2 Information and Beliefs

Each player observes a private signal $s_i \in \Omega$. The joint distribution of (θ, s_1, s_2) is common knowledge and symmetric across players. We impose two reduced-form monotonicity conditions: one on beliefs about the opponent's signal, and one on posteriors about the fundamental.

Assumption 3 (Monotone beliefs about others)**.** *For any cutoff $x \in \Omega$, the map $s_i \mapsto \Pr(s_j \geq x \mid s_i)$ is weakly increasing.*

Assumption 3 is a reduced-form restriction on *beliefs about the opponent*: a higher private signal makes player i weakly more optimistic that player j's signal exceeds any given threshold. Many standard signal structures imply this property, including common additive-noise models under familiar likelihood-ratio or TP_2 conditions; see Appendix A.2. The assumption is stated directly at the level of beliefs, not at the level of a specific signal-generating process. This is substantively weaker: many different primitives can induce the same monotone strategic belief comparison, and the analysis only needs that reduced-form ordering.

Assumption 4 (Monotone posteriors about fundamentals)**.** *For any weakly increasing function $g : \Omega \to \mathbb{R}$, the map $s_i \mapsto \mathbb{E}[g(\theta) \mid s_i]$ is weakly increasing.*

Assumption 4 is a reduced-form restriction on *posteriors*: higher private signals shift beliefs toward higher fundamentals. Its role is to deliver the single-crossing property that underlies cutoff strategies. By Assumption 2, for each a_j the payoff advantage $d(a_j, \theta)$ is weakly increasing in θ; applying Assumption 4 with $g(\theta) = d(a_j, \theta)$ yields that, for each a_j,

$$s_i \;\longmapsto\; \mathbb{E}[d(a_j, \theta) \mid s_i]$$

is weakly increasing. So higher signals make action $a_i = 1$ weakly more attractive for any fixed a_j, which is the single-crossing property we need.

The two assumptions discipline different channels. Assumption 4 concerns how s_i moves beliefs about θ (the *fundamental* channel), while Assumption 3 concerns how s_i moves beliefs about s_j (the *strategic* channel). In general joint distributions, neither assumption implies the other: learning that θ is likely higher need not translate into learning that the opponent's signal is likely higher, and vice versa. However, in common benchmark structures with conditional independence of signals given θ (i.e., $s_i \perp s_j \mid \theta$) and in which $s_j \mid \theta$ is stochastically increasing in θ, Assumption 3 follows from Assumption 4 by iterated expectations. A formal proof is provided in Appendix A.1.

Remark 1 (Additive signals as a benchmark special case). In Section 4 we specialise to the additive private-signal model

$$s_i = \theta + \varepsilon_i,$$

where $(\varepsilon_1, \varepsilon_2)$ are i.i.d. on Ω and independent of θ. In this benchmark, Assumption 4 is satisfied under familiar regularity conditions on the signal structure (e.g. a monotone likelihood ratio property); see Appendix A.2 for discrete sufficient conditions. Moreover, once Assumption 4 holds, Assumption 3 follows in this benchmark.

3.3 A Canonical Payoff Structure

The belief assumptions tell us how signals move posteriors and strategic beliefs. To translate those restrictions into a clean equilibrium characterization, we now impose a separable payoff structure standard in binary-action coordination models. This makes the interim payoff gain depend on the posterior only through a scalar fundamental term and a scalar coordination term, and it makes cutoff strategies natural on the ordered signal space Ω.

Assumption 5 (Separable payoff structure)**.** *There exist a function $v : \Omega \to \mathbb{R}$ and constants $b, c \in \mathbb{R}$ such that for all $\theta \in \Omega$ and $a_j \in \{0, 1\}$,*

$$d(a_j, \theta) = v(\theta) - c + b\, a_j.$$

Remark 2. Given Assumption 5, Assumptions 1 and 2 imply $b \geq 0$ and that $v(\cdot)$ is weakly increasing. We therefore keep b and v unrestricted in Assumption 5 and impose the economic content via Assumptions 1–2.

Assumption 5 says that the fundamental affects the net return to the high action through a single index $v(\theta)$, while the opponent's action adds a constant *coordination premium* b. In the usual investment/adoption/compliance interpretation, $v(\theta)$ captures direct profitability, c is the baseline cost or threshold for taking the high action, and b is the extra payoff from taking the high action when the opponent also takes the high action.

The payoff consequence is simple. For any opponent cutoff x and signal s, define the interim gain from choosing action 1 rather than action 0 as

$$\Delta(s; x) \equiv \mathbb{E}[v(\theta) \mid s_i = s] - c + b \Pr(s_j \geq x \mid s_i = s).$$

Interim incentives therefore depend on beliefs through two objects only: a posterior moment about fundamentals and the probability that the opponent is in the action region. That is exactly what we need to prove monotone best responses without committing to a specific signal-generating process.

3.4 Monotone Bayesian Equilibria

Definition 1 (Cutoff strategy)**.** *A (pure) strategy $a_i(s_i) : \Omega \to \{0, 1\}$ is a cutoff strategy if there exists $x \in \Omega$ such that*

$$a_i(s_i) = \mathbf{1}\{s_i \geq x\}.$$

Fix $x \in \Omega$ and suppose player j plays $a_j = \mathbf{1}\{s_j \geq x\}$. Given signal $s_i = s$, the interim payoff gain from choosing $a_i = 1$ rather than $a_i = 0$ is $\Delta(s; x)$. A best response is a cutoff if $s \mapsto \Delta(s; x)$ is weakly increasing, so that the set $\{s : \Delta(s; x) \geq 0\}$ is an upper set of Ω.

Lemma 1 (Monotone best response)**.** *Under Assumptions 1, 2, 3, 4, and 5, for any opponent cutoff $x \in \Omega$ the function $s \mapsto \Delta(s; x)$ is weakly increasing. Hence player i admits a (pure) cutoff best response.*

Proof. Fix $x \in \Omega$. Under Assumption 5,

$$\Delta(s; x) = \mathbb{E}[v(\theta) \mid s_i = s] - c + b \Pr(s_j \geq x \mid s_i = s).$$

By Assumption 2 applied to $d(0,\theta) = v(\theta) - c$, the index $v(\cdot)$ is weakly increasing. Hence Assumption 4 implies $s \mapsto \mathbb{E}[v(\theta) \mid s_i = s]$ is weakly increasing. By Assumption 1 and $d(1,\theta) - d(0,\theta) = b$, we have $b \geq 0$, and Assumption 3 implies $s \mapsto \Pr(s_j \geq x \mid s_i = s)$ is weakly increasing. Therefore $\Delta(s;x)$ is weakly increasing in s, so the best-response set $\{s : \Delta(s;x) \geq 0\}$ is an upper set of Ω, i.e. a cutoff. □

Lemma 1 gives the closure step: against a cutoff strategy, a best response can also be chosen as a cutoff. The next theorem adds the fixed-point step. It shows that within the cutoff class a symmetric equilibrium exists, and that the set of equilibrium cutoffs has extremal elements.

Theorem 1 (Existence of a symmetric cutoff equilibrium)**.** *Maintain Assumptions 1, 2, 3, 4, and 5. Consider the class of cutoff strategies $a^x(s) = \mathbf{1}\{s \geq x\}$ indexed by $x \in \Omega$. Then there exists a symmetric cutoff Bayesian Nash equilibrium, i.e., there exists $x^* \in \Omega$ such that a^{x^*} is a best response to itself. Moreover, among symmetric cutoff equilibria there exist a smallest and a largest equilibrium cutoff, denoted $\underline{x}$ and $\overline{x}$.*

Proof sketch. Lemma 1 implies that against any cutoff a^x, a best response can be chosen to be a cutoff; hence $BR(x)$ is nonempty for every x. Strategic complementarities (Assumption 1) imply that the best-response operator is isotone under the pointwise order on strategies; restricted to cutoff strategies, this yields monotonicity of the cutoff best-response correspondence.[1] Tarski's fixed-point theorem then yields existence, and moreover delivers the extremal cutoff equilibria. For a complete proof, see Appendix A.3. □

Theorem 1 justifies the paper's focus on symmetric cutoff equilibria indexed by a single scalar $x \in \Omega$. The existence claim is what we need immediately; the extremal cutoffs will matter later when we compare equilibrium multiplicity and welfare across information environments.

4 Benchmark: Posterior Ties and the Equilibrium Set

Section 3 established a general monotone-cutoff logic: under monotone beliefs and a canonical payoff structure (Assumption 5), best responses are cutoff rules and symmetric cutoff equilibria exist (Theorem 1). This section has a narrower goal. We specialise to a translation-invariant additive benchmark in which the full set of symmetric cutoff equilibria can be solved in closed form, and we isolate the statistic that governs its size.

[1] Formally, the cutoff strategy set is a complete lattice under the pointwise order, and the best-response operator is isotone; Tarski's fixed-point theorem yields existence and extremal fixed points. See, e.g., Topkis (1998, Ch. 2) or Vives (1990).

The key benchmark fact is that on a discrete state space $\Omega = \Omega_\Delta$ players can be *exactly tied* in posterior-relevant objects with positive probability. Those ties drive multiplicity here, and in the additive benchmark they are summarized by a simple tie probability.

4.1 Finite Resolution and Posterior Ties

In this section we specialise to an additive benchmark,

$$s_i = \theta + \varepsilon_i, \qquad i \in \{1, 2\},$$

with i.i.d. idiosyncratic shocks $\varepsilon_1, \varepsilon_2$ that are independent of θ. All random variables take values in a discrete, totally ordered set $\Omega = \Omega_\Delta = \{k\Delta : k \in \mathbb{Z}\}$, where $\Delta > 0$ is the *resolution*.

We use Δ to capture a common institutional feature: fundamentals and signals are recorded and communicated in a finite alphabet, through rounding rules, reporting conventions, rating buckets, or discretised scores. The benchmark treats this grid as part of the environment and asks how it reshapes higher-order beliefs and equilibrium multiplicity.

A distinctive feature of discrete environments is that two independent noise draws can coincide *exactly* on the grid with positive probability. That matters because cutoff incentives compare neighboring grid points, and exact ties are precisely what prevent those boundary comparisons from being strict.

Let p_ε denote the common pmf of ε_i on Ω_Δ. We therefore define the key statistic:

Definition 2 (Tie probability)**.** *The* tie probability *in the posterior-relevant noise is*

$$\kappa \;\equiv\; \Pr(\varepsilon_1 = \varepsilon_2) \;=\; \sum_{u \in \Omega_\Delta} p_\varepsilon(u)^2. \tag{1}$$

The interpretation is immediate: κ is the probability that two agents' posterior-relevant shocks are *exactly tied*. In information-theoretic terms, κ is the collision probability, and the associated Rényi collision entropy is $H_2(\varepsilon) \equiv -\log \kappa$ (Rényi, 1961). Our comparative statics for κ can therefore also be read as comparative statics for collision entropy: more tie risk means lower H_2, and $H_2(\varepsilon) = 0$ iff the posterior-relevant shock is a point mass, so the environment has collapsed to complete information. We work directly with κ because it enters the equilibrium formulas linearly.

4.2 Translation Invariance and Normalisation

To obtain closed-form expressions, we impose a translation-invariance condition standard in the global-games benchmark. After observing s_i, player i learns about the opponent only through *signal differences*: shifting the realised signal level does not change the conditional distribution of $s_j - s_i$.

Assumption 6 (Translation invariance)**.** *For every measurable set $A \subseteq \Omega$, the probability $\Pr(s_j - s_i \in A \mid s_i = s)$ is independent of s.*

A convenient sufficient condition for Assumption 6 is the canonical "diffuse prior + i.i.d. additive noise" benchmark used in the global-games tradition. Here "diffuse" simply means globally flat: the prior places equal weight on all locations, so only signal differences matter for higher-order beliefs.[2] Under this benchmark, $s_j - s_i = \varepsilon_j - \varepsilon_i$, so the conditional law of $s_j - s_i$ is independent of s_i and Assumption 6 holds automatically.

Translation invariance is a simplifying benchmark, not the source of the mechanism. It holds exactly under diffuse priors, and it can also be a good approximation when the prior is sufficiently flat over the region where posterior mass concentrates. Technically, Assumption 6 makes beliefs about the opponent depend only on relative positions: for any cutoff $x \in \Omega$, $\Pr(s_j \geq x \mid s_i = s)$ is a function of $x - s$. This removes level effects from the prior, reduces equilibrium conditions to local boundary comparisons, and makes the closed form depend on a single tie statistic. Section 6 shows what survives once this simplification is dropped.

Next, recall from Section 3 that under the separable payoff structure (Assumption 5), the fundamental enters interim incentives through the scalar

$$m(s) \equiv \mathbb{E}[v(\theta) \mid s_i = s].$$

In the canonical investment benchmark with $v(\theta) = \theta$, additive signals $s_i = \theta + \varepsilon_i$, and a diffuse prior on θ, the posterior distribution satisfies $(\theta \mid s_i = s) \overset{d}{=} s - \varepsilon_i$.[3] If, in addition, the noise is unbiased, $\mathbb{E}[\varepsilon_i] = 0$, then

$$m(s) = \mathbb{E}[\theta \mid s_i = s] = \mathbb{E}[s - \varepsilon_i] = s. \tag{2}$$

[2] A diffuse prior is a translation-invariant prior that assigns equal weight to all locations on the real line (or on the grid Ω). Economically, it captures a situation in which agents share no informative common prior anchor about the level of the fundamental, so only relative signal differences matter for higher-order beliefs. This is a standard device in global games to isolate the role of higher-order uncertainty; see, e.g., Morris and Shin (2001); Hellwig (2002).

[3] This is the standard translation-invariance property of the diffuse prior: conditioning on s shifts the posterior one-for-one.

Equation (2) is a convenient normalisation, not an extra economic assumption. For any strictly increasing m, one can work in transformed-signal space $\tilde{s} \equiv m(s)$ and then map cutoffs back via m^{-1}. We keep (2) only to obtain closed-form endpoints and keep the algebra focused on ties.

When the noise is not unbiased, let $\chi \equiv \mathbb{E}[\varepsilon_i]$. Under the same translation-invariant posterior,

$$m(s) = \mathbb{E}[\theta \mid s_i = s] = \mathbb{E}[s - \varepsilon_i] = s - \chi.$$

Relative to (2), this changes only the *location* of the relevant index: the mapping $s \mapsto m(s)$ is a pure translation by χ. As a result, any cutoff interval characterised in m-space has the same *width* as under unbiased noise; only its position on the raw signal scale shifts.

More importantly, this translation does not change the substantive equilibrium behaviour. A cutoff rule stated in terms of $m(s)$ is invariant to relabelling of the raw signal by a constant. Concretely, suppose two environments differ only in the noise mean, with $\mathbb{E}[\varepsilon_i] = \chi$ in one case and $\mathbb{E}[\varepsilon_i] = 0$ in the other. Then a cutoff at signal level $s^* + \chi$ in the first environment is equivalent to a cutoff at s^* in the second one. This shift leaves agents' actions and all equilibrium event probabilities unchanged; it merely expresses the same decision rule in a translated signal language. In this sense, relaxing unbiasedness does not affect the coordination mechanism driven by discrete ties; it only re-centres the signal index, without rescaling and without altering the equilibrium set up to this trivial translation.

4.3 Closed Form and the Role of Ties

We now characterise *all* symmetric cutoff equilibria in the benchmark. Given a cutoff $x \in \Omega$, write $a^x(s) = \mathbf{1}\{s \geq x\}$ for the corresponding cutoff strategy. Define the set of symmetric cutoff equilibria by

$$\mathcal{E} \equiv \{\, x \in \Omega : \ (a^x, a^x) \text{ is a Bayesian Nash equilibrium} \,\}.$$

Theorem 2 (Closed-form equilibrium set in the additive benchmark)**.** *Suppose Assumptions 1, 2, 4, 5, and 6 hold, as well as the additive signal structure, and the benchmark normalisation $m(s) = s$ applies.*[4] *Then the set of symmetric cutoff equilibria $\mathcal{E}$ is exactly*

$$\mathcal{E} = \left[c - \frac{b}{2} - \frac{b\kappa}{2}, \quad c + \Delta - \frac{b}{2} + \frac{b\kappa}{2} \right) \cap \Omega.$$

Hence its width is $\Delta + b\kappa$, which is increasing in the tie probability κ.

[4] In the additive benchmark, Assumption 3 is redundant: it follows from Assumption 4 under conditional independence (Lemma 9).

Remark 3 (Mapping back via m^{-1}). The normalisation $m(s) = s$ is adopted only for expositional convenience. If m is strictly increasing, the same closed-form interval applies in transformed-signal space $\tilde{s} = m(s)$, and the corresponding cutoffs in raw signal space are obtained by applying m^{-1} to the two endpoints from Theorem 2.

Proof sketch. Fix a candidate cutoff $x \in \Omega$ and suppose player j plays $a_j = \mathbf{1}\{s_j \geq x\}$. Under Assumption 5, player i's interim gain from choosing 1 rather than 0 at signal s is

$$\Delta(s; x) = m(s) - c + b \Pr(s_j \geq x \mid s_i = s),$$

which is weakly increasing in s by Lemma 1. Hence x is a symmetric cutoff equilibrium iff the boundary conditions $\Delta(x; x) \geq 0$ and $\Delta(x - \Delta; x) < 0$ hold.

With the benchmark normalisation $m(s) = s$ and translation invariance, these boundary conditions depend only on the distribution of the signal difference $s_j - s_i = \varepsilon_j - \varepsilon_i$. By symmetry,

$$\Pr(\varepsilon_j - \varepsilon_i \geq 0) = \frac{1}{2} + \frac{1}{2} \Pr(\varepsilon_i = \varepsilon_j) = \frac{1}{2} + \frac{\kappa}{2},$$

where the tie probability $\Pr(\varepsilon_i = \varepsilon_j) = \kappa$ is defined in (1). Moreover, because the grid has no support points in $(0, \Delta)$,

$$\Pr(\varepsilon_j - \varepsilon_i \geq \Delta) = \Pr(\varepsilon_j - \varepsilon_i > 0) = \frac{1}{2} - \frac{\kappa}{2}.$$

Substituting these expressions into the two boundary inequalities yields the stated interval for $\mathcal{E}$, whose width is $\Delta + b\kappa$. The full proof is given in Appendix A.4. □

Theorem 2 shows why reporting resolution matters in its own right. In this benchmark, multiplicity is governed by one discrete object: the tie probability $\kappa = \Pr(\varepsilon_1 = \varepsilon_2)$ from Definition 2. Figure 1 summarises the logic visually: discretisation creates an atom at zero in the signal-difference distribution, and that atom widens the admissible cutoff region symmetrically around the no-tie benchmark. The reason is local. On a grid, cutoff equilibrium is pinned down by the two boundary comparisons at $s = x$ and $s = x - \Delta$. At those signals, the only source of non-strictness is an exact tie in the posterior-relevant shock. When $\kappa > 0$, the two inequalities can both hold over an interval of cutoffs, so the set $\mathcal{E}$ becomes non-degenerate. A finer reporting language is therefore not merely a more detailed description of the same noisy measurement. It changes the local boundary comparisons that select equilibrium.

The formula also separates resolution from accuracy. Holding the underlying measurement environment fixed, a finer grid spreads probability mass across more reported values,

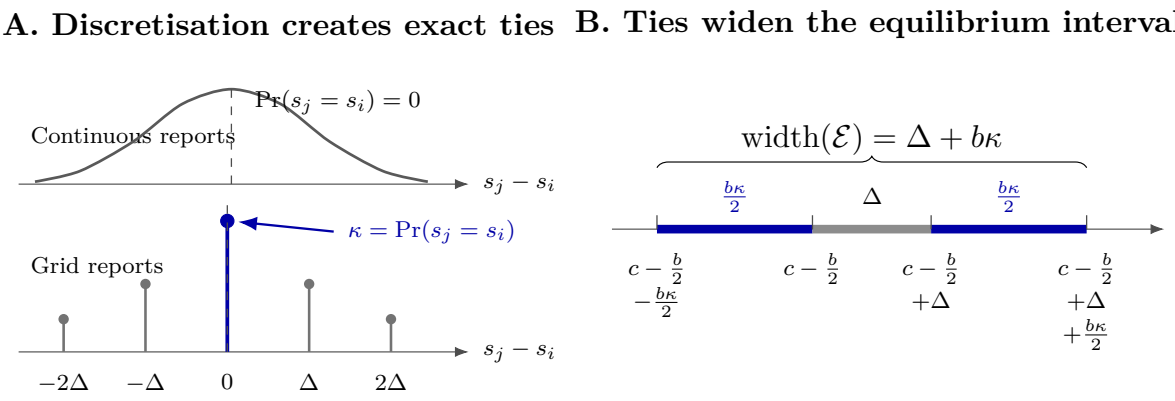

Figure 1: Mechanism in the additive benchmark. Left: discretisation places positive mass at zero in the signal-difference distribution, creating the tie probability $\kappa = \Pr(s_j = s_i)$. Right: ties expand the equilibrium cutoff interval symmetrically by $b\kappa/2$ on each side, so $\text{width}(\mathcal{E}) = \Delta + b\kappa$ instead of the no-tie width Δ.

lowers tie probability, and shrinks the equilibrium interval. Holding Δ fixed, more accurate measurements can instead concentrate mass on a few grid points, raise tie probability, and widen the interval. Resolution and accuracy therefore need not move together, even though both concern the information environment. Section 5 turns these heuristics into formal results.

These comparative statics clarify what "improving information" means in applications. Accuracy improves each agent's measurement of the fundamental, but on a coarse grid it can also increase ties and expand the set of self-fulfilling coordination outcomes. Resolution works more directly on equilibrium selection: it refines the reporting language and reduces ties without changing measurement accuracy. The policy question is therefore not only whether to improve information, but on which margin.

This perspective also has an implication for the familiar discontinuity in the original global-games limit. In the canonical Morris–Shin environment with continuous signals, equilibrium remains essentially unique at every strictly positive noise level, yet multiplicity returns once idiosyncratic noise vanishes (Carlsson and van Damme, 1993; Morris and Shin, 1998). Theorem 2 shows why: atomlessness keeps $\kappa = 0$ for every strictly positive noise level. At zero noise, by contrast, both agents observe the fundamental itself, so ties return with probability one and $\kappa = 1$. Once resolution is finite, tie probability can remain positive even when noise is arbitrarily small, so whether multiplicity disappears depends on the relative speed at which *accuracy* improves ($\sigma \downarrow 0$) and *resolution* improves ($\Delta \downarrow 0$). Section 5.2 makes that joint-limit dependence precise. It also anticipates our later argument that in rich information structures noise variance does not fully capture the uncertainty relevant for

coordination, so the apparent jump near complete information partly reflects the limitation of that metric rather than a genuine discontinuity in equilibrium responses.

5 Comparative Statics: Resolution and Accuracy

Section 4 shows that in the translation-invariant additive benchmark the equilibrium set is an interval whose width equals $\Delta + b\kappa$ (Theorem 2). This section makes precise how the two primitives —*resolution* Δ and *accuracy* σ—shape posterior ties through the tie statistic κ, and therefore the size of the equilibrium set. The point is to keep distinct two margins that are often conflated: the accuracy of the underlying measurement and the fineness of the reporting language.

To keep reporting resolution separate from measurement accuracy, we model reported signals as discretisations of an underlying continuous measurement error ε^c on $\mathbb{R}$. Changes in the measurement environment correspond to changes in the continuous distribution of ε^c. By contrast, resolution is a property of the reporting rule: it specifies a grid $\Omega_\Delta = \{k\Delta : k \in \mathbb{Z}\}$ according to which the continuous scale is grouped into finitely or countably many categories. Holding fixed the underlying continuous scale, changing Δ alters only how probability mass is aggregated across categories. Even when the latent measurement environment is held fixed, that regrouping changes how posterior mass is clustered across reported values and therefore changes tie incidence.

To discretise the continuous scale, one must also specify how category boundaries are aligned relative to the grid points. We capture this alignment by a *phase* (grid-origin) parameter $\phi \in \mathbb{R}$.[5] For a given (Δ, ϕ), define the induced pmf $p_\varepsilon^{\Delta,\phi}$ on Ω_Δ by[6]

$$p_\varepsilon^{\Delta,\phi}(u) \equiv \Pr\big(\varepsilon^c \in [u+\phi, u+\phi+\Delta)\big), \qquad u \in \Omega_\Delta. \tag{3}$$

where ε^c denotes the underlying continuous error.

Three points need to be separated. Refining the reporting grid lowers κ by a purely algebraic splitting argument and needs no further shape restriction on the underlying continuous distribution. By contrast, monotonicity in *accuracy* does not follow without additional structure. We therefore proceed in four steps: first, how nested refinement of the grid affects ties; second, accuracy in a location–scale family together with the small-Δ approximation

[5]For implementation-specific questions, ϕ is part of the reporting rule. For comparative statics in the mesh width Δ, however, ϕ is an alignment parameter rather than the primitive of interest; Section 5.4 integrates it out to obtain a phase-robust tie index.

[6]Since ε^c is continuously distributed, bin endpoints have zero probability. Hence the choice of half-open versus closed bins is immaterial. If F_{ε^c} denotes the cdf of ε^c, then $p_\varepsilon^{\Delta,\phi}(u) = F_{\varepsilon^c}(u+\phi+\Delta) - F_{\varepsilon^c}(u+\phi)$.

behind the joint-limit results; third, global monotonicity under symmetry and log-concavity; and fourth, a phase-robust summary, obtained by averaging over alignment, which yields a single-ratio description in (Δ, σ).

5.1 Resolution: Refinement Lowers Tie Probability

Fix an underlying continuously distributed measurement error ε^c and an alignment parameter ϕ, and construct $p_\varepsilon^{\Delta,\phi}$ via (3). We hold ϕ fixed throughout so that comparative statics in Δ capture genuine refinements of the grid rather than shifts in its alignment. Define the associated tie probability

$$\kappa_\Delta \;\equiv\; \sum_{u\in\Omega_\Delta} \big(p_\varepsilon^{\Delta,\phi}(u)\big)^2.$$

For any integer $n \geq 2$, let $\Delta' = \Delta/n$ denote an n-fold refinement of the representational grid.

Proposition 1 (Refinement lowers tie probability)**.** *For $\Delta' = \Delta/n$ with $n \geq 2$,*

$$\kappa_{\Delta'} \;\leq\; \kappa_\Delta,$$

with strict inequality whenever ε^c assigns positive probability mass to at least two distinct subcells of some coarse cell $[u+\phi, u+\phi+\Delta)$. Consequently, in the benchmark of Theorem 2, the equilibrium interval width $\Delta + b\kappa_\Delta$ is weakly decreasing under refinement.

Proof. For each coarse grid point $u \in \Omega_\Delta$, the coarse-cell probability is

$$p(u) \equiv p_\varepsilon^{\Delta,\phi}(u) = \Pr\big(\varepsilon^c \in [u+\phi, u+\phi+\Delta)\big).$$

Under refinement, the same coarse cell $[u+\phi, u+\phi+\Delta)$ is partitioned into n subcells of length Δ', namely $[u+\phi+r\Delta',\, u+\phi+(r+1)\Delta')$ for $r = 0, \ldots, n-1$, with probabilities

$$p_r(u) \;\equiv\; p_\varepsilon^{\Delta',\phi}(u+r\Delta') = \Pr\big(\varepsilon^c \in [u+\phi+r\Delta',\, u+\phi+(r+1)\Delta')\big), \qquad r = 0, \ldots, n-1,$$

so that $\sum_{r=0}^{n-1} p_r(u) = p(u)$. Since $p_r(u) \geq 0$,

$$p(u)^2 = \Big(\sum_{r=0}^{n-1} p_r(u)\Big)^2 = \sum_{r=0}^{n-1} p_r(u)^2 + 2 \sum_{0\leq r<q\leq n-1} p_r(u)\, p_q(u) \;\geq\; \sum_{r=0}^{n-1} p_r(u)^2,$$

with strict inequality if at least two $p_r(u)$ are strictly positive. Summing over coarse cells

yields

$$\kappa_{\Delta'} = \sum_{u\in\Omega_\Delta} \sum_{r=0}^{n-1} p_r(u)^2 \;\le\; \sum_{u\in\Omega_\Delta} p(u)^2 = \kappa_\Delta.$$

Finally, Theorem 2 gives width $\Delta + b\kappa_\Delta$, which weakly decreases when both Δ and κ_Δ weakly decrease. □

Proposition 1 formalises the heuristic in Section 4: holding fixed the underlying measurement environment, a finer representational grid spreads probability mass over more grid points and lowers the probability that two independent draws land on the same point. This nested-refinement comparative static is entirely mechanical: no likelihood-ratio or log-concavity assumption is needed. In the benchmark of Theorem 2, the equilibrium interval has width $\Delta + b\,\kappa_\Delta$, so refinement shrinks the multiplicity region through both terms at once. In entropy terms, refinement weakly raises collision entropy $H_2 = -\log \kappa_\Delta$, while coarsening lowers it.

5.2 Joint Limits: Accuracy Versus Resolution

We next relate tie probability to *accuracy* through a convenient one-parameter family that scales the dispersion of the underlying continuous noise. Specifically, we parameterise measurement error as a location–scale family,

$$\varepsilon_\sigma^c \;=\; \sigma\,\xi, \qquad \sigma > 0,$$

where ξ has a continuous cdf F and density f, and its distribution has nonempty interior support. The parameter σ governs the overall spread of the noise: smaller σ makes the distribution more concentrated, and the limit $\sigma \downarrow 0$ captures the near-complete-information regime. For a given representational resolution Δ and phase ϕ, discretisation on Ω_Δ yields the pmf

$$p_{\varepsilon,\sigma}^{\Delta,\phi}(u) \equiv \Pr\left(\varepsilon_\sigma^c \in [u+\phi, u+\phi+\Delta)\right) = F\!\left(\frac{u+\phi+\Delta}{\sigma}\right) - F\!\left(\frac{u+\phi}{\sigma}\right), \qquad u \in \Omega_\Delta, \quad (4)$$

and tie probability

$$\kappa_\Delta(\sigma) \;\equiv\; \sum_{u\in\Omega_\Delta} \left(p_{\varepsilon,\sigma}^{\Delta,\phi}(u)\right)^2.$$

In what follows we fix Δ and ϕ and study how $\kappa_\Delta(\sigma)$ varies with the accuracy parameter σ.

For general F, $\kappa_\Delta(\sigma)$ need not be globally monotone in σ without additional restrictions. Nevertheless, the fine-grid regime admits a clean and robust characterisation, which is the

relevant input for the joint-limit analysis below. Moreover, under standard shape restrictions (e.g. symmetry and log-concavity), one can strengthen the accuracy comparative static to a global monotonicity result for any fixed Δ; see Section 5.3.

Lemma 2 (Small-Δ approximation)**.** *Assume $f \in L^2(\mathbb{R})$. Fix $\sigma > 0$ and let $\kappa_\Delta(\sigma)$ be defined above. As $\Delta \downarrow 0$,*

$$\frac{\kappa_\Delta(\sigma)}{\Delta} \;\longrightarrow\; \int_{\mathbb{R}} f_\sigma(x)^2\,dx \;=\; \frac{1}{\sigma}\int_{\mathbb{R}} f(u)^2\,du, \qquad f_\sigma(x) \equiv \frac{1}{\sigma} f(x/\sigma).$$

Proof. By (4), for each $u \in \Omega_\Delta$,

$$p^{\Delta,\phi}_{\varepsilon,\sigma}(u) = \int_{u+\phi}^{u+\phi+\Delta} f_\sigma(x)\,dx, \qquad \frac{\kappa_\Delta(\sigma)}{\Delta} = \sum_{u\in\Omega_\Delta} \Delta\left(\frac{1}{\Delta}\int_{u+\phi}^{u+\phi+\Delta} f_\sigma(x)\,dx\right)^2.$$

The term in parentheses is the cell average of f_σ on $[u+\phi, u+\phi+\Delta)$. As $\Delta \downarrow 0$, these averages converge to f_σ pointwise a.e.; since $f_\sigma \in L^2(\mathbb{R})$, the Riemann sums above converge to $\int_{\mathbb{R}} f_\sigma(x)^2\,dx$. Finally, a change of variables gives $\int f_\sigma^2 = (1/\sigma)\int f^2$. □

Lemma 2 gives the local approximation that links accuracy and resolution:

$$\kappa_\Delta(\sigma) \;\approx\; \frac{\Delta}{\sigma}\int_{\mathbb{R}} f(u)^2\,du.$$

This is a fine-grid statement: it is accurate when Δ/σ is small, not when almost all mass lies in a single cell. In that regime, improving accuracy (smaller σ) raises tie probability in proportion to Δ/σ. That scaling is exactly what drives the near-complete-information limit.

We now use this fine-grid approximation to formalise the joint-limit perspective highlighted in Section 4. In the fine-grid regime (formally, when Δ/σ is small),

$$\text{width}(\mathcal{E}) \;\approx\; \Delta + b\,\frac{\Delta}{\sigma}\int_{\mathbb{R}} f(u)^2\,du.$$

Thus, whether near-complete information yields uniqueness or multiplicity depends on the relative speed at which accuracy improves ($\sigma \downarrow 0$) and resolution improves ($\Delta \downarrow 0$).

Corollary 1 (Near-complete-information limits)**.** *Assume the benchmark of Theorem 2 and the conditions of Lemma 2. Fix a phase ϕ and consider a sequence (Δ_n, σ_n) with $\Delta_n \downarrow 0$ and $\sigma_n \downarrow 0$.*

(i) If $\Delta_n/\sigma_n \to 0$, then $\kappa_{\Delta_n}(\sigma_n) \to 0$ and the equilibrium set shrinks to a singleton in the limit.

(ii) If $\liminf_{n\to\infty} \Delta_n/\sigma_n > 0$*, then* $\liminf_{n\to\infty} \kappa_{\Delta_n}(\sigma_n) > 0$*, so the equilibrium interval does not collapse.*

(iii) If $\Delta_n/\sigma_n \to \infty$*, then* $\liminf_{n\to\infty} \kappa_{\Delta_n}(\sigma_n) \geq 1/3$*. Moreover, if in addition*

$$\sum_{n=1}^{\infty} \frac{\sigma_n}{\Delta_n} \;<\; \infty, \tag{5}$$

then $\kappa_{\Delta_n}(\sigma_n) \to 1$ *for Lebesgue-a.e. phase* ϕ *(i.e. except for a measure-zero set of phases).*

Proof sketch. Part (i) follows from the small-Δ approximation: when the grid is much finer than the noise scale ($\Delta_n/\sigma_n \to 0$), tie probability is of order Δ_n/σ_n and therefore vanishes, so the equilibrium interval shrinks to a point.

For part (ii), if Δ_n/σ_n is bounded away from 0, the bin that contains 0 has width at least a fixed multiple of σ_n. Since it straddles 0, it must contain a fixed-size neighbourhood of 0 in σ_n-units, so it captures a uniformly positive probability mass. This yields a uniform positive lower bound on tie probability, hence the equilibrium interval cannot collapse.

For part (iii), when $\Delta_n/\sigma_n \to \infty$ the distribution becomes extremely concentrated relative to the bin width. In particular, almost all probability mass lies in a narrow neighbourhood of 0 of width $2\Delta_n$, which can intersect at most three adjacent bins on a Δ_n-grid. Therefore the discretised distribution is asymptotically supported on at most three grid points, implying $\liminf \kappa_{\Delta_n}(\sigma_n) \geq 1/3$. Under the summability condition (5), for Lebesgue-a.e. phase the bin boundary is eventually far from 0 relative to σ_n, so almost all mass falls into a single bin and $\kappa_{\Delta_n}(\sigma_n) \to 1$. Appendix A.5 provides the details. □

Corollary 1 partitions near-complete-information limits by the single ratio Δ_n/σ_n. If refinement outpaces accuracy, ties vanish and asymptotic uniqueness obtains. If Δ_n/σ_n stays bounded away from zero, some cell keeps positive mass and multiplicity persists. If $\Delta_n/\sigma_n \to \infty$, noise is negligible relative to the reporting grid, so almost all mass is forced into the few cells adjacent to zero. Without further control on phase this gives only the universal lower bound $\liminf \kappa_{\Delta_n}(\sigma_n) \geq 1/3$, because in the worst alignment the vanishing mass can still straddle up to three neighbouring bins. Under the summability condition in part (iii), that boundary pathology occurs only on a measure-zero set of phases, so for almost every phase all mass eventually falls into a single bin and $\kappa_{\Delta_n}(\sigma_n) \to 1$.

When part (iii) yields $\kappa_{\Delta_n}(\sigma_n) \to 1$, Theorem 2 implies that the equilibrium interval

$$\left[c - \frac{b}{2} - \frac{b\kappa_{\Delta_n}(\sigma_n)}{2},\; c + \Delta_n - \frac{b}{2} + \frac{b\kappa_{\Delta_n}(\sigma_n)}{2}\right)$$

converges to $[c-b, c]$ once $\Delta_n \to 0$. In the benchmark normalisation, $[c-b, c]$ is exactly

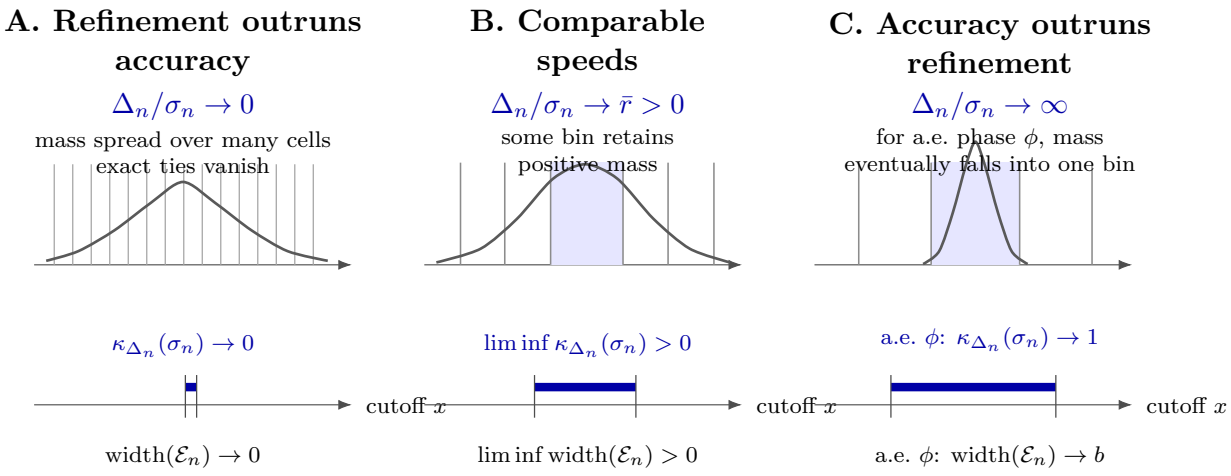

Figure 2: Joint limits in accuracy and resolution. Left: refinement outruns accuracy, so ties and the equilibrium interval vanish. Middle: comparable speeds leave a positive limiting interval. Right: under the summability condition in Corollary 1, accuracy outruns refinement, and for a.e. phase ties converge to one while the equilibrium interval converges to $[c-b, c]$.

the complete-information multiplicity region: if $s > c$, action 1 is strictly dominant, whereas if $s < c - b$, action 0 is strictly dominant, so strategic coordination matters only on that closed interval, whose width is exactly b. For every threshold $\theta \in [c-b, c]$, the cutoff strategy that chooses action 1 if and only if the fundamental is at least θ is a complete-information equilibrium. Thus, when accuracy outruns refinement, the model approaches complete-information multiplicity rather than uniqueness.

Figure 2 summarises the three regimes. The relevant comparison is not whether noise vanishes in isolation, but whether the shrinking noise scale is dominated by, matched to, or dominates the shrinking bin width.

This comparison also clarifies the Morris–Shin benchmark. With continuous reports, atomlessness forces $\kappa = 0$ at every strictly positive noise level, so vanishing noise alone selects uniqueness. With finite reporting resolution, ties need not vanish as noise vanishes. What matters is whether resolution improves fast enough relative to accuracy. The joint-limit analysis gives a simpler reading of the Morris–Shin jump. It is not best viewed as a discontinuous equilibrium response to uncertainty itself. In richer information structures, an apparent jump can arise because noise variance is an incomplete measure of uncertainty: it does not record how posterior mass is concentrated across reports. A natural alternative is to borrow entropy from information theory as the measure of uncertainty. In the present benchmark, we already have κ, the collision probability of the posterior-relevant report

distribution, so the corresponding object is the collision entropy

$$H_2(p) \equiv -\log \sum_u p(u)^2 = -\log \kappa.$$

Lower collision entropy corresponds exactly to more tie risk. Measured this way, the apparent jump disappears in the relevant near-complete-information limit. Section 8 returns to this interpretation and places it in the broader Rényi family.

We now leave the joint-limit comparison and return to fixed-Δ comparative statics. The next subsection asks when higher accuracy has a globally monotone effect on tie probability under standard shape restrictions.

5.3 Accuracy: Global Monotonicity under Log-Concavity

The preceding subsection characterised the joint limit. Here we keep Δ fixed and ask when higher accuracy has a global effect on tie probability. Without further restrictions, monotonicity need not hold. Under standard shape conditions, however, it does: for a symmetric log-concave noise distribution, shrinking the scale parameter systematically pulls mass toward the centre, which is enough to raise the sum-of-squares functional $\sum_k p_k^2$.

To make the ordering argument transparent, we use the *symmetric* binning, i.e. we choose the phase $\phi = -\Delta/2$ so that the bins are centred at 0:

$$[k\Delta + \phi,\ k\Delta + \phi + \Delta) = \Big[\Big(k - \tfrac{1}{2}\Big)\Delta,\ \Big(k + \tfrac{1}{2}\Big)\Delta\Big), \qquad k \in \mathbb{Z}.$$

The point of this alignment is that, when the underlying density f is symmetric about 0, the resulting binned pmf $\big(p_k(\sigma)\big)_{k\in\mathbb{Z}}$ is *even* (i.e. $p_k(\sigma) = p_{-k}(\sigma)$), so comparisons can be made in the distance from the centre (as a function of $|k|$). This symmetry is what streamlines the single-crossing/central-mass-dominance argument used below. Accordingly, fix $\Delta > 0$, let $\varepsilon_\sigma^c = \sigma\xi$, and define

$$p_k(\sigma) \ \equiv\ \Pr\Big(\varepsilon_\sigma^c \in \Big[\Big(k - \tfrac{1}{2}\Big)\Delta,\ \Big(k + \tfrac{1}{2}\Big)\Delta\Big)\Big), \qquad \kappa_\Delta(\sigma) \ \equiv\ \sum_{k\in\mathbb{Z}} p_k(\sigma)^2.$$

Proposition 2 (Accuracy and ties under log-concavity)**.** *Suppose ξ has a density f that is symmetric about 0 and log-concave. Fix $\Delta > 0$ and consider the symmetric binning defined above. Then $\kappa_\Delta(\sigma)$ is weakly decreasing in σ (equivalently, weakly increasing in accuracy $1/\sigma$):*

$$0 < \sigma_1 < \sigma_2 \quad \Longrightarrow \quad \kappa_\Delta(\sigma_1) \ \geq\ \kappa_\Delta(\sigma_2).$$

Proof sketch. Under the symmetric binning, $(p_k(\sigma))_{k\in\mathbb{Z}}$ is an even and unimodal probability sequence. The key step is an ordering induced by log-concavity: for $\sigma_1 < \sigma_2$, the binwise likelihood ratio $p_k(\sigma_1)/p_k(\sigma_2)$ is weakly decreasing in $|k|$ (an MLR-type ordering across bins). This implies a single-crossing property for the centred partial sums, hence a *central-mass dominance* ordering of the binned vectors: relative to $p(\sigma_2)$, the vector $p(\sigma_1)$ places (weakly) more mass on bins near the origin at every centred truncation level. A summation-by-parts argument then shows that the sum-of-squares functional $\sum_k p_k^2$ increases under such a centralisation, yielding $\sum_k p_k(\sigma_1)^2 \geq \sum_k p_k(\sigma_2)^2$. Appendix A.6 provides the full argument. □

Taken together, the fine-grid approximation and Proposition 2 show how accuracy maps into posterior ties. At small Δ, the relevant object is the ratio Δ/σ. For a fixed and possibly coarse Δ, log-concavity restores a clean global monotonicity result. We now turn to the phase-robust tie object, obtained by integrating out alignment, which gives a global single-index description before Section 6 relaxes translation invariance and additivity.

5.4 A Phase-Robust Tie Index

Section 5 studies how *resolution* Δ and *accuracy* σ shape the tie channel, summarised by the tie probability κ_Δ. The preceding subsections deliver two directional messages: nested refinement lowers ties (Proposition 1), and under shape restrictions accuracy can raise ties at fixed Δ (Proposition 2). What is still missing for policy design is a *quantitative* mapping from (Δ, σ) to tie risk that is informative beyond the asymptotic regimes of Section 5.2. In the limit analysis we obtain a clean scale and a ratio-type dependence (e.g., leading terms governed by Δ relative to σ), but those expressions need not approximate κ_Δ well for empirically relevant, moderate values of (Δ, σ). In particular, away from small-Δ asymptotics, κ_Δ at a fixed phase ϕ can be non-monotone in Δ because changing the bin width also changes how probability mass is sliced by misaligned boundaries.

This subsection introduces a *phase-robust* tie index, obtained by averaging over the alignment parameter ϕ. A phase-specific tie probability is the right object for one particular reporting implementation with fixed known thresholds. But if the question is how the bin width Δ changes tie risk in a common underlying continuous environment, then κ_Δ^ϕ loads on two margins at once: coarseness and the placement of the boundaries within a Δ-period. Averaging over ϕ integrates out that alignment effect and leaves the component of tie risk attributable to width itself. Appendix A.7 briefly contrasts the phase-specific and phase-robust objects. Accordingly, instead of fixing one alignment, we average κ_Δ^ϕ over $\phi \in [0, \Delta)$.

This phase-robust construction is not a new behavioural assumption. It is a different

description of the same tie event after integrating out the alignment nuisance parameter. In the original formulation the grid is fixed and the question is whether two draws fall in the same bin. After integrating out phase, the draws are fixed and the question is whether a random boundary falls between them. This converts ties into a pure distance comparison.

Fix $\Delta > 0$ and define the phase-robust tie probability by

$$\bar{\kappa}_\Delta \;\equiv\; \frac{1}{\Delta}\int_0^\Delta \kappa_\Delta^\phi \, d\phi, \qquad \kappa_\Delta^\phi \equiv \sum_{u\in\Omega_\Delta} \big(p_\varepsilon^{\Delta,\phi}(u)\big)^2,$$

where $p_\varepsilon^{\Delta,\phi}$ is defined in (3). Equivalently, $\bar{\kappa}_\Delta$ is the expected tie probability when the phase Φ is drawn uniformly on $[0,\Delta)$ and is independent of the underlying continuous measurement error.

Proposition 3 (Phase-robust ties: monotonicity in Δ)**.** *Let X and Y be i.i.d. with the continuous distribution of the underlying measurement error. For each $\Delta > 0$, let $\Phi \sim \mathrm{Unif}[0,\Delta)$ be independent of (X,Y), and let $\bar{\kappa}_\Delta$ be defined above. Then*

(i) Representation. *One has*

$$\bar{\kappa}_\Delta \;=\; \mathbb{E}\left[\left(1 - \frac{|X-Y|}{\Delta}\right)_+\right], \qquad (t)_+ \equiv \max\{t,0\}. \tag{6}$$

(ii) Monotonicity. *For any $0 < \Delta' < \Delta$,*

$$\bar{\kappa}_{\Delta'} \;\le\; \bar{\kappa}_\Delta,$$

with strict inequality whenever $\Pr(|X-Y| < \Delta) > 0$.

Proof. We first prove (6). For a fixed $\Delta > 0$ and phase $\phi \in [0,\Delta)$, let $q_{\Delta,\phi} : \mathbb{R} \to \Omega_\Delta$ be the binning map

$$q_{\Delta,\phi}(x) \;\equiv\; \Delta\left\lfloor \frac{x-\phi}{\Delta} \right\rfloor,$$

so that $x \in [q_{\Delta,\phi}(x)+\phi,\; q_{\Delta,\phi}(x)+\phi+\Delta)$ and $p_\varepsilon^{\Delta,\phi}(u) = \Pr\big(q_{\Delta,\phi}(X) = u\big)$. For i.i.d. X, Y, conditioning on ϕ and using independence yields

$$\kappa_\Delta^\phi = \sum_{u\in\Omega_\Delta} \Pr\big(q_{\Delta,\phi}(X)=u\big)\Pr\big(q_{\Delta,\phi}(Y)=u\big) = \Pr\big(q_{\Delta,\phi}(X) = q_{\Delta,\phi}(Y)\big).$$

Hence, with $\Phi \sim \mathrm{Unif}[0,\Delta)$ independent of (X,Y),

$$\bar{\kappa}_\Delta = \frac{1}{\Delta}\int_0^\Delta \kappa_\Delta^\phi \, d\phi = \Pr\big(q_{\Delta,\Phi}(X) = q_{\Delta,\Phi}(Y)\big) = \mathbb{E}\Big[\Pr\big(q_{\Delta,\Phi}(X) = q_{\Delta,\Phi}(Y) \,\big|\, X, Y\big)\Big].$$

Fix realisations (x, y) and let $d \equiv |x - y|$. If $d \geq \Delta$, then x and y cannot lie in the same length-Δ bin for any phase, so the conditional probability is 0. Suppose $d < \Delta$ and, without loss, $x < y$. The event $q_{\Delta,\phi}(x) = q_{\Delta,\phi}(y)$ fails if and only if a bin boundary of the grid with phase ϕ falls in the open interval $(x, y]$. Since bin boundaries are located at $\phi + k\Delta$, $k \in \mathbb{Z}$, and ϕ is uniform on $[0, \Delta)$, the location of the boundary in a given Δ-period is uniform. Exactly one boundary can fall in an interval of length $d < \Delta$ within a period, so

$$\Pr\left(q_{\Delta,\Phi}(x) = q_{\Delta,\Phi}(y) \,\middle|\, X = x, Y = y\right) = 1 - \frac{d}{\Delta}.$$

Combining the cases gives

$$\Pr\left(q_{\Delta,\Phi}(X) = q_{\Delta,\Phi}(Y) \,\middle|\, X, Y\right) = \left(1 - \frac{|X - Y|}{\Delta}\right)_+,$$

and taking expectations yields (6).

We next prove monotonicity. For each fixed $d \geq 0$, the function $g_\Delta(d) \equiv (1 - d/\Delta)_+$ is weakly increasing in Δ. Therefore, for $0 < \Delta' < \Delta$,

$$\bar{\kappa}_{\Delta'} = \mathbb{E}[g_{\Delta'}(|X - Y|)] \;\leq\; \mathbb{E}[g_\Delta(|X - Y|)] = \bar{\kappa}_\Delta.$$

If $\Pr(|X - Y| < \Delta) > 0$, then $g_{\Delta'}(d) < g_\Delta(d)$ for all $d \in (0, \Delta)$, so the inequality is strict. □

Proposition 3 gives a global comparative static in resolution that does not require nested refinements and is robust to sub-Δ alignment. The representation (6) is the key point: ties become a pure distance comparison, so only the ratio $|X - Y|/\Delta$ matters. Conditional on two latent draws X and Y, integrating over phase asks whether a random grid boundary lands between them; when $|X - Y| < \Delta$, that happens with probability $|X - Y|/\Delta$, so the tie probability is exactly $1 - |X - Y|/\Delta$. With scale noise $\varepsilon^c_\sigma = \sigma\xi$, this becomes a single-index relation in σ/Δ for tie-based equilibrium selection. The next proposition turns this geometry into the corresponding accuracy comparative static.

Proposition 4 (Phase-robust ties and the scale ratio)**.** *Fix $\Delta > 0$ and let $\bar{\kappa}_\Delta(\sigma)$ denote the phase-robust tie probability generated by $\varepsilon^c_\sigma = \sigma\xi$, where ξ has continuous cdf F. Then*

(i) Ratio form. *Let ξ_1, ξ_2 be i.i.d. copies of ξ and set $Z \equiv |\xi_1 - \xi_2|$. Then*

$$\bar{\kappa}_\Delta(\sigma) \;=\; \mathbb{E}\left[\left(1 - \frac{\sigma Z}{\Delta}\right)_+\right], \tag{7}$$

so $\bar{\kappa}_\Delta(\sigma)$ is a function of the single ratio σ/Δ.

(ii) Accuracy. $\bar{\kappa}_\Delta(\sigma)$ *is weakly decreasing in* σ*:*

$$0 < \sigma_1 < \sigma_2 \quad \Longrightarrow \quad \bar{\kappa}_\Delta(\sigma_1) \;\geq\; \bar{\kappa}_\Delta(\sigma_2),$$

with strict inequality whenever $\Pr(|\xi_1 - \xi_2| < \Delta/\sigma_2) > 0$.

Proof. Apply (6) to $X_\sigma = \sigma\xi_1$ and $Y_\sigma = \sigma\xi_2$:

$$\bar{\kappa}_\Delta(\sigma) = \mathbb{E}\left[\left(1 - \frac{|X_\sigma - Y_\sigma|}{\Delta}\right)_+\right] = \mathbb{E}\left[\left(1 - \frac{\sigma|\xi_1 - \xi_2|}{\Delta}\right)_+\right],$$

which is (7). If $\sigma_1 < \sigma_2$, then for each $z \geq 0$ the map $\sigma \mapsto (1 - \sigma z/\Delta)_+$ is weakly decreasing, so

$$\left(1 - \frac{\sigma_1 Z}{\Delta}\right)_+ \;\geq\; \left(1 - \frac{\sigma_2 Z}{\Delta}\right)_+.$$

Moreover, for every $z \in (0, \Delta/\sigma_2)$ one has

$$\left(1 - \frac{\sigma_1 z}{\Delta}\right)_+ \;>\; \left(1 - \frac{\sigma_2 z}{\Delta}\right)_+.$$

Hence, if $\Pr(0 < Z < \Delta/\sigma_2) > 0$, then the inequality is strict after taking expectations. Under our standing continuity assumptions, $\Pr(Z = 0) = 0$, so $\Pr(Z < \Delta/\sigma_2) > 0$ implies $\Pr(0 < Z < \Delta/\sigma_2) > 0$, and therefore

$$\mathbb{E}\left[\left(1 - \frac{\sigma_1 Z}{\Delta}\right)_+\right] \;>\; \mathbb{E}\left[\left(1 - \frac{\sigma_2 Z}{\Delta}\right)_+\right].$$

□

Proposition 4 sharpens the role of accuracy in a way that is useful for policy. Under the phase-robust summary, accuracy enters only through the scale of the underlying measurement error: shrinking σ compresses the distribution of pairwise gaps $|X - Y|$ and therefore increases the mass of pairs that fall within a Δ-neighbourhood. This gives an unambiguous accuracy comparative static without the shape restrictions used in Section 5.3. At the same time, the representation (7) shows that the joint effect of accuracy and resolution on tie incidence is one-dimensional: under the phase-robust summary, only the ratio σ/Δ matters for this equilibrium-selection effect.

This ratio form echoes the limit analysis of Section 5.2: there, too, resolution and accuracy interact through their relative scale. The next result makes the connection literal. Taking the same small-Δ limit as in Lemma 2, the phase-robust tie probability admits the same leading term, governed by the L^2 quantity $\int f_\sigma^2$.

Lemma 3 (Small-Δ limit for the phase-robust tie probability)**.** *Assume $f \in L^2(\mathbb{R})$ and let $\varepsilon_\sigma^c = \sigma\xi$ with density $f_\sigma(x) \equiv \sigma^{-1} f(x/\sigma)$. Then, for each fixed $\sigma > 0$,*

$$\frac{\bar{\kappa}_\Delta(\sigma)}{\Delta} \;\longrightarrow\; \int_{\mathbb{R}} f_\sigma(x)^2\, dx \;=\; \frac{1}{\sigma}\int_{\mathbb{R}} f(u)^2\, du, \qquad \text{as } \Delta \downarrow 0.$$

Proof. Let $D \equiv X_\sigma - Y_\sigma$ with density $h_\sigma(d)$. Then h_σ is continuous and $h_\sigma(0) = \int f_\sigma^2$. By (6),

$$\bar{\kappa}_\Delta(\sigma) = \mathbb{E}\left[\left(1 - \frac{|D|}{\Delta}\right)_+\right] = \int_{\mathbb{R}} \left(1 - \frac{|d|}{\Delta}\right)_+ h_\sigma(d)\, dd = \int_{-\Delta}^{\Delta} \left(1 - \frac{|d|}{\Delta}\right) h_\sigma(d)\, dd.$$

Divide by Δ and substitute $d = \Delta t$:

$$\frac{\bar{\kappa}_\Delta(\sigma)}{\Delta} = \int_{-1}^{1} (1 - |t|)\, h_\sigma(\Delta t)\, dt.$$

Since $0 \le (1 - |t|) \le 1$ on $[-1, 1]$ and h_σ is continuous at 0, dominated convergence yields

$$\frac{\bar{\kappa}_\Delta(\sigma)}{\Delta} \;\longrightarrow\; \int_{-1}^{1} (1 - |t|)\, h_\sigma(0)\, dt = h_\sigma(0) \int_{-1}^{1} (1 - |t|)\, dt = h_\sigma(0) = \int_{\mathbb{R}} f_\sigma(x)^2\, dx.$$

The final equality $\int f_\sigma^2 = (1/\sigma)\int f^2$ is the standard change of variables. □

Lemma 3 shows that integrating out phase preserves the leading-order small-Δ economics. It removes alignment-driven irregularities, yields a well-behaved single-index description in σ/Δ, and nests the same L^2 object that appears in Lemma 2. This ratio form is the main input into the policy analysis in Section 7, where it indexes the policy-relevant equilibrium-selection effect.

6 Tie–Tilt Theory Beyond Translation Invariance

Section 4 gave a sharp closed form under translation invariance: the equilibrium set is interval-valued, and its width decomposes into the grid term Δ plus a strategic tie term proportional to the tie probability. We now drop translation invariance and allow an arbitrary common prior π on the fundamental. Once we do so, boundary beliefs become signal-dependent and the benchmark formula no longer survives automatically.

The extra complexity is structured, not arbitrary. Even before imposing the additive signal structure, the boundary conditions admit a useful decomposition. A *tie* component captures exact equality of signals and creates local *thickening*: neighboring cutoff values may

satisfy the boundary conditions together. A *tilt* component captures directional asymmetry in the boundary belief and creates a *location* effect, shifting where those conditions bind. The additive structure is useful only for the next step, where we study the sign, shape, and magnitude of tilt and ask when the benchmark logic is approximately or exactly recovered.

6.1 Cutoff-Relevant Region and Boundary Conditions

The translation-invariant benchmark of Section 4 is analytically convenient, but it relies on a strong prior restriction: the environment is invariant to shifts of the fundamental over the whole line. The economic problem is more local than that. Outside a bounded part of the signal space, one of the two actions is already pinned down by dominance, so multiplicity can arise only within a cutoff-relevant region.

Recall that when the opponent uses cutoff x, the gain from taking action 1 at signal s is

$$\Delta(s; x) = m(s) - c + b\, q(s; x).$$

If $m(s) < c - b$, then $\Delta(s; x) < 0$ for every x, so action 0 is strictly dominant; if $m(s) > c$, then $\Delta(s; x) > 0$ for every x, so action 1 is strictly dominant. Hence any equilibrium cutoff must lie in the closed interval

$$c - b \leq m(s) \leq c.$$

Under Assumptions 2 and 4, $m(s)$ is increasing in s, so this set is an interval. Define the cutoff-relevant region by

$$X^{\text{rel}} \equiv \{\, s \in \mathbb{R} :\ c - b \leq m(s) \leq c \,\}.$$

Under the normalisation $m(s) = s$, this is exactly the interval $[c - b, c]$, whose width is b. By construction, any symmetric equilibrium cutoff must lie in $X^{\text{rel}} \cap \Omega$. By Lemma 1, verifying whether a candidate $x \in X^{\text{rel}} \cap \Omega$ is an equilibrium cutoff requires checking only the two adjacent boundary signals. The player with signal x must be willing to take action 1 when the opponent uses cutoff x, whereas the player with signal $x - \Delta$ must not.

To write these two adjacent boundary conditions compactly, define for $s \in X^{\text{rel}} \cap \Omega$

$$q^L(s) \equiv \Pr(s_j \geq s \mid s_i = s), \qquad q^U(s) \equiv \Pr(s_j \geq s + \Delta \mid s_i = s),$$

which are the relevant opponent-tail probabilities at the lower and upper adjacent boundaries.

Under the normalisation $m(s) = s$, it is then convenient to define[7]

$$A(s) \equiv c - b\,q^L(s), \qquad B(s) \equiv c + \Delta - b\,q^U(s).$$

These two objects summarise the lower and upper adjacent boundary comparisons, respectively.

Lemma 4 (Adjacent boundary characterisation)**.** *Under Assumptions 1, 2, 3, 4, and 5, and under the normalisation $m(s) = s$, a grid point $x \in X^{\mathrm{rel}} \cap \Omega$ is a symmetric cutoff equilibrium if and only if*

$$x \geq A(x) \qquad \textit{and} \qquad x < B(x - \Delta).$$

Proof. By Lemma 1, best responses are monotone, so it is enough to check the two adjacent boundary signals around a candidate cutoff x. At signal x, taking action 1 is optimal if and only if

$$\Delta(x; x) = x - c + b\,q^L(x) \geq 0,$$

which is equivalent to $x \geq A(x)$. At signal $x - \Delta$, taking action 1 must be suboptimal, so

$$\Delta(x - \Delta; x) = x - \Delta - c + b\,q^U(x - \Delta) < 0,$$

which is equivalent to $x < B(x - \Delta)$. These two inequalities are therefore necessary and sufficient. □

Lemma 4 reduces the equilibrium cutoff problem to two adjacent boundary objects, $A(s)$ and $B(s)$. The next step is to decompose the boundary belief that enters them.

6.2 Tie–Tilt Decomposition

The next lemma decomposes the underlying boundary belief into two parts: one associated with exact ties in signals, and one associated with directional asymmetry around the boundary signal.

For each $s \in X^{\mathrm{rel}}$, define the *tie* term

$$\kappa(s) \equiv \Pr(s_j = s \mid s_i = s),$$

[7]As discussed following (2) and in Remark 3, this normalisation is adopted only for expositional convenience. For any strictly increasing m, one can solve the cutoff conditions in transformed-signal space $\tilde{s} \equiv m(s)$ and map the resulting cutoffs back via m^{-1}. This leaves the underlying equilibrium mechanism unchanged and is used here only to keep the formulas transparent and comparable to the closed-form expressions in Section 4.

and the *tilt* term

$$\delta(s) \equiv \Pr(s_j > s \mid s_i = s) - \Pr(s_j < s \mid s_i = s).$$

The tie term records the conditional mass on exact equality of signals. The tilt term records whether, conditional on $s_i = s$, the opponent is more likely to lie above or below the boundary signal.

Lemma 5 (Tie–tilt decomposition)**.** *For every $s \in X^{\text{rel}}$,*

$$q^L(s) = \frac{1 + \kappa(s) + \delta(s)}{2}, \qquad q^U(s) = \frac{1 - \kappa(s) + \delta(s)}{2}.$$

In particular,

$$q^L(s) - q^U(s) = \kappa(s).$$

Proof Sketch. Conditional on $s_i = s$, partition the opponent's signal into the three events $s_j > s$, $s_j = s$, and $s_j < s$. The tie term $\kappa(s)$ is the mass on the middle event, while the tilt term $\delta(s)$ is the difference between the upper and lower masses. Solving the resulting two linear equations yields the stated formulas for $q^L(s)$ and $q^U(s)$. Appendix A.8 gives the full proof. □

Substituting Lemma 5 into the definitions of $A(s)$ and $B(s)$ gives

$$A(s) = c - \frac{b}{2}\big(1 + \kappa(s) + \delta(s)\big), \qquad B(s) = c + \Delta - \frac{b}{2}\big(1 - \kappa(s) + \delta(s)\big).$$

The two components play different roles. The tie term $\kappa(s)$ enters with opposite signs: it lowers $A(s)$ and raises $B(s)$. That creates slack between the two boundary conditions, so neighboring grid points become easier to sustain together. This is the *thickening* effect. By contrast, the tilt term $\delta(s)$ enters both objects with the same sign. It shifts the pair of boundary conditions together rather than creating the local slack. This is a *location* effect. Figure 3 illustrates the distinction: higher tie mass thickens the admissible region, whereas tilt mainly moves where that region sits.

This distinction also separates two forms of multiplicity. *Interval (connected) multiplicity* means a nondegenerate block of consecutive grid points. *Multi-branch (disconnected) multiplicity* means separated branches with gaps in between. The next theorem formalises the first part of the message: ties are necessary for the connected form. Tilt, by contrast, can still shift where the boundary conditions bind and thereby generate disconnected branches; the subsequent subsections develop that channel.[8]

[8]This terminology is also consistent with the complete-information benchmark. In the complete-

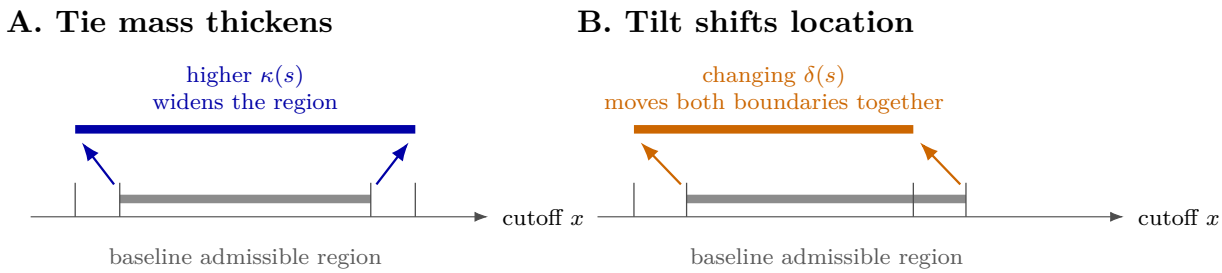

Figure 3: Local roles of ties and tilt. Left: higher $\kappa(s)$ thickens the admissible cutoff region. Right: changing $\delta(s)$ shifts its location.

Theorem 3 (Ties are necessary for connected multiplicity)**.** *Under Assumptions 1, 2, 3, 4, and 5, and under the normalisation $m(s) = s$, suppose*

$$\kappa(s) = 0 \qquad \textit{for all } s \in X^{\mathrm{rel}} \cap \Omega.$$

Then no two adjacent grid points in $X^{\mathrm{rel}} \cap \Omega$ can both be symmetric cutoff equilibria. Equivalently, interval (connected) multiplicity cannot arise. Any multiplicity, if it occurs, must therefore be multi-branch (disconnected).

Proof. Suppose, to the contrary, that both x and $x + \Delta$ are symmetric cutoff equilibria for some $x, x + \Delta \in X^{\mathrm{rel}} \cap \Omega$. By Lemma 4,

$$x \geq A(x) \qquad \text{and} \qquad x + \Delta < B(x).$$

Hence

$$\Delta < B(x) - A(x).$$

By Lemma 5,

$$B(x) - A(x) = \Delta + b\,\kappa(x).$$

If $\kappa(x) = 0$, then $B(x) - A(x) = \Delta$, contradicting $\Delta < B(x) - A(x)$. Therefore no two adjacent grid points can both be equilibrium cutoffs. This rules out interval (connected) multiplicity. □

Theorem 3 isolates the role of ties cleanly. Without exact equality of signals, the local thickening margin disappears, so adjacent grid points cannot both be equilibrium cutoffs. Ties are therefore the unique local source of connected multiplicity on a grid. Tilt still

information game, for every $\theta \in [c - b, c]$, the cutoff strategy that chooses action 1 if and only if the fundamental is at least θ is an equilibrium. Viewed as a set of cutoff equilibria, the equilibrium set is therefore interval-valued, with width b.

matters, but through a different channel: it shifts where the boundary conditions bind and can thereby shape the global pattern of equilibrium branches.

6.3 Additive Structure and Tilt

The tie–tilt decomposition itself did not require additivity. To study tilt more sharply in this and the following subsection, however, it is useful to impose an additive signal structure. We assume that signals take the form

$$s_i = \theta + \varepsilon_i,$$

where the noise terms are independent and identically distributed across players. This additional structure makes it possible to study the sign, shape, and magnitude of tilt directly.

Assumption 7 (Symmetric noise)**.** *The noise pmf p_ε is symmetric:*

$$p_\varepsilon(u) = p_\varepsilon(-u) \qquad \text{for all } u \in \Omega.$$

Under the additive signal structure and symmetric noise, tilt reflects local prior shape rather than mechanical asymmetry from the noise itself.

Lemma 6 (Sign of tilt under local prior monotonicity)**.** *Maintain the additive signal structure and Assumption 7. Fix $s \in X^{\text{rel}}$. Let π denote the common prior on θ, let $\mu_s(\theta) \equiv \Pr(\theta \mid s_i = s)$ be the posterior pmf of θ, and write $\operatorname{supp}(\mu_s) \equiv \{\theta \in \Omega : \mu_s(\theta) > 0\}$ for its support.*

1. *If π is constant on $\operatorname{supp}(\mu_s)$, then $\delta(s) = 0$.*
2. *If π is weakly increasing on $\operatorname{supp}(\mu_s)$, then $\delta(s) \geq 0$.*
3. *If π is weakly decreasing on $\operatorname{supp}(\mu_s)$, then $\delta(s) \leq 0$.*

Proof Sketch. Under the additive signal structure,

$$\delta(s) = \sum_{\theta \in \Omega} \mu_s(\theta)\, D(s - \theta),$$

where $D(u) \equiv \Pr(\varepsilon > u) - \Pr(\varepsilon < u)$ is odd under symmetric noise. Pairing posterior mass at θ with mass at its reflection $\theta^\sharp \equiv 2s - \theta$, symmetry of the noise implies that the sign of each paired contribution is governed by the posterior comparison between θ and $\theta^\sharp$, and hence by the local monotonicity of the prior. Appendix A.9 gives the full proof. □

The intuition is straightforward. If the prior is locally increasing, then a relatively low signal is more naturally read as a negative noise realisation than as evidence that the fundamental is truly low. Conditional on such a signal, the opponent is then more likely to lie above than below the player's own signal, which yields positive tilt. Since positive tilt lowers both $A(s)$ and $B(s)$, it shifts the equilibrium cutoff region to the left and makes action 1 easier to sustain. The reverse holds under a locally decreasing prior.

To isolate the directional-asymmetry channel on its own, it is useful to look at the corresponding continuous additive model. Tilt continues to matter in a continuous environment and therefore has a natural continuous analogue. By contrast, the tie channel relies on exact equality of signals; once signals are continuous and such ties carry no mass, that channel disappears.

Lemma 7 (Log-concavity implies decreasing continuous tilt)**.** *Consider the corresponding continuous additive model with prior density π. If π is positive and log-concave, then the continuous tilt is weakly decreasing.*

Proof. See Appendix A.10. □

A decreasing continuous tilt means that the location effect varies monotonically with the signal: lower signals are associated with weakly larger tilt values, and higher signals with weakly smaller ones. Under a log-concave prior, tilt therefore shifts the boundary conditions in a disciplined directional way across the signal space.

Another question is whether branching can still arise when the tie channel is absent. If $\kappa \equiv 0$, then the two adjacent boundary conditions collapse to a single indifference condition,

$$x = c - \frac{b}{2} - \frac{b}{2}\,\delta(x) \qquad \Longleftrightarrow \qquad \delta(x) = \frac{2(c-x)}{b} - 1.$$

Equilibrium cutoffs are then determined by intersections between the tilt curve and an affine function of x. If the location effect bends sufficiently, several separated crossings may still arise, producing multi-branch multiplicity even without local thickening.

Figure 4 illustrates this branching mechanism.[9] An S-shaped or otherwise sufficiently nonlinear tilt function can produce several separated fixed points. This is the familiar branching logic of continuous global games; see, for example, Morris and Shin (2001). In the discrete model studied here, the same branching force remains present, but ties add a second margin: once a branch exists, local tie mass can thicken it into a connected interval of grid cutoffs.

[9]The bottom-right example is not monotone because the bimodal Gaussian-mixture prior is not log-concave. Lemma 7 applies only under log-concavity of the prior.

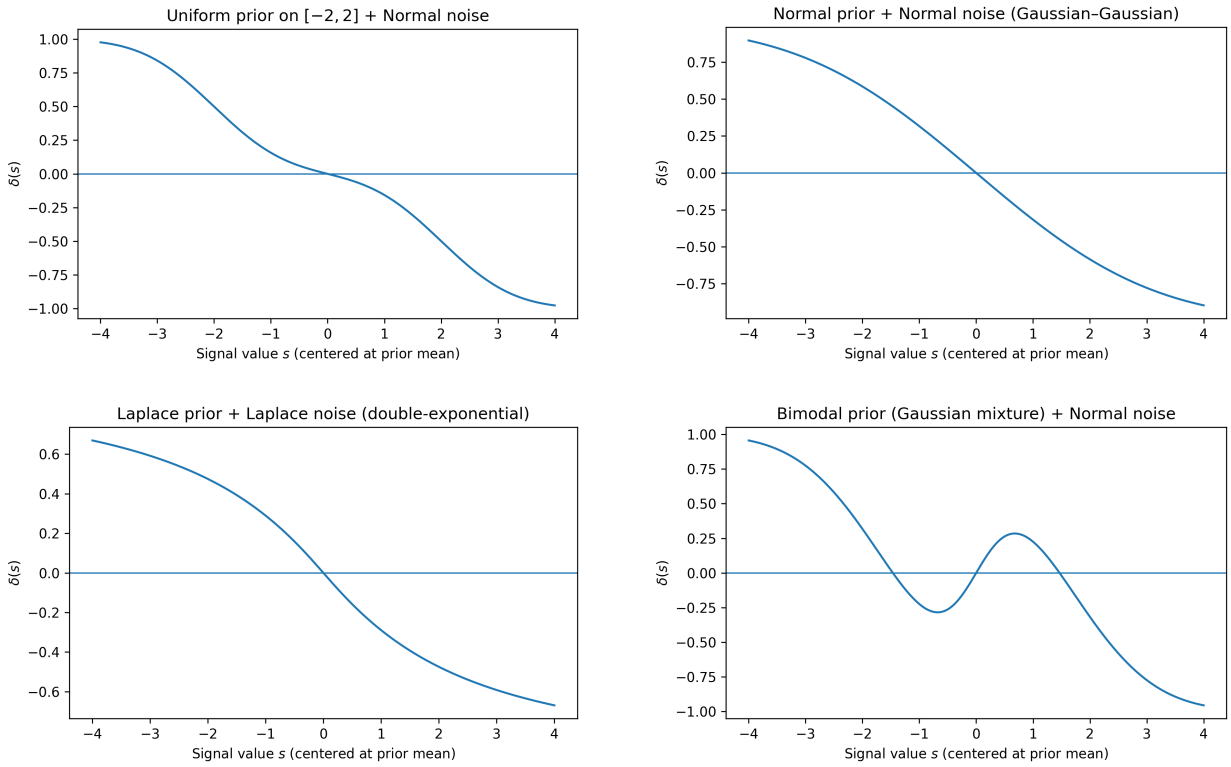

Figure 4: Examples of the tilt function $\delta(s)$ (signal s centered at the prior mean). Top left: Uniform prior on $[-2, 2]$ with Normal noise. Top right: Normal prior with Normal noise. Bottom left: Laplace prior with Laplace noise. Bottom right: Bimodal prior (Gaussian mixture) with Normal noise.

Lemma 6 suggests that tilt is governed by local prior shape. It is therefore natural to expect the magnitude of tilt to depend on how flat the prior remains around the relevant signal region. To study the size of tilt, we introduce a local log-slope restriction on the prior. Fix a radius $R > 0$ and define the relevant band

$$\Theta^{\mathrm{rel}}(R) \equiv X^{\mathrm{rel}} \oplus [-R, R] = \{\theta : \ \theta = x + u, \ x \in X^{\mathrm{rel}}, \ u \in [-R, R] \cap \Omega\}.$$

Assumption 8 (Local log-slope bound on the relevant band)**.** *There exists* $L_{\mathrm{loc}}(R) < \infty$ *such that for all* $\theta, \theta + \Delta \in \Theta^{\mathrm{rel}}(R)$*,*

$$\big|\log \pi(\theta + \Delta) - \log \pi(\theta)\big| \leq L_{\mathrm{loc}}(R)\, \Delta.$$

Proposition 5 (Bound on tilt)**.** *Maintain the additive signal structure and Assumption 7.*

Fix $R > 0$ and suppose Assumption 8 holds at this R. Then for every $s \in X^{\mathrm{rel}}$,

$$|\delta(s)| \;\le\; \tanh\big(L_{\mathrm{loc}}(R)\, R\big) \;+\; \Pr\big(|s-\theta| > R \mid s_i = s\big). \tag{8}$$

Consequently, if

$$T(R) \equiv \sup_{s \in X^{\mathrm{rel}}} \Pr\big(|s-\theta| > R \mid s_i = s\big),$$

then

$$\sup_{s \in X^{\mathrm{rel}}} |\delta(s)| \;\le\; \tanh\big(L_{\mathrm{loc}}(R)R\big) + T(R). \tag{9}$$

Proof Sketch. Write tilt as a posterior average of the directional comparison between the opponent's signal being above or below the player's own signal. Split posterior mass into a local part $|s-\theta| \le R$ and a tail part $|s-\theta| > R$. The tail contributes at most the posterior tail probability. On the local part, compare each θ with its reflection $2s-\theta$. Symmetry of the noise implies that the posterior ratio across reflected points is determined entirely by the prior ratio, and Assumption 8 controls that ratio uniformly on the local window. This yields the stated bound on the local asymmetry term, and hence on $|\delta(s)|$. Appendix A.11 gives the full proof. □

The bound has a simple reading. The first term captures local prior shape, and the second captures posterior tail leakage. Tilt is therefore not a free residual. When the prior is locally flat and private noise is small enough that posterior mass stays local, tilt is uniformly small.

The interaction between $L_{\mathrm{loc}}(R)$ and R is the main comparative-static message. The radius R is the scale on which posterior concentration is required: noisier signals require a larger neighborhood to make the tail term small. The quantity $L_{\mathrm{loc}}(R)$ measures how sharply the prior varies on that same scale. So controlling tilt requires a joint restriction across scales: if noise is large and one must take R large, then the prior must stay flatter over a wider band. This is the same comparative-static tension emphasized by Morris and Shin (2001): noisier private information forces one to control beliefs over a wider region, so keeping multiplicity in check requires a flatter prior on that wider scale.[10] When that

[10] A central message of Morris and Shin (2001) is that multiplicity is easier to control when private information is precise and the prior is diffuse, and harder to control when private information is weak and the prior is concentrated. In their continuous setting, the relevant multiplicity takes the form of multiple separated cutoff equilibria—that is, what we call multi-branch multiplicity here. In the present language, this is precisely the form of multiplicity associated with a sufficiently strong tilt channel. Our bound on tilt delivers a discrete counterpart of the same logic: if noise is large and one must take R large to control the tail term, then preventing large tilt requires the prior to remain flat over a wider band, meaning a smaller $L_{\mathrm{loc}}(R)$ on that scale.

restriction holds, the role of tilt remains limited, and the tie channel becomes the dominant source of multiplicity.

6.4 Bounds and Benchmark Recovery

Once tilt is controlled, the tie–tilt decomposition yields useful quantitative bounds on the equilibrium set even away from translation invariance. To state these bounds, define the worst-case tie and tilt magnitudes on the cutoff-relevant region by

$$\overline{\kappa}^{\mathrm{rel}} \equiv \sup_{s \in X^{\mathrm{rel}}} \kappa(s), \qquad \overline{\delta}^{\mathrm{rel}} \equiv \sup_{s \in X^{\mathrm{rel}}} |\delta(s)|.$$

Theorem 4 (Robust interval bounds on the cutoff-relevant region)**.** *Maintain Assumptions 1, 2, 4, and 5, the additive signal structure, Assumption 7, and the normalisation $m(s) = s$. Then any symmetric cutoff equilibrium $x \in X^{\mathrm{rel}}$ must satisfy*

$$x \in \Big[c - \frac{b}{2} - \frac{b}{2}\,\overline{\kappa}^{\mathrm{rel}} - \frac{b}{2}\,\overline{\delta}^{\mathrm{rel}}, \quad c + \Delta - \frac{b}{2} + \frac{b}{2}\,\overline{\kappa}^{\mathrm{rel}} + \frac{b}{2}\,\overline{\delta}^{\mathrm{rel}}\Big) \cap \Omega.$$

In particular, the equilibrium set has width at most

$$\Delta + b\,\overline{\kappa}^{\mathrm{rel}} + b\,\overline{\delta}^{\mathrm{rel}}.$$

Moreover, if Assumption 8 holds at some radius R, then

$$\overline{\delta}^{\mathrm{rel}} \le \tanh\big(L_{\mathrm{loc}}(R)R\big) + T(R).$$

Proof. By Lemma 4, an equilibrium cutoff must satisfy the two adjacent boundary inequalities. Substituting the tie–tilt representation of $A(s)$ and $B(s)$ and then taking worst-case bounds over $s \in X^{\mathrm{rel}}$ yields the stated interval restriction. The final inequality follows directly from Proposition 5. □

Theorem 4 is the general non-translation-invariant bound. It separates the discrete thickening channel, captured by $\overline{\kappa}^{\mathrm{rel}}$, from the location and branching channel, captured by $\overline{\delta}^{\mathrm{rel}}$. The tilt bound shows that the latter is uniformly small when the prior is locally diffuse on the relevant scale and posterior tails are thin, so the first-order effect of discretisation still works through ties.

Under local uniformity and bounded symmetric support, that worst-case interval restriction sharpens to an exact closed-form characterization: the tilt term vanishes pointwise and the benchmark interval is recovered on X^{rel}.

Theorem 5 (Benchmark recovery under local uniformity)**.** *Maintain Assumptions 1, 2, 4, and 5, the additive signal structure, Assumption 7, and the normalisation $m(s) = s$. Suppose the noise has bounded support contained in $[-R, R] \cap \Omega$ for some $R > 0$, and the prior is locally uniform on $\Theta^{\mathrm{rel}}(R)$. Then the set of symmetric cutoff equilibria in $X^{\mathrm{rel}} \cap \Omega$ is exactly*

$$\Big[c - \frac{b}{2} - \frac{b}{2}\kappa, \quad c + \Delta - \frac{b}{2} + \frac{b}{2}\kappa\Big) \cap X^{\mathrm{rel}} \cap \Omega.$$

Proof. Fix $s \in X^{\mathrm{rel}}$. Because the noise support is contained in $[-R, R] \cap \Omega$, observing $s_i = s$ implies that the posterior support of θ is contained in

$$s + [-R, R] \subseteq \Theta^{\mathrm{rel}}(R).$$

Since the prior is constant on $\Theta^{\mathrm{rel}}(R)$, Lemma 6 gives $\delta(s) = 0$ for all $s \in X^{\mathrm{rel}}$.

For the tie term,

$$\kappa(s) = \Pr(s_j = s \mid s_i = s) = \Pr(\varepsilon_j = \varepsilon_i \mid s_i = s).$$

Under local uniformity, conditioning on $s_i = s$ leaves the posterior distribution of ε_i equal to its unconditional distribution p_ε. Therefore

$$\kappa(s) = \sum_{u \in [-R,R] \cap \Omega} p_\varepsilon(u)^2 = \Pr(\varepsilon_j = \varepsilon_i) = \kappa \qquad \text{for all } s \in X^{\mathrm{rel}}.$$

Substituting $\delta(s) = 0$ and $\kappa(s) = \kappa$ into Lemma 4 yields exactly the stated interval. □

This result should be read as a sharpening of Theorem 4. Section 4 obtained the same equilibrium interval under global translation invariance. Here, global uniformity is unnecessary. It is enough that the prior be locally uniform on the part of the fundamental space reached by posterior updating from X^{rel}, together with bounded noise support to keep posterior comparisons inside that flat region. Under those primitive conditions, the worst-case bound tightens to the benchmark closed form on X^{rel}, and the comparative statics and policy logic built on it carry over there.

6.5 Sufficient Conditions for Interval Widening

Theorem 3 gave the necessary side: without ties, there is no non-degenerate connected equilibrium region. This subsection asks when the same tie channel also yields the expected *sufficient* comparative static: coarsening widens that region, while refinement shrinks it.

The remaining issue is that adjacent-boundary comparisons depend on within-cell position. At a fixed phase, changing the mesh affects not only the local overlap created by ties but also which predecessor boundary enters the comparison. For the comparative static of interest, however, the relevant design margin is mesh width rather than boundary origin. We therefore work, as in Section 5.4, with the phase-robust connected cutoff region $\bar{\mathcal{E}}_\Delta$ obtained by integrating over phase. This isolates the effect of mesh width Δ from arbitrary boundary placement, so the tie-driven thickening margin can be compared cleanly across meshes.

As in Section 5, we compare different reporting resolutions by embedding them in a common underlying continuous signal space: different values of Δ are interpreted as alternative discretisations of the same continuous variable. Let y denote this underlying continuous signal, of which the reported discrete signal s is a discretised version.

Assumption 9 (Regularity of the conditional difference law)**.** *For all $y, y' \in X^{\mathrm{rel}}$,*

$$\left\|\mathrm{Law}(D_y) - \mathrm{Law}(D_{y'})\right\|_{\mathrm{TV}} \leq L_D\,|y - y'|, \qquad D_y \equiv y_j - y_i \mid y_i = y,$$

for some constant $L_D < 1/b$.

This assumption requires the conditional difference law to change only gradually with the underlying signal. Economically, it says that nearby types face nearby local strategic environments: a small change in a player's signal should not drastically alter the conditional distribution of the opponent's signal relative to her own. In that sense, the assumption is a regularity condition rather than a substantive restriction on the direction of beliefs or incentives. The threshold $L_D < 1/b$ also has a natural interpretation. A larger b means stronger strategic complementarity, so the comparison becomes more sensitive to residual shifts in conditional beliefs; correspondingly, the conditional difference law must vary more smoothly. By contrast, when b is smaller, the requirement is weaker. In the translation-invariant benchmark, the law of D_y does not depend on y at all, so one has $L_D = 0$. Thus the benchmark case satisfies the condition automatically, and the present assumption can be read as a smooth departure from that idealised environment.

Theorem 6 (Coarsening widens, refinement shrinks)**.** *Maintain Assumptions 1, 2, 4, 5, and the normalisation $m(s) = s$. Fix $0 < \Delta' < \Delta$. Let $\bar{\mathcal{E}}_\Delta$ denote the phase-robust cutoff region defined in Appendix A.12. If Assumption 9 holds, then*

$$\bar{\mathcal{E}}_{\Delta'} \subseteq \bar{\mathcal{E}}_\Delta.$$

Equivalently, coarsening weakly widens and refinement weakly shrinks the phase-robust cutoff

region. In particular, if these regions are intervals, their endpoints satisfy

$$\underline{x}_{\Delta} \leq \underline{x}_{\Delta'}, \qquad \overline{x}_{\Delta} \geq \overline{x}_{\Delta'}.$$

Proof Sketch. The proof proceeds in two steps. First, after integrating out phase, coarsening has a direct monotone effect on the boundary comparison: the lower adjacent-boundary condition becomes more permissive, while the upper one becomes less restrictive. Second, the upper condition still involves a predecessor correction, because the coarser region compares a cutoff with the phase-robust upper boundary evaluated one coarse step below, rather than one fine step below. Assumption 9 controls exactly that residual term. The inequality $L_D < 1/b$ ensures that the predecessor correction is always dominated by the direct grid gain $\Delta - \Delta'$, so every cutoff that survives under the finer mesh also survives under the coarser one. Appendix A.13 gives the full proof. □

This theorem is the sufficient-side counterpart to Theorem 3. Once the conditional difference law is regular enough, the same tie channel that is necessary for connected multiplicity also delivers the expected monotone comparative static: coarsening widens the connected phase-robust region, while refinement shrinks it.

The role of the phase-robust formulation is the same as in Section 5.4: it isolates the mesh-width effect from boundary placement. A phase-specific statement would combine the same economic force with within-cell position and would therefore require additional bookkeeping about predecessor shifts. The theorem shows that once this alignment margin is integrated out and the conditional difference law is regular, coarsening widens the connected region for the substantive reason identified in the benchmark: more ties relax the boundary comparisons.

The section's message is therefore straightforward. Once translation invariance is relaxed, reporting resolution still matters through a tie-driven thickening force, while local prior shape matters through tilt. The benchmark closed form is not fully robust, yet its core economic lesson survives: resolution remains a distinct design margin for equilibrium selection, and finer reporting can still make coordination outcomes more selective even when accuracy is unchanged.

7 Policy Guidance: Resolution and Accuracy

This section uses the equilibrium characterisations of Sections 4–5 to translate the theory into policy guidance about two distinct informational margins: *resolution*, captured by the bin

width Δ, and *accuracy*, captured by the dispersion parameter σ of the underlying continuous measurement error.

The analysis distinguishes two effects of policy changes. Holding the cutoff fixed isolates a learning effect, which works through the latent measurement distribution. Changing the set of self-consistent cutoffs isolates an equilibrium-selection effect, which works through multiplicity.

The goal is not to solve a universal design problem. Optimal policy depends on the application, including the direct learning gains from greater accuracy, the welfare criterion, and the implementation costs of refining the reporting language or improving the measurement technology. What the theory does deliver is a way to isolate how resolution and accuracy load onto these two effects. Even when measurement accuracy is limited, finer reporting can still reduce posterior ties and narrow the set of self-consistent cutoffs, whereas accuracy directly improves learning and can also change equilibrium selection through tie incidence. With continuous reports this second effect is hidden because exact ties occur with probability zero; with discrete reports it becomes observable because tie incidence varies systematically with Δ and σ.

7.1 Objective and Two Channels

We evaluate policies under a general measurable welfare criterion. Let $a_i \in \{0, 1\}$ be agent i's action and let $G(\theta, a_1, a_2)$ be any measurable planner objective (e.g. the success indicator $\mathbf{1}\{a_1 = a_2 = 1\}$ or a surplus function). The induced welfare is

$$W \;\equiv\; \mathbb{E}\big[G(\theta, a_1, a_2)\big].$$

Policy affects W only through the induced joint distribution of actions (a_1, a_2) conditional on θ. This subsection separates the effect of policy holding the cutoff fixed from the effect working through equilibrium selection.

Following Section 5.2, write the observed report as a discretisation of an underlying continuous measurement:

$$y_i = \theta + \varepsilon^c_{\sigma,i}, \qquad \varepsilon^c_{\sigma,i} = \sigma\xi_i,$$

with (ξ_i) i.i.d. and independent of θ. The reported signal is

$$s_i = q_{\Delta,\phi}(y_i) \in \Omega_\Delta, \qquad q_{\Delta,\phi}(y) = \Delta\Big\lfloor \frac{y - \phi}{\Delta} \Big\rfloor,$$

where $\Delta > 0$ is the bin width, ϕ is grid alignment, and actions follow an equilibrium cutoff

rule

$$a_i = \mathbf{1}\{s_i \geq x\} \qquad \text{for some equilibrium cutoff } x \in \Omega_\Delta,$$

where x may be set-valued under multiplicity.

We begin by holding the cutoff fixed. The next lemma rewrites the event $\{s_i \geq x\}$ on the latent measurement scale.

Lemma 8 (Grid cutoffs are latent thresholds)**.** *Fix (Δ, ϕ) and let $x \in \Omega_\Delta$. For any latent measurement $y \in \mathbb{R}$, if the reported signal is $s_i = q_{\Delta,\phi}(y)$, then*

$$\{s_i \geq x\} \iff \{y \geq x + \phi\}.$$

Consequently, under the cutoff rule $a_i = \mathbf{1}\{s_i \geq x\}$ the realised action can be written as

$$a_i \;=\; \mathbf{1}\{y_i \geq x + \phi\}.$$

Proof. Write $x = \Delta k$ for some $k \in \mathbb{Z}$ since $x \in \Omega_\Delta$. Then

$$\begin{aligned} q_{\Delta,\phi}(y) \geq x &\iff \Delta\Big\lfloor \frac{y-\phi}{\Delta} \Big\rfloor \geq \Delta k \iff \Big\lfloor \frac{y-\phi}{\Delta} \Big\rfloor \geq k, \\ &\iff \frac{y-\phi}{\Delta} \geq k \iff y \geq \phi + \Delta k = x + \phi. \end{aligned}$$

□

This isolates the *learning channel*: once the equilibrium cutoff is fixed, accuracy matters only through the latent measurement distribution, that is, through how well the latent measurement tracks the underlying state. The next corollary records the corresponding invariance with respect to the bin width Δ.

Corollary 2 (Δ-invariance of the learning channel)**.** *Fix a cutoff $x \in \Omega_\Delta$ and phase ϕ, and define*

$$W(\sigma; x, \phi) \;\equiv\; \mathbb{E}\Big[G\Big(\theta, \mathbf{1}\{s_1 \geq x\}, \mathbf{1}\{s_2 \geq x\}\Big)\Big].$$

Then $W(\sigma; x, \phi)$ is independent of Δ.

Proof. By Lemma 8, the action indicators can be rewritten as $\mathbf{1}\{y_i \geq x + \phi\}$. The joint law of (θ, y_1, y_2) is generated by $y_i = \theta + \sigma\xi_i$ and therefore depends on σ but not on Δ, once the cutoff x and phase ϕ are held fixed. □

The role of Δ is different. Once the cutoff is fixed, it does not affect the learning channel. Its effect on realised outcomes works instead through the set of possible equilibrium cutoffs:

with discrete reports, ties make more cutoffs self-consistent and thereby make equilibrium selection more consequential.

Proposition 6 (Learning and equilibrium-selection channels)**.** *Let $X(\Delta, \sigma)$ denote the set of equilibrium cutoffs induced by the reporting environment with reporting granularity Δ and accuracy σ. Let $\mathcal{S}(\Delta, \sigma)$ be any selection rule, and write $x = \mathcal{S}(\Delta, \sigma) \in X(\Delta, \sigma)$. Define realised welfare by*

$$W^{\mathrm{eq}}(\Delta, \sigma) \;\equiv\; W\big(\sigma;\, \mathcal{S}(\Delta, \sigma), \phi\big).$$

Then:

*(i) (*Learning channel*) Holding the cutoff fixed isolates the effect of accuracy on behaviour through the latent measurement distribution. In particular, the mapping $\sigma \mapsto W(\sigma; x, \phi)$ is independent of Δ (Corollary 2).*

*(ii) (*Equilibrium-selection channel*) Reporting granularity Δ affects $W^{\mathrm{eq}}(\Delta, \sigma)$ only through its effect on the equilibrium cutoff set $X(\Delta, \sigma)$, and hence on the selected cutoff $\mathcal{S}(\Delta, \sigma)$.*

*(iii) (*Accuracy affects both channels*) Accuracy σ affects realised welfare through both channels: it changes $W(\sigma; x, \phi)$ for any fixed cutoff x, and it can also change equilibrium-selection risk through its effect on ties, and hence on the set $X(\Delta, \sigma)$.*

Proof. Part (i) is Corollary 2. For (ii), once the cutoff is fixed, welfare is Δ-invariant, so resolution can affect $W^{\mathrm{eq}}(\Delta, \sigma)$ only by changing which cutoffs are self-consistent and hence which cutoff is selected. Finally, (iii) follows because changing σ changes both the latent law $y_i = \theta + \sigma\xi_i$ and the incidence of ties. □

Proposition 6 gives the basic division of labour between the two instruments. Holding the cutoff fixed isolates the learning effect of accuracy. Reporting granularity matters instead through equilibrium selection, by changing which cutoffs are self-consistent. Accuracy is the subtler instrument because it moves both channels at once. This is why refinement of resolution need not be redundant even when the underlying measurement remains noisy.

Principle 1 (Refinement targets selection risk). *To reduce equilibrium-selection risk created by discretisation, refine the reporting grid. This reduces ties directly while leaving the learning channel unchanged.*

Principle 2 (Accuracy cuts both ways). *Accuracy upgrades strengthen learning, but with coarse reports they can also raise tie incidence and thereby amplify selection risk.*

Principles 1–2 summarise the two channels. Reporting granularity acts directly on ties and therefore on equilibrium-selection risk. Accuracy affects learning, but it can also affect selection by changing tie incidence. The rest of the section quantifies that second effect.

7.2 The Phase-Robust σ/Δ Rule

This subsection applies the phase-robust tie object from Section 5.4 to policy comparisons. When the policy margin is reporting width rather than threshold origin, the relevant statistic is $\bar{\kappa}_\Delta(\sigma)$, not a phase-specific tie probability. This is exactly the case when threshold origin is uncontrolled, heterogeneous, or institutionally incidental. Once phase is integrated out, tie risk depends on policy only through the scale comparison between measurement noise and bin width.

Applying the phase-robust representation from Section 5.4 to the tie probability, and writing $Z \equiv |\xi_1 - \xi_2|$, yields the following single-index form.

Proposition 7 (A single-index representation and its marginal statistic)**.** *Assume* $\mathbb{E}[Z] < \infty$*, and define* $\Psi(r) \equiv \mathbb{E}[(1 - rZ)_+]$ *for* $r > 0$*. Then under the phase-robust summary,*

$$\bar{\kappa}_\Delta(\sigma) = \Psi(\sigma/\Delta).$$

Moreover, Ψ *is differentiable with* $\Psi'(r) = -\mathbb{E}[Z\,\mathbf{1}\{Z < 1/r\}]$*. Hence, with* $\mathcal{M}(t) \equiv \mathbb{E}[Z\,\mathbf{1}\{Z < t\}]$ *for* $t > 0$*,*

$$\frac{\partial \bar{\kappa}_\Delta(\sigma)}{\partial \sigma} = -\frac{1}{\Delta}\,\mathcal{M}(\Delta/\sigma), \quad \frac{\partial \bar{\kappa}_\Delta(\sigma)}{\partial \Delta} = \frac{\sigma}{\Delta^2}\,\mathcal{M}(\Delta/\sigma).$$

Proof. By the phase-robust representation (6) in Proposition 3, applied to $X = \sigma\xi_1$ and $Y = \sigma\xi_2$,

$$\bar{\kappa}_\Delta(\sigma) = \mathbb{E}\left[\left(1 - \frac{|X - Y|}{\Delta}\right)_+\right] = \mathbb{E}\big[(1 - (\sigma/\Delta)Z)_+\big] = \Psi(\sigma/\Delta),$$

where $Z = |\xi_1 - \xi_2|$. Since $\mathbb{E}[Z] < \infty$, dominated convergence justifies differentiation under the expectation and yields the expression for $\Psi'(r)$. Substituting $r = \sigma/\Delta$ and writing $\mathcal{M}(t) = \mathbb{E}[Z\,\mathbf{1}\{Z < t\}]$ gives the stated marginal formulas. □

The implication is immediate. Under the phase-robust summary, tie risk depends on the two policy instruments only through the ratio σ/Δ. Any policy reform that keeps σ/Δ fixed also keeps $\bar{\kappa}_\Delta(\sigma)$ fixed. The phase-robust formulation therefore compresses the equilibrium-selection side into a one-dimensional design index.

Principle 3 (Accuracy upgrades should be matched by finer reporting). *If policy raises accuracy, it should also refine reporting whenever the goal is to avoid making equilibrium selection more fragile.*

Once multiplicity is present, welfare also depends on how equilibrium is selected. Different institutions therefore act like different selection rules: some place more weight on

avoiding adverse coordination outcomes and hence care more about shrinking the range of self-consistent cutoffs. In our setting, this means placing more weight on reducing tie risk.

Principle 4 (More conservative selection favours a larger σ/Δ). *Relative to more optimistic selection criteria, more conservative selection criteria place greater weight on avoiding adverse coordination outcomes and therefore favour policy configurations with a larger σ/Δ.*

Proposition 7 also yields a directly estimable object for the selection side: the truncated first moment

$$\mathcal{M}(t) = \mathbb{E}[Z\,\mathbf{1}\{Z < t\}], \qquad t = \Delta/\sigma.$$

Because $\bar{\kappa}_\Delta(\sigma) = \Psi(\sigma/\Delta)$, this same object measures how sensitive tie risk is to changes in the single index σ/Δ. Equivalently, when Δ is held fixed, it gives the normalized marginal effect of an accuracy upgrade. The empirical problem is therefore not to estimate tie incidence itself, but to trace the normalized sensitivity curve $t \mapsto \mathcal{M}(t)$.

In experiments, one can expose pairs of subjects to the same underlying state (or to matched states) and record repeated measurements or reported signals. The resulting within-pair differences provide observations of Z, from which $\mathcal{M}(t)$ is obtained by averaging Z over pairs with $Z < t$. In observational settings, the same construction applies whenever two units can reasonably be treated as facing a common fundamental, for instance repeated readings of the same object by different inspectors, duplicate sensor readings, parallel audits, or matched agents operating under the same realised shock. Given a sample $\{Z_\ell\}_{\ell=1}^{L}$ of such pairwise disagreements, the plug-in estimator

$$\widehat{\mathcal{M}}(t) = \frac{1}{L}\sum_{\ell=1}^{L} Z_\ell\,\mathbf{1}\{Z_\ell < t\}$$

is immediate. The resulting curve measures local normalized sensitivity and provides a convenient way to classify when, in the fixed-Δ comparison relevant for accuracy upgrades, the tie-risk response is likely to be large or small.

Corollary 3 (Three regimes of normalized sensitivity)**.** *Let $r = \sigma/\Delta$ and $t = \Delta/\sigma = 1/r$. The marginal effect of the phase-robust tie risk $\Psi(r) = \bar{\kappa}_\Delta(\sigma)$ with respect to the single index r is*

$$-\frac{d\Psi(r)}{dr} = \mathcal{M}(t) \in \Big[0, \mathbb{E}Z\Big].$$

*(i) (*High-sensitivity region*) If t is large, then $\mathcal{M}(t)$ is close to $\mathbb{E}Z$, so normalized sensitivity is close to $\mathbb{E}Z$.*

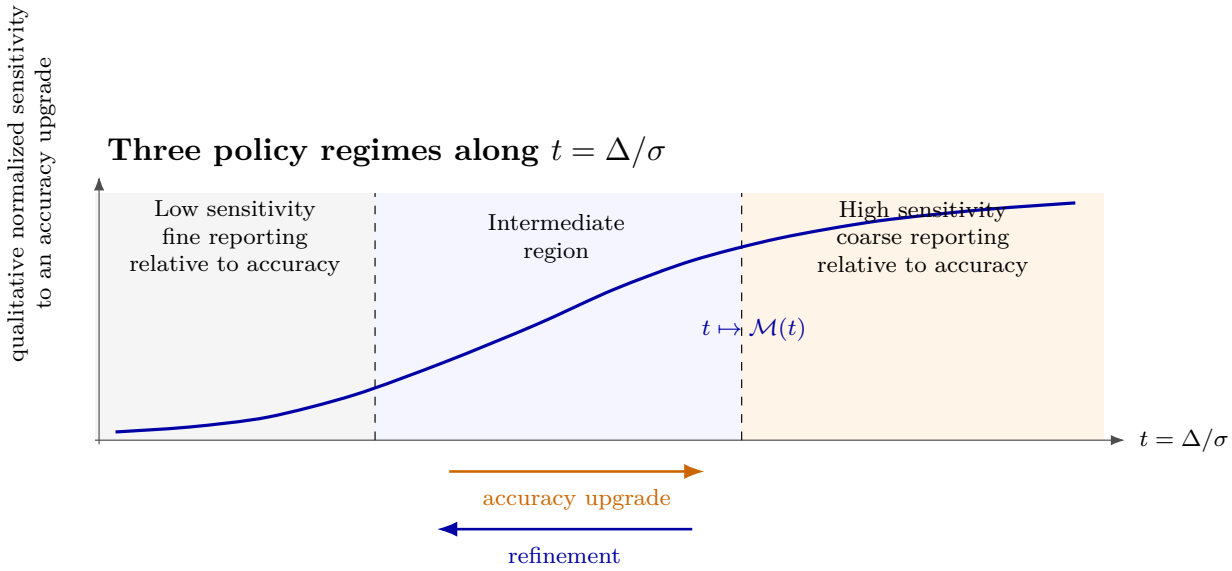

Figure 5: Three policy regimes along $t = \Delta/\sigma$. Moving right makes reporting coarser relative to accuracy and increases normalized sensitivity to an accuracy upgrade.

(ii) (Low-sensitivity region) *If t is small, then $\mathcal{M}(t)$ is close to 0, so normalized sensitivity is close to 0.*

(iii) (Intermediate region) *For intermediate t, normalized sensitivity varies smoothly with $\mathcal{M}(t)$.*

Proof. Immediate from Proposition 7: with $r = \sigma/\Delta$ and $t = \Delta/\sigma = 1/r$,

$$-\Psi'(r) = \mathcal{M}(t).$$

The regime statements then follow from monotonicity of $\mathcal{M}(t)$ in t. □

Corollary 3 turns the qualitative trade-off into a quantitative one. The relevant state variable is $t = \Delta/\sigma$, and the curve $t \mapsto \mathcal{M}(t)$ gives the scale-free local strength of the selection channel. Formally, the corollary concerns the marginal effect with respect to the ratio $r = \sigma/\Delta$. When Δ is held fixed, the same object is the normalized marginal effect of an accuracy upgrade. Figure 5 translates that one-dimensional regime map into a compact policy diagram.

Principle 5 (When accuracy helps, when it backfires). *When reporting is coarse relative to accuracy, an accuracy upgrade has a large normalized selection effect, so matching refinement becomes more important. When reporting is already fine relative to accuracy, that normalized effect is small, and the case for higher accuracy is driven mainly by learning.*

Principle 5 translates the three regimes into policy language. When reporting is coarse relative to accuracy, the selection channel is more sensitive, so improving accuracy is more

likely to expand the set of self-consistent cutoffs through ties and therefore to call for accompanying refinement. When reporting is already fine relative to accuracy, that normalized sensitivity is weak and the main effect of higher accuracy is to improve learning. Corollary 3 is useful because $\mathcal{M}(t)$ is directly estimable, so the magnitude of that selection margin can be measured rather than inferred indirectly.

8 Conclusion and Discussion

This paper studies coordination when signals are represented on a finite grid and treats reporting resolution as a primitive distinct from measurement accuracy. In the additive benchmark, the equilibrium set has width $\Delta + b\kappa$, so posterior ties are the sufficient statistic linking information structure to multiplicity. That observation separates two margins that are often conflated: finer reporting narrows the equilibrium set mechanically, whereas higher accuracy improves learning but can also enlarge multiplicity when reports remain coarse.

Beyond the benchmark, the tie–tilt decomposition shows which parts of the equilibrium geometry survive once translation invariance is relaxed, and the phase-robust analysis delivers a practical single-index summary for policy work. The broader lesson is that once representation is coarse, one should ask not only how noisy information is, but also how that information is encoded.

That distinction is also practical. In applications, *resolution* is the granularity of the reporting language: the number of categories in a rating system, the width of disclosure bands, or the rounding rule (e.g. nearest integer versus one-decimal reporting). By contrast, *accuracy* is the quality of the underlying measurement environment: audit intensity, sensor maintenance, evaluator training and standardisation, or the precision of measurement protocols. These margins can often be adjusted separately, which is why separating the learning channel from the equilibrium-selection channel is practically useful.

The joint-limit analysis also clarifies how to read the familiar Morris–Shin discontinuity. In the canonical continuous-signal model, every strictly positive noise level yields an atomless report distribution and therefore $\kappa = 0$, while complete information yields $\kappa = 1$. The jump at zero noise is therefore a property of the continuous representation. Once one allows richer information structures, a second point appears: noise variance does not fully measure the uncertainty relevant for coordination, because it does not record how posterior mass is distributed across report categories. A jump seen through variance can therefore reflect the weakness of the measure rather than a genuine discontinuity in equilibrium responses.

This suggests using entropy to measure uncertainty over posterior-relevant reports. Entropy is an information-theoretic measure of uncertainty defined on a distribution (Shannon,

1948). Here the relevant distribution is the posterior-relevant distribution of reports, because equilibrium selection depends on how strongly that distribution clusters on a few report values. In our benchmark, $\kappa = \sum_u p(u)^2$ is the collision probability, so the natural choice is collision entropy $H_2(p) = -\log \kappa$, the order-two Rényi entropy (Rényi, 1961). More generally, the same continuity conclusion extends across the broader Rényi family, and therefore also across Tsallis entropies (Tsallis, 1988):

Proposition 8 (Entropy continuity of the equilibrium interval)**.** *Consider a sequence of benchmark environments with posterior-relevant report pmf p_n and grid width Δ_n. Fix any $\alpha \in [0, \infty]$, and let $H_\alpha(p_n)$ be the Rényi entropy of order α. If $H_\alpha(p_n) \to 0$ and $\Delta_n \to 0$, then the equilibrium interval from Theorem 2 converges to the complete-information interval $[c-b, c]$. The same conclusion holds for Tsallis entropies, since they are monotone transforms of the Rényi family.*

Appendix A.14 proves the Rényi claim and the Tsallis corollary. The proposition captures the key point: once uncertainty is measured by entropy, vanishing uncertainty means concentration on a single posterior-relevant report, and the equilibrium set approaches its complete-information limit continuously. On this margin, entropy captures something that variance misses and that is central in coordination problems.

More broadly, variance and entropy answer different questions. Variance tracks the scale of latent measurement noise. Entropy tracks the concentration of the induced report distribution. It is also naturally suited to discrete objects, which is conceptually attractive here: coarse information need not be treated as a messy approximation to a continuous benchmark. With devices such as phase-robust averaging, discrete information structures can often be analysed cleanly on their own terms.

The framework also suggests a complementary perspective on the classic debate—initiated by Morris and Shin (2002) and developed in subsequent contributions such as Svensson (2006), Morris et al. (2006), Angeletos and Pavan (2007), and Hellwig (2005)—about when "public information" is socially beneficial or harmful in coordination games. The point here is not to revisit that debate, but to note that tie statistics provide one concrete reduced-form measure of effective commonality in agents' information: when posteriors cluster more tightly, ties become more frequent.

More frequent ties enlarge the set of self-consistent cutoffs and make equilibrium selection more salient. This mechanism is distinct from the standard variance-decomposition language, but it points in a related direction: when information structures generate more effective commonality, they may also generate more selection risk. Because κ is the collision probability, the same point also has an entropy formulation: lower entropy means less residual idiosyncratic uncertainty in the posterior-relevant report distribution, that is, a more

effectively public information structure. Agents are then more likely to receive the same report, ties become more frequent, and the set of self-consistent cutoffs widens. In that sense, tie statistics connect "publicness" to an observable object that can be measured from paired disagreements and used in policy analysis.

A Omitted Proofs and Additional Results

This appendix collects proofs and extensions referenced in the main text.

A.1 Monotone Posteriors and Beliefs About Others

Lemma 9 (Monotone posteriors imply MLR beliefs)**.** *Suppose that signals are conditionally independent given the fundamental, $s_i \perp s_j \mid \theta$. Suppose moreover that for each cutoff $x \in \Omega$, the function*

$$h_x(\theta) := \Pr(s_j \geq x \mid \theta)$$

is weakly increasing in θ. Then Assumption 4 implies Assumption 3.

Proof. Fix any cutoff $x \in \Omega$. By the law of iterated expectations,

$$\begin{aligned}\Pr(s_j \geq x \mid s_i) &= \mathbb{E}[\mathbf{1}\{s_j \geq x\} \mid s_i] \\ &= \mathbb{E}[\mathbb{E}[\mathbf{1}\{s_j \geq x\} \mid \theta, s_i] \mid s_i] \\ &= \mathbb{E}[\mathbb{E}[\mathbf{1}\{s_j \geq x\} \mid \theta] \mid s_i] = \mathbb{E}[h_x(\theta) \mid s_i]\,,\end{aligned}$$

where the third equality uses conditional independence $s_j \perp s_i \mid \theta$. Since $h_x(\theta)$ is weakly increasing in θ by assumption, Assumption 4 (applied with $g = h_x$) implies that $s_i \mapsto \mathbb{E}[h_x(\theta) \mid s_i] = \Pr(s_j \geq x \mid s_i)$ is weakly increasing. As x was arbitrary, Assumption 3 follows. □

A.2 Sufficient Conditions for Monotone Posteriors and Beliefs

This appendix collects simple sufficient conditions for Assumptions 4 and 3. We first record a general result for conditionally independent signals. We then specialise to the additive benchmark used in Section 4.

Proposition 9 (General sufficient conditions)**.** *Suppose $s_i \perp s_j \mid \theta$. If the family $\{f(\cdot \mid \theta)\}_{\theta \in \Omega}$ satisfies MLRP (equivalently TP_2) in (s, θ), then Assumption 4 holds. Moreover, if $\theta \mapsto \Pr(s_j \geq x \mid \theta)$ is weakly increasing for every cutoff x, then Assumption 3 holds as well.*

Proof. We first show Assumption 4. Under MLRP (equivalently TP_2), the joint kernel

$$k(s, \theta) := \pi(\theta) f(s \mid \theta)$$

is TP_2 in (s, θ) for any prior π with full support on Ω.[11] A standard implication of TP_2 is that the posterior family $\{\mathcal{L}(\theta \mid s)\}_{s \in \Omega}$ is increasing in s in the monotone likelihood ratio order (hence in the sense of first-order stochastic dominance). Therefore, for every weakly increasing function g on Ω, $s \mapsto \mathbb{E}[g(\theta) \mid s_i = s]$ is weakly increasing, which is Assumption 4.

For the second claim, suppose in addition that $\theta \mapsto \Pr(s_j \geq x \mid \theta)$ is weakly increasing for every cutoff x. Since we also assume $s_i \perp s_j \mid \theta$, Lemma 9 applies and yields Assumption 3. □

Proposition 10 (Additive benchmark: a discrete sufficient condition)**.** *Suppose $s_i = \theta + \varepsilon_i$, where $(\varepsilon_1, \varepsilon_2)$ are i.i.d. on Ω and independent of θ. Assume the noise pmf p_ε has full support on Ω and satisfies the following (discrete) monotone likelihood ratio condition: for every $\delta \in \Omega$ with $\delta > 0$,*[12]

$$e \;\mapsto\; \frac{p_\varepsilon(e - \delta)}{p_\varepsilon(e)} \quad \textit{is weakly increasing in } e.$$

Then Assumption 4 holds. Consequently, Assumption 3 holds as well.

Proof. **Step 1: the additive structure implies MLRP.** Under additivity, the conditional pmf of s_i given θ is

$$f(s \mid \theta) = p_\varepsilon(s - \theta), \qquad s, \theta \in \Omega.$$

Fix $\theta' > \theta$ and let $\delta := \theta' - \theta \in \Omega$ (so $\delta > 0$). Then for any $s \in \Omega$,

$$\frac{f(s \mid \theta')}{f(s \mid \theta)} = \frac{p_\varepsilon(s - \theta')}{p_\varepsilon(s - \theta)} = \frac{p_\varepsilon\big((s - \theta) - \delta\big)}{p_\varepsilon(s - \theta)}.$$

By the assumed discrete likelihood-ratio condition, the right-hand side is weakly increasing in $e := s - \theta$, hence weakly increasing in s. Therefore the family $\{f(\cdot \mid \theta)\}$ satisfies MLRP in (s, θ).

Step 2: MLRP implies monotone posteriors (Assumption 4). Since $(\varepsilon_1, \varepsilon_2)$ are i.i.d. and independent of θ, we have $s_i \perp s_j \mid \theta$. Applying Proposition 9 to the signal family just shown to satisfy MLRP yields Assumption 4.

[11] See, e.g., Milgrom (1981).

[12] Although imposed for all $\delta > 0$, this is simply the discrete monotone-likelihood-ratio (equivalently, TP_2) condition for the translation family $f(s \mid \theta) = p_\varepsilon(s - \theta)$. It is satisfied by a broad class of "well-behaved" unimodal noise distributions, including many common log-concave families (and their discretisations).

Step 3: monotone posteriors imply monotone beliefs (Assumption 3). Fix any cutoff $x \in \Omega$ and define $h_x(\theta) := \Pr(s_j \geq x \mid \theta)$. Under additivity and independence,

$$h_x(\theta) = \Pr(\theta + \varepsilon_j \geq x \mid \theta) = 1 - F_\varepsilon(x - \theta),$$

which is weakly increasing in θ. Hence Lemma 9 implies Assumption 3. □

A.3 Proof of Theorem 1

Maintain Assumptions 1, 2, 3, 4, and 5.

Step 0: The cutoff lattice. Let $\overline{\Omega} := \Omega \cup \{-\infty, +\infty\}$ ordered by the usual $\leq$ (with $-\infty \leq x \leq +\infty$ for all x). Then $(\overline{\Omega}, \leq)$ is a complete lattice.

For each cutoff $x \in \overline{\Omega}$, define the cutoff strategy $a^x(s) := \mathbf{1}\{s \geq x\}$. For cutoffs $x \leq x'$, we have $a^{x'} \preceq a^x$ pointwise.

Step 1: Interval-Valued Best Response. Fix an opponent cutoff $x \in \overline{\Omega}$. Let

$$\Delta(s; x) \equiv \mathbb{E}[d(\mathbf{1}\{s_j \geq x\}, \theta) \mid s_i = s]$$

be the interim payoff gain from choosing action 1 at signal s when the opponent plays a^x. By Lemma 1, $s \mapsto \Delta(s; x)$ is weakly increasing on Ω. Hence the set $\{s \in \Omega : \Delta(s; x) \geq 0\}$ is an upper set of Ω. Therefore the set of cutoff best responses is nonempty and forms an interval in $\overline{\Omega}$:

$$BR(x) := \{y \in \overline{\Omega} : \ a^y \text{ is a best response to } a^x\} \neq \emptyset,$$

and define the endpoint selections

$$\underline{b}(x) := \inf BR(x), \qquad \overline{b}(x) := \sup BR(x).$$

In our grid setting (augmented by $\pm\infty$), the interval $BR(x)$ attains its endpoints, so $\underline{b}(x) \in BR(x)$ and $\overline{b}(x) \in BR(x)$.

Step 2: Isotonicity of the extremal selections. Take $x \leq x'$. Then for every signal s, $\mathbf{1}\{s_j \geq x\} \geq \mathbf{1}\{s_j \geq x'\}$ pointwise in s_j, hence $\Pr(s_j \geq x \mid s_i = s) \geq \Pr(s_j \geq x' \mid s_i = s)$.

Under Assumption 5, for any opponent cutoff x we can write

$$\Delta(s; x) = \mathbb{E}[v(\theta) \mid s_i = s] - c + b \Pr(s_j \geq x \mid s_i = s),$$

where $b \geq 0$ by strategic complementarity (Assumption 1). Therefore $\Delta(s; x) \geq \Delta(s; x')$ for all s. Since each $\Delta(\cdot; x)$ is weakly increasing in s, this implies that when the opponent cutoff increases from x to x', any cutoff best response becomes (weakly) more conservative, i.e.,

$$x \leq x' \implies \underline{b}(x) \leq \underline{b}(x') \quad \text{and} \quad \overline{b}(x) \leq \overline{b}(x').$$

Thus $\underline{b}$ and $\overline{b}$ are isotone self-maps on the complete lattice $\overline{\Omega}$.

Step 3: Existence and extremal equilibria (Tarski). A symmetric cutoff equilibrium is an $x^* \in \overline{\Omega}$ such that $x^* \in BR(x^*)$.

Existence. Since $\underline{b}$ is isotone and $\overline{\Omega}$ is a complete lattice, Tarski's fixed-point theorem implies that $\underline{b}$ admits at least one fixed point. Let $\underline{x}$ be any fixed point of $\underline{b}$. Then $\underline{x} = \underline{b}(\underline{x}) \in BR(\underline{x})$ by Step 1, so $a^{\underline{x}}$ is a symmetric cutoff equilibrium. This establishes existence.

Extremality. By Tarski, the set of fixed points of $\underline{b}$ is a nonempty complete lattice and hence has a least element; denote it by $\underline{x}$. Similarly, the fixed points of $\overline{b}$ form a nonempty complete lattice and hence have a greatest element; denote it by $\overline{x}$.

Let $\mathcal{E} := \{x \in \overline{\Omega} : \ x \in BR(x)\}$ be the equilibrium cutoff set. If $x \in \mathcal{E}$, then $x \in BR(x)$ implies $\underline{b}(x) \leq x \leq \overline{b}(x)$. In particular, x is a post-fixed point of $\underline{b}$ (i.e. $\underline{b}(x) \leq x$) and a pre-fixed point of $\overline{b}$ (i.e. $x \leq \overline{b}(x)$). Tarski's theorem implies that the least fixed point of $\underline{b}$ is the infimum of its post-fixed points, and the greatest fixed point of $\overline{b}$ is the supremum of its pre-fixed points. Hence

$$\underline{x} \leq x \leq \overline{x} \qquad \text{for all } x \in \mathcal{E},$$

so $\underline{x}$ and $\overline{x}$ are respectively the smallest and largest equilibrium cutoffs. □

A.4 Proof of Theorem 2

Proof. Fix a candidate cutoff $x \in \Omega$ and suppose player j plays $a_j = \mathbf{1}\{s_j \geq x\}$. Under Assumption 5, player i's interim gain from choosing $a_i = 1$ rather than $a_i = 0$ at signal s is

$$\Delta(s; x) = m(s) - c + b \Pr(s_j \geq x \mid s_i = s).$$

With (2), this becomes $\Delta(s; x) = s - c + b \Pr(s_j \geq x \mid s_i = s)$.

A symmetric cutoff equilibrium at x requires two boundary conditions: (i) at $s = x$, acting must be (weakly) optimal; and (ii) at the previous grid point $s = x - \Delta$, not acting must be strictly optimal. Because $\Delta(s; x)$ is weakly increasing in s (Lemma 1), these two inequalities are sufficient.

Step 1: the boundary belief at $s_i = x$. Under translation invariance,

$$\Pr(s_j \geq x \mid s_i = x) = \Pr(s_j - s_i \geq 0 \mid s_i = x) = \Pr(\varepsilon_j - \varepsilon_i \geq 0).$$

By symmetry of $\varepsilon_j - \varepsilon_i$, we have

$$\Pr(\varepsilon_j - \varepsilon_i > 0) = \Pr(\varepsilon_j - \varepsilon_i < 0),$$

and therefore

$$\Pr(\varepsilon_j - \varepsilon_i \geq 0) = \frac{1}{2} + \frac{1}{2}\Pr(\varepsilon_j - \varepsilon_i = 0) = \frac{1}{2} + \frac{1}{2}\Pr(\varepsilon_1 = \varepsilon_2) = \frac{1}{2} + \frac{\kappa}{2}.$$

Thus $\Delta(x; x) \geq 0$ is equivalent to

$$x - c + b\left(\frac{1}{2} + \frac{\kappa}{2}\right) \geq 0 \quad \Longleftrightarrow \quad x \geq c - \frac{b}{2} - \frac{b\kappa}{2}.$$

Step 2: the boundary belief at $s_i = x - \Delta$. Similarly,

$$\Pr(s_j \geq x \mid s_i = x - \Delta) = \Pr(s_j - s_i \geq \Delta \mid s_i = x - \Delta) = \Pr(\varepsilon_j - \varepsilon_i \geq \Delta).$$

Again by symmetry,

$$\Pr(\varepsilon_j - \varepsilon_i \geq \Delta) = \Pr(\varepsilon_j - \varepsilon_i > 0) - \Pr(0 < \varepsilon_j - \varepsilon_i < \Delta).$$

On the grid, the interval $(0, \Delta)$ contains no support points, so $\Pr(0 < \varepsilon_j - \varepsilon_i < \Delta) = 0$, hence

$$\Pr(\varepsilon_j - \varepsilon_i \geq \Delta) = \Pr(\varepsilon_j - \varepsilon_i > 0).$$

Moreover,

$$1 = \Pr(\varepsilon_j - \varepsilon_i < 0) + \Pr(\varepsilon_j - \varepsilon_i = 0) + \Pr(\varepsilon_j - \varepsilon_i > 0),$$

and symmetry implies $\Pr(\varepsilon_j - \varepsilon_i > 0) = \Pr(\varepsilon_j - \varepsilon_i < 0) = \frac{1}{2}\big(1 - \Pr(\varepsilon_1 = \varepsilon_2)\big) = \frac{1}{2} - \frac{\kappa}{2}$. Therefore,

$$\Pr(s_j \geq x \mid s_i = x - \Delta) = \frac{1}{2} - \frac{\kappa}{2}.$$

The strict inequality $\Delta(x - \Delta; x) < 0$ is equivalent to

$$x - \Delta - c + b\left(\frac{1}{2} - \frac{\kappa}{2}\right) < 0 \quad \Longleftrightarrow \quad x < c + \Delta - \frac{b}{2} + \frac{b\kappa}{2}.$$

Combining the two boundary conditions yields the stated interval. Its width is $\Delta + b\kappa$. $\square$

A.5 Proof of Corollary 1

Proof. For each n, let

$$p_n(u) \;\equiv\; p_{\varepsilon,\sigma_n}^{\Delta_n,\phi}(u) = \Pr\left(\varepsilon_{\sigma_n}^{c} \in [u+\phi, u+\phi+\Delta_n)\right), \qquad \kappa_n \;\equiv\; \sum_{u\in\Omega_{\Delta_n}} p_n(u)^2.$$

Then $(p_n(u))_{u\in\Omega_{\Delta_n}}$ is a probability mass function, so $\sum_u p_n(u) = 1$ and $0 \le \kappa_n \le 1$.

Part (i). Assume $\Delta_n/\sigma_n \to 0$. Fix $\varepsilon > 0$. By Lemma 2, for each n there exists $\bar{\Delta}(\sigma_n, \varepsilon) > 0$ such that whenever $\Delta \le \bar{\Delta}(\sigma_n, \varepsilon)$,

$$\frac{\kappa_\Delta(\sigma_n)}{\Delta} \;\le\; \int_{\mathbb{R}} f_{\sigma_n}(x)^2\,dx + \varepsilon \;=\; \frac{1}{\sigma_n}\int_{\mathbb{R}} f(u)^2\,du + \varepsilon.$$

Since $\Delta_n \downarrow 0$, for all sufficiently large n we have $\Delta_n \le \bar{\Delta}(\sigma_n, \varepsilon)$. Therefore, for all large n,

$$\kappa_n = \kappa_{\Delta_n}(\sigma_n) \le \Delta_n\Big(\frac{1}{\sigma_n}\int f^2 + \varepsilon\Big) = \frac{\Delta_n}{\sigma_n}\int f^2 + \varepsilon\Delta_n \;\longrightarrow\; 0,$$

where the convergence uses $\Delta_n/\sigma_n \to 0$ and $\Delta_n \to 0$. Hence $\mathrm{width}(\mathcal{E}) = \Delta_n + b\,\kappa_n \to 0$, so the equilibrium interval collapses.

Part (ii). Assume $\liminf_{n\to\infty} \Delta_n/\sigma_n > 0$. Then there exist $\eta > 0$ and N such that $\Delta_n \ge \eta\sigma_n$ for all $n \ge N$.

By the nonempty-interior-support assumption on F, let x_0 be an interior point of its support. Fix $\delta \equiv \eta/2$, and define

$$m(\eta; x_0) \;\equiv\; \min\{F(x_0) - F(x_0 - \eta/2),\; F(x_0 + \eta/2) - F(x_0)\}.$$

By the interior-support property at x_0 (i.e., $F(x_0 - \delta) < F(x_0) < F(x_0 + \delta)$ for all $\delta > 0$), we have $m(\eta; x_0) > 0$.

Fix any $n \ge N$. Let $u_n \in \Omega_{\Delta_n}$ be such that $\sigma_n x_0 \in [u_n + \phi, u_n + \phi + \Delta_n)$, and write $a_n \equiv u_n + \phi$. Then $a_n \le \sigma_n x_0 < a_n + \Delta_n$. We claim that the cell $[a_n, a_n + \Delta_n)$ must contain at least one of the two subintervals $[\sigma_n(x_0 - \eta/2),\, \sigma_n x_0]$ or $[\sigma_n x_0,\, \sigma_n(x_0 + \eta/2)]$.

- If $a_n \le \sigma_n(x_0 - \eta/2)$, then $[\sigma_n(x_0 - \eta/2),\, \sigma_n x_0] \subseteq [a_n, \sigma_n x_0] \subseteq [a_n, a_n + \Delta_n)$.
- Otherwise $a_n > \sigma_n(x_0 - \eta/2)$, and since $\Delta_n \ge \eta\sigma_n$ we have $a_n + \Delta_n > \sigma_n(x_0 - \eta/2) + \eta\sigma_n = \sigma_n(x_0 + \eta/2)$, so $[\sigma_n x_0,\, \sigma_n(x_0 + \eta/2)] \subseteq [a_n, a_n + \Delta_n)$.

Therefore,

$$p_n(u_n) = \Pr\left(\varepsilon^c_{\sigma_n} \in [a_n, a_n + \Delta_n)\right) \;\geq\; \min\Big\{ \Pr\left(\varepsilon^c_{\sigma_n} \in [\sigma_n(x_0 - \eta/2),\, \sigma_n x_0]\right), \\ \Pr\left(\varepsilon^c_{\sigma_n} \in [\sigma_n x_0,\, \sigma_n(x_0 + \eta/2)]\right)\Big\}.$$

Using $\varepsilon^c_{\sigma_n} = \sigma_n \xi$,

$$\Pr\left(\varepsilon^c_{\sigma_n} \in [\sigma_n(x_0 - \eta/2),\, \sigma_n x_0]\right) = F(x_0) - F(x_0 - \eta/2),$$
$$\Pr\left(\varepsilon^c_{\sigma_n} \in [\sigma_n x_0,\, \sigma_n(x_0 + \eta/2)]\right) = F(x_0 + \eta/2) - F(x_0).$$

Hence $p_n(u_n) \geq m(\eta; x_0)$ for all $n \geq N$, and consequently

$$\kappa_n \;\geq\; p_n(u_n)^2 \;\geq\; m(\eta; x_0)^2 \;>\; 0 \quad \text{for all } n \geq N.$$

This implies $\liminf_{n\to\infty} \kappa_n > 0$, so $\mathrm{width}(\mathcal{E}) = \Delta_n + b\kappa_n$ does not converge to 0 and the equilibrium interval does not collapse.

Part (iii): phase-robust lower bound. Assume $\Delta_n/\sigma_n \to \infty$. Consider the deterministic interval $J_n \equiv [-\Delta_n, \Delta_n]$, which has length $2\Delta_n$ and contains 0. Then

$$\Pr\left(\varepsilon^c_{\sigma_n} \in J_n\right) = \Pr\left(|\xi| \leq \Delta_n/\sigma_n\right) = F\left(\frac{\Delta_n}{\sigma_n}\right) - F\left(-\frac{\Delta_n}{\sigma_n}\right) \;\longrightarrow\; 1.$$

Next we use a geometric fact: any interval of length $2\Delta_n$ can intersect at most three adjacent cells of the Δ_n-grid (regardless of phase). Therefore, J_n is contained in the union of at most three cells of the discretisation, say those corresponding to grid points $v_{n,1}, v_{n,2}, v_{n,3} \in \Omega_{\Delta_n}$. Hence

$$p_n(v_{n,1}) + p_n(v_{n,2}) + p_n(v_{n,3}) \;\geq\; \Pr\left(\varepsilon^c_{\sigma_n} \in J_n\right) \;\longrightarrow\; 1.$$

Let $s_{n,i} \equiv p_n(v_{n,i})$ for $i = 1, 2, 3$. Since $(s_{n,1}, s_{n,2}, s_{n,3})$ are nonnegative and sum to $S_n \equiv s_{n,1} + s_{n,2} + s_{n,3} \to 1$, the inequality

$$s_{n,1}^2 + s_{n,2}^2 + s_{n,3}^2 \;\geq\; \frac{1}{3}\left(s_{n,1} + s_{n,2} + s_{n,3}\right)^2 = \frac{1}{3}\, S_n^2$$

follows from Cauchy–Schwarz. Therefore,

$$\kappa_n = \sum_u p_n(u)^2 \;\geq\; s_{n,1}^2 + s_{n,2}^2 + s_{n,3}^2 \;\geq\; \frac{1}{3}\, S_n^2 \;\longrightarrow\; \frac{1}{3},$$

which proves $\liminf_{n\to\infty} \kappa_n \geq 1/3$. Since $\Delta_n \to 0$ and $\kappa_n \geq 1/3 + o(1)$,

$$\frac{\Delta_n}{b\,\kappa_n} \;\leq\; \frac{3\Delta_n}{b} \;\longrightarrow\; 0.$$

Part (iii): Lebesgue-a.e. strengthening under (5). To state a "Lebesgue-a.e. in ϕ" conclusion, note that the Δ_n-grid is periodic: shifting ϕ by an integer multiple of Δ_n only relabels grid points and does not change the induced partition. Consequently, for a given n it suffices to consider ϕ on any interval of length at least one full grid period Δ_n. Since $\Delta_n \downarrow 0$, we have $\Delta_n < 1$ for all sufficiently large n, so the fixed window $[0, 1]$ covers an entire period for all large n. We therefore work with $\phi \in [0, 1]$ below; this is without loss for the Lebesgue-a.e. statement.

Fix an integer $M \geq 1$. For each n define the set of phases for which 0 lies within distance $M\sigma_n$ of some bin boundary:

$$E_n(M) \;\equiv\; \{\phi \in [0, 1] : \operatorname{dist}(\phi, \Delta_n\mathbb{Z}) \leq M\sigma_n\}.$$

We bound the Lebesgue measure $\lambda(E_n(M))$. The set $\Delta_n\mathbb{Z} \cap [0, 1]$ has at most $\lceil 1/\Delta_n \rceil + 1$ points. Around each boundary point, the phases within distance $M\sigma_n$ form an interval of length at most $2M\sigma_n$. Hence

$$\lambda(E_n(M)) \;\leq\; 2M\sigma_n\big(\lceil 1/\Delta_n \rceil + 1\big) \;\leq\; \frac{4M\sigma_n}{\Delta_n} \quad \text{for all sufficiently large } n,$$

using $\Delta_n \downarrow 0$.

Under (5), we have $\sum_n \lambda(E_n(M)) < \infty$ for each fixed M. We now use the following elementary fact: if $\sum_n \lambda(E_n) < \infty$, then $\lambda(\limsup_n E_n) = 0$, where $\limsup_n E_n \equiv \cap_{n\geq 1} \cup_{k\geq n} E_k$ is the set of points that belong to E_n infinitely often. Indeed, for each N,

$$\lambda\Big(\bigcup_{k\geq n} E_k\Big) \;\leq\; \sum_{k\geq n} \lambda(E_k) \;\longrightarrow\; 0 \qquad \text{as } n \to \infty,$$

so $\lambda(\cap_n \cup_{k\geq n} E_k) = 0$.

Applying this to $E_n(M)$, we obtain: for each fixed M, for Lebesgue-a.e. ϕ there exists $N_M(\phi)$ such that $\phi \notin E_n(M)$ for all $n \geq N_M(\phi)$, i.e. $\operatorname{dist}(\phi, \Delta_n\mathbb{Z}) > M\sigma_n$ eventually. Taking the intersection over $M = 1, 2, \ldots$ (a countable intersection), we conclude that for Lebesgue-a.e. ϕ,

$$\frac{\operatorname{dist}(\phi, \Delta_n\mathbb{Z})}{\sigma_n} \;\longrightarrow\; \infty.$$

Fix such a phase ϕ. Let u_n denote the grid point whose cell $[u_n + \phi, u_n + \phi + \Delta_n)$ contains

0, and let

$$d_n \;\equiv\; \min\{-u_n - \phi,\; u_n + \phi + \Delta_n\}$$

be the distance from 0 to the nearest boundary of that cell. By construction, $d_n = \operatorname{dist}(\phi, \Delta_n\mathbb{Z})$. Therefore $d_n/\sigma_n \to \infty$. Finally,

$$1 - p_n(u_n) = \Pr\left(\varepsilon^c_{\sigma_n} \notin [u_n + \phi, u_n + \phi + \Delta_n)\right) \le \Pr\left(|\varepsilon^c_{\sigma_n}| \ge d_n\right) = \Pr\left(|\xi| \ge d_n/\sigma_n\right) \;\longrightarrow\; 0,$$

so $p_n(u_n) \to 1$. Hence $\kappa_n \ge p_n(u_n)^2 \to 1$, and since $\kappa_n \le 1$ we conclude $\kappa_n \to 1$. □

A.6 Global Monotonicity under Log-Concavity

Throughout this appendix we fix $\Delta > 0$ and work with the *symmetric* binning (phase $\phi = -\Delta/2$),

$$C_k \;\equiv\; \Big[\Big(k - \tfrac{1}{2}\Big)\Delta,\; \Big(k + \tfrac{1}{2}\Big)\Delta\Big), \qquad k \in \mathbb{Z}.$$

Let $\varepsilon^c_\sigma = \sigma\xi$ where ξ has density f. Define

$$p_k(\sigma) \;\equiv\; \Pr(\varepsilon^c_\sigma \in C_k), \qquad \kappa_\Delta(\sigma) \;\equiv\; \sum_{k\in\mathbb{Z}} p_k(\sigma)^2.$$

We prove Proposition 2. The argument has three steps: (i) log-concavity implies a pointwise likelihood-ratio monotonicity in $|x|$ and hence an MLR ordering across bins; (ii) the binned MLR ordering implies central-mass dominance for symmetric unimodal sequences; and (iii) central-mass dominance implies a larger tie probability via summation by parts.

Step 1: Log-concavity implies an MLR ordering across bins

Assumption 10 (Symmetric log-concavity)**.** *The density f is symmetric about 0 and log-concave on $\mathbb{R}$.*

Write $f_\sigma(x) = (1/\sigma)f(x/\sigma)$ for the density of ε^c_σ.

Lemma 10 (Pointwise likelihood-ratio monotonicity in $|x|$)**.** *Suppose Assumption 10 holds. For $0 < \sigma_1 < \sigma_2$, the pointwise likelihood ratio*

$$\rho(x) \;\equiv\; \frac{f_{\sigma_1}(x)}{f_{\sigma_2}(x)}$$

is (weakly) decreasing in $|x|$.

Proof. By symmetry it suffices to consider $x \geq 0$. Let $\varphi = \log f$, which is concave by log-concavity. Fix $a \equiv \sigma_2/\sigma_1 > 1$ and write $u \equiv x/\sigma_1$. Up to an additive constant,

$$\log \rho(x) = \varphi(x/\sigma_1) - \varphi(x/\sigma_2) = \varphi(u) - \varphi(u/a).$$

Define $g(u) \equiv \varphi(u) - \varphi(u/a)$ for $u \geq 0$. Since φ is concave, φ' is weakly decreasing wherever it exists; moreover, symmetry and log-concavity imply that φ is weakly decreasing on $[0, \infty)$, hence $\varphi'(t) \leq 0$ for a.e. $t > 0$. For $u > 0$,

$$g'(u) = \varphi'(u) - \frac{1}{a}\,\varphi'(u/a) \;\leq\; \varphi'(u/a) - \frac{1}{a}\,\varphi'(u/a) = \Big(1 - \frac{1}{a}\Big)\varphi'(u/a) \;\leq\; 0 \qquad \text{for a.e. } u > 0,$$

where we used $u/a < u$ and the monotonicity of φ'. Thus g is weakly decreasing on $[0, \infty)$, which implies that $\log \rho(x)$ (and hence $\rho(x)$) is weakly decreasing in $x \geq 0$. By symmetry, $\rho(x)$ is weakly decreasing in $|x|$. □

Lemma 11 (MLR across symmetric bins)**.** *Suppose Assumption 10 holds and fix $0 < \sigma_1 < \sigma_2$. Then the ratio*

$$r_k \;\equiv\; \frac{p_k(\sigma_1)}{p_k(\sigma_2)}$$

is (weakly) decreasing in $|k|$.

Proof. By symmetry, it suffices to consider $k = 0, 1, 2, \ldots$. Let $\rho(x) = f_{\sigma_1}(x)/f_{\sigma_2}(x)$, which is decreasing on $[0, \infty)$ by Lemma 10. For $k \geq 0$, write

$$p_k(\sigma_1) = \int_{C_k} f_{\sigma_1}(x)\,dx = \int_{C_k} \rho(x)\, f_{\sigma_2}(x)\,dx.$$

Hence

$$r_k = \frac{p_k(\sigma_1)}{p_k(\sigma_2)} = \frac{\int_{C_k} \rho(x)\, f_{\sigma_2}(x)\,dx}{\int_{C_k} f_{\sigma_2}(x)\,dx} = \mathbb{E}\big[\rho(X) \mid X \in C_k\big],$$

where X has density proportional to f_{σ_2}. Since ρ is decreasing and the sets $C_k \cap [0, \infty)$ shift right as k increases, the conditional expectation $\mathbb{E}[\rho(X) \mid X \in C_k]$ is weakly decreasing in k. Therefore r_k is weakly decreasing in $k \geq 0$, and by symmetry in $|k|$. □

Step 2: MLR implies central-mass dominance For a symmetric probability sequence $p = (\ldots, p_{-1}, p_0, p_1, \ldots)$ with $p_{-m} = p_m$ and $p_0 \geq p_1 \geq p_2 \geq \cdots$, define central partial sums

$$S_M(p) \equiv p_0 + 2\sum_{m=1}^{M} p_m, \qquad M = 0, 1, 2, \ldots$$

(i.e., total mass in the $2M+1$ central bins).

Lemma 12 (MLR implies central-mass dominance)**.** *Let p and q be symmetric unimodal probability sequences on $\mathbb{Z}$ with full support on their indices where defined. If p_m/q_m is weakly decreasing in $|m|$, then*

$$S_M(p) \;\geq\; S_M(q) \qquad \textit{for all } M \geq 0.$$

Proof. By symmetry it suffices to work with $m \geq 0$. Consider the difference sequence $d_m \equiv p_m - q_m$ for $m \geq 0$. The condition that p_m/q_m is decreasing implies a single-crossing property: there exists an index $m^\star \in \{0, 1, 2, \ldots, \infty\}$ such that $d_m \geq 0$ for $m \leq m^\star$ and $d_m \leq 0$ for $m > m^\star$. (Indeed, if $p_{m_1} \geq q_{m_1}$ at some m_1 then $p_m/q_m \geq 1$ for all $m \leq m_1$; similarly, if $p_{m_2} \leq q_{m_2}$ then $p_m/q_m \leq 1$ for all $m \geq m_2$.) Because $\sum_{m\geq 0}(2-\mathbf{1}\{m=0\})\, d_m = 0$ (both are probability sequences on $\mathbb{Z}$), the cumulative sums $\sum_{m=0}^{M}(2-\mathbf{1}\{m=0\})\, d_m$ are weakly nonnegative for every M, which is exactly $S_M(p) - S_M(q) \geq 0$. □

Step 3: Central-Mass Dominance

Lemma 13 (Summation by parts (Abel))**.** *Let $(b_m)_{m\geq 0}$ be any real sequence such that the partial sums*

$$B_M \;\equiv\; \sum_{m=0}^{M} b_m$$

are well-defined, and let $(c_m)_{m\geq 0}$ be a real sequence. Then for every $N \geq 0$,

$$\sum_{m=0}^{N} b_m c_m = B_N c_N \;+\; \sum_{M=0}^{N-1} B_M\,(c_M - c_{M+1}).$$

Proof. Write $b_m = B_m - B_{m-1}$ with the convention $B_{-1} = 0$. Then

$$\sum_{m=0}^{N} b_m c_m = \sum_{m=0}^{N}(B_m - B_{m-1})c_m = \sum_{m=0}^{N} B_m c_m - \sum_{m=0}^{N} B_{m-1}c_m = \sum_{m=0}^{N} B_m c_m - \sum_{m=-1}^{N-1} B_m c_{m+1}.$$

Since $B_{-1} = 0$, this becomes

$$\sum_{m=0}^{N} b_m c_m = B_N c_N + \sum_{m=0}^{N-1} B_m(c_m - c_{m+1}),$$

as claimed. □

Lemma 14 (Central-mass dominance and tie probability)**.** *Let p and q be symmetric unimodal probability sequences on $\mathbb{Z}$, i.e., $p_{-m} = p_m$, $q_{-m} = q_m$, and $p_0 \geq p_1 \geq p_2 \geq \cdots$,*

$q_0 \geq q_1 \geq q_2 \geq \cdots$. *If* $S_M(p) \geq S_M(q)$ *for all* $M \geq 0$, *then*

$$\sum_{k\in\mathbb{Z}} p_k^2 \;\geq\; \sum_{k\in\mathbb{Z}} q_k^2.$$

Proof. By symmetry,

$$\sum_{k\in\mathbb{Z}} (p_k^2 - q_k^2) = (p_0^2 - q_0^2) + 2\sum_{m\geq 1}(p_m^2 - q_m^2) = \sum_{m\geq 0} w_m\,(p_m - q_m)\,(p_m + q_m),$$

where $w_0 = 1$ and $w_m = 2$ for $m \geq 1$.

Define

$$b_m \equiv w_m\,(p_m - q_m), \qquad c_m \equiv p_m + q_m, \qquad B_M \equiv \sum_{m=0}^{M} b_m.$$

Then $B_M = S_M(p) - S_M(q) \geq 0$ by assumption.

Moreover, unimodality of p and q implies $c_m = p_m + q_m$ is weakly decreasing in $m \geq 0$, hence $c_M - c_{M+1} \geq 0$ for all $M \geq 0$.

Apply Lemma 13 to $\sum_{m=0}^{N} b_m c_m$:

$$\sum_{m=0}^{N} b_m c_m = B_N c_N \;+\; \sum_{M=0}^{N-1} B_M\,(c_M - c_{M+1}).$$

Every term on the right-hand side is weakly nonnegative: $B_M \geq 0$ and $c_M - c_{M+1} \geq 0$, and also $B_N c_N \geq 0$ since $c_N \geq 0$. Therefore $\sum_{m=0}^{N} b_m c_m \geq 0$ for each N.

Since p and q are probability sequences, we have $\sum_{m\geq 0} |b_m c_m| < \infty$: indeed, $|b_m c_m| \leq w_m(p_m + q_m)^2$ and

$$\sum_{m\geq 0} w_m (p_m + q_m)^2 \leq 2\sum_{m\geq 0} w_m (p_m^2 + q_m^2) = 2\sum_{k\in\mathbb{Z}} (p_k^2 + q_k^2) \leq 4.$$

Hence the partial sums $\sum_{m=0}^{N} b_m c_m$ converge to $\sum_{m\geq 0} b_m c_m$. Since each partial sum is nonnegative, the limit is nonnegative as well, i.e. $\sum_{m\geq 0} b_m c_m \geq 0$. This proves $\sum_k p_k^2 \geq \sum_k q_k^2$. □

Conclusion

Proof of Proposition 2. Fix $\Delta > 0$ and $0 < \sigma_1 < \sigma_2$. Lemma 11 implies that $p_k(\sigma_1)/p_k(\sigma_2)$ is decreasing in $|k|$. Since each $p(\sigma)$ is symmetric and unimodal under Assumption 10,

Lemma 12 gives $S_M(p(\sigma_1)) \geq S_M(p(\sigma_2))$ for all M. Lemma 14 then yields

$$\kappa_\Delta(\sigma_1) = \sum_k p_k(\sigma_1)^2 \;\geq\; \sum_k p_k(\sigma_2)^2 = \kappa_\Delta(\sigma_2).$$

□

A.7 Phase-Specific and Phase-Robust Tie Objects

This appendix clarifies the division of labour between the two tie objects used in the paper. For a fixed reporting scheme with known thresholds, the phase-specific tie probability κ_Δ^ϕ is the correct object: it describes that particular implementation. The phase-robust object $\bar{\kappa}_\Delta$ serves a different purpose. It is the relevant summary when the comparative-static question concerns the effect of the bin width Δ in a common underlying continuous environment rather than the consequences of one particular grid origin.

Proposition 3 already gives the mathematical reason for using $\bar{\kappa}_\Delta$: once phase is integrated out, tie risk becomes a pure distance comparison. The present appendix adds interpretive context.

First, in latent-variable environments the continuous scale itself often has no canonical origin. Ordered-response models are formulated on a latent index whose location and scale are fixed only after normalization or other identifying restrictions (Greene and Hensher, 2010; Wooldridge, 2010). In that class of settings, the absolute placement of discretisation thresholds partly inherits the arbitrariness of the coordinate system. Treating one particular phase as structurally privileged would therefore over-interpret a normalization choice.

Second, even when a reporting rule is institutionally fixed, threshold placement is often partly conventional or heterogeneous across uses. Work on response-category incomparability and anchoring vignettes documents that respondents need not apply the same verbal categories in the same way (King et al., 2003; Rice et al., 2010). Relatedly, Gideon et al. (2017) show that salient round-number thresholds can generate mechanical clustering in quantitative reports. These observations do not imply that the phase-robust object is always the right one. They do show that boundary placement can contain implementation-specific variation separate from the coarseness margin itself.

A nearby statistical analogy makes the same point. The averaged shifted histogram of Scott (1985) keeps bin width fixed while smoothing away origin sensitivity. The present construction plays an analogous role: it holds Δ fixed, integrates over translations, and isolates the part of tie risk attributable to reporting resolution rather than to where the cell edges happen to fall.

For that reason, the paper states its general monotonicity and single-index results in terms of $\bar{\kappa}_\Delta$. Those results concern comparative statics in reporting resolution on a common latent scale. By contrast, κ_Δ^ϕ remains the right object when the application is one particular implementation with fixed known thresholds.

A.8 Proof of Lemma 5

Proof of Lemma 5. Fix $s \in X^{\text{rel}}$ and work conditional on $s_i = s$. Let

$$p_+ \equiv \Pr(s_j > s \mid s_i = s), \qquad p_0 \equiv \Pr(s_j = s \mid s_i = s) = \kappa(s), \qquad p_- \equiv \Pr(s_j < s \mid s_i = s).$$

Then

$$p_+ + p_0 + p_- = 1, \qquad \delta(s) = p_+ - p_-.$$

Solving these two linear equations for p_+ and p_- gives

$$p_+ = \frac{1 - \kappa(s) + \delta(s)}{2}, \qquad p_- = \frac{1 - \kappa(s) - \delta(s)}{2}.$$

Now

$$q^L(s) = \Pr(s_j \geq s \mid s_i = s) = p_0 + p_+ = \kappa(s) + \frac{1 - \kappa(s) + \delta(s)}{2} = \frac{1 + \kappa(s) + \delta(s)}{2}.$$

Because Ω is a grid with mesh Δ, the event $s_j > s$ is equivalent to $s_j \geq s + \Delta$. Hence

$$q^U(s) = \Pr(s_j \geq s + \Delta \mid s_i = s) = \Pr(s_j > s \mid s_i = s) = p_+ = \frac{1 - \kappa(s) + \delta(s)}{2}.$$

This proves the two identities in the lemma, and the final display follows by subtraction. □

A.9 Proof of Lemma 6

Proof. Fix $s \in X^{\text{rel}}$. Under the additive signal structure and independence, the posterior pmf of θ conditional on $s_i = s$ is

$$\mu_s(\theta) \equiv \Pr(\theta \mid s_i = s) = \frac{\pi(\theta) p_\varepsilon(s - \theta)}{\sum_{\theta' \in \Omega} \pi(\theta') p_\varepsilon(s - \theta')}.$$

Define

$$D(u) \equiv \Pr(\varepsilon > u) - \Pr(\varepsilon < u), \qquad u \in \Omega.$$

Then

$$\begin{aligned}\delta(s) &= \Pr(s_j > s \mid s_i = s) - \Pr(s_j < s \mid s_i = s) \\ &= \sum_{\theta \in \Omega} \mu_s(\theta) \Big(\Pr(\theta + \varepsilon > s) - \Pr(\theta + \varepsilon < s) \Big) \\ &= \sum_{\theta \in \Omega} \mu_s(\theta)\, D(s - \theta). \end{aligned} \tag{10}$$

Under Assumption 7, the noise distribution is symmetric, so D is odd:

$$D(-u) = -D(u) \qquad \text{for all } u \in \Omega. \tag{11}$$

Now define the reflection of θ around s by

$$\theta^\sharp \equiv 2s - \theta.$$

Because $s, \theta \in \Omega$, we also have $\theta^\sharp \in \Omega$. By symmetry of the noise,

$$p_\varepsilon(s - \theta) = p_\varepsilon(\theta - s) = p_\varepsilon(s - \theta^\sharp),$$

and therefore

$$\frac{\mu_s(\theta)}{\mu_s(\theta^\sharp)} = \frac{\pi(\theta) p_\varepsilon(s - \theta)}{\pi(\theta^\sharp) p_\varepsilon(s - \theta^\sharp)} = \frac{\pi(\theta)}{\pi(\theta^\sharp)}, \tag{12}$$

whenever both θ and $\theta^\sharp$ lie in $\mathrm{supp}(\mu_s)$.

Using (11), we have

$$D(s - \theta^\sharp) = D(\theta - s) = -D(s - \theta).$$

Hence the terms in (10) can be paired at θ and $\theta^\sharp$:

$$\mu_s(\theta) D(s - \theta) + \mu_s(\theta^\sharp) D(s - \theta^\sharp) = \big(\mu_s(\theta) - \mu_s(\theta^\sharp) \big) D(s - \theta).$$

Summing over $\theta \in \Omega$ and dividing by 2 to avoid double counting gives

$$\delta(s) = \frac{1}{2} \sum_{\theta \in \Omega} \big(\mu_s(\theta) - \mu_s(\theta^\sharp) \big) D(s - \theta). \tag{13}$$

If π is weakly increasing on $\mathrm{supp}(\mu_s)$, then for every reflected pair,

$$\mu_s(\theta) \geq \mu_s(\theta^\sharp) \qquad \text{whenever } \theta \geq s.$$

If $\theta \geq s$, then $s - \theta \leq 0$, so $D(s-\theta) \geq 0$ by oddness and symmetry. Hence

$$\big(\mu_s(\theta) - \mu_s(\theta^\sharp)\big) D(s-\theta) \geq 0.$$

If instead $\theta \leq s$, then $\theta^\sharp \geq s$, so

$$\mu_s(\theta) - \mu_s(\theta^\sharp) \leq 0 \qquad \text{and} \qquad D(s-\theta) \leq 0,$$

and the same product is again weakly nonnegative. Therefore every paired term in (13) is weakly nonnegative, and so $\delta(s) \geq 0$.

Finally, if π is constant on $\operatorname{supp}(\mu_s)$, then $\pi(\theta) = \pi(\theta^\sharp)$ for every reflected pair in the posterior support. By (12), this implies

$$\mu_s(\theta) = \mu_s(\theta^\sharp),$$

so every paired term in (13) is zero, and therefore $\delta(s) = 0$. □

A.10 Proof of Lemma 7

Proof. Let

$$D(u) \equiv \Pr(\varepsilon > u) - \Pr(\varepsilon < u) = 1 - 2F_\varepsilon(u),$$

and let

$$\lambda(\theta) \equiv (\log \pi)'(\theta).$$

Since F_ε is weakly increasing, D is weakly decreasing in u. Under log-concavity of π, λ is weakly decreasing in θ.

For a given signal s, write the posterior density of θ conditional on $s_i = s$ as

$$\mu_s(\theta) \equiv \frac{\pi(\theta) f_\varepsilon(s-\theta)}{\int_{\mathbb{R}} \pi(t) f_\varepsilon(s-t)\, dt}.$$

Then the continuous tilt can be written as

$$\delta^c(s) = \int_{\mathbb{R}} D(s-\theta)\, \mu_s(\theta)\, d\theta.$$

Using the change of variables $u = s - \theta$, this becomes

$$\delta^c(s) = \int_{\mathbb{R}} D(u)\, w_s(u)\, du, \qquad w_s(u) \equiv \frac{\pi(s-u) f_\varepsilon(u)}{Z(s)}, \qquad Z(s) \equiv \int_{\mathbb{R}} \pi(s-v) f_\varepsilon(v)\, dv.$$

Differentiating $w_s(u)$ with respect to s gives

$$\frac{\partial}{\partial s} w_s(u) = w_s(u) \left(\frac{\pi'(s-u)}{\pi(s-u)} - \frac{Z'(s)}{Z(s)} \right).$$

Moreover,

$$\frac{Z'(s)}{Z(s)} = \frac{\int_{\mathbb{R}} \pi'(s-u) f_\varepsilon(u)\, du}{\int_{\mathbb{R}} \pi(s-u) f_\varepsilon(u)\, du} = \int_{\mathbb{R}} \lambda(s-u)\, w_s(u)\, du.$$

Hence

$$\frac{\partial}{\partial s} w_s(u) = w_s(u) \big(\lambda(s-u) - \mathbb{E}_{w_s}[\lambda(s-U)] \big),$$

where U has density w_s. Therefore

$$\begin{aligned}
\delta^{c\prime}(s) &= \frac{d}{ds} \int D(u)\, w_s(u)\, du \\
&= \int D(u)\, \frac{\partial}{\partial s} w_s(u)\, du \\
&= \int D(u)\, w_s(u) \big(\lambda(s-u) - \mathbb{E}_{w_s}[\lambda(s-U)] \big)\, du \\
&= \mathbb{E}_{w_s}[D(U)\lambda(s-U)] - \mathbb{E}_{w_s}[D(U)]\, \mathbb{E}_{w_s}[\lambda(s-U)] \\
&= \mathrm{Cov}_{w_s} \big(D(U),\, \lambda(s-U) \big).
\end{aligned}$$

Now $D(U)$ is weakly decreasing in U, while $\lambda(s-U)$ is weakly increasing in U, because λ is weakly decreasing in its argument. Hence the covariance of these two functions under the same probability law is weakly negative. Therefore

$$\delta^{c\prime}(s) \leq 0 \qquad \text{for all } s.$$

So the continuous tilt is weakly decreasing. □

A.11 Proof of Proposition 5

Proof. Fix $s \in X^{\mathrm{rel}}$ and $R > 0$. Under the additive signal structure and independence, the posterior pmf of θ conditional on $s_i = s$ is

$$\mu_s(\theta) \equiv \Pr(\theta \mid s_i = s) = \frac{\pi(\theta) p_\varepsilon(s-\theta)}{\sum_{\theta' \in \Omega} \pi(\theta') p_\varepsilon(s-\theta')}.$$

Define

$$D(u) \equiv \Pr(\varepsilon > u) - \Pr(\varepsilon < u), \qquad u \in \Omega.$$

Under Assumption 7, D is odd and satisfies $|D(u)| \le 1$ for all u.

As in the proof of Lemma 6,

$$\delta(s) = \sum_{\theta \in \Omega} \mu_s(\theta)\, D(s-\theta).$$

Split the posterior into a local part and a tail part:

$$\delta(s) = \sum_{|s-\theta| \le R} \mu_s(\theta)\, D(s-\theta) + \sum_{|s-\theta| > R} \mu_s(\theta)\, D(s-\theta).$$

Hence

$$|\delta(s)| \le \left| \sum_{|s-\theta| \le R} \mu_s(\theta)\, D(s-\theta) \right| + \Pr(|s-\theta| > R \mid s_i = s), \tag{14}$$

since $|D| \le 1$.

It remains to bound the local term. Let

$$\mu_s^R(\theta) \equiv \Pr(\theta \mid s_i = s,\ |s-\theta| \le R) = \frac{\mu_s(\theta)\mathbf{1}\{|s-\theta| \le R\}}{\Pr(|s-\theta| \le R \mid s_i = s)}$$

be the posterior truncated to the local window, and write

$$\delta_R(s) \equiv \mathbb{E}_{\mu_s^R}[D(s-\theta)].$$

Then the first term in (14) equals

$$\Pr(|s-\theta| \le R \mid s_i = s)\, |\delta_R(s)| \le |\delta_R(s)|.$$

Now define the reflection of θ around s by

$$\theta^\sharp \equiv 2s - \theta.$$

Because $|s-\theta| \le R$ implies $|s-\theta^\sharp| \le R$, and because $\theta, \theta^\sharp \in \Theta^{\mathrm{rel}}(R)$, Assumption 8 gives

$$\left| \log \pi(\theta) - \log \pi(\theta^\sharp) \right| \le L_{\mathrm{loc}}(R)\, |\theta - \theta^\sharp| = 2L_{\mathrm{loc}}(R)\, |s-\theta| \le 2L_{\mathrm{loc}}(R)R.$$

By symmetry of the noise,

$$p_\varepsilon(s-\theta) = p_\varepsilon(\theta - s) = p_\varepsilon(s - \theta^\sharp),$$

so

$$\frac{\mu_s(\theta)}{\mu_s(\theta^\sharp)} = \frac{\pi(\theta)p_\varepsilon(s-\theta)}{\pi(\theta^\sharp)p_\varepsilon(s-\theta^\sharp)} = \frac{\pi(\theta)}{\pi(\theta^\sharp)}.$$

The same equality holds for the truncated posterior μ_s^R, because truncation only rescales both numerator and denominator by the same conditioning probability. Therefore

$$e^{-2L_{\mathrm{loc}}(R)R} \le \frac{\mu_s^R(\theta)}{\mu_s^R(\theta^\sharp)} \le e^{2L_{\mathrm{loc}}(R)R} \qquad \text{for all } |s-\theta| \le R.$$

A standard likelihood-ratio bound then implies

$$\|\mu_s^R - (\mu_s^R)^\sharp\|_{\mathrm{TV}} \le \tanh\big(L_{\mathrm{loc}}(R)R\big), \tag{15}$$

where $(\mu_s^R)^\sharp(\theta) \equiv \mu_s^R(\theta^\sharp)$.[13]

Next, using the oddness of D, we compute

$$\begin{aligned}
\mathbb{E}_{(\mu_s^R)^\sharp}[D(s-\theta)] &= \sum_\theta (\mu_s^R)^\sharp(\theta)\, D(s-\theta) \\
&= \sum_{\theta^\sharp} \mu_s^R(\theta^\sharp)\, D\big(-(s-\theta^\sharp)\big) \\
&= -\sum_{\theta^\sharp} \mu_s^R(\theta^\sharp)\, D(s-\theta^\sharp) \\
&= -\mathbb{E}_{\mu_s^R}[D(s-\theta)] \\
&= -\delta_R(s),
\end{aligned}$$

where the second equality is the change of variables $\theta^\sharp = 2s-\theta$, and the third uses $D(-u) = -D(u)$.

Therefore,

$$2|\delta_R(s)| = \Big|\mathbb{E}_{\mu_s^R}[D(s-\theta)] - \mathbb{E}_{(\mu_s^R)^\sharp}[D(s-\theta)]\Big| \le 2\|\mu_s^R - (\mu_s^R)^\sharp\|_{\mathrm{TV}},$$

where the inequality uses $|D| \le 1$.[14] Combining with (15) yields

$$|\delta_R(s)| \le \tanh\big(L_{\mathrm{loc}}(R)R\big).$$

[13] If two probability measures ν and $\nu^\sharp$ satisfy $e^{-a} \le \nu(\theta)/\nu^\sharp(\theta) \le e^a$ pointwise, then $\|\nu - \nu^\sharp\|_{\mathrm{TV}} \le \tanh(a/2)$. Here $a = 2L_{\mathrm{loc}}(R)R$, which gives (15).

[14] For any bounded measurable function f with $|f| \le 1$, $|\mathbb{E}_\nu[f] - \mathbb{E}_{\nu'}[f]| \le 2\|\nu - \nu'\|_{\mathrm{TV}}$. We apply this with $f(\theta) = D(s-\theta)$.

Substituting this into (14) gives

$$|\delta(s)| \le \tanh\big(L_{\mathrm{loc}}(R)R\big) + \Pr(|s-\theta| > R \mid s_i = s),$$

which proves (8).

Taking the supremum over $s \in X^{\mathrm{rel}}$ and using the definition of $T(R)$ gives (9). □

A.12 Explicit Phase-Robust Boundary Kernels

This appendix records the formal construction used in Section 6.5. Its purpose is twofold: first, to define the phase-robust cutoff region $\bar{\mathcal{E}}_\Delta$ rigorously; second, to collect the explicit phase-robust boundary kernels used in the proof of Theorem 6.

As in Section 5, fix an underlying continuous signal and interpret different values of Δ as alternative discretisations of that same variable. For each mesh Δ and phase $\phi \in [0, \Delta)$, let

$$q_{\Delta,\phi} : \mathbb{R} \to \Omega_\Delta$$

denote the binning map

$$q_{\Delta,\phi}(x) \equiv \Delta \left\lfloor \frac{x-\phi}{\Delta} \right\rfloor.$$

Thus $x \in [q_{\Delta,\phi}(x) + \phi,\ q_{\Delta,\phi}(x) + \phi + \Delta)$, and integrating out phase amounts to averaging over the within-cell position, equivalently over the phase ϕ of the discretisation grid.

For a continuous signal difference $d \in \mathbb{R}$, define the phase-robust lower-boundary, upper-boundary, and tie kernels by

$$\rho_\Delta(d) \equiv \Pr_\Phi\big(q_{\Delta,\Phi}(y_j) \ge q_{\Delta,\Phi}(y_i) \mid y_j - y_i = d\big),$$

$$\eta_\Delta(d) \equiv \Pr_\Phi\big(q_{\Delta,\Phi}(y_j) \ge q_{\Delta,\Phi}(y_i) + \Delta \mid y_j - y_i = d\big),$$

$$\tau_\Delta(d) \equiv \Pr_\Phi\big(q_{\Delta,\Phi}(y_j) = q_{\Delta,\Phi}(y_i) \mid y_j - y_i = d\big).$$

Given these kernels, define for each $y \in X^{\mathrm{rel}}$ the conditional difference law

$$D_y \equiv y_j - y_i \mid y_i = y,$$

the corresponding phase-robust boundary beliefs

$$\bar{q}^L_\Delta(y) \equiv \mathbb{E}[\rho_\Delta(D_y)], \qquad \bar{q}^U_\Delta(y) \equiv \mathbb{E}[\eta_\Delta(D_y)], \qquad \bar{\kappa}_\Delta(y) \equiv \mathbb{E}[\tau_\Delta(D_y)],$$

and the induced boundary maps

$$\bar{A}_\Delta(y) \equiv c - b\,\bar{q}^L_\Delta(y), \qquad \bar{B}_\Delta(y) \equiv c + \Delta - b\,\bar{q}^U_\Delta(y).$$

The phase-robust cutoff region is then defined by

$$\bar{\mathcal{E}}_\Delta \equiv \left\{x \in X^{\text{rel}} :\ x \geq \bar{A}_\Delta(x) \text{ and } x < \bar{B}_\Delta(x-\Delta)\right\}.$$

Thus $\bar{\mathcal{E}}_\Delta$ is the set of cutoffs in the cutoff-relevant region that satisfy the lower and upper adjacent-boundary conditions after integrating out phase.

Lemma 15 (Explicit phase-robust boundary kernels)**.** *For every $d \in \mathbb{R}$,*

$$\rho_\Delta(d) = \begin{cases} 1, & d \geq 0, \\ \left(1 + \dfrac{d}{\Delta}\right)_+, & d < 0, \end{cases} \qquad \eta_\Delta(d) = \begin{cases} 0, & d \leq 0, \\ \left(\dfrac{d}{\Delta}\right) \wedge 1, & d > 0, \end{cases}$$

and

$$\tau_\Delta(d) = \left(1 - \frac{|d|}{\Delta}\right)_+.$$

Moreover,

$$\rho_\Delta(d) = \eta_\Delta(d) + \tau_\Delta(d),$$

$\rho_\Delta(d)$ is weakly increasing in Δ, and $\eta_\Delta(d)$ is weakly decreasing in Δ.

Proof. Fix $d = y_j - y_i$.

If $d \geq 0$, then the opponent's continuous signal is weakly above the player's, so after discretisation the opponent's bin is always weakly above the player's bin. Hence $\rho_\Delta(d) = 1$. If instead $d < 0$, the lower-boundary event fails precisely when a grid boundary falls in the interval $(y_j, y_i]$, whose length is $|d|$. Under a uniform phase draw, the boundary location is uniform on a Δ-period, so the failure probability is $(|d|/\Delta) \wedge 1$. Therefore

$$\rho_\Delta(d) = 1 - \left(\frac{|d|}{\Delta} \wedge 1\right) = \left(1 + \frac{d}{\Delta}\right)_+.$$

This gives the stated formula for ρ_Δ.

The same argument yields η_Δ. If $d \leq 0$, then the opponent cannot be at least one bin above the player, so $\eta_\Delta(d) = 0$. If $0 < d < \Delta$, the upper-boundary event occurs exactly when a grid boundary falls in the interval $(y_i, y_j]$, whose length is d, so $\eta_\Delta(d) = d/\Delta$. If $d \geq \Delta$, the opponent is at least one full bin above the player for every phase, so $\eta_\Delta(d) = 1$.

Finally, the tie event occurs exactly when no grid boundary falls between the two continuous signals. Thus its phase-robust probability is

$$\tau_\Delta(d) = \left(1 - \frac{|d|}{\Delta}\right)_+ .$$

Because the lower-boundary event is the disjoint union of the upper-boundary event and the tie event, we have $\rho_\Delta = \eta_\Delta + \tau_\Delta$.

The monotonicity statements are immediate from the explicit formulas: for fixed d, enlarging Δ weakly raises $\rho_\Delta(d)$ and weakly lowers $\eta_\Delta(d)$. □

As a consequence, for each $y \in X^{\mathrm{rel}}$,

$$\bar{q}^L_\Delta(y) = \bar{q}^U_\Delta(y) + \bar{\kappa}_\Delta(y)$$

and therefore

$$\bar{B}_\Delta(y) - \bar{A}_\Delta(y) = \Delta + b\,\bar{\kappa}_\Delta(y).$$

A.13 Proof of Theorem 6

Proof. Fix $0 < \Delta' < \Delta$. We show that every $x \in \bar{\mathcal{E}}_{\Delta'}$ also belongs to $\bar{\mathcal{E}}_\Delta$.

By Lemma 15, for every $d \in \mathbb{R}$,

$$\rho_\Delta(d) \geq \rho_{\Delta'}(d), \qquad \eta_\Delta(d) \leq \eta_{\Delta'}(d).$$

Hence, for every $y \in X^{\mathrm{rel}}$,

$$\bar{q}^L_\Delta(y) \geq \bar{q}^L_{\Delta'}(y) \qquad \text{and} \qquad \bar{q}^U_\Delta(y) \leq \bar{q}^U_{\Delta'}(y).$$

Equivalently,

$$\bar{A}_\Delta(y) \leq \bar{A}_{\Delta'}(y) \qquad \text{for all } y \in X^{\mathrm{rel}}.$$

Now let $x \in \bar{\mathcal{E}}_{\Delta'}$. By definition,

$$x \geq \bar{A}_{\Delta'}(x) \qquad \text{and} \qquad x < \bar{B}_{\Delta'}(x - \Delta').$$

The lower-boundary inequality is immediate from $\bar{A}_\Delta(x) \leq \bar{A}_{\Delta'}(x)$, so it remains to show

$$\bar{B}_\Delta(x - \Delta) \geq \bar{B}_{\Delta'}(x - \Delta').$$

Write

$$\bar{B}_\Delta(x-\Delta) - \bar{B}_{\Delta'}(x-\Delta') = (\Delta - \Delta') - b\Big(\bar{q}^U_\Delta(x-\Delta) - \bar{q}^U_{\Delta'}(x-\Delta')\Big).$$

Add and subtract $\bar{q}^U_{\Delta'}(x-\Delta)$ to obtain

$$\begin{aligned}\bar{B}_\Delta(x-\Delta) - \bar{B}_{\Delta'}(x-\Delta') = (\Delta - \Delta') &- b\Big(\bar{q}^U_\Delta(x-\Delta) - \bar{q}^U_{\Delta'}(x-\Delta)\Big) \\ &- b\Big(\bar{q}^U_{\Delta'}(x-\Delta) - \bar{q}^U_{\Delta'}(x-\Delta')\Big).\end{aligned}$$

The first bracketed term is weakly nonpositive because $\bar{q}^U_\Delta \le \bar{q}^U_{\Delta'}$ pointwise. Therefore

$$\bar{B}_\Delta(x-\Delta) - \bar{B}_{\Delta'}(x-\Delta') \ge (\Delta - \Delta') - b\Big|\bar{q}^U_{\Delta'}(x-\Delta) - \bar{q}^U_{\Delta'}(x-\Delta')\Big|. \tag{16}$$

It remains to bound the predecessor correction. Since $0 \le \eta_{\Delta'} \le 1$,

$$\Big|\bar{q}^U_{\Delta'}(y) - \bar{q}^U_{\Delta'}(y')\Big| = \Big|\mathbb{E}[\eta_{\Delta'}(D_y)] - \mathbb{E}[\eta_{\Delta'}(D_{y'})]\Big| \le \big\|\mathrm{Law}(D_y) - \mathrm{Law}(D_{y'})\big\|_{\mathrm{TV}},$$

where the inequality uses $0 \le \eta_{\Delta'} \le 1$.[15] By Assumption 9,

$$\Big|\bar{q}^U_{\Delta'}(y) - \bar{q}^U_{\Delta'}(y')\Big| \le L_D\,|y - y'|.$$

Applying this with $y = x - \Delta$ and $y' = x - \Delta'$ gives

$$\Big|\bar{q}^U_{\Delta'}(x-\Delta) - \bar{q}^U_{\Delta'}(x-\Delta')\Big| \le L_D(\Delta - \Delta').$$

Substituting into (16),

$$\bar{B}_\Delta(x-\Delta) - \bar{B}_{\Delta'}(x-\Delta') \ge (1 - bL_D)(\Delta - \Delta').$$

Because $L_D < 1/b$, the right-hand side is strictly positive. Since $x < \bar{B}_{\Delta'}(x-\Delta')$, it follows that

$$x < \bar{B}_\Delta(x-\Delta).$$

Together with $x \ge \bar{A}_\Delta(x)$, this shows $x \in \bar{\mathcal{E}}_\Delta$. Therefore

$$\bar{\mathcal{E}}_{\Delta'} \subseteq \bar{\mathcal{E}}_\Delta.$$

[15] For any bounded measurable function f with $0 \le f \le 1$, $|\mathbb{E}_\nu[f] - \mathbb{E}_{\nu'}[f]| \le \|\nu - \nu'\|_{\mathrm{TV}}$. We apply this with $f(d) = \eta_{\Delta'}(d)$.

The endpoint conclusion for interval regions follows immediately from set inclusion. □

A.14 Proof of Proposition 8

We first prove the Rényi claim in Proposition 8; the Tsallis statement then follows by monotone transformation. Consider a sequence of benchmark environments with posterior-relevant report pmf p_n and grid width Δ_n. Define

$$\kappa_n \equiv \sum_u p_n(u)^2, \qquad L_n \equiv c - \frac{b}{2} - \frac{b\kappa_n}{2}, \qquad U_n \equiv c + \Delta_n - \frac{b}{2} + \frac{b\kappa_n}{2}.$$

For each n, let

$$p_n^* \equiv \max_{u \in \Omega_{\Delta_n}} p_n(u).$$

For $\alpha \in (0, \infty)$, define the Rényi entropy of order α by

$$H_\alpha(p_n) \equiv \frac{1}{1-\alpha} \log\Big(\sum_u p_n(u)^\alpha \Big),$$

with $H_1(p_n)$ given by the Shannon limit and $H_\infty(p_n) \equiv -\log p_n^*$. Since $\kappa_n = \sum_u p_n(u)^2$, we have

$$\kappa_n \geq (p_n^*)^2.$$

Therefore

$$1 - \kappa_n \leq 1 - (p_n^*)^2 \leq 2(1 - p_n^*).$$

Hence

$$1 - \kappa_n \leq 2\big(1 - e^{-H_\infty(p_n)}\big) \leq 2H_\infty(p_n),$$

where the last inequality uses $1 - e^{-x} \leq x$ for $x \geq 0$.

It remains to show that $H_\infty(p_n) \leq H_\alpha(p_n)$ for every $\alpha \in (0, \infty)$. If $\alpha > 1$, then

$$\sum_u p_n(u)^\alpha = \sum_u p_n(u)\, p_n(u)^{\alpha-1} \leq (p_n^*)^{\alpha-1} \sum_u p_n(u) = (p_n^*)^{\alpha-1}.$$

Since $1 - \alpha < 0$,

$$H_\alpha(p_n) = \frac{1}{1-\alpha} \log\Big(\sum_u p_n(u)^\alpha \Big) \geq \frac{1}{1-\alpha} \log\big((p_n^*)^{\alpha-1}\big) = -\log p_n^* = H_\infty(p_n).$$

If $0 < \alpha < 1$, then $\alpha - 1 < 0$, so $p_n(u)^{\alpha-1} \geq (p_n^*)^{\alpha-1}$ for every u, and hence

$$\sum_u p_n(u)^\alpha = \sum_u p_n(u)\, p_n(u)^{\alpha-1} \geq (p_n^*)^{\alpha-1} \sum_u p_n(u) = (p_n^*)^{\alpha-1}.$$

Since $1-\alpha > 0$, the same calculation gives $H_\alpha(p_n) \geq H_\infty(p_n)$. The case $\alpha = \infty$ is tautological, and $\alpha = 1$ follows by continuity of the Rényi family at $\alpha = 1$. Therefore

$$\frac{b}{2}(1 - \kappa_n) \leq b\, H_\alpha(p_n).$$

Thus, for every $\alpha \in (0, \infty]$, $H_\alpha(p_n) \to 0$ implies $1 - \kappa_n \to 0$, and hence $\kappa_n \to 1$.

For Hartley entropy,

$$H_0(p_n) = \log |\mathrm{supp}(p_n)|.$$

If $H_0(p_n) \to 0$, then $|\mathrm{supp}(p_n)| \to 1$. Since support size is an integer at least 1, this means $|\mathrm{supp}(p_n)| = 1$ eventually, so $\kappa_n = 1$ eventually.

In either case, $\kappa_n \to 1$. The endpoints of the equilibrium interval therefore satisfy

$$0 \leq L_n - (c - b) = \frac{b}{2}(1 - \kappa_n) \to 0$$

and

$$|U_n - c| = \left| \Delta_n - \frac{b}{2}(1 - \kappa_n) \right| \leq \Delta_n + \frac{b}{2}(1 - \kappa_n) \to 0,$$

whenever $\Delta_n \to 0$. Hence $L_n \to c - b$ and $U_n \to c$, which proves convergence of the equilibrium interval to $[c - b, c]$.

For Tsallis entropy, let

$$T_q(p_n) \equiv \frac{1 - \sum_u p_n(u)^q}{q - 1}, \qquad q > 0,\ q \neq 1,$$

with $T_1(p_n) = H_1(p_n)$. For $q \neq 1$,

$$H_q(p_n) = \frac{1}{1 - q} \log\bigl(1 + (1 - q) T_q(p_n)\bigr),$$

which is a continuous strictly increasing transform on the nonnegative domain. Hence $T_q(p_n) \to 0$ if and only if $H_q(p_n) \to 0$, and the Rényi case applies.

References

Anctil, R. M., Dickhaut, J., Johnson, C., and Kanodia, C. (2010). Does information transparency decrease coordination failure? *Games and Economic Behavior*, 70(2):228–241.

Angeletos, G.-M. and Pavan, A. (2004). Transparency of information and coordination in economies with investment complementarities. *American Economic Review*, 94(2):91–98.

Angeletos, G.-M. and Pavan, A. (2007). Efficient use of information and social value of information. *Econometrica*, 75(4):1103–1142.

Bergemann, D. and Morris, S. (2019). Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95.

Carlsson, H. and van Damme, E. (1993). Global games and equilibrium selection. *Econometrica*, 61(5):989–1018.

Frankel, D. M., Morris, S., and Pauzner, A. (2003). Equilibrium selection in global games with strategic complementarities. *Journal of Economic Theory*, 108(1):1–44.

Gideon, M., Hsu, J., and Helppie-McFall, B. (2017). Heaping at round numbers on financial questions: The role of satisficing. *Survey Research Methods*, 11(2):189–214.

Greene, W. H. and Hensher, D. A. (2010). *Modeling Ordered Choices: A Primer*. Cambridge University Press, Cambridge.

Heinemann, F., Nagel, R., and Ockenfels, P. (2004). The theory of global games on test: Experimental analysis of coordination games with public and private information. *Econometrica*, 72(5):1583–1599.

Hellwig, C. (2002). Public information, private information, and the multiplicity of equilibria in coordination games. *Journal of Economic Theory*, 107(2):191–222.

Hellwig, C. (2005). Heterogeneous information and the benefits of public information disclosures. UCLA Economics Online Paper No. 283, October 2005 version.

Kajii, A. and Morris, S. (1997). The robustness of equilibria to incomplete information. *Econometrica*, 65(6):1283–1309.

Kamenica, E. (2019). Bayesian persuasion and information design. *Annual Review of Economics*, 11(1):249–272.

Kamenica, E. and Gentzkow, M. (2011). Bayesian persuasion. *American Economic Review*, 101(6):2590–2615.

King, G., Murray, C. J. L., Salomon, J. A., and Tandon, A. (2003). Enhancing the validity and cross-cultural comparability of measurement in survey research. *American Political Science Review*, 97(4):567–583.

Lizzeri, A. (1999). Information revelation and certification intermediaries. *RAND Journal of Economics*, 30(2):214–231.

Lyu, Q., Suen, W., and Zhang, Y. (2023). Coarse information design. *arXiv preprint*. arXiv:2305.18020, revised June 2025.

Milgrom, P. R. (1981). Rational expectations, information acquisition, and competitive bidding. *Econometrica*, 49(4):921–943.

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

Morris, S. and Shin, H. S. (1997). Approximate common knowledge and co-ordination: Recent lessons from game theory. *Journal of Logic, Language and Information*, 6(2):171–190.

Morris, S. and Shin, H. S. (1998). Unique equilibrium in a model of self-fulfilling currency attacks. *American Economic Review*, 88(3):587–597.

Morris, S. and Shin, H. S. (2001). Global games: Theory and applications. Cowles Foundation Discussion Paper No. 1275R (revised August 2001), Yale University.

Morris, S. and Shin, H. S. (2002). Social value of public information. *American Economic Review*, 92(5):1521–1534.

Morris, S., Shin, H. S., and Tong, H. (2006). Social value of public information: Morris and shin (2002) is actually pro-transparency, not con: Reply. *American Economic Review*, 96(1):453–455.

Morris, S., Shin, H. S., and Yildiz, M. (2016). Common belief foundations of global games. *Journal of Economic Theory*, 163:826–848.

Rayo, L. and Segal, I. (2010). Optimal information disclosure. *Journal of Political Economy*, 118(5):949–987.

Rényi, A. (1961). On measures of entropy and information. In *Proceedings of the Fourth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Contributions to the Theory of Statistics*, pages 547–561, Berkeley and Los Angeles. University of California Press.

Rice, N., Robone, S., and Smith, P. C. (2010). International comparison of public sector performance: The use of anchoring vignettes to adjust self-reported data. *Evaluation*, 16(1):81–101.

Scott, D. W. (1985). Averaged shifted histograms: Effective nonparametric density estimators in several dimensions. *The Annals of Statistics*, 13(3):1024–1040.

Shannon, C. E. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423.

Svensson, L. E. O. (2006). Social value of public information: Comment: Morris and shin (2002) is actually pro-transparency, not con. *American Economic Review*, 96(1):448–452.

Topkis, D. M. (1998). *Supermodularity and Complementarity*. Princeton University Press, Princeton, NJ.

Tsallis, C. (1988). Possible generalization of Boltzmann–Gibbs statistics. *Journal of Statistical Physics*, 52(1–2):479–487.

Vives, X. (1990). Nash equilibrium with strategic complementarities. *Journal of Mathematical Economics*, 19(3):305–321.

Weinstein, J. and Yildiz, M. (2007). Impact of higher-order uncertainty. *Games and Economic Behavior*, 60(1):200–212.

Weinstein, J. and Yildiz, M. (2011). Sensitivity of equilibrium behavior to higher-order beliefs in nice games. *Games and Economic Behavior*, 72(1):288–300.

Wooldridge, J. M. (2010). *Econometric Analysis of Cross Section and Panel Data*. MIT Press, Cambridge, MA, 2 edition.